



Craig S. Mullins

[Return to Home Page](#)

September 1995



An Introduction to the Architecture of DB2 for OS/2

By Craig S. Mullins

The generic name for the DB2 workstation products is now DB2 for common servers. This change occurred with the general availability of Version 2 of the DB2 for common servers product. Version 2 of both DB2 for OS/2 and DB2 for AIX is available, with V2 code for Windows NT, Sun Solaris, Siemens Nixdorf SINIX, and HP-UX, soon to follow. This article will focus on one of these products-DB2 for OS/2-although most of the information covered in this article will apply to the other platforms as well.

Database structures

To effectively administer DB2 for OS/2 one must first understand the difference between the DB2 for OS/2 Database Manager and a DB2 for OS/2 database. The Database Manager is the actual DB2 for OS/2 product

in which one or more databases can be created. An installation of the Database Manager is also referred to as an instance.

Each database is composed of multiple files. First, every database has its own set of system catalog tables containing the meta data description of the objects created in that database. Table spaces, new to Version 2 of DB2 for OS/2, can be used to partition data storage based upon the type of data. Up to three table spaces can be utilized per table (one each for normal data, index, and large object fields). Table spaces can span multiple volumes and can be extended without stopping the database. However, it is crucial to note that the table space concept implemented by DB2 for OS/2 is not the same as that used by DB2 for MVS.

Additionally, each database has an associated set of log files used to record database changes. Additionally, log files will be archived by DB2 for OS/2 when transaction activity ceases and the log file is closed or when the log is full. Recovery can then be requested from the active and archived logs using the RESTORE and ROLLFORWARD utilities.

Directories

DB2 for OS/2 utilizes a set of directories for establishing an environment for enabling data access. The directories used by DB2 for OS/2 are as follows:

- The **System Database Directory** identifies the databases that can be accessed from the workstation. It contains an entry for each local and remote database that can be accessed from the workstation. Each entry contains the database name, alias, entry type, and location.
- One **Volume Database Directory** is allocated per disk drive that contains a workstation database. Each entry identifies the location of a specific database on the drive.
- The **Workstation Directory** is used to make a connection to a remote database server. It is used in conjunction with the DCS Directory to make a connection to a remote host server.
- The **Database Connection Services Directory** is used by DDCS for OS/2 to make a connection to a remote host server. An entry is required for each host database specifying the local database name, target database name, location of DRDA code, LU 6.2 parameters, and SQLCODE mapping requirements.

These directories define the overall database manager environment. Without the proper information recorded in these directories, DB2 for OS/2 will not function in the desired manner. For local databases, population of these directories is done by default during CREATE DATABASE processing. Figure 1 offers a synopsis of these directories.

Memory structures

RDBMS products love memory - DB2 for OS/2 is no

different. DB2 employs several different types of memory structures to more efficiently process data.

Heaps are groupings of memory segments utilized for specific types of processing. Some of the more important heaps used by DB2 for OS/2 are:

- The *agent heap* is used for external agents that communicate with DB2 for OS/2
- The *application heap* is an area of memory used by individual applications as they execute
- The *database heap* is a block of memory containing the internal control blocks for SQL statement execution
- The *sort heap* is used to hold temporary results when the database manager performs a sort
- The *statement heap* is used by DB2 when binding static and dynamic SQL
- The *statistics heap* is used by DB2 when gathering database statistics

DB2 for OS/2 uses several other heaps, including the utility heap and the agent heap. DB2 for OS/2 also requires memory for storing locks. Every lock handled by the database manager requires a block of memory. As locks accumulate, memory is consumed. A section of memory, known as the lock list, must be configured for this purpose.

Every DB2 for OS/2 database also has a buffer pool associated with it. The buffer pool is a memory block set aside as a staging area for database reads and

writes. It contains the most recently used pages of data read from the database. Additionally, a DB2 for OS/2 employs a directory cache for staging reads and writes from the directories covered in the previous section.

The actual size of the heaps, lock list, caches, and buffers may be modified by changing database configuration parameters.

System and database configuration

Many parameters for tuning and setting up your DBMS are available to the user of DB2 for OS/2. These are known as configuration parameters; they are changed using a GUI-based configuration tool. Two levels of configuration are provided: system and database. System-level configuration, usually referred to as database manager configuration, applies to the entire DB2 for OS/2 installation. This is where basic database manager specifications such as defaults, maximums, and minimums are made. For example, this is where the following parameters can be set:

- Maximum number of open files for a single process (*maxtotfilop*)
- Maximum number of concurrently active databases (*numdb*)
- The backup buffer size (*backbufsiz*)
- Default database system monitor switches (*dft_monswitches*)

Database configuration, on the other hand, impacts a single database only, not the entire DB2 for OS/2 database manager installation. Database configuration parameters are typically used for setup and performance tuning and are modified more frequently than system configuration parameters. This is the level at which buffers, logging, locking, and heap sizes are specified. For example:

- Buffer pool size (*buffpage*)
- Sort heap size (*sortheap*)
- Time-out specification for locks (*locktimeout*)
- Number of primary log files (*logprimary*)

Refer to the DB2 for OS/2 Planning Guide (The guide may be ordered from IBM by calling 1-800-879-2755) for a description of the configuration parameters.

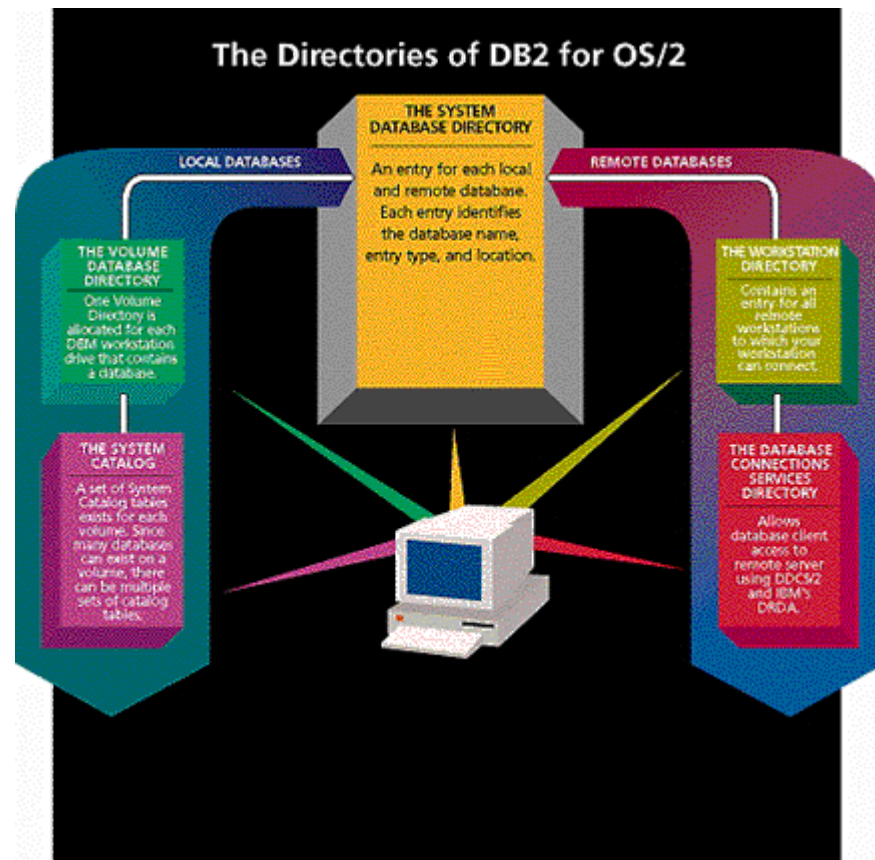


Figure 1

Threads vs. processes

Version 2 of DB2 for OS/2 runs in a multithreaded environment. It is designed to use OS/2 operating system threads for execution instead of processes. The DB2 database server process creates threads for all application agents. By using threads, substantially less memory is required to support multiple users, thereby freeing additional memory for more suitable uses (such as larger buffer pool specifications).

User profile management

Access to DB2 for OS/2 is validated outside of Database Manager by the User Profile Management (UPM) feature: every user must provide a valid userid to UPM before accessing data. This feature is not necessary on the other platforms (including DB2 for AIX). UPM provides userid validation, user and group management, and log in/log out facilities. User log in is performed using alternative methods by DB2 on other platforms.

Distributed support

To access data across a network from the server version of DB2 for OS/2, IBM provides Client Application Enabler (CAE) software. CAEs enable client workstations running disparate operating systems to access the DB2 for OS/2 server. The CAE also requires a communication protocol such as APPC, NetBIOS, TCP/IP, or NetWare IPX/SPX.

By contrast, applications running on the workstation that require access to a DRDA server other than DB2 for OS/2, will require the Distributed Database Connection Services (DDCS) product to enable the distributed access. DB2 for common servers can also be accessed as a DRDA Application Server as of Version 2.

Version 2 highlights

Version 2 of DB2 for OS/2 is a significant new release that propels DB2 into the ranks of the leading RDBMS products. Support for user-defined functions and data

types, triggers, check constraints, recursive SQL, compound SQL, large objects, table spaces, and more efficient utility processing merely scrapes the surface of the new features in this release.

IBM is the undisputed leader in terms of SQL optimization; the DB2 for OS/2 optimizer was completely rewritten to utilize the Starburst optimization technique created at IBM's Almaden research lab. The new optimizer uses advanced query rewrite rules to automatically transform a complex query into a simpler query that is easier to optimize. It also examines a greater number of alternative access paths as it searches for the best query execution plan. These advances will make it less important for users to understand the mechanics of writing efficient SQL queries, enabling them to concentrate on formulating accurate queries instead of accurate and efficient ones.

The new Visual Explain tool enables users to more easily decipher the access path chosen by the optimizer. It depicts the access path in a graphical format instead of an encoded or tabular format used by other RDBMS products.

Additional performance techniques have been incorporated into DB2 for OS/2 making applications and data access much more efficient:

- Based on the nature of each specific query, if DB2 can anticipate that additional data will be

required, *data prefetch* can be enabled to asynchronously request data pages be moved to the buffer pool before they are explicitly requested by the query

- Many input/output operations can be performed in parallel on behalf of a single query
- *Big block* read capability enables multiple pages to be read from disk using a single I/O operation
- Dirty buffer pool pages can be written back to disk asynchronously

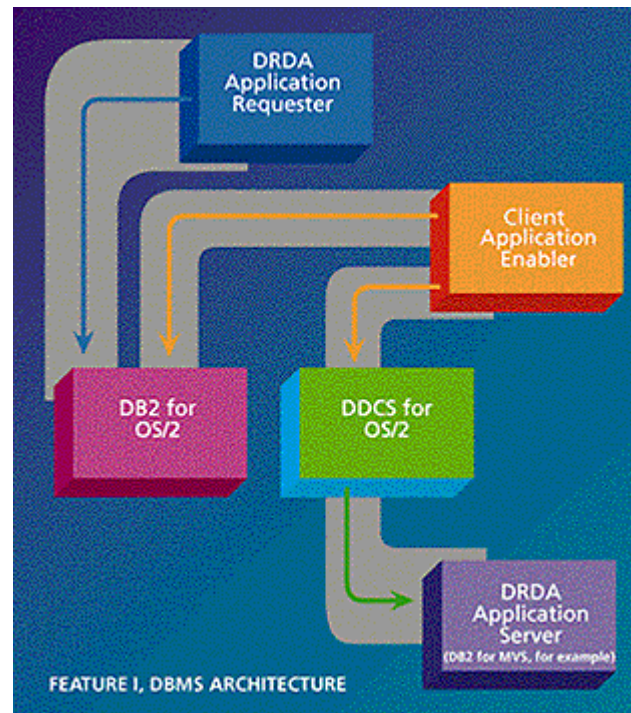


Fig. 2 - Distributed Connections with DB2 for OS/2

Synopsis

Version 2 of DB2 for OS/2 is a robust database server

that compares favorably to any other popular database server. In terms of performance, availability, technical features, and support, DB2 V2 provides a most effective and useful relational database management system.

EDGE Magazine Online, published by PLATINUM *technology, inc.*

© 1999 Craig S. Mullins & Associates, Inc. All rights reserved.
[Home](#). Phone: 630-393-0870 Fax: 630-393-2811