

June 7, 1983

NUMBER 283-160

DATE 830607

TYPE Programming

TITLE IBM DATABASE 2 (DB2) IS ANNOUNCED WITH AVAILABILITY PLANNED FOR
THIRD QUARTER 1984
5740-XYR

OVERVIEW IBM DATABASE 2 (DB2) IS ANNOUNCED WITH AVAILABILITY PLANNED FOR
THIRD QUARTER 1984
5740-XYR

IBM Database 2 (DB2) is IBM's relational data base management system licensed program for the MVS/XA and MVS/370 environments. It may coexist with and complements Information Management System/Virtual Storage (IMS/VS) Data Base System in these environments.

DB2 supports a relational data model. A DB2 data base can be thought of as a collection of tables. Data is defined in terms of tables and accessed through operations on tables. Data definition, retrieval, manipulation, and control operations are supported by the Structured Query Language (SQL). SQL is a high-level, data language available to users through an interactive terminal and through application programs written in COBOL, FORTRAN, PL/I, or Assembler language.

IBM Database 2 may be accessed concurrently by the IMS/VS Data Communication Feature, by CICS/OS/VS, by TSO users, and by batch jobs. The figure below shows the configurations that are possible using IMS/VS Data Base System and/or DB2 as data base managers. Application programs running under control of the IMS/VS Data Communication Feature or CICS/OS/VS may access DB2 data, IMS/VS Data Base data, or both.

DB2 operates as an MVS subsystem. DB2 is designed to utilize the XA architecture, including 31-bit virtual addressing and large real storage. DB2's architecture provides for very large data bases (up to 64 billion $.64 \times 10^9$ bytes per table). It is supported by a comprehensive set of data base utilities which operate online. Its security and authorization mechanism offers field content security and allows various levels of authority to be delegated to users as appropriate.

To assist IBM in introducing the significant new IBM Database 2 product, DB2 is currently being tested in a number of customer and IBM locations. Early experience with a limited number of users will assist IBM in evaluating the product in varied environments and in developing additional support material prior to general availability. An early support program (ESP) is being planned for DB2. This ESP will be conducted by the NAD Dallas Systems Center.

See Announcement Letter for chart.

Highlights

IBM Database 2 is an easy-to-install, easy-to-use licensed program supporting the relational data model with a high-level language to access DB2 data. DB2 offers many functions in support of both traditional data base application areas and the information center. Specific highlights include:

- Relational data model: DB2 provides a basic tabular data structure. Data is viewed by the user as a series of rows and columns.
- Structured Query Language (SQL): SQL is an easy-to-use, high-level language used for data manipulation, data definition, and control. It is not dependent on data paths, placement, or order when accessing the relational data bases.
- Continuous operations: Data definitions may be dynamically changed without stopping DB2. Also, data base utilities can be run without stopping DB2.
- Concurrent access to data: DB2 may be accessed concurrently by the IMS/VS Data Communication Feature, by CICS/OS/VS, by TSO users, and by batch jobs. Additionally, application programs

running under control of the IMS/VS Data Communication Feature or CICS/OS/VS may access DB2 data, IMS/VS Data Base data, or both.

- Application programming: Users may access and manipulate DB2 data using SQL through application programs written in COBOL, FORTRAN, PL/I, and Assembler language.
- DB2 Interactive (DB2I): DB2I is an interactive facility based on Interactive System Productivity Facility (ISPF). Users may access and manipulate DB2 data interactively through DB2I. Additionally, other operations, such as the DB2 data base utilities, may be invoked through DB2I.
- Base for Query Management Facility (QMF): QMF operates with DB2 data in the MVS/XA and MVS/370 environments. See the QMF Announcement Letter, dated June 7, 1983, for details on this product.
- Large processor support: DB2 operates in the MVS/XA and MVS/370 environments, thus providing support for the full range of processors supported by those operating systems.

Description

Relational data model: DB2's relational data model is easy to learn and easy to use.

Simple data structure: The basic data structures in DB2 are tables. A DB2 data base may contain many tables; the simple tabular data format protects the user from having to know and cope with more complicated representations of the data in storage. All relationships are represented by data in tables so users do not have to "navigate" through data structures. The user may experience an increase in productivity, because there is no need to concentrate on how the data is to be accessed, but only on what data is to be accessed.

Structured Query Language (SQL): SQL is a high-level nonnavigational language. SQL statements reference data by name and/or value and not by data access paths, data placement, or order.

SQL, in a single statement, can access:

- One or more columns from a row of a table
- One or more columns from multiple rows of a single table
- One or more columns from multiple rows of multiple tables
- One or more columns calculated from other columns

SQL can be used for retrieval, insertion, deletion, updating, data definition, commitment of changes, and granting and revoking of DB2 authorization. SQL is consistent; it provides essentially the same features and syntax whether it is used as an interactive language or invoked from a COBOL, FORTRAN, PL/I, or Assembler language program.

High productivity: Both the data structures and the data language are simple. As a result:

- The amount of data processing knowledge, training, and experience required to become productive with DB2 is minimal.
- Data base administrators may design, define, and modify data bases quickly.
- Application programmers may develop programs quickly, and hence may be highly productive.
- Using SQL interactively, some users may create and maintain their own data bases without traditional application programming or data administration. Therefore, the application program backlog in many data base installations may be reduced.

Continuous operation: DB2 may operate for long periods without interruption. Many of the activities that may require other systems to stop running do not stop DB2's operation.

- DB2 is a separate MVS subsystem and may be available even when IMS/VS, CICS/OS/VS, or TSO is unavailable.
- Data definitions may be added, deleted, or changed interactively without stopping DB2.
- DB2 utilities can perform their functions without stopping activity on other data bases. In some cases, concurrent activity on the same data base is permitted.

These capabilities allow user data bases to evolve without disrupting ongoing processes and tables that have already been established.

New data base applications and users: In addition to DB2 use in the traditional application areas, its ease-of-use characteristics will attract many new data base users. DB2 will provide an excellent foundation for decision support systems and information centers. Application areas where DB2 is particularly appropriate include:

- Decision support systems.
- Environments in which changing requirements make preplanning difficult or impossible.
- Applications involving exhaustive searches or unanticipated queries.
- Situations in which demanding time constraints necessitate getting the application "on the air" quickly.
- Applications with modeling characteristics -- the need to answer "what if" questions.
- Analysis of data extracted from existing systems.
- Short-term data bases when an extensive design and development cycle is not warranted.
- In combination with QMF, DB2 provides new opportunities for end users to gain interactive access to data through query (SQL and QBE) and reporting capabilities.

Access to data: IBM Database 2 data may be accessed concurrently by the IMS/VS Data Communication Feature, by CICS/OS/VS, and by TSO users. TSO programs may be invoked interactively or by using JCL. Application programs running under the control of the IMS/VS Data Communication Feature or CICS/OS/VS may access DB2 data, IMS/VS Data Base data, or both. Many users will initially install DB2 for interactive access and will add IMS/VS and/or CICS/OS/VS transaction access as volumes increase.

Users can operate on DB2 data interactively and/or through applications written in COBOL, FORTRAN (for TSO users only), PL/I, and Assembler language. In all cases, they access and manipulate the data using the SQL language. DB2 interactive (DB2I) is an ISPF-based interactive facility. Using DB2I, an authorized user can:

- Enter most SQL statements and view the results at the terminal
- Enter DB2 commands
- Perform other programming-related DB2 functions, such as precompilation
- Run TSO-based applications
- Invoke the DB2 data base utilities

For example, using DB2I, application programmers can create and edit most SQL statements they plan to include in their programs, execute those statements, and receive and review the results at their terminals. They can also create, load, and delete test tables. They can use DB2I to check the contents of existing tables. DB2I is a powerful tool for application programmers, data base administrators, and system administrators.

The Query Management Facility (QMF) licensed program (5668-972) provides query and report-writing capability for DB2 data bases. It allows end users to create, save, modify, and execute queries and produce reports. QMF provides these functions in an easy-to-use way. Queries may be created using either SQL or Query-By-Example (QBE) syntax. Subject to DB2 authorization, QMF allows creation and maintenance of DB2 data bases, tables, and indexes.

DB2I and QMF are complementary. DB2I will be used primarily by DP professionals in the application development process. QMF will be used primarily by end users as an easy-to-use alternative to writing application programs.

Technology extended

Utilization of current technologies: DB2 operates in the MVS/XA and MVS/370 environments. DB2 takes advantage of many new and existing hardware and software technologies. It operates with IMS/VS and CICS/OS/VS, so users familiar with these systems don't have to learn a new data communication (online) system. At the same time, DB2 provides a simple data base system for new users who may not be

concerned with data communications. DB2 is designed to utilize the XA architecture, including 31-bit virtual addressing and large real storage.

- Large tables: DB2's architecture allows tables to be very large, up to 64 billion (64×10^9) bytes in size. Large tables can be divided into smaller parts that can be reorganized and recovered individually to improve their manageability.
- Integrated recovery: DB2 has a comprehensive and integrated recovery mechanism. It has disk logging and automatic log archiving capabilities, automatic recovery on restart, and utilities to perform commonly needed recovery functions, such as image copy. DB2 recovery is fully coordinated with the recovery of IMS/VS and CICS/OS/VS with which it may operate. DB2 communicates with both IMS/VS and CICS/OS/VS so that potential inconsistencies can be resolved without loss of data.
- Security: DB2 has a very flexible and comprehensive authorization mechanism. The flexibility and power of the authorization mechanism allow installations to choose to centralize control, decentralize control, or use these in combination. For example, the authority to control particular data bases can be delegated to individual users or to groups of users outside the central organization.

Additional DB2 details

Relational data model

- Data storage
 - Tables and table spaces: Tables are logical data structures consisting of rows and columns. One or more tables are stored in a table space. In DB2, table spaces, not data bases, are the recoverable structure. A table space, and therefore the tables it contains, can be very large, up to approximately 64 billion (64×10^9) bytes.
 - Partitioned table spaces: To make large table spaces manageable, DB2 supports partitioned table spaces. A table space can be divided into partitions on the basis of ranges of data values. Partitions can also be reorganized or recovered independently; reducing the time the table space is unavailable. Partitioning also allows active data to be stored on fast devices and inactive data on slower devices.
 - Views: A DB2 view is an alternative representation of the data in one or more tables. Once a view has been defined, it can be used in most SQL statements just as a table would be. A view has two important advantages as a means of representing data in DB2: views protect sensitive data and views further reduce complexity.
 - Storage groups: With DB2, users need not be concerned with defining MVS data sets. One or more storage groups can be defined indicating which DASD volumes are available for use. DB2 then invokes access method services functions of the Data Facility Product to define, extend, and delete data sets within the storage groups as required.
 - Data bases: A DB2 data base is a collection of table spaces (and, by extension, tables and indexes). The data base is an operational tool. Data bases may be started and stopped independently. A data base can be considered an entity for authorization purposes.
 - Indexes: Indexes can be defined by the user to provide for faster access to data in a table. The index key is made up of one or more columns. Indexes can be defined interactively and created and deleted online.
 - DB2 catalog: The DB2 system catalog is a set of maintained tables that contain information about the data and objects that DB2 manages. Authorized users can retrieve the data in these tables in the same way as any other DB2 table: with SQL statements. DB2 catalog tables contain information about, for example, tables, columns, indexes, storage, and authorization.
- Structured Query Language (SQL): SQL is much more than a query language. It is also an easy-to-use data manipulation, data definition, and control language. Administrators, analysts, programmers, and interactive users all use SQL. Application programmers may embed SQL statements in COBOL, FORTRAN, PL/I, and Assembler language programs to define a complete application. They can also test the "data base" portions of these applications by trying out their SQL statements using the DB2 interactive facility (DB2I) before embedding them in their programs. SQL has similar capabilities and syntax whether entered interactively or embedded in an application program. Data base administrators use SQL to define or change data base descriptors. They can also define and delete tables, indexes, and entire data bases. System administrators use SQL to grant and revoke authorization of all other users to data. Properly authorized interactive users can use SQL to retrieve, create, and modify data from a terminal.

- Automatic access path selection: DB2 uses information maintained within the DB2 system to "navigate" to requested data. DB2 users don't have to know how data is represented in storage in order to retrieve and use it. DB2's ability to select access paths is important in minimizing program maintenance. For example, an administrator might decide to eliminate a seldom-used index. Programs that previously used the index would have their access paths automatically and dynamically reestablished by DB2 the next time they were executed. The logic of the application programs themselves is unaffected and no program maintenance is required to cope with this change. Set-level operation: The SQL language is based on mathematical set-level operations. Therefore, when DB2 processes data, it returns, updates, or deletes the entire set of data that meets the conditions specified by the user. This can reduce the application programming necessary to accomplish a given task.

Access to data

DB2I: DB2 provides users with an interactive online tool. Called DB2I for DB2 Interactive, it uses the Interactive System Productivity Facility (ISPF) to help users perform tasks interactively. To use DB2I, users sign on to TSO, enter ISPF, and select the menu for DB2I. That menu supports several activities. For example:

- SQL processor using file input (SPUFI) is an online function through which DB2 users can submit SQL statements. They can execute those statements and receive the results at the terminal. Using SPUFI, for example, application programmers can create their own test tables and can test SQL statements prior to inclusion in programs, data base administrators can use SPUFI to define and administer data and to grant and revoke the authority to access data and programs, and system administrators can use it to grant and revoke privileges or to query the DB2 system catalog.

Online help: To help insure ease of use, DB2 provides online help information. By pressing a program function key on the display terminal, DB2I users can retrieve panels of information that explain DB2 functions in detail. Over 500 panels are available.

TSO help text is also provided to aid the application programmer or data base administrator who is entering DB2 commands, precompiling, or using utilities directly from TSO.

The DB2 precompiler: Before a DB2 COBOL, FORTRAN, PL/I, or Assembler language program is compiled (or assembled), the DB2 precompiler processes it to extract the SQL statements for subsequent processing within DB2.

Since the precompiler executes independently of DB2, programmers may precompile programs without access or reference to the rest of the DB2 system.

Security, recovery, utilities, and installation

Security and authorization: DB2 provides two mechanisms that allow installations to protect data:

- Authorization statements: Authorization statements in SQL are used to grant and revoke authority for all levels of data, command, and program access in DB2. These authorization statements permit an installation to ensure that only properly authorized users perform data base operations.
- Views: Views play an important role in controlling access to sensitive data. A view that presents only certain rows or columns in a table can be defined. The user or application accesses only the data defined in the view; other data in the table is not available.

The ability to define views combined with the ability to grant and revoke authorization to access those views and the tables on which they are based gives DB2 installations the control they need to keep valuable and sensitive data secure from both unauthorized and inadvertent activities.

Recovery: DB2 protects data from three types of failure: system, media, and application. Applications and transactions that use DB2 data operate across subsystem boundaries; DB2 supports fully synchronized recovery across these boundaries.

The points at which changes to the data base are committed are synchronized so that recovery can be coordinated if a failure occurs at any point.

On system failure, a restart of DB2 automatically restores data to a consistent state by backing out uncommitted changes and completing the processing of the committed changes.

Restart can be speeded and simplified, since not all the data bases that were online at the time of failure need to be available when DB2 is restarted. Some data bases can be started later if the installation chooses.

DB2 makes provisions for media recovery (for example, failure of a disk device or failure of a read or write to disk) by providing disk logging and, optionally, dual logging. DB2 tracks log data sets and image copies so that recovery from media failures is simplified.

If an application program fails, DB2 isolates the work associated with the failing program. It then backs out all uncommitted data changes dynamically, without interfering with other system activities.

Utilities: DB2 provides a number of fully integrated data base utility programs. They are executed online. The long-running ones are restartable. These utilities may be initiated by JCL, using supplied TSO CLISTS (command list, a set of TSO commands that perform a complete task when executed), or through the use of ISPF panels supplied with DB2. The major utilities provided with DB2 are load, image copy, recovery, reorganize, and statistics.

Installation: DB2 has been designed for easy installation.

No system generation process is required. Tailoring a DB2 system is accomplished by supplying parameters to supplied TSO CLISTS. Tailoring is simplified through use of supplied ISPF dialogs.

Planned availability: IBM Database 2 general availability is planned for third quarter 1984.

Technical information

Specified operating environment

Hardware requirements: IBM Database 2 has been designed to operate on any processor supported by MVS/SP Version 2 Release 1.1 or MVS/SP Version 1 Release 3. For the 3033 processor, it is recommended that the hardware cross-memory extension feature (6850) be used. The processors must have sufficient real storage to satisfy the combined requirements of IBM Database 2, MVS/XA or MVS/370, appropriate Data Facility Product, the appropriate access methods, batch requirements, and other customer-required applications. The configuration must include sufficient I/O devices to support the requirements for system output, system residence, and system data sets. Sufficient direct access storage must be available to satisfy the user information storage requirements and may consist of any direct access facility supported by the system configuration and the programming system.

External storage: DB2 is independent of both DASD and tape device type. Any DASD or tape device supported by Data Facility Product may be used. The following data sets are supported by the device types shown:

Active recovery log data sets	DASD
Archive recovery log data sets	DASD, tape, MSS
Image copy data sets	DASD, tape, MSS
Bootstrap data set	DASD
Data base data sets	DASD, MSS
DB2 catalog data sets	DASD
Work data sets (for utilities)	DASD, tape, MSS

Note: DB2 may be used with MSS (IBM 3850 Mass Storage Subsystem) and with appropriate prerequisite program support. The DB2 data base data set internal format precludes the use of VSAM utilities, such as IMPORT, EXPORT, and REPRO, as well as programs that use them, such as DASD Migration Aid (5668-002) and the MVS Hierarchical Storage Manager (HSM) (5740-XRB).

Data communications devices: DB2 uses the system console.

An installation can control DB2 operations from:

The system console

Authorized IMS/VS terminals

Authorized CICS/OS/VS terminals

TSO terminals operated by authorized users

For the data communication devices supported by IMS/VS, CICS/OS/VS, and TSO, see the appropriate documentation.

Software requirements: IBM Database 2 requires the functions of the following licensed programs.

For MVS/Extended Architecture (MVS/XA) environment (with the appropriate prerequisites for each product):

- MVS/System Product-JES2 (5740-XC6) or -JES3 (5665-291) Version 2.1
 - MVS/XA Data Facility Product (5665-284) Release 1.1
 - MVS TSO Extensions (TSO/E) Release 1 (5665-285)
- For MVS/370 environment (with the appropriate prerequisites for each product):
- MVS/System Product-JES2 (5740-XYS) or -JES3 (5740-XYN) Version 1
 - MVS/370 Data Facility Product (5665-295) Release 1
 - MVS TSO Command Package Release 1.1 (5740-XT6) or
 - MVS TSO Extensions (TSO/E) Release 1 (5665-285)
- For both MVS/370 and MVS/XA environments:
 - OS/VS Sort/Merge Release 5 (5740-SM1)
 - For functional ease of use: Interactive System Productivity Facility (ISPF) (5668-960) and ISPF/Program Development Facility (ISPF/PDF) (5665-268)

The other programs that may be used with IBM Database 2 are:

- Query Management Facility (QMF) Release 1 (5668-972)
- Data Extract (DXT) Release 1 (5668-973)
- Information Management System/Virtual Storage (IMS/VS) Version 1 Release 3 (5740-XX2)
- Customer Information Control System/OS/Virtual Storage (CICS/OS/VS) Version 1 Release 6 (5740-XX1)
- Resource Access Control Facility (RACF) Release 5 (5740-XXH)
- OS/VS COBOL Compiler and Library (5740-CB1)
- TSO Assembler Prompter (5734-CP2)
- TSO COBOL Prompter (5734-CP1)
- OS PL/I Optimizing Compiler and Libraries (5734-PL1, 5734-LM4, 5734-LM5, or composite 5734-PL3)
- VS FORTRAN Compiler and Library (5748-FO3)
- OS Assembler H (5734-AS1)
- Assembler H Version 2 (5668-962)

Compatibility/conversion

SQL/DS: The SQL language used by IBM's relational data base management system for DOS/VSE and VM systems, the SQL/DS licensed program (5748-XXJ), is broadly compatible with DB2's SQL. It is possible to move most data from SQL/DS to DB2 by unloading it from SQL/DS tables with the SQL/DS data base services utility and then loading it into DB2 tables with the DB2 load utility.

Information on conversion of SQL/DS programs and data to DB2 will be provided in the IBM Database 2 Application Programming Guides and the IBM Database 2 Data Base Planning and Administration Guide.

Virtual storage considerations: The DB2 storage requirements are dependent on various configurations and work-load parameters including the number of users and their usage patterns, the number and size of data

bases, and the number and complexity of applications. The following estimates of virtual storage requirements are provided for general guidance.

- For MVS/XA systems:
Common storage area (CSA) - Most of the DB2-required areas reside in the extended CSA. The residual CSA requirement is expected to be less than 100K bytes.
Private address space - Most modules, control blocks, and buffers reside in the extended private area. The minimum residual private address space requirement is approximately 2.5M bytes.
- For MVS/370 systems:
CSA - The minimum CSA is approximately 512K bytes. The typical CSA requirement is expected to be 512-532K bytes.
Private address space - For planning purposes, the practical minimum virtual storage requirements for private address space should be considered to be 4.5-5.0M bytes. However, some users with minimum requirements may have DB2 systems requiring less than 4.5M bytes. Virtual storage requirements grow as the number of concurrent users and/or defined tables grows. Thus, the user should consider the maximum address space size that will be available to DB2 as part of his overall capacity planning procedures.

Performance considerations: The performance of IBM Database 2 in a virtual storage environment is highly dependent on the system resources available, the programs that operate concurrently and their relative priorities, and system and application data set placement. Performance also depends on the paging characteristics and storage reference patterns of DB2 and its application programs, the allocation of data sets to particular devices, table sizes, indexes created, and many other factors.

Performance of DB2 is also dependent on the amount of concurrent demand for processor time by DB2 and other users of the system.

It is IBM's intent to provide improved DB2 utilization of 308X class processors through the facilities of MVS/XA.

Data security and auditability: The IBM Database 2 function is subject to all the controls in its environment, such as those provided by the operating system, access methods, and subsystems with which it is communicating. Customer management is responsible for the selection, application, and adequacy of those controls. For further description of the security available in DB2, see "Security, recovery, utilities, and installation."

IMS/VS and CICS/OS/VS also have security capabilities. (See those products' descriptions for details.) Also, through the security access facility (SAF) router of the operating system, Resource Access Control Facility (RACF) (5740-XXH) may be used to control access at the data-set level.

MVS system integrity: IBM will accept APARs describing situations in which the installation of IBM Database 2 introduces an exposure to the system integrity of MVS. This program is intended to run authorized. Refer to the programming announcement dated October 21, 1981.

Programming services: Central service, including the IBM support center, will be available until is continued by IBM upon twelve months' written notice under the terms and conditions of the Agreement for IBM Licensed Programs. For DSLO licenses, central service, including the IBM support center, will be provided only through the customer location designated for the basic license. Local licensed program support is available under the terms and conditions of the Agreement for Local Licensed Program Support for IBM Licensed Programs at the monthly licensed program support charge, monthly multiple licensed program support charge, or at the applicable IBM hourly service rate. Local licensed program support under this agreement will be provided by IBM Field Engineering.