



DB2 Family Newsletter

July 1993 - No 3.

DB2 GOES DESKTOP

In this issue

1. In The News (pg 2)
2. DB2/2 V1 and DDCS/2 V2 (pg 4)
3. DB2/6000 and DDCS/6000 (pg 19)
4. DB2 V3 for MVS (pg 27)
5. Other:
 - Software Spotlight (pg 30)
 - Calendar of Events - Where you can catch us! (pg 37)
 - Submitting Suggestions to IBM (pg 37)
 - Newsletter Subscription (pg 38)

Thanks so much for all the feedback submitted to us. Please continue to stay in touch by completing the Reader's Reply Form and returning it to us by FAX or mail.

I N T H E N E W S

Here is the latest news since the April 1993 newsletter:

DB2/2 and DDCS/2:

Hurry...Introductory Offer ends soon!: The DB2/2 and DDCS/2 introductory price ends soon. Hurry and take advantage of this special offer:

1. In the United States:
 - DB2/2 V1.0: \$199 US (single-user), offer ends August 31,1993
 - DDCS/2 V2.0: \$229 US (single-user), offer ends August 31,1993
 - To order, call 1-800-3IBMOS2 today!
2. In Canada:
 - DB2/2 V1.0:
 - Single-user: \$249 CDN
 - Multi-user: \$1,500 CDN
 - DDCS/2 V2.0:
 - Single-user: \$299 CDN
 - Multi-user: \$3,000 CDN
 - Orders must be placed by September 17,1993 and delivered by September 30,1993.
 - To order, call 1-800-465-7999 today!

DB2/2 Demo Diskette: We've developed an exciting demo diskette for DB2/2 and DDCS/2. To order your free demo, call 1-800-342-6672 (USA) and 1-800-465-7999, ext 850 (Canada).

Downsizing EXPO Conference & Exposition (September 13-15,1993): The Downsizing EXPO Conference & Exposition will be held in Toronto, Canada from September 13-15,1993. On Wednesday September 15,1993 (2:00pm - 3:00pm), Janet Perna, Director of Database Technology for IBM Programming Systems Toronto Laboratory will give the keynote address for the topic **Client/Server Strategy for Relational Database Development**.

For additional conference information, contact:

Digital Consulting, Inc.
204 Andover Street
MA 01810 U.S.A

508-470-3880 (9:00am - 6:00pm EDT)

DB2/2 Developer Assistance Program: The DB2 Developer Assistance Program (DAP) is designed to secure feedback from professional developers and to provide those developers with an opportunity to enhance their products with support of DB2/2. If you are developing software for resale and would like to provide support for DB2/2, call 1-800-627-8363 and ask for the enrollment package for the DB2 DAP. Outside the US and Canada, send a fax to 1-817-961-8410.

The following is a representative list of developers in the DB2/2 Developer Assistance Program. It includes developers who have indicated to us that they intend to support DB2/2 with at least some of their products and who have given us permission to use their names.

ACCESS, a communications company
Alphasoft Marketing International
Bridge Technology, Inc.
BTB GmbH
CGI Informatique
CIRRUS Technology, Inc.
Computer Associates International
Compuware Corporation
Comshare, Inc.
DataSix
DBopen, Inc.
Easel Corporation
Gupta Technologies, Inc.
IBM Corporation
Information Builders, Inc.
Integrated Systems Solutions Corp.
Intelligent Environments
Intelligent Office Company Inc.
KASEWORKS, INC.
Metaphor, Inc.
Micro Focus, Inc.
N Systems
OpenBooks Software, Inc.
PeopleSoft, Inc.
Policy Management Systems Corp.
Q+E Software Inc.
Repository Technologies Inc.
SAS Institute Inc.
Syntelligence Systems, Inc.
TechGnosis, Inc.
Trinzic Corporation
Uniface Corporation
XDB Systems, Inc.

Summary information on the products developed by the above companies is available in a document called the **DB2/2 Solutions Directory**. Customers may ask their IBM Marketing Representative to retrieve this document as DB22VEND TRS3820, in the WORKDBM PACKAGE on MKTTOOLS. We are also expecting to ship a hardcopy of the Directory with the next edition of this Newsletter. And if you just can't wait that long, you can request a copy using the Reader's Reply Form contained herein.

DB2/6000 and DDCS/6000

DB2/6000 Beta Program: The PRGS Toronto lab is pleased to announce the availability of the DB2/6000 beta code for IBM customers. The terms and conditions are included in the package that can be obtained through your local IBM representative. Support for this program will be through your local IBM representative or a special customer forum that is available through IBMLink. The initial beta release of DB2/6000 includes the database manager, support for AIX clients via TCP/IP and OS/2 clients via Extended Services APPC. A complete list of the functional content of the beta driver is included in the package.

You can receive the package by asking your IBM representative to request DB2BETA from the AIXTOOLS conferencing disk. Publications can be requested as well (DB26PUBS from AIXTOOLS).

Section 1: DB2/2 V1 and DDCS/2 V2

DB2 Goes Desktop at DB/EXPO

By Jay Lennox and Liwen Yeow

DB/EXPO was an exciting event for the PRGS Toronto Lab Database Technology Group. With a theme of DB2 Goes Desktop, the show marked the general availability of our OS/2 relational database, DB2/2 Release Version 1.0. The show was held May 4-6, 1993 in San Francisco, California.

To really show the power of DB2/2 as a client/server database, we decided to implement a major LAN in the IBM booth, and to connect this LAN via Distributed Database Connection Services/2 (DDCS/2 V2) to IBM's host relational databases, SQL/DS and DB2. DDCS/2 is the IBM implementation of Distributed Relational Database Architecture (DRDA) for OS/2 and allows an OS/2 user transparent access to relational data in SQL/DS, DB2 and OS/400 databases.

We also wanted to demonstrate solutions to day to day business problems that utilized DB2/2 and the recently announced DB2/6000. Products from 14 Independent Software Vendors (ISVs) and from IBM were utilized in this effort. The ISVs demonstrated query and report writing solutions, connectivity solutions (Macintosh to DB2/2), 3rd generation languages, 4th generation languages, and office automation solutions. When combined with database administration and application generation solutions from IBM, the strength of DB2/2 to address customers business needs was very evident.

The DB/EXPO setup is schematically represented in the accompanying diagram, Figure 1 on page 5. From a networking and connectivity perspective, the LAN demonstrated:

- the mixed use of Netbios, SNA/APPC and TCP/IP on the Token Ring LAN to access four DB2/2 database servers and a DB2/6000 database server
- the connectivity of client workstations running OS/2 DB2/2 client code, DOS/WIN DB2/2 client code, and a Macintosh System 7 running code from TechGnosis Inc.
- an X-Station attached via the ring to DB2/6000

All of our DB2/2 servers also provided gateway support via DDCS/2 to host relational databases:

- 3 of them were connected to SQL/DS and DB2 systems running in Dallas via SDLC dial-up lines
- the 4th gateway was connected to a DB2 system running in IBM's Santa Teresa Laboratory in San Jose, CA, via a TCP/IP SLIP line

We managed the entire LAN and host environment using DataHub and LAN Network Manager, both from IBM.

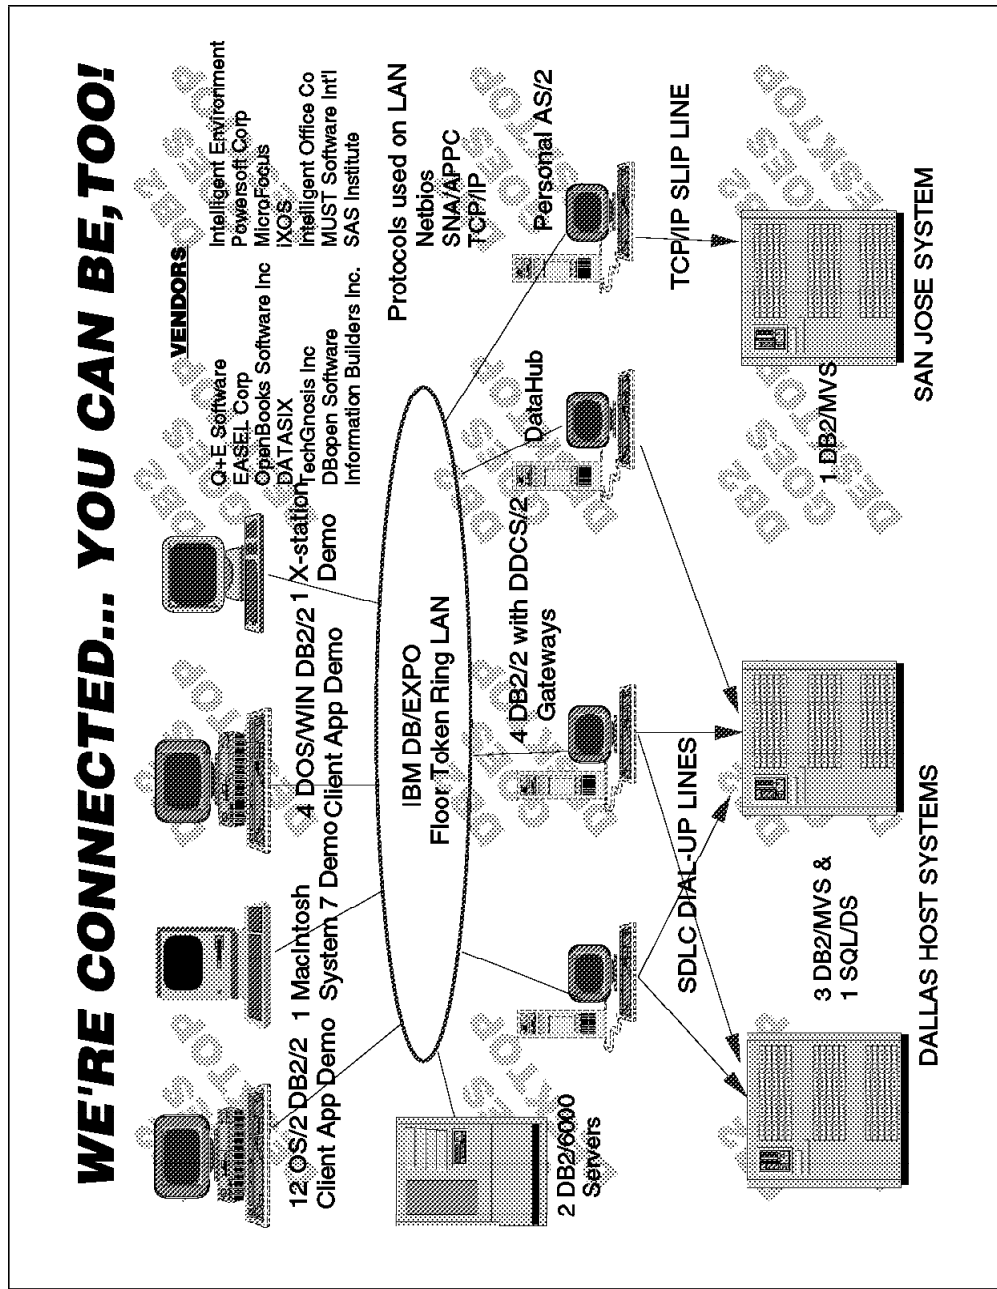


Figure 1. DB/EXPO Client/Server Configuration

Everything was done live; no story boards were allowed. If we expected customers to buy DB2/2, then we felt that we needed to show that it was ready to go. And ready it is. Throughout the course of the show, we did not experience one problem with software, even with applications from 14 ISVs and IBM running live against the DB2/2 and DB2/6000 servers, and through DDCS/2 gateway code to SQL/DS and DB2. We did have a couple of problems; a dial-up phone line 'hung-up' a few times, and one workstation needed servicing.

For us DB/EXPO was a great success. We achieved all of our objectives - we proved that DB2/2 is customer ready, that it is a true client/server database, and that we have many quality products from ISVs and IBM available to complement it. Since the DB/EXPO show in early May, the number of ISV and IBM products supporting DB2/2 has continued to grow. These products are now included in the

DB2/2 Solutions Directory. Refer to **DB2/2 Developer Assistance Program** on page 2 of this document for additional information.

We would like to take just a quick moment to thank everyone for their efforts at DB/EXPO, especially the customers who came to see us at the show, and the ISVs who demonstrated their products with us at DB/EXPO.

DB2/2 V1 Function/Feature Matrix

Feature	Description for DB2/2 V1
Precompilers	
C	Yes
C++	No
Cobol	Yes
Ada	No
Fortran	Yes
Pascal	No
PL/I	No
Basic	No
Call-level interface	
C	No
Callable APIs	
C	Database manager routines are available as function calls: environment routines, utility routines, configuration routines, client-server routines, run-time routines, and additional development routines. This is not SQL Access Group CLI nor ODBC.
Cobol	Yes - refer to C
Fortran	Yes - refer to C
Pascal	No
REXX	Yes - refer to C
Command Line Processor	Yes - can issue SQL statements and database manager commands that perform the same functions as the routines listed above from the operating system command line or from a command file.
API features	
Positioned UPDATE/DELETE	Yes
Cursor Context Preservation	Yes. Allows database cursors to continue to be used after a commit thus avoiding the overhead of reestablishing each cursor in the application after a commit. (via DECLARE CURSOR WITH HOLD)
Scrollable Cursor	No
Procedural SQL	No
Stored Proc. /w Proc. SQL	No
Stored Proc. /w 3GL	Yes - Database Application Remote Interface (DARI). To use DARI, a calling procedure is contained in an application running on the client. A server procedure executes against the database at the database server. The client application declares, allocates, and initializes any data structures, connects to the database and invokes the server procedure by calling the DARI routine. The server procedure may be written in C, COBOL or REXX and accepts optional input data structures from the client application. It executes under the same transaction as the client application and returns the SQLCA and optional output data structure to the client application upon completion.
Bind Files	Yes
Array Inserts	No
Interfaces (Vendor's Own)	

Feature	Description for DB2/2 V1
Programmer tools	QM. Refer to the DB2/2 Solutions Directory for more information.
End-user tools	QM. Refer to the DB2/2 Solutions Directory for more information.
Applications	Personal AS, Data Interpretation System. Refer to the DB2/2 Solutions Directory for more information.
DBA Tools	Configuration Tool, Recovery Tool, Command Line Interface, Directory Tool, SQLQMF, Reorgchk, Status commands.
4GLs	
Interfaces (Third Party)	
Programmer tools	Microfocus COBOL Workbench - Host compatibility option, Advanced Revelations, Application Manager (AM), Building Blocks, Choreographer, Clarion Database Developer, EASEL Workbench, Enfin/2, Enfin/3, ERwin/ERX, 4D SQL Server, 4th Dimension DAL, Object/1, Q+E DB Library, SequeLink, Smalltalk/SQL, JAM. Refer to the DB2/2 Solutions Directory for more information.
End-user tools	DataPrism, Easy SQR, Q+E DB Editor, SmartLeaf, SmartLeaf/Compare, SmartLeaf/Librarian, Snow Report Writer, SQR, TeleMagic, WinList, ClearAccess, Quest, InfoAlliance, LightShip. Refer to the DB2/2 Solutions Directory for more information.
Applications	Refer to the DB2/2 Solutions Directory for more information.
DBA Tools	Bridge/FASTLOAD for DB2/2
4GLs	Refer to the DB2/2 Solutions Directory for more information.
Standards and Compatibility	
FIPS 127.1 (SQL89 Level 2)	No
Entry SQL2	No
DB2 Compatibility	partial - high level. New this release: <ul style="list-style-type: none"> • NOT NULL WITH DEFAULT • CONNECT statement • implicit connection • CURRENT SERVER special register • scalar functions: FLOAT, DECIMAL, INTEGER • RUN as synonym for EXECUTE on GRANT and REVOKE
SAA Compatibility	partial - high level. See new items added above - all SAA compliant.
MIA	Yes
NAS	No
ACE	No
CA90	No
SQL Function (Extensions)	
Scalar Functions	CHAR, DATE, DAY, DAYS, DECIMAL, FLOAT, HOUR, INTEGER, LENGTH, MICROSECOND, MINUTE, MONTH, SECOND, SUBSTR, TIME, TIMESTAMP, TRANSLATE, YEAR
Column Functions	AVG, COUNT, MAX, MIN, SUM
Multi-table view update	No
Update with SELECT INTO	No
SELECT as expression	No
Outer Join	No
Recursion	No
User Defined Functions	No
Very Long Fields	No - max is 32,700.
Event Alerters	No
Integrity Constraints	
Declarative RI	Yes - RESTRICT, CASCADE, SET NULL options on DELETE, RESTRICT on UPDATE.

Feature	Description for DB2/2 V1
Triggers	No
Base datatypes	Char(254), varchar(4000), long varchar(32700), smallint, integer, float, decimal/numeric, date, time , timestamp.
User-defined types	No
View check option	Yes
Column constraints	No
Column defaults	Yes - same support as for DB2 - the system will assign a pre-defined default.
NULL support	Yes
Security	
C2 level	No
B1/B2 level	No
User Group security	Yes - defined using UPM
Audit trail	No
Authorization	Yes, SYSADM, DBADM, bindadd, connect, createtab, control, alter, delete, index, insert, references, select, update, bind, execute.
Roles	No; however, groups may be created that correspond to a required role - groups could be used to simulate roles.
Passwords	Validated outside of the database: on OS/2: managed by UPM, on DOS and Windows: encrypted and stored in a hidden profile on client and sent to server for validation.
Distributed Database	
Remote Unit of Work	Yes - application requestor to DB2/MVS, SQL/DS, SQL/400 via application program and DB2/2 command line processor using DDCS/2.
Distributed Unit of Work	No
Distributed Request	No
Distributed Request /w Global Optimization	No
Snapshots	No
Transparency	Yes - the database is catalogued once and referred to, from then on, by its alias.
IBM RDB Connectivity	Read/write access to DB2/MVS, SQL/DS, SQL/400
OEM RDB Connectivity	No
Non-relational DB Connectivity	No
Communication Protocols	
NetBIOS	Yes
APPC/APPN	Yes
TCP/IP	No
IPX/SPX	No
Named Pipes	No
Portability (server)	
DOS	No
Windows	No
OS/2	Yes
AIX	Yes
SunOs	No
Ultrix	No
HP-UX	No
Pyramid	No
Sequent	No

Feature	Description for DB2/2 V1
Decstation (MIPS)	No
Macintosh	No
Novell Netware	No
Portability (client)	
DOS	Yes
Windows	Yes
OS/2	Yes - EE 1.3, ES 1.0 & DB2/2 1.0
AIX	Yes
SunOs	No
Ultrix	No
HP-UX	No
Macintosh	No
Decstation (MIPS)	No
System Integration	
Single O/S	Yes
Multiple O/S	No
DCE/OSI Support	No
Performance (Access Paths)	
Cost-based Optimizer	Yes
Query transformations	Yes - in some cases. E.g. if there's a suitable index on a column, an IN list might be transformed into a set of equality predicates connected by OR so that index-ORing may be done; there are a few more cases that are too long to explain here.
EXPLAIN	Yes - applet on diskette #1, not installed with the DB2/2. Documentation can be found in EXPLAIN package on OS2TOOLS and may be distributed to customers.
Index ORing	Yes
Index ANDing	No
Index-only access	Yes
Recreate Index option on database startup	Yes
Performance (Architecture)	
Static SQL	Yes
Process Model	Multiserver architecture - process per user with 32-bit exploitation
Symmetric Multiprocessing	No
Parallelism	No
Data Organization	
Clustering Index	No - rows physically stored in index order when reorganized but will become unclustered with inserts and updates.
Storage methods	B ⁺ tree
Inter-table cluster	No
Write-once media (optical)	No
Performance (I/O Management)	
Group Commits	Yes
Deferred writes	Yes
Data page buffering	Yes
Chained writes	No
Sequential Prefetch	No
Index Prefetch	No

Feature	Description for DB2/2 V1
Performance (Tuning)	
Performance monitor	System, database and user status interface for a point in time
Governor	No
Tuning parameters	Configuration parameters for the database manager and each database, ARI, record blocking, isolation parm at bind
Transaction monitor	No
Recovery	
Forward Recovery	Yes and point-in-time recovery is possible
Log Archive	Yes
Incremental Backup	Yes, at the page level. Not available with forward recovery
Granular Recovery	No
Set logging off	No
Tape Management	No
Availability	
Separate Logs and DB	Yes
Dual Logging	No
Mirrored Disks (OS Level)	Yes
Mirrored Disks (SW Level)	No
Online Backup/Restore	No
Online Configuration	No
Online Reorganization	Yes - table level reorg - only the table and its indexes are unavailable during a reorg.
Run with Partial Database	No
Locking	
Row-level locks	Yes
Page-level locks	No
Table-level locks	Yes
Lock Escalation	Yes, row to table
Wait/nowait lock	No, always waits for a lock to be freed.
Deadlock detection	Aborts transaction that caused deadlock and on timeout.
Isolation Level	
Repeatable Read	Yes
Cursor Stability	Yes
Dirty Read	Yes
Multi-version reads	No
Committed read	No
Autonomous/Remote Operations	
Remote Installation	Yes
Remote Administration	No
Autonomous Reorganization	No
Autonomous Statistics	No
Manageability	
Truncate Table	No
Analyze Table	No
Fast dataload	Yes, via Bridge/FASTLOAD for DB2/2
Object Features	
Object IDs	No

Feature	Description for DB2/2 V1
Smalltalk Pre-Compiler	No
Multiple inheritance	No
Late binding	No
Overloading	No
Class Libraries	No
Links	No
Encapsulation	No
Object caching	No
Pointer swizzling	No
Object-level locking	No
Check-in/check-out locks	No
Miscellaneous	
National Language Support	Yes, sorting/collating, translated error messages and documentation for twelve languages
Utilities	Import/Export- can import from these file formats: non-delimited ASCII, delimited ASCII, WSF (Lotus worksheet), PC IXF; can export in these file formats: delimited ASCII, WSF & PC IXF; Configuration Tool; Recovery Tool; Directory Tool.
Diagnostics	Corruption detection, trace facility, first failure support technology (FFST/2) exploitation
Installation	Config/Install/Distribute (C/I/D) Enablement

Performance Tips & Techniques

Database management requires that you set performance objectives. The applications and system must be designed and tailored to meet those objectives. The following tips and techniques provide an introduction to performance management. For additional information, consult the following publications.

Table 1. Performance Related Publications	
Publication Number	Publication Title
S62G-3663	DB2/2 Guide
G362-0012	Comprehensive Database Performance for OS/2 2.0 Extended Services by Bruce Tate, Tim Malkemus, and Terry Gray

The following information is reprinted from the book "Comprehensive Database Performance for OS/2 2.0 Extended Services", Appendix B, with permission from the publisher VAN NOSTRAND REINHOLD.

The Comprehensive Plan

Tip 1: Develop a comprehensive performance plan. Include requirements collection, database design, development, benchmarking, and tuning.

Configuration

Tip 2: Ensure adequate hardware for your performance requirements.

Tip 3: There must be adequate RAM installed to achieve good performance.

Tip 4: Tune at least the ROW BUFFER and LOCK LIST parameters of OS/2 Extended Edition 1.3 and before. The default values for these configuration parameters on the above releases are much too low for most applications.

Tip 5: Where possible, place all data in a single database to conserve resources.

Tip 6: When possible, place the database files and log files on different physical disks to decrease the lateral movement of the disk head.

Tip 7: When Forward Recovery is used, the size of the logs should be convenient for the backup devices. In all other systems, it is best to have a small number (3-5) of primary logs which are large enough to handle all mainline processing.

Tip 8: For applications that are experiencing concurrency problems manifested in deadlocks or long waits, try increasing the LOCK LIST parameter of the database configuration file.

Tip 9: Begin tuning efforts with the database configuration parameters ROW BUFFER, LOCK LIST, and SORT HEAP (for applications that sort extensively). These are likely to have the biggest impact.

Database Design

Tip 10: Spend enough time up front to get a good database design, or performance and integrity problems will result.

Tip 11: Use the comment statement to describe columns, tables, views, and indexes.

Tip 12: Document indexes that are used only by dynamic queries, or they may be inadvertently removed during periodic maintenance.

Efficient Data Access

Tip 13: Dynamic SQL may outperform static SQL if the dynamic query contains a literal pattern for a LIKE predicate that does not begin with "%" or "_" such as LIKE 'ABC%'.

Tip 14: Though indexes can dramatically improve performance, consider their cost: decreased performance impact to adds, inserts, and deletes and increased disk requirements.

Tip 15: Do not create redundant indexes, for example, i=(a,b), and j=(a,b,c) are redundant. Index is unnecessary.

Tip 16: Choose a "clustering" index and specify it when you use REORG.

Tip 17: Specify index columns in the ordering of decreasing cardinality or example (last_name, sex) rather than (sex, last_name) in the absence of other considerations.

Tip 18: Avoid correlated subqueries and quantified subqueries where possible.

Tip 19: Create an index on each table on the columns that are used in join predicates.

Tip 20: Draw a simple graph depicting join predicates to avoid unqualified Cartesian products.

Tip 21: Specify indexes where appropriate to assist GROUP BY and ORDER BY clauses.

Tip 22: Use views that will be merged rather than materialized.

Tip 23: When lexical ordering is not required, specify an appropriate collating sequence that gives each character a unique weight to improve the performance of comparisons. The identity weight table is best.

Concurrency and Integrity

Tip 24: Put each separate concurrent task in a separate process.

Tip 25: When locking is not important, consider table locks or the UNCOMMITTED READ level of isolation to reduce locking overhead.

Tip 26: Choose the weakest isolation level that maintains data integrity to get the best concurrency.

Tip 27: Never commit in the middle of a logical unit of work.

Tip 28: Tune commit frequency to get the best performance. Make commits more frequent to improve concurrency, less frequent to boost performance within a process.

Tip 29: If all else fails, consider replication or fragmentation of the database to improve concurrency.

Remote Connections

Tip 30: For remote connections, use record blocking for applications that retrieve many database rows with the same cursor.

Tip 31: Do not use dynamic SQL in applications that will record-block.

Tip 32: For remote connections, use Database Application Remote Interface for applications that issue several SQL calls at once or SQL within loops.

Tip 33: When using Database Application Remote Interface, use a single SQLDA entry to pass a structure. This cuts down on the amount of data passed.

DB2/2 - At the Max

The following table describes certain maximums inherent in DB2/2. Adhering to the most restrictive case can help the programmer design application programs that are easily portable.

Table 2. DB2/2 Maximums	
DB2/2 Maximums	Limit
Most columns in a table	255
Most columns in a view	255
Maximum length of a row including all overhead	4005
Maximum size of a table ^a	69 x 10 ⁹
Maximum size of an index	2 x 10 ⁹
Most rows in a table	4 x 10 ⁹
Longest index key including all overhead	255
Most columns in an index key	16
Most indexes on a table	32 767
Most tables referenced in an SQL statement or a view	15
Most host variable declarations in a precompiled program	880
Most host variable references in an SQL statement	1489
Longest host variable value used for insert or update	32 700
Longest SQL statement	32 765
Most elements in a select list	255
Most predicates in a WHERE or HAVING clause	300
Maximum number of columns in a GROUP BY clause	255
Maximum total length of columns in a GROUP BY clause	4005
Maximum number of columns in an ORDER BY clause	255
Maximum total length of columns in an ORDER BY clause	4005
Maximum size of an SQLDA	65 532
Most declared cursors in a program	400
Maximum number of cursors opened at one time	storage
Most tables in a relational database	32767
Maximum number of prepared statements	400
Maximum number of packages	storage
Note:	
a The numbers shown are architectural limits and approximations. The practical limits may be less.	

Question & Answer

Question & Answer
<p>Table 3 (Page 1 of 4). DB2/2 Q&As</p>
<p>Q: Does DB2/2 have a multi-process architecture?</p> <p>A: DB2/2 allows application programmers to design and write their applications using a multi-threaded design. It is to be noted that a COMMIT on one thread COMMITs the entire PROCESS connected to the database.</p> <p>REMINDER : Only one database can be accessed per process with DB2/2.</p>
<p>Q: I'm migrating from ES 1.0 to DB2/2 and received message "SQL5122N: Access to the database was invalid because of a machine dependent check". What is the cause?</p> <p>A: The database and database configuration file cannot be accessed because of copy protection. To correct this error, return to the original database and modify the configuration file to turn the copy protection off, then create a new backup which may be used to restore the database. This must be done by an individual with SYSADMIN authority.</p>
<p>Q: What options are required to compile and link a DB2/2 program using C Set/2?</p> <p>A: There are no REQUIRED options, the defaults are sufficient. There is one exception, to build a .DLL file for DARI Server applications the Ge- parameter must be specified.</p>
<p>Q: How do I get BRIDGE/FASTLOAD?</p> <p>A: Contact Bridge Technology at 617-424-6266 (FAX 617-424-6621).</p>
<p>Q: Can I place my database logs on a VDISK?</p> <p>A: This is not recommended for data integrity reasons.</p>
<p>Q: If I purchase DB2/2 (multi-user), does any communications software come with DB2/2? If no, what communications software is required?</p> <p>A: DB2/2 provides the database software. Communications software is available as follows:</p> <ul style="list-style-type: none"> • You require CM/2 version 1.0 or ES 1.0 for APPC communications on the server machine. • If you decide to use the NetBIOS protocol the required communication software is Network Transport Services 2 version 1.0 NTS/2 or LAN Adapter and Protocol Support (LAPS). LAPS is provided with the following products : ES 1.0, LAN Server 2.0, or LAN Server 3.0.
<p>Q: Specifically, how is DB2/2 compatible with DB2 for MVS?</p> <p>A: The following DB2 Compatibility and SQL Language Extensions are supported by DB2/2:</p> <ol style="list-style-type: none"> 1. NOT NULL WITH DEFAULT support for column definitions (DB2 MVS compat.) 2. SAA SQL Level 2 compliant CONNECT statement (replacing START USING/ STOP USING APIs, though they will continue to be supported). 3. Implicit CONNECT (no longer need to do START USING, but can define default database name to connect to). 4. CURRENT SERVER clause (SAA SQL Level 2 CPI) to find out the name of the database you are connected to. 5. Support for RUN as a synonym for EXECUTE (DB2 MVS compatibility) 6. Support for DECIMAL, INTEGER, and FLOAT scalar functions (part of SAA SQL Level 2 CPI)
<p>Q: I'm looking for a database that will handle images and textual information. How can DB2/2 help?</p> <p>A: IBM recently announced the Ultimedia Workplace/2 product which brings to your desktop an easy way to organize and manage multimedia (image, audio, video) files and story objects. Put briefly, Workplace/2 offers:</p> <ul style="list-style-type: none"> • Simple, visual management of multimedia objects • A means of adding descriptive information to multimedia objects • A means of integrating multimedia objects with DB2/2 databases • Tools that help you define, sequence, and run searches on multimedia objects • Help with finding and managing objects, on line and off line.
<p>Q: What do I have to do to run my 16-bit application with DB2/2?</p> <p>A: A 16-bit application does not require re-compiling or linking to work with DB2/2. Please be careful to run SQLPREP with DB2/2 if you are planning to make any application changes (ie. don't simply re-link old .OBJ files using the DB2/2 libraries).</p>

Table 3 (Page 2 of 4). DB2/2 Q&As

Question & Answer													
<p>Q: What are my connectivity options with IBM's OS/2 databases?</p> <p>A: Here are your options:</p> <table border="1"> <thead> <tr> <th>Client Type</th> <th>DB2/2 Server</th> </tr> </thead> <tbody> <tr> <td>EE 1.3 OS/2</td> <td>APPC</td> </tr> <tr> <td>ES 1.0 OS/2</td> <td>APPC, APPN, NETBIOS</td> </tr> <tr> <td>ES 1.0 DOS/Windows</td> <td>NETBIOS</td> </tr> <tr> <td>DB2/2 OS/2</td> <td>APPC, APPN, NETBIOS</td> </tr> <tr> <td>DB2/2 DOS/Windows</td> <td>NETBIOS</td> </tr> </tbody> </table>		Client Type	DB2/2 Server	EE 1.3 OS/2	APPC	ES 1.0 OS/2	APPC, APPN, NETBIOS	ES 1.0 DOS/Windows	NETBIOS	DB2/2 OS/2	APPC, APPN, NETBIOS	DB2/2 DOS/Windows	NETBIOS
Client Type	DB2/2 Server												
EE 1.3 OS/2	APPC												
ES 1.0 OS/2	APPC, APPN, NETBIOS												
ES 1.0 DOS/Windows	NETBIOS												
DB2/2 OS/2	APPC, APPN, NETBIOS												
DB2/2 DOS/Windows	NETBIOS												
<p>Q: Can I use the Borland C++ compiler to develop DB2/2 applications?</p> <p>A: No, currently there is no support for the C++ precompiler. Since DB2/2 provides a C precompiler, the only supported 32-bit compilers are IBM C Set/2 version 1.0 and WATCOM FORTRAN77/386 version 9.01 (Patch Level D).</p>													
<p>Q: Can DDCS/2 V2 work with ES 1.0?</p> <p>A: No. DDCS/2 V2 works with DB2/2 V1. DDCS/2 V1 works with ES 1.0. (Note: DDCS/2 V2 may function with Communications Manager from ES 1.0).</p>													
<p>Q: How can I keep my passwords synchronized?</p> <p>A: IBM Network SignON Coordinator/2 Version 1.1 provides DOS and OS/2 users a simplified method of logging on and off a variety of systems. It also keeps passwords synchronized across multiple locally-attached workstations and central site hosts.</p> <p>IBM Network SignON Coordinator/2 extends its support to provide single signon and password coordination for Novell (1) NetWare (1) v3.11 systems and local DOS host operations. Users can benefit from extensive usability enhancements such as:</p> <ul style="list-style-type: none"> • A new graphical, icon-based user interface for OS/2 (R) enables users to select single systems for signon operations • Increased flexibility in the host communication script language, including dynamic selection of terminal emulation sessions • Easy administration and maintenance of configuration file information • Access to configuration files from a server instead of storing them on the workstations <p>Network SignON Coordinator supports DOS, OS/2 and NetWare clients, IBM OS/2 LAN Servers, NetWare v3.11 servers, remote OS/2 workstations (for example, Extended Services (TM) database servers and DATABASE 2 (TM) OS/2 servers) and VM, MVS and OS/400 (R) hosts.</p>													
<p>Q: Do the precompilers handle the C Set/2 preprocessor conditional compile functions line #ifdef and #undef?</p> <p>A: SQLPREP will ignore #ifdef and #undef contained in the application. The only things that are preprocessed by the precompiler are statements beginning with EXEC SQL and host variables declared between the EXEC SQL BEGIN DECLARE SECTION and EXEC SQL END DECLARE SECTION statements. Everything else is left to the host language compiler.</p>													
<p>Q: What is the best SQLDS data type to use if the C Set/2 data type to be passed into the database is a double float?</p> <p>A: Use Float. In DB2/2, a float column is an 8-byte float, same format as a double in C.</p>													

Table 3 (Page 3 of 4). DB2/2 Q&As

Question & Answer

Q: Does DB2/2 support data format MM/DD/YYYY for compatibility with DB2?

A:

Yes, DB2/2 will support dates in MM/DD/YYYY format. When you precompile your applications, you can add the /F=<format> to the SQLPREP command line. This sets the date/time format as follows:

<format>	DATE	TIME
DEF	Use the default format for the database codepage (If no /F= is given, this will be used)	
USA	MM/DD/YYYY	HH:MM AM/PM
EUR	DD.MM.YYYY	HH.MM.SS
ISO	YYYY-MM-DD	HH.MM.SS
JIS	YYYY-MM-DD	HH:MM:SS
LOC	Special date and time formats depending on country code associated with database.	

This parameter is also supported by SQLBIND and can also be set through the options parameter of sqlabndr().

Q: Does DB2/2 provide a sample database so that I can test my application?

A: SQLSAMPL is an executable file that installs a sample database. To install the database, you must have SYSADM authority. At the command prompt, type **SQLSAMPL x** where x is the drive letter. For a hardcopy version of the sample database, refer to Appendix D of the DB2/2 Guide.

Q: Can I use PL/I to develop DB2/2 database applications?

A: No, PL/I is not a support programming interface language for DB2/2.

Q: Can I use DDCS/2 with non-IBM databases? How is this done?

A: The host database must adhere to Distributed Relational Database Architecture (DRDA) protocol. Below is a summary of DRDA supported by Vendors:

VENDOR NAME:	STATUS:					
	(1)	(2)	(3)	(4)	(5)	(6)
Borland International, Inc.	x					
Cincom Systems, Inc.	x					
Cross Access Corporation	x					
Information Builders Inc. (IBI)	x	x				
Informix Software, Inc.	x	x	x			x
Ingres Corporation	x					
Micro Decisionware, Inc. (MDI)	x	x	x			x
Novell	x					
Object Technology Intl., Inc. (OTI)	x	x	x			
Oracle Corporation	x	x				
Progress Software Corporation	x					
Sybase, Inc.	x					x
Wall Data, Inc.	x					
XDB Systems, Inc.	x	x	x	x		x

Notes:

- (1) Have LICENSED the DRDA Architecture
- (2) Have publicly "ANNOUNCED" DRDA AR implementation(s)
- (3) Have SHIPPED DRDA AR implementation(s)
- (4) Have publicly "ANNOUNCED" DRDA Application Server (AS) implementation(s)
- (5) Have SHIPPED DRDA AS implementation(s)
- (6) Have stated intention to implement DRDA Level 2

Q: Can Query Manager work with DDCS/2?

A: QM cannot access host databases using DDCS/2.

Q: What are the common pitfalls and errors using DDCS/2?

A: Communication and database configurations problems seem to be the most common pitfalls with DDCS/2.

Table 3 (Page 4 of 4). DB2/2 Q&As
Question & Answer
<p>Q: What do I look for if my application ABENDED while connected to a DB?</p> <p>A: Look for a .TRP file in x:\SQLLIB directory, where x: is the installation drive for DB2/2. This file is created during an ABEND while using DB2/2. The file contains the user's stack and current register information. Using this information you should be able to find the failing area of code, if you are unable to resolve the condition phone IBM support.</p>
<p>Q: What are the recommended trace settings for DB2/2?</p> <p>A:</p> <pre>>SQLTRC ON -1 1000000 -e 10 >SQLTRC DUMP TRACE.DMP >SQLTRC FMT TRACE.DMP TRACE.FMT</pre> <p>This procedure will create a 1MB trace file which is useful by IBM support to help resolve any DB2/2 problems. The TRACE.FMT file will be an ASCII formatted file with DB2/2 activity recorded during the problem recreation.</p>
<p>Q: When attempting to use DB2/2 with Lotus Datalens, I receive the following message when attempting to connect to a DB2/2 database:</p> <p style="padding-left: 40px;">ERROR CODE SQL0804N is returned and you are unable to connect.</p> <p>What is wrong?</p> <p>A: Lotus 1-2-3 for OS/2 has a driver called DATALENS that allows you to bring data from a database into your spreadsheet.</p> <p>This is a known problem with the Lotus Datalens product. The fix should be available from your local Lotus Technical Support centre in a SLIPSTREAM Driver package by the end of July 1993.</p>
<p>Q: I'm having a problem when installing OS/2 2.1 after ES 1.0 has been installed and vice versa. How can I correct this?</p> <p>A: If you are installing OS/2 2.1 on a system with ES 1.0 on already replace the existing ESSTART.COMD file with the ESSTART.BAK file from the OS2\INSTALL directory.</p> <p>If you are installing ES 1.0 over OS/2 2.1 then rename the existing ESSTART.COMD file from the ES 1.0 diskette 1. Then, copy the file ESSTART.BAK file from the \OS2\INSTALL directory to ES 1.0 diskette 1. To complete the procedure, rename the ESSTART.BAK file to ESSTART.COMD and begin the installation of ES 1.0.</p>
<p>Q: I have printing problems after installing DB2/2 on OS/2 2.1. I lose the ability to print any file to a printer after installing DB2/2 on a computer running with OS/2 2.1 operating system. The file goes to the PM print spooler. What is wrong?</p> <p>A: As documented in the OS/2 2.1 readme file, you must REMark the RUN=C:\OS2\EPW.EXE and RUN=C:\OS2\EPWRROUT.EXE 1 statements out of the CONFIG.SYS files and DETACH C:\OS2\EPW.EXE and DETACH C:\OS2\EPWRROUT.EXE -1 in the STARTUP.COMD file. DB2/2 will install the FFST/2 code (EPW) and update your config.sys file.</p>
<p>Q: The message SQL0818 - Timestamp conflict occurs when a DB2/2 V1.0 server is used with mixed clients (DB2/2 client and ES 1.0 client). Even though the user has the proper authority, the user is not able to access the database. If the user is a DB Administrator, they can access the database. What gives?</p> <p>A: A fix is available from IBM support - APAR JR07179.</p>

Section 2: DB2/6000 and DDCS/6000

Index of Publications

The following publications will be provided with DB2/6000 and DDCS/6000 at General Availability:

Table 4. Index of Publications	
Publication Number	Publication Title
SC09-1569	DB2/6000 Information and Planning Guide
SC09-1571	DB2/6000 Guide
SC09-1570	DB2/6000 Installation Guide
SC09-1572	DB2/6000 Programming Guide
SC09-1573	DB2/6000 Programming Reference
SC09-1574	DB2/6000 SQL Reference
SC09-1577	DB2/6000 Messages & Problem Determination Guide
SC09-1578	DB2/6000 Master Index and Glossary
SC09-1575	DB2/6000 Command Reference
SC09-1526	DB2/6000 Call Level Interface Reference
SC09-1568	DDCS/6000 Guide

DB2/6000 Function/Feature Matrix

Feature	Description for DB2/6000
Precompilers	
C	Yes
C+ +	No
Cobol	Yes - only from Micro Focus COBOL V3 compiler
Ada	No
Fortran	Yes
Pascal	No
PL/I	No
Basic	No
Call-level interface	
C	Yes - ODBC/CLI: the Core ODBC Functions plus 3 Level 1 ODBC functions plus 1 Level 2 ODBC function for a total of 28 functions. Windows applications would use the ODBC specifications and DOS and OS/2 applications would use the X/Open CLI specifications.
Callable APIs	
C	Database manager routines are available as function calls: environment routines, utility routines, configuration routines, client-server routines, run-time routines, additional development routines and database monitor calls
Cobol	Database manager routine function calls as above
Fortran	Database manager routine function calls as above
Pascal	No
REXX	No

Feature	Description for DB2/6000
Command Line Processor	Yes - can issue SQL statements and database manager commands that perform the same functions as the routines listed above from the operating system command line or from a shell script.
API features	
Positioned UPDATE/DELETE	Yes
Cursor Context Preservation	Yes. Allows database cursors to continue to be used after a commit thus avoiding the overhead of reestablishing each cursor in the application after a commit. (via DECLARE CURSOR WITH HOLD)
Scrollable Cursor	No
Procedural SQL	No
Stored Proc. /w Proc. SQL	No
Stored Proc. /w 3GL	Yes - Database Application Remote Interface (DARI). To use DARI, a calling procedure is contained in an application running on the client. A server procedure executes against the database at the database server. The client application declares, allocates, and initializes any data structures, connects to the database and invokes the server procedure by calling the ARI routine. The server procedure may be written in C or FORTRAN and may be stored in a static or shared library. It executes in a separate process space than its associated Database Agent. It is executed by a special ARI process running on behalf of the client. (The client must have sufficient authorities and privileges to execute the statements within the stored procedure.) It accepts optional input data structures from the client application and returns the SQLCA and optional output data structure to the client application upon completion.
Bind Files	Yes
Array Inserts	No
Compound SQL	Yes - groups several SQL statements into a single executable block. The statements execute at the server with no response back to the application until either all statements are successfully executed or an error is raised by any one of the statements.
Force Users	Yes - mechanism to force remote or local users selectively from the database connection.
Interfaces (Vendor's Own)	
Programmer tools	TBA
End-user tools	Query/6000
Applications	TBA
DBA Tools	Configuration Tool, Recovery Tool, Command Line Processor, Directory Tool, Reorgchk, Status commands, Explain.
Interfaces (Third Party)	
Programmer tools	TBA
End-user tools	TBA
Applications	TBA
DBA Tools	TBA
4GLs	TBA
Standards and Compatibility	
FIPS 127.1 (SQL89 Level 2)	No
Entry SQL2	No
DB2 Compatibility	partial - high level. Same as DB2/2 plus new this release: <ul style="list-style-type: none"> • VALUE scalar function • String concatenation with 'CONCAT' and ' '
SAA Compatibility	partial - high level. See new items added above - all SAA compliant.
MIA	Yes
NAS	No
ACE	No

Feature	Description for DB2/6000
CA90	No
X/Open XA Compliance	As a resource manager of a database resource, DB2/6000 can participate in a global transaction that is being coordinated by an XA-compliant transaction manager. DB2/6000 supports CICS/6000, Encina and Tuxedo as the XA-compliant transaction monitor.
SQL Function (Extensions)	
Scalar Functions	CHAR, DATE, DAY, DAYS, DECIMAL, FLOAT, HOUR, INTEGER, LENGTH, MICROSECOND, MINUTE, MONTH, SECOND, SUBSTR, TIME, TIMESTAMP, TRANSLATE, VALUE, YEAR
Column Functions	AVG, COUNT, MAX, MIN, SUM
Multi-table view update	No
Update with SELECT	No
SELECT as expression	No
Outer Join	No
Recursion	No
User Defined Functions	No
Very Long Fields	No - max is 32,700.
Event Alerters	No
Integrity Constraints	
Declarative RI	Yes - RESTRICT, CASCADE, SET NULL options on DELETE, RESTRICT on UPDATE.
Triggers	No
Base datatypes	Char(254), varchar(4000), long varchar(32700), smallint, integer, float, decimal/numeric, date, time, timestamp.
User-defined types	No
View check option	Yes
Column constraints	No
Column defaults	Yes - same support as for DB2 - the system will assign a pre-defined default.
NULL support	Yes
Security	
C2 level	No
B1/B2 level	No
User Group security	Yes - uses AIX group facilities
Audit trail	No
Authorization	Yes, SYSADM, DBADM, bindadd, connect, createtab, control, alter, delete, index, insert, references, select, update, bind, execute.
Roles	No; however, groups may be created that correspond to a required role - groups could be used to simulate roles.
Passwords	User authentication can be done on the client side or be passed on to the server. When the user is verified locally (at the client node), no passwords are sent across the network. The user is expected to be verified at the location they first signed on to. For server authentication, the password is sent over the network. DB2/6000 must receive the password in clear text ASCII, but the communication mechanism may encrypt it before it is sent across the network and decrypt it before passing it on to DB2/6000.
Distributed Database	
Remote Unit of Work	Yes - application requestor to DB2 on MVS, SQL/DS, SQL/400 via application program and DB2/6000 command line interface using DDCS/6000.
Distributed Unit of Work	No
Distributed Request	No

Feature	Description for DB2/6000
Distributed Request /w Global Optimization	No
Snapshots	No
Transparency	Yes - the database is catalogued once and referred to, from then on, by its alias.
IBM RDB Connectivity	Read/write access to DB2 on MVS, SQL/DS, SQL/400
OEM RDB Connectivity	No
Non-relational DB Connectivity	No
Communication Protocols	
NetBIOS	No
APPC/APPN	Yes - for OS/2 and AIX clients to DB2/6000 server
TCP/IP	Yes - for DOS, Windows, OS/2 and AIX clients to DB2/6000 server
IPX/SPX	No
Named Pipes	No
Portability (server)	
DOS	No
Windows	No
OS/2	Yes
AIX	Yes
SunOs	No
Ultrix	No
HP-UX	No
Pyramid	No
Sequent	No
Decstation (MIPS)	No
Macintosh	No
Novell Netware	No
Portability (client)	
DOS	Yes
Windows	Yes
OS/2	Yes - ES 1.0 & DB2/2 1.0
AIX	Yes
SunOs	No
Ultrix	No
HP-UX	No
Macintosh	No
Decstation (MIPS)	No
System Integration	
Single O/S	No
Multiple O/S	Yes - the system-specific calls are tuned for the particular operating system
DCE/OSI Support	No
Performance (Access Paths)	
Cost-based Optimizer	Yes
Query transformations	Yes - in some cases. E.g. if there's a suitable index on a column, an IN list might be transformed into a set of equality predicates connected by OR so that index-ORing may be done; there are a few more cases that are too long to explain here.
EXPLAIN	Yes - this tool is installed in the SAMPLE directory.

Feature	Description for DB2/6000
Index ORing	Yes
Index ANDing	No
Index-only access	Yes
Recreate Index option on database startup	Yes
Performance (Architecture)	
Static SQL	Yes
Process Model	Multi-server architecture: n - n process model with 32-bit exploitation. Application programs are isolated from internal database manager shared data structures and data files. SQL data file handles are owned by database manager Agents and therefore inaccessible to user application code. User applications will be unable to accidentally compromise the integrity of their own or another user's data.
Symmetric Multiprocessing	No
Parallelism	No
Data Organization	
Clustering Index	No - rows physically stored in index order when reorganized but will become unclustered with inserts and updates.
Storage methods	B ⁺ tree
Inter-table cluster	No
Write-once media (optical)	No
Performance (I/O Management)	
Group Commits	Yes
Deferred writes	Yes
Data page buffering	Yes
Chained writes	No
Sequential Prefetch	No
Index Prefetch	No
Performance (Tuning)	
Performance monitor	API to provide snapshot information about the current state of the database manager. Type of information collected: database connections, locks, buffer pool activity, SQL statement execution, sorting, cursors, deadlocks, table activity, unit of work status, database status, and communications activity.
Governor	No
Tuning parameters	Configuration parameters for the database manager and each database, ARI, record blocking, isolation parm at bind
Transaction monitor	Yes - support for CICS/6000, Encina and Tuxedo with XA compliance as a resource manager
Recovery	
Forward Recovery	Yes and point-in-time recovery is possible
Log Archive	Yes
Incremental Backup	No
Granular Recovery	No
Set logging off	No
Tape Management	No
Availability	
Separate Logs and DB	Yes
Dual Logging	No
Mirrored Disks (OS Level)	Yes - using Logical Volume Manager
Mirrored Disks (SW Level)	No

Feature	Description for DB2/6000
Online Backup/Restore	Yes - online backup to fixed disk, tape or to the DFDSM utility. Can be invoked from an application, via the Command Line Processor, via the recovery tool or via the DFDSM utility.
Online Configuration	No
Online Reorganization	Yes - table level reorg - only the table and its indexes are unavailable during a reorg.
Run with Partial Database	No
Locking	
Row-level locks	Yes
Page-level locks	No
Table-level locks	Yes
Lock Escalation	Yes, row to table
Wait/nowait lock	No, always waits for a lock to be freed.
Deadlock detection	Aborts transaction that caused deadlock and on timeout.
Isolation Level	
Repeatable Read	Yes
Cursor Stability	Yes
Dirty Read	Yes
Multi-version reads	No
Committed read	No
Autonomous/Remote Operations	
Remote Installation	Yes - may use netinstall or additional products such as NetView DM and SOFTDIST/6000.
Remote Administration	No
Autonomous Reorganization	No
Autonomous Statistics	No
Manageability	
Truncate Table	No
Analyze Table	No
Fast dataload	Yes, via Bridge/FASTLOAD for DB2/6000
Object Features	
Object IDs	No
Smalltalk Pre-Compiler	No
Multiple inheritance	No
Late binding	No
Overloading	No
Class Libraries	No
Links	No
Encapsulation	No
Object caching	No
Pointer swizzling	No
Object-level locking	No
Check-in/check-out locks	No
Miscellaneous	
National Language Support	Yes, sorting/collating, translated error messages and documentation for seven languages

Feature	Description for DB2/6000
Utilities	Import/Export- can import from these file formats: non-delimited ASCII, delimited ASCII, WSF (Lotus worksheet), PC IXF and can commit at a regular designated interval so that following a crash the import would not have to start at the beginning; can export in these file formats: delimited ASCII, WSF & PC IXF; Configuration Tool; Recovery Tool; Directory Tool.
Diagnostics	Corruption detection, trace facility
Installation	Follows established AIX installation guidelines using installp from the command line or via SMIT menu interface.

Question & Answer

Question & Answer												
<p>Table 5 (Page 1 of 2). DB2/6000 Q&As</p>												
<p>Q: What is the maximum size for a database?</p> <p>A: The table and database size limits are being increased from the DB2/2 version of the product. Maximum table size is 64GB, not including indexes or long field data. The maximum database size would be measured in petabytes (peta is the next size after tera), so the actual number is essentially irrelevant.</p>												
<p>Q: What database objects (databases, tables, views, indexes, and packages) can or cannot have the same name?</p> <p>A: The only namespace restriction is that a table cannot have the same name as a view.</p>												
<p>Q: What are my connectivity options with DB2/6000?</p> <p>A: Here are your options:</p> <table border="1"> <thead> <tr> <th>Client Type</th> <th>DB2/6000 Server</th> </tr> </thead> <tbody> <tr> <td>ES 1.0 OS/2</td> <td>APPC</td> </tr> <tr> <td>DB2/2 V1.0</td> <td>APPC</td> </tr> <tr> <td>DB2 Client Enabler/DOS</td> <td>TCP/IP</td> </tr> <tr> <td>DB2 Client Enabler/2</td> <td>APPC, TCP/IP</td> </tr> <tr> <td>DB2 Client Enabler/6000</td> <td>APPC, TCP/IP</td> </tr> </tbody> </table>	Client Type	DB2/6000 Server	ES 1.0 OS/2	APPC	DB2/2 V1.0	APPC	DB2 Client Enabler/DOS	TCP/IP	DB2 Client Enabler/2	APPC, TCP/IP	DB2 Client Enabler/6000	APPC, TCP/IP
Client Type	DB2/6000 Server											
ES 1.0 OS/2	APPC											
DB2/2 V1.0	APPC											
DB2 Client Enabler/DOS	TCP/IP											
DB2 Client Enabler/2	APPC, TCP/IP											
DB2 Client Enabler/6000	APPC, TCP/IP											
<p>Q: When a record is locked and another user tries to access it, the other user cannot do anything until the lock is released. Is there a parameter or something that can be set to allow the other user's request to timeout or get some sort of error message in return instead? The other user would like to be able to continue with other type of processing instead of waiting for the release of the lock.</p> <p>A: DB2/6000 does not provide that capability</p>												
<p>Q: Why does this statement not work - db2 select * from table_name?</p> <p>A: The select command is interpreted without the quotes, and * is taken as the entire list of file names in the current directory. Double quotes are needed around the command (i.e. db2 "select * from table_name").</p>												
<p>Q: How can I do the same using Call Level Interface for this statement: EXEC SQL FETCH c1 INTO :lMessageNumber:N1, :zFunctionalArea:N2; The null indicator variable is important.</p> <p>A:</p> <pre>SQLINTEGER lMessageNumber; SQLCHAR zFunctionalArea*...+1"; SQLINTEGER N1,N2; SQLBindCol (hstmt,1,SQL_C_LONG,&lMessageNumber,4,&N1); SQLBindCol (hstmt,2,SQL_C_CHAR,zFunctionalArea,sizeof(zFunctionalArea,&N2); /* SQLBindCol calls only have to be performed once. They are valid */ /* until SQLFreeStmt(hstmt, SQL_UNBIND) is called. */ SQLFetch(hstmt); /* After SQLFetch call, N? will either contain the number of */ /* bytes in corresponding variable, or SQL_NULL_DATA (-1) which */ /* indicate that a NULL value for that column. */</pre>												

Table 5 (Page 2 of 2). DB2/6000 Q&As										
Question & Answer										
Q: What is the replacement function for sqlabndr() that was on OS/2?	A: The replacement is sqlabndx(). Basically, it has exactly the same function as sqlabndr(), but requires you to be connected to a database (via the EXEC SQL CONNECT TO statement) before calling the API. This is to allow the BIND operation to use the database username and password specified in the CONNECT.									
Q: DB2START returns error code SQL5042N. What happened?	A: Ensure that you have the environment variable DB2COMM=APPC defined.									
Q: DB2START returns error code SQL5041N What happened?	A: Setting the environment variable DB2COMM=APPC indicates that you want to start up the APPC server support. Hence, you must define the TP name in the Database Manager Configuration file, if you intend to use APPC or remove APPC from the DB2COMM variable.									
Q: Receiving SQL0969N stating that there is no message text corresponding to SQL error "-1403" in the message file when connecting to a DB2/6000 server from an OS/2 client. How do I correct this?	A: SQL error code -1403 is a DB2/6000 return code which means invalid user name and/or password during connection. Ensure that the userid that you enter on the client has a corresponding lower case definition on the 6000. The password on the 6000 for this userid will have to be in upper case since userids are folded to UPPER case but passwords are not. For example: <table border="0"> <tr> <td>OBJECT</td> <td>OS/2 UPM Entry</td> <td>AIX/6000 User id Entry</td> </tr> <tr> <td>userid:</td> <td>MYUSER</td> <td>myuser</td> </tr> <tr> <td>password:</td> <td>MYPASSWD</td> <td>MYPASSWD</td> </tr> </table>	OBJECT	OS/2 UPM Entry	AIX/6000 User id Entry	userid:	MYUSER	myuser	password:	MYPASSWD	MYPASSWD
OBJECT	OS/2 UPM Entry	AIX/6000 User id Entry								
userid:	MYUSER	myuser								
password:	MYPASSWD	MYPASSWD								
Q: Can OS/2 database applications developed to access Extended Services 1.0 Database Manager or DB2/2 database servers also now access a DB2/6000 AIX database server?	A: Yes. Customers using or planning to use OS/2 Extended Services 1.0 Database Manager or the recently-announced DB2/2 database products will be able to also get access to data stored in a DB2 database via the APPC communications protocol. Applications running on the OS/2 platform can act as APPC database clients to a DB2/6000 database server. The DB2/6000 database server must also have the SNA Support Feature installed which provides the APPC communications protocol required to support these clients.									
Q: Does DB2/6000 provide a sample database so that I can test my application?	A: SQLSAMPL is an executable file that installs a sample database. To install the database, you must have SYSADM authority. At the command prompt, type SQLSAMPL x where x is the drive letter. For a hardcopy version of the sample database, refer to the DB2/6000 Guide.									
Q: Does DB2/6000 support Compound SQL? If so, how? Can you provide a code sample?	A: Yes, Compound SQL is supported. Compound SQL is useful for minimizing communication between the application and the database engine (either local or remote). Basically, it removes the (non-trivial) pathlength involved in doing Inter-process communication for each sub-statement. Depending on workload, this can significantly improve your performance. Each sub-statement of your compound SQL statement is then run in order - there's a loop near the top of your agent processing which drives each sub-statement to completion and manages the aggregation of the sqlca information. Here is some sample code to do this: <pre>exec sql begin compound atomic static UPDATE accounts SET Abalance = Abalance + :delta WHERE Aid = :Aid; SELECT Abalance INTO :Abalance FROM accounts WHERE Aid = :Aid; UPDATE tellers SET Tbalance = Tbalance + :delta WHERE Tid = :Tid; UPDATE branches SET Bbalance = Bbalance + :delta WHERE Bid = :Bid; INSERT INTO history(Tid,Bid,Aid,delta,ptime) VALUES(:Tid,:Bid,:Aid,:delta,CURRENT TIMESTAMP); COMMIT WORK; end compound;</pre>									
Q: Does DB2/6000 support 2-phase commit?	A: Support for 2-phase commit is provided through a transaction manager such as CICS/6000.									

Section 3: IBM DATABASE 2 Version 3 for MVS

Question & Answer

Table 6 (Page 1 of 3). DB2 for MVS Q&As
Question & Answer
<p>Q: Can you give me a brief summary of the DB2 Version 3 announcement?</p> <p>A: Key improvements are being made to DB2 in response to customer and market requirements in three broad areas: availability, client/server, and performance.</p> <ul style="list-style-type: none">• Recognizing the importance of pursuing continuous availability goals, DB2 Version 3 minimizes planned outages and helps provide faster recovery from unplanned outages with the new support of partition independence and data compression.• DB2's value for open systems and client/server is enhanced by multi-site update. These changes are based on Distributed Relational Database Architecture (DRDA), and strengthen DB2's role as the premier enterprise data server.• Compression can also improve the price/performance of customers' systems.• Query performance is enhanced by the support of query input/output parallelism. Buffer management enhancements and hiperpool support allow customers to improve the use of processor storage, and to take advantage of improved DB2 performance.• DB2 Version 3 will also meet the increasing quality needs of its customers, participating in the Quality Partnership Program (QPP), a development process initiated by IBM to further ensure high-quality software delivery. Once product quality is confirmed by use in a defined production environment, IBM will provide availability and ordering information.
<p>Q: What are the key messages of the DB2 Verion 3 announcement?</p> <p>A:</p> <ol style="list-style-type: none">1. Chief among the themes of DB2 Version 3 is IBM's continuing commitment to the client/server environment, and to making DB2 the premier relational data server in the industry.<ul style="list-style-type: none">• "DB2 has become one of the most vital pieces of software technology in corporate America," wrote Corporate Computing in its January 1993 edition, awarding DB2 Version 2 Release 3 a "Best Buy for 1993" award.• "If you're developing an information-anywhere approach to connectivity, DB2 is the place to start", wrote Corporate Computing in its January 1993 edition, awarding DB2 Version 2 Release 3 a "Best Buy for 1993" award.<p>Version 3 extends DB2's distributed capabilities significantly with its Multi-site Update feature, and by extending the number of concurrently active distributed users. IBM is signalling its long-term commitment to connectivity and openness with this new version.</p>2. Version 3 also represents a giant step in DB2's continuing progress toward the goal of 24X7 operations:<ul style="list-style-type: none">• Enabling the concurrent execution of multiple batch utilities.• The use of parallel I/O processing, and a resulting reduction in the size of the batch window.<p>Availability has always been a primary goal for DB2, and Version 3 validates that goal once again with significant improvements.</p>3. Improved performance continues to be a major theme for DB2, and Version 3 offers several enhancements. Large queries, for example, can have improved response time via:<ul style="list-style-type: none">• Increased I/O bandwidth associated with I/O parallelism.• Reduced I/O operations through compression and expanded buffer pools.
<p>Q: Is DB2 is a mature technology?</p> <p>A: Relational Technology although mature is still evolving. DB2 has been, and continues to be, the best of breed large scale relational database. Customers can continue investing in DB2 knowing that we will provide the support necessary for their continued growth as they move towards the world of open Client/Server computing and new application types involving objects and multimedia.</p>

Table 6 (Page 2 of 3). DB2 for MVS Q&As
Question & Answer
<p>Q: With this new Version can DB2 really be considered a viable solution for people who want to take advantage of Client/Server technology?</p> <p>A: Absolutely. Many large enterprises will find the small-scale, LAN-based approach to client/server technology both costly and restrictive. They will require better performance and availability than can be offered by these solutions. DB2, with its robust implementation of DRDA, and its proven record for superior performance, availability, and recoverability, will be a very logical answer to this dilemma. Offering, as it does, up to 2000 concurrently active distributed connections, Version 3 is a natural choice for enterprises with large scale client/server needs.</p>
<p>Q: When is the GA? Why won't IBM reveal the GA at this time? Customers need to know this in order to do the appropriate planning.</p> <p>A: We have worked hard in Version 3 to ensure significantly improved quality. First Customer Shipment of DB2 V3 will be in June '93. This signals the beginning of a validation program with our QPP customers to validate this quality, and to determine the date of general availability.</p>
<p>Q: What kind of services are available in the United States?</p> <p>A: The IBM Application Enabling Support Center has extensive expertise in database design, performance analysis and tuning, database recovery and repair, client-server database services, distributed application and database design as well as Information Warehouse database design.</p>
<p>Q: Why was the Catalog Visibility function removed ?</p> <p>A: The Information Warehouse strategy is to provide customers with products that best meet their needs in the areas in which those products are strongest. Catalog Visibility provided a series of interactive panels that displayed information from the DB2 catalog. Similar function is provided not only by DataHub, but also by our International Alliance Members and development partners, such as Platinum, Candle and Legent, with their respective products.</p>
<p>Q: What is the DB2 Version 3 support for DFSMS Concurrent Copy?</p> <p>A: The DB2 Version 3 Recover Utility has been enhanced to provide a Log-only Recovery function. This enhancement permits:</p> <ul style="list-style-type: none"> • Bypassing the Recover utility restore phase. Users can therefore restore DB2 datasets from non-DB2 backups, such as those provided by the DFSMS Concurrent Copy. • Performing forward recovery with only log records. <p>This enhancement will make the task of recovering from non-DB2 copies more feasible, thereby permitting customers to take advantage of the of DFSMS Concurrent Copy's availability and performance.</p>
<p>Q: Will I have to compile and rebind any of my applications if I move to this new version of DB2?</p> <p>A: In general, the answer is no. There are circumstances that will require rebinding to take advantage of new functions. The use of query I/O parallelism by your application is a good example. Others include data compression and distribution.</p>
<p>Q: I am running a previous release of DB2. Can I migrate directly to Version 3?</p> <p>A: If you are running Version 2 Release 3, you can migrate directly. If you are running a prior release, you must first migrate to Version 2 Release 3.</p>
<p>Q: Can I have the same logical unit of work requests update non-DB2 data (IMS, VSAM)?</p> <p>A: Yes. Multi-site updates extend the unit of work to include remove updates. IMS and CICS two-phase commit is synchronized with multi-site update.</p>
<p>Q: Do I need the ESA/9000 hardware in order to take advantage of the data compression function ?</p> <p>A: The ESA/9000 data compression hardware will reduce the CPU load associated with the compression and decompression of data as well as provide superior performance. Version 3 provides software data compression as an alternative in the event the hardware is not available. Software data compression performance is similar to that provided by our DSN8HUFF (Huffman) sample, which was provided in DB2 Version 2 Release 3.</p>
<p>Q: Will there be tools to help me understand what data I should compress, how much compression I can achieve, and how much CPU time compression will require?</p> <p>A: Each customer workload is unique. Data compression can provide significant reduction in DASD space required for large table spaces by compressing the data without information loss. The customer compression ratio is customer-data dependent and there will be some variation. For the typical case (text, integer, decimal data; row size 100-300 bytes), we would expect good compression and good query performance. Row size of 100-300 bytes are typical of many query applications.</p> <p>Version 3 provides a new stand alone utility, DSN1COMP, that estimates how much space can be saved without actually compressing the data so that customers can determine where data compression will be effective for them.</p>

Table 6 (Page 3 of 3). DB2 for MVS Q&As
Question & Answer
<p>Q: My installation doesn't have large partitioned tables spaces. Will we still derive benefits from Query I/O parallelism?</p> <p>A: There is no direct support for simple or segmented tables. The primary benefits are provided for very large tables, which generally need to be partitioned for a variety of reasons. Version 3 adds two more reasons for partitioning: query I/O parallelism, and partition independence. If your table sizes are measured in gigabytes, we strongly encourage partitioned tables.</p>
<p>Q: Will partition independence provide me with complete, continuous availability for my mission critical applications?</p> <p>A: We believe partition independence to be the most important current requirement for continuous availability, but there are many other needs. Continuous availability is a critical area for DB2, and more improvements are planned. Continuous availability requires careful management. Key examples are to keep commit intervals short, to manage copy and recovery processes, to examine performance and concurrency, and to practice recovery.</p>
<p>Q: We are not planning to upgrade our processor to acquire the necessary hardware for support of hiperpools. Without this, what benefits will we see with the new buffer pool management support?</p> <p>A: Improvements in buffer management include improved availability, improved tuning, and better flexibility and control.</p> <ul style="list-style-type: none"> • Continuous Availability: Buffer pool sizes may be configured while DB2 is up to satisfy different workloads. This allows system administrators to improve transaction, query, and batch response time by increasing database caching space when more central and expanded storage is available. • Buffer pool display and tuning: Buffer pool tuning information can be obtained with the new display, and the parameters can be changed dynamically. • Better flexibility and control over buffer pools: The number of DB2 buffer pools has been increased from 4 to 60. This enables database administrators to have more flexibility to give preferential performance to some data.
<p>Q: How will this version of DB2 affect the performance of my decision support applications?</p> <p>A: By providing Query I/O Parallelism, new buffer management support, and compression, customer applications that take advantage of these functions may see significant response time and cost improvements.</p>
<p>Q: One of my customers' biggest concerns is quick recoverability from unplanned outages. Does this new version of DB2 help with this?</p> <p>A: DB2 Version 3 provides a number of enhancements to facilitate improved recoverability:</p> <ul style="list-style-type: none"> • Partition Independence enables a single RECOVER job to be broken up into separate jobs, each processing a different partition of a table space. The jobs are submitted in parallel. The total elapsed time required to recover the partitions of the same tablespace in parallel is significantly less than a for single recover job processing the entire tablespace. • Version 3 also improves archive log processing and reduces recovery time by using multi-tasking for increased parallelism, and by facilitating the way tapes are handled and mounted.

Software Spotlight

Software spotlight serves to highlight new products that have come to the marketplace that are designed to work with DB2/2 and DB2/6000.

IBM Personal AS/2 V3

This section deals with the IBM Personal AS/2 V3 (Personal AS/2 V3) and IBM Personal AS/2 Builder/2 (Personal AS/2 Builder/2).

Overview

Personal AS/2 V3 is IBM's 32-bit OS/2 solution for business professionals working in an office environment who have requirements for data access, query, data analysis, presentation and communication of business results.

Personal AS/2 provides extensive reach to relational and non-relational data at an Enterprise, Departmental or personal level via

- Exploitation of Distributed Relational Database Architecture(DRDA)
- Use of Dynamic Data Exchange (DDE)
- Support of Enterprise Data Access (EDA) technology
- Linkage with AS/400 via PC Support as well as DRDA
- Use of the Application System (AS) Intelligent Server
- Native ability to read and write various data formats

Personal AS/2 contains a wide range of decision making tools that transform data into valuable information. These integrated facilities include

- Table Edit and Browsing
- Query
- Data Analysis and Manipulation
- Report writing and usage
- Chart creation and usage
- Task automation via iconic procedures

Personal AS/2 is designed for use by non computer professionals and provides exceptional levels of ease of use through its exploitation of the OS/2 Workplace Shell.

Integration with IBM's OS/2 office systems will be provided enabling files and data to be sent to both IBM and non IBM applications.

Users who develop customized solutions for others will benefit greatly from the functions of Personal AS/2 Builder/2. Facilities include:

- a highly productive event driven development language
- comprehensive SQL access to the OS/2 Client Server Database
- support of DDE as both a client and a server
- programmable interface to the standard mail APIs.
- Powerful editors for the visual layout of windows and the placement of controls
- Intelligent program editor to enable rapid logic creation
- Interface to 'C'
- Support for clipboard
- An interactive debug capability
- Application Program Interfaces (APIs) to objects such as Report and Chart

Personal AS/2 Builder/2 should be viewed as a rapid high level development tool with a smaller learning curve when compared with languages such as 'C'. It possesses unique capability when combining customisation with other powerful business tools via its APIs and DDE linkages.

Personal AS/2 and Personal AS/2 Builder/2 both work in Client Server mode to take full advantage of the data manipulation and processing power of Application System.

Personal AS/2 and Personal AS/2 Builder/2 are significant product enhancements to IBM Personal Application System/2 V2 and its Development/2 option. As such customer solutions developed with either Personal AS/2 V2 or Personal AS/DOS for Windows will execute in the Personal AS/2 V3 environment.

Note: Direct movement between operating systems is provided but the ease of moving depends on whether the application has used operating system dependent facilities.

Highlights

- Key position within the Information Warehouse framework
- IBM's 32-bit client server front end to data
- Exploits the OS/2 workplace shell
- Iconic user interface including Drag+Drop
- Powerful Personal AS/2 Builder/2 functions for rapid customisation of OS/2 based client server applications
- Enabled for communication and receipt of information through standard mail functions
- Interoperability with leading workstation products via DDE
- AS/400 exploitive applications via Builder/2
- 32 Bit Client to the ASV3.2 Intelligent Server

Personal AS/2 (5622-103 and 5622-104) Supported Environments

Hardware

- PS/2 386 25Mhz or above, PS/2 486 recommended
- recommended 8MB Memory additional to OS/2
- 8MB disk space for product files
- Mouse, Printer ,Color display, Math Co processor

Software

- OS/2 V2.0 (Servicepak XR06055 or later) plus OS/2 Extended Services
- Distributed Database Connection Services/2 for DRDA Links
- PC/Support/400 Version 2.1.1 for connection to AS/400
- ASV2.2 or 3.X for S/390 Connection to AS via SRPI
- ASV3.2 for S/390 Client Connection to AS via CPI-C

DataHub

Summary

IBM SystemView Information Warehouse DataHub is designed to meet the challenges of managing complex database environments while increasing productivity and saving costs.

DataHub provides a set of integrated database management functions for relational database across four Systems Application Architecture (SAA) platforms. These functions significantly simplify the tasks of database administrators, system administrators, help desk personnel, system programmers, and application programmers in distributed as well as centralized database environments. DataHub functions address database systems management tasks in all areas of operations.

DataHub support the achievement of customer business objectives through:

- Designated workstation control point
- Automatic handling of many of the differences between relational database management systems
- Seamless invocation of tools using object-action orientation.

Highlights

DataHub is designed to help increase the productivity of the systems management staff and simplify database management tasks. DataHub offers a workstation-based control point that provides a consistent object-based graphical user interface that associates data with function. From this designated control point, users can:

- Distribute data across the enterprise by copying database objects from one relational database management system to another
- Perform database management tasks across multiple homogeneous or heterogeneous database management systems
- Perform database management tasks across multiple sites
- Reduce skill requirements by masking most of the differences among IBM relational database management systems
- Speed up problem detection and minimize downtime.

DataHub Release 1 supports these SAA relational database management systems:

- DB2/2 or OS/2 Extended Services database manager for the OS/2 environment
- OS/400 database for the OS/400 environment
- IBM DATABASE 2 (DB2) for the MVS/ESA or MVS/XA environment
- SQL/DS for the VM/ESA or VSE environment

DataHub extends the advantages of its database management functions to the IBM Information Warehouse framework. Administrators can use DataHub to create and manage a customer's implementation of the Information Warehouse framework. DataHub provides the ability to copy database objects from one relational database to another.

DataHub addresses the database category within SystemView. It provides tools in five of the SystemView disciplines: business management, change management, operations management, problem management, and configuration management. DataHub conforms to level 1 SystemView integration providing a user interface consistent with other SystemView conforming products.

DataHub is designed to promote the integration of database systems management tools developed by IBM, other software vendors, and customers. Its integration platform provides developers with a set of common services on the workstation and the host. These services help reduce development efforts and promote consistency among products.

The DataHub Tool Builder's Guide and Reference (SC26-3046) describes how to utilize DataHub's common services when developing database systems management products. The DataHub integration platform is used by all DataHub functions.

ProductManager/6000

Summary

The IBM ProductManager licensed programs are IBM's application solutions for Product Data Management, Product Definition and Release, and Product Change Management. These comprehensive solutions provide robust function and state-of-the-art object-oriented technology which, when combined, enable the implementation of a comprehensive Engineering Data Management system in full support of concurrent engineering practices in a centralized or distributed open systems environment.

With this announcement, the ProductManager programs are now available for operation on the IBM RISC System/6000 family of advanced workstations. These products, already announced and available for the MVS environment (refer to Announcement 293-009, dated January 26, 1993), have been extended to provide consistent functions across the MVS and RISC System/6000 AIX platforms and the flexibility of choice in operating environment and configuration. They are being announced as Version 2 because they have the advanced function provided by ProductManager/MVS Version 2, announced January 26, 1993.

Running under IBM's RISC System/6000 Advanced Interactive Executive (AIX) operating system, ProductManager/6000 provides the same functionality as is available in the MVS/TSO environment, and adds an enhanced Graphical User Interface, to provide an advanced function, low cost, distributed engineering management solution. This brings ProductManager function to smaller enterprises and

allows more flexible choices for medium and large enterprises, enabling an enterprise to rapidly right-size their product data management environment.

The availability of the ProductManager applications in the RISC System/6000 AIX environment is particularly suited for those accounts that have large investment in workstations and/or multiple departmental environments within the manufacturing, engineering, process, utilities and construction industries. Each department or project can be autonomous having its own data and ProductManager/6000 installation, yet distributing information for review to other departments or projects to facilitate concurrent engineering, thus reducing product development and change cycle time.

Banking, Insurance, Health, and Distribution/Retail industries may benefit through the use of the Application Services Manager/6000 electronic folder facility to meet their industry-specific electronic foldering requirements.

When coupled with the optional use of IBM AIX EXPRESSight/6000, and IBM Distributed Application Environment's Entry Communications System/6000, these programs provide a tightly integrated solution for the creation, distribution, review and approval processes of electronic work folders and their associated graphics or image contents. In addition, with optional use of IBM Common Data Facility (CDF/MVS), mainframe users can integrate drawings and images stored in host data bases.

ProductManager/6000 includes three products:

- Application Services Manager/6000 - an advanced function, stand-alone application enabler which provides common services such as an electronic folder, administrative services and productivity enablers for rapid customization of the ProductManager/6000 applications.
- Product Change Manager/6000 - which provides the ability to manage the design and release process within an enterprise, for engineering and manufacturing, and facilitates concurrent engineering practices through its Engineering Change (EC) functions, as well as the provision of a Request for Engineering Action (REA) function which facilitates an automated process for communication and approval of engineering change requests.
- Product Structure Manager/6000 - which provides the ability to define and maintain multiple versions of Product Definition Data and Product Structure Data, such as Bills of Material, for both engineering and manufacturing, within a department, project or throughout an enterprise.

Highlights

In addition to comprehensive product data management functionality, ProductManager/6000 provides additional function and technology available with the RISC System/6000 AIX and advanced workstation environment. These additional functions include:

- An advanced Graphical User Interface (GUI), architected for ease of use and aligned with OSF/Motif guidelines.

Through the exploitation of AIXwindows Environment/6000, an icon-based mouse-driven point and click user interface enhances usability, user productivity, application integration and provides for effective presentation of data to the ProductManager user's desktop. ProductManager/6000, running on the RS/6000, can be accessed by any local area network attached workstation, PS2 or PC with X-Server (X Version 11 Release 4 compliant) capabilities, preserving a customer's current workstation investments.

- Exploitation of IBM's Client/Server Architecture through the use of IBM Database 2/6000 (DB2/6000).

ProductManager/6000 facilitates distributed management of an enterprise's engineering management business processes across multiple connected local area networks, with specific support for distributed management of Electronic Folders, Engineering Change Reviews and Deliverables, as well as Requests for Engineering Action. These powerful functions provide support for both small and medium installations or provide a large enterprise with the flexibility of choice to install and configure one or more ProductManager/6000 instances in support of their implementation requirements and ensure the availability and synchronization of these processes across the multiple environments.

- Enhanced network integration capability for supporting heterogeneous data.

IBM AIX Entry Communication System/6000 (A Distributed Application Environment Program Product) may be optionally used with ProductManager's Distributed Resource Manager to provide access to heterogeneous data stores referenced through Logical Data Reference Objects (LDROs) in the ProductManager/6000 database. These referenced objects may include such items as drawings, specifications, and images. Through the optional use of IBM AIX EXPRESSight/6000, the drawings may be viewed and annotated as part of a ProductManager process. Using IBM Distributed Application Environment's Entry Communications Systems/6000 ProductManager/6000 also supports movement of referenced objects to and from CATIA Data Management (CDM) and AIX files, enabling drawings to be brought into the ProductManager session for viewing. Linkage to IBM Common Data Facility

(CDF/MVS), when resident in a network, is also supported to send and receive Logical Data Objects (LDOs) from MVS data stores. Customization capability and published procedures are provided to extend the default support to other CAD data stores.

Description

ProductManager helps plan, document, distribute, and track the release of new products and engineering changes while controlling the flow of product and process information. Such information ranges from changes in product design to the type of material used in a part.

ProductManager/6000 offers fast, easy access to product information. Utilizing one or more of ProductManager/6000 systems, an enterprise can:

- Increase early manufacturing awareness and involvement in the initial design of products and changes to these products. This supports the implementation of concurrent engineering.
- Automate control of the distribution, review, and approval process for engineering change and manufacturing engineering change management.
- Automate work folder definition, distribution, and approval.
- Improve control over engineering change cycle time.
- Improve efficiency in evaluating design introduction and design changes.
- Implement manageable and orderly changes to the enterprise's business processes.

Investment Protection

To protect the investment of ProductManager/MVS customers:

- Customisation capabilities provided with ProductManager/6000 are consistent with those currently available in ProductManager/MVS.
- ProductManager/6000 dialog tags, used to define panels, are consistent with ProductManager/MVS dialog tags.
- Product Data Interface Files (PDIF) that were used to perform initial load of data into ProductManager/MVS can also be used to perform initial load of data into ProductManager/6000. In addition, through customer code at an installation-wide exit or through a request from outside ProductManager, a ProductManager/MVS customer can export selected data to an MVS data set. He can then move this data to a RISC System/6000 AIX file and invoke ProductManager/6000 to load this data through the initial data load capabilities of ProductManager, provided the data is in the appropriate status. The same scenario can be used to move selected data from ProductManager/6000 to ProductManager/MVS.

Application Services Manager/6000

Application Services Manager is a stand alone application enabler which provides common user services such as folder management, system administration functions and system services. It also provides the base services for the Product Change Manager and Product Structure Manager applications. Application Services Manager provides common user and application services required to:

- Enable users to package information in an "electronic file folder" and distribute the folder for review and action by other local or remote users.
- Support ProductManager's Administrative Services, such as:
 - User profile management
 - Security control
 - Definitions of location, company and supplier data
- Customize (or tailor) the ProductManager system to match the business processes established within an enterprise.
- Extend or add ProductManager functions to meet the unique needs of the enterprise.
- Import data into ProductManager using ProductManager's published product data interface format without requiring knowledge of its object-oriented environment.
- Export information through a customer developed program using a ProductManager installation-wide exit or via a request originating outside ProductManager.
- Provide an Event Log of significant activities within the major ProductManager processes. This can be accessed by customer written applications to trigger events in downstream systems, send data to downstream systems, or request ProductManager to export objects that were involved in the event.
- Support an object oriented development environment.

Product Structure Manager/6000

Product Structure Manager provides the ability to define and maintain product definition data and product structure data, such as bills of material, for both engineering and manufacturing. Product Structure Manager provides the capability to:

- Create and manage multiple levels or versions of item and product structure data.
- Migrate existing engineering and manufacturing product information into the ProductManager system.
- Define initial product definition data in a structured, yet flexible, development environment without engineering change (EC) control.
- Convert non EC-controlled product data to EC control for further review and release to manufacturing organizations and locations.
- Establish new part numbers to authorized groups or departments for automatic assignment when creating new items.
- Establish a link between product item data and reference documentation such as drawings and specifications.
- Establish and maintain multiple views of a product bill of material.

Product Change Manager/6000

Product Change Manager/6000 provides the capability to manage the design and release process within a department, project, or an enterprise through its Engineering Change (EC) functions, for engineering and manufacturing. There is also the provision of a Request for Engineering Action (REA) function which facilitates an automated process for communication and approval of engineering change requests. With Product Change Manager the enterprise can:

- Create, maintain, and release engineering changes (EC) for developing new or existing products.
- Create and maintain manufacturing engineering changes (MEC) for released products.
- Implement an automated problem identification process, allowing any area of the enterprise to define problems and pursue solutions.
- Assign approved Request for Engineering Actions (REA) to responsible engineers and automatically create engineering change data from these REAs.
- Define and track the tasks required to prepare an engineering change for release and introduction to manufacturing.
- Establish an automated engineering change review process that supports automated creation, distribution, review, and approval processes.
- Manage complex engineering changes and ensure an orderly release of these changes.
- Establish an integrated product development and release control process that supports multiple enterprise-wide design and manufacturing departments or projects.

Bridge Technology - BRIDGE/FASTLOAD

The following product description applies to both the DB2/2 and DB2/6000 versions of BRIDGE/FASTLOAD. The DB2/2 version is currently in production. The DB2/6000 version will ship at the same time as IBM's first production ship of DB2/6000.

BRIDGE/FASTLOAD significantly reduces the time and disk space you need to load data into IBM's DATABASE 2 for OS/2 (DB2/2) and DATABASE 2 AIX/6000 (DB2/6000). You can take advantage of the latest enhancements in IBM's client server technology by using BRIDGE/FASTLOAD to move corporate data from your mainframe to OS/2 or AIX LANs, or from one LAN to another. BRIDGE/FASTLOAD's powerful toolset allows you to move large quantities of data faster than conventional methods, and with more efficient use of hardware resources. Time savings of up to 10:1 can be achieved.

BRIDGE/FASTLOAD uses a proprietary data loading procedure to produce a DB2/2 or DB2/6000 DAT retrieval file. This file is then placed directly into DB2/2 or DB2/6000 without invoking conventional logging procedures. The result is a dramatic reduction in the elapsed time required to complete the data load.

Eliminating the logging process not only reduces load time, but also reduces your disk space requirements. BRIDGE/FASTLOAD does not need expansion space for logging, so disk utilization is optimized and your overall cost is reduced.

BRIDGE/FASTLOAD can build the DB2/2 or DB2/6000 DAT file on the mainframe from data stored in DB2. Or, on the OS/2 or AIX LAN, BRIDGE/FASTLOAD will accept data from any delimited or non-delimited text file, or SQL statement.

BRIDGE/FASTLOAD is available in a mainframe version, or in stand-alone LAN versions.

BRIDGE/FASTLOAD MAINFRAME VERSION

BRIDGE/FASTLOAD for the mainframe consists of one component residing on the mainframe, and a second component on your OS/2 or AIX server. This version uses the power of the mainframe to retrieve data directly from DB2 for MVS and build the DB2/2 or DB2/6000 DAT file on the host in preparation for placement into DB2/2 or DB2/6000 on the LAN.

Once the DB2/2 or DB2/6000 DAT files are prepared, you can transmit them to a single copy of the database, or to multiple copies of the database at different locations. When the data arrives at the server, BRIDGE/FASTLOAD places the data directly into the database.

To transfer the data from the mainframe to the LAN, you can select the communications method that is best suited to your specific installation. Use BRIDGE/FASTLOAD to send the data, or use other tools that you already have in place (NETVIEW/DM, DDCS/2, XCOM, tape, etc.).

Whichever method you choose, as soon as the data is available at the server, BRIDGE/FASTLOAD takes the DB2/2 or DB2/6000 DAT file and places it directly into the database.

BRIDGE/FASTLOAD OS/2 AND AIX LAN VERSIONS

BRIDGE/FASTLOAD is available in a server edition and a single-user edition. Both editions employ technology similar to the mainframe version, but they cannot connect to the host.

The server and single-user editions reside exclusively on the OS/2 or AIX LAN and will accept data from any delimited or non-delimited file or SQL statement. Just like the mainframe version, BRIDGE/FASTLOAD builds the DB2/2 or DB2/6000 DAT file on the LAN and places the data into the database without invoking any logging procedures.

TIME SAVINGS

For the mainframe version, speeds of 10 Megabytes per minute or more are not uncommon (depending on CPU speed) for BRIDGE/FASTLOAD to extract and prepare DB2 data for transmission to the LAN. Actual transmission time will of course depend on the method used. Once the data is on the LAN, the amount of time it takes to place the data into the database on the server is negligible because the DAT file is placed directly into the database.

In a typical example, a load of 500 megabytes on a 486-50 machine may take over 10 hours using standard loading procedures. BRIDGE/FASTLOAD will accomplish the same task in less than 1 hour!

Load times are also reduced with the OS/2 and AIX LAN editions. On the LAN, BRIDGE/FASTLOAD builds the DB2/2 DAT file at speeds of 3 Megabytes per minute or better. Tests of 600 Megabyte loads have shown a reduction of time up to 8:1 when compared to other available methods.

DISK SAVINGS

The amount of disk space required to accommodate data downloads is often underestimated or completely overlooked. To handle the logging procedures using conventional methods, you need available expansion space of up to 4 times the amount of data to be downloaded. So for a 250 Megabyte download, you need 1 Gigabyte drive available on the server - at significant cost.

The mainframe version of BRIDGE/FASTLOAD eliminates the need for expansion space beyond the size of the data file itself. For a 250 Megabyte download, all you need on the server is 250 Megabytes of space - a dramatic savings compared to conventional methods.

The LAN editions of BRIDGE/FASTLOAD will reduce your space requirements by 33% or more. Beyond the data file itself, some additional space is required on the LAN to process the DAT file because the mainframe is not involved. However, the elimination of logging still results in significant overall disk space savings.

BRIDGE/FASTLOAD is designed for your applications that move sizeable amounts of read-only data into IBM's DATABASE 2 OS/2 (DB2/2) or DATABASE 2 AIX/6000 (DB2/6000). BRIDGE/FASTLOAD enables the Data Delivery Component of IBM's Information Warehouse Framework by providing a high speed method of migrating data to the server platform. Mainframe query workloads can be downsized to the server. Corporate data can be accessed and processed locally to support the business function that needs it. The speed, reliability, and low resource utilization of BRIDGE/FASTLOAD make it the premier data loading tool for DB2/2 and DB2/6000.

PRICES: The mainframe version of BRIDGE/FASTLOAD ranges from \$15,000 - \$25,000 (US) depending on the size of the CPU. the LAN version of BRIDGE/FASTLOAD for DB2/2 is available in a server edition for \$999 (US), and a single-user edition for \$99 (US). The price for BRIDGE/FASTLOAD for DB2/6000 have not yet been announced.

For additional information, please contact:

Bridge Technology
 419 Boylston St.
 Boston, MA 02116
 Phone: 617-424-6266
 FAX: 617-424-6621

Calendar of Events - Where you can catch us!

The following table provides a list of upcoming shows and conferences through to November 1993.

Table 7. Calendar of Events		
Show	Dates	Location
SHARE	August 15 - August 21	Washington
OS/2 Technical Interchange	August 30 - September 3	Orlando
Downsizing EXPO	September 13 - September 15	Toronto
UNIX EXPO	September 21 - September 23	New York
DB2 Technical Conference	September 26 - October 1	Anaheim
COMMON	September 27 - October 1	New Orleans
OS/2 Professional Interchange	October 17 - October 20	Palm Desert, California
European IDUG	October 17 - October 21	Copenhagen
PC EXPO	October 19 - October 22	Chicago
COMDEX	November 15 - November 20	Las Vegas
GUIDE	November 15 - November 20	New Orleans

Submitting Suggestions to IBM

Do you have any suggestions to enhance our products?

If so, you have two vehicles to communicate your requirements with IBM.

First of all, if your organization has an IBM Representative, discuss your requirements with your IBM Representative, and have them submit your requirements to the **Database Technology Planning** group at **TOROLAB2(DBMREQ)**.

Secondly, send your detailed requirements to IBM by using the Reader's Reply Form (at the back of this Newsletter). If you need more room, include additional paper as needed.

Subscribing to this Newsletter

If you are interested in subscribing to the DB2 Family Newsletter, indicate this on the Reader's Reply Form and mail/fax it back to us.

Requesting Back Issues

To date, the following issues have been released:

1. January 1993 Edition
2. April 1993 Edition
3. July 1993 Edition

If you are interested in receiving back issues, indicate your request on the Reader's Reply Form and mail/fax it back to us.

Publication Frequency

This newsletter will be published on a quarterly basis. The next issue will become available September 1993.

This newsletter was produced by the Database Technology Planning department at the IBM PRGS Toronto Lab. For further information on any of the products mentioned, please contact your local IBM office, or an authorized IBM Business Partner.

Numerous product references in this publication are registered trademarks or trademarks of International Business Machines Corporation. IBM Canada Ltd., a related company, is a registered user.

Sample code is provided for information purposes only, and is used by readers at their risk. IBM makes an effort to provide accurate and safe code examples, but does not warrant their correctness. user.

This newsletter was created and marked for processing using IBM BookMaster (Program Number 5688-015) and IBM Document Composition Facility (DCF)TM (Program Number 5748-XX9).

The final copy was printed on an IBM 3825 Page Printer, an Advanced Function Printer.

This newsletter is © IBM Corporation 1993.

DB2 Family Newsletter No. 3, July 1993 Edition

Reader's Reply Form

1. *Did you find this newsletter useful?*
2. *Is there any way you think we could improve this newsletter?*
3. *Is there any database-related subject you would like to see addressed in this newsletter?*
4. *Would you like to receive the DB2/2 Solutions Directory? Y N*

Please note:

- *If you need additional writing space, write your comments on additional paper and include it with your reply.*
- *IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you, and all such information will be considered non-confidential.*
- *Do not use this form to report database problems or to request copies of publications. Instead, contact your IBM representative or an authorized IBM Business Partner.*
- *If you wish, you may include your name, address, and company name if applicable, and your phone number.*

Date: _____

Your Name: _____

Company Name: _____

Address: _____

Telephone: _____ *FAX:* _____

Thank you for your cooperation and help. You can either mail this form to us, or hand it into an IBM office for forwarding.

You can also FAX the form to us. Our FAX is 416-448-2114. Please mark your FAX to the attention of L.G. Pedron.

DB2 Family Newsletter No. 3, July 1993 Edition

Reader's Reply Form

Fold here and tape.....fold here and tape.....fold here and tape.....fold here and tape.....fold here and tape.....fold here and tape.....

IBM

*IBM Canada Ltd
Database Technology Planning
31/110/895/TOR
895 Don Mills Road
North York, Ontario
Canada M3C 1W3*

ATTN: L.G. PEDRON

Fold here and tape.....fold here and tape.....fold here and tape.....fold here and tape.....fold here and tape.....fold here and tape.....