WHY SOFTWARE IS IBM'S MOST IMPORTANT BUSINESS

Its development tools, Repository and SQL will play a major role for next 25 years

By Dr. George Schussel

With 1988 software and services sales of approximately $20 billion, IBM is by far the largest software company in the world.

IBM's dominance in software is actually greater than in hardware. Large, competing independent software vendors like Cullinet Software, Inc., Westwood, Mass., and Oracle Corp., Belmont, Calif., have sales of a few hundred million dollars per annum. Computer Associates International, Inc., Garden City, N.Y., the largest independent, has an annual revenue of about $1 billion.

IBM needs to do well in software because software sales drive the sales of hardware, and the profit margins and growth prospects for software are greater than the hardware business. Based on current projections, it seems likely that IBM's attention to its software businesses is only likely to increase (see Table 1).

Systems Application Architecture (SAA) plays the central role in IBM's 1990s software strategies. The SAA set of common user interfaces, languages and communications will provide IBM's customers with consistency, portability, migration and connectivity over IBM's three principal hardware platforms (PS/2, AS/400 and 370).

Table 1

IBM's Software and Services Revenue

Year	Percent of Total Revenue
1985	22%
1986	25%
1987	27%
1988	30%
1989	34%*
1990	38%*
1991	42%*
1992	46%*

*estimated

SOURCE: YANKEE GROUP AND DIGITAL CONSULTING INC.

When delivered, SAA will provide a scalable software architecture and the flexibility/portability benefits that IBM first introduced to the industry with the 360 hardware architecture. IBM needs something like SAA to compete in the '90s with the DEC VAX/VMS architecture and the emerging clout of Unix.

To the extent that IBM's SAA is able to increase the number and diversity of computers that run compatible software, then SAA will make the IBM environment even more attractive to independent software vendors than it has been in the past. Because software availability has become more important than hardware price-performance for many users, the result of this is more hardware sales for IBM.

Another important goal of SAA is to simplify the migration of PC users upward into the world of IBM mainframes. The decade of the 1980s has witnessed the education of at least 25 million personal computer users. Because SAA's user interface will come from the PS/2's Presentation Manager, it is clear that this strategy has a good chance of succeeding.

CAUTION SURROUNDS SAA

SAA's goals will be difficult to achieve, will probably be only partially reached, and will take years to be delivered. It is unlikely that a reasonably complete complement of SAA will be available much before the turn of the century. IBM staffers have made statements to the effect that they expect only 20% or so of their own applications to be compliant with SAA by 1993-1994.

SAA's definition is not likely to remain static over time. I expect continuing enhancements to its definition. Missing pieces will be announced this year or in 1990: Repository; security across networks; distributed DBMS; and common screen access.

Commitment by users to the SAA

Schussel is president and principal consultant of Digital Consulting Inc., Andover, Mass. He is a well-known lecturer and futurist. Conferences that he has created and chairs include: Database '89, Software Futures, IBM Software Strategies, Application Development in the '90s and Unix Futures.

protocols means a willingness to accept a significant performance penalty risk. IBM will deliver SAA as another software layer on top of its existing products. The underlying architectures of the PS/2 and 370 will remain fundamentally incompatible. The software layer that is required to translate the source code into compatible object code is likely to consume a lot of cycles.

SAA may redefine the software business. Until now, this market has offered rather well-defined niches—applications, micro/PC software, mainframe systems software and minicomputer software.

Companies like Microsoft Corp., Redmond, Wash., and Lotus Development Corp., Cambridge, Mass., have competed on PCs. Relational Technology Inc., Alameda, Calif., Oracle and Cognos, Inc., Peabody, Mass., have sold minicomputer software, while Software AG, Reston, Va., Cullinet and CAI have largely targeted mainframes.

SAA, by offering portability, will completely mix up this neat order. Lotus is likely to compete with Cullinet, and all of them will both complement and compete with IBM's SAA world.

At this time SAA does not incorporate an architecture. SAA is a set of proposed protocols and standards but it is not really an architecture—architecture is defined as a grand design for computing, which defines various functions and how the pieces fit together.

For example, Systems Networking Architecture (SNA) is an architecture. There is an architecture (unannounced) in SAA; it is being developed by the Data Systems Architecture staff at IBM's Santa Teresa, Calif., laboratory. That architecture, when it is announced, is likely to be called the IBM Application Development Environment (see Table 2). It will be a repository-based architecture, offering a sophisticated software development environment.

The structure of this architecture is shown in Table 2. It will offer a global conceptual view to the data administrator, a logical view to the analyst and a relational view to the DBMS. The conceptual and logical views will utilize entity/relationship models.

IBM's goal is to successfully define the standard environment and repository format, and thus establish IBM's products as defining these functions. A myriad of independent vendors can then offer their toolsets and application building languages for the IBM ADE.

IBM has already succeeded in establishing SQL as a standard for database access. It is also likely that the IBM Repository will define a standard in this area.

IBM researchers talk about an initial delivery of the ADE products for 370 and PS/2 for late 1990. That may sound like a long time, but IBM's recent record for meeting its commitments for on-time delivery of new software has been much better than average.

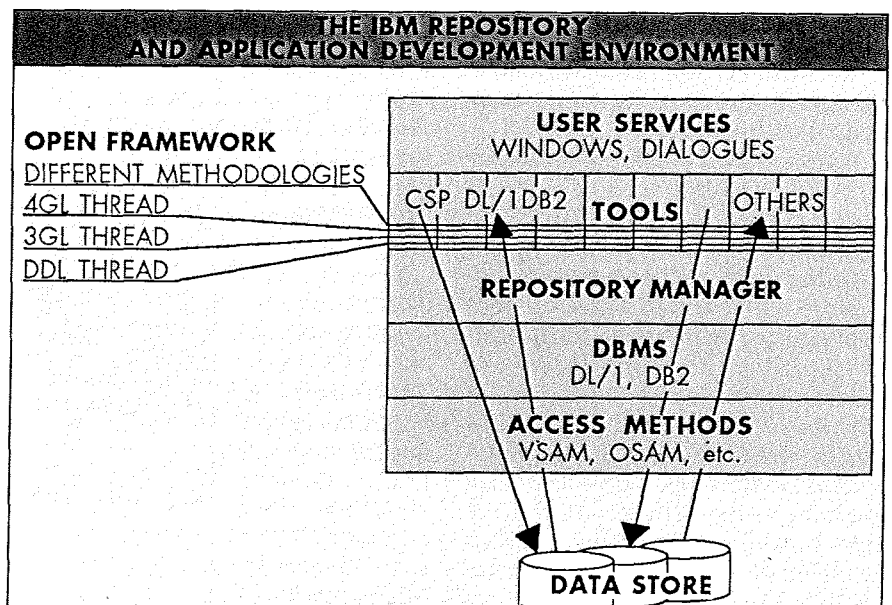

Table 2

SOURCE: DIGITAL CONSULTING, INC.

While IBM wants a strong independent software product community, it would prefer that that community not offer alternative DBMS products. IBM would like to be your only supplier for DBMS software.

This goal is already a fact of life on the AS/400, which has a hardware price that includes the OS/400 operating system and its internal relational database management system. Although, on the PS/2 platform the data manager (OS/2 Extended Edition) is separately priced.

In the 370 line, IBM has been careful not to talk about a bundling of DB2. However, it is likely that as the next generation Summit hardware series is delivered, and IBM's architecture develops over the 1990s, DB2 will evolve into a common subsystem with MVS, ultimately to be installed by most customers using large IBM mainframes.

This does not necessarily mean that the market for alternative database managers on large mainframes will go away. In fact, over the 1990s the DBMS choice for most companies will become more tactical than strategic.

Emergence of SQL as a standard database access language for all DBMS vendors will allow more portability of applications over different DBMSs. Most large shops are likely to have several DBMS products in the 1990s, with DB2 being one of them.

In the PC arena, competition from alternative suppliers for LAN and DBMS managers will probably deny IBM a position of hegemony in systems software. Companies like Sybase Inc., Emeryville, Calif.; Gupta Technologies, Menlo Park, Calif.; Novell, Orem, Utah; Microsoft; Lotus, and others have the market credibility and technology to give IBM a tough run.

A principal SAA goal is to allow

mainframe and minicomputers to become network servers in a world where small machines run most of the mips. IBM's future architectures stress the use of intelligent workstations rather than 327X-type terminals. In this environment a "system" becomes a network where all processors have common access to repository-managed databases.

Another goal for SAA is to provide the connectivity and database consistency across IBM's diverse platforms. IBM knows that in the 1990s the mainframe will become a repository and network server to a large variety of mid-sized and small machines where most processing will occur. Cooperative processing and distributed database will be the new technologies that allow this.

In the past, IBM's support for connectivity across different IBM architectures has been poor. In the future, IBM plans to offer distributed database software solutions that span the SAA world.

A distributed database manager offers the advantages of a single logical view of data with the physical implementations spread over a number of heterogeneous computers. Changes in the physical reorganization, or breaks in the network, will be hidden from the application developer. The future of database management lies in the field of distributed databases.

IBM has developed a multiphased plan for bringing distributed database capability to its customers. Four principal development laboratories are participating in the evolution of this plan.

Laboratory	Product
Toronto	SQL/DS
Santa Teresa	DB2
Rochester	AS/400 SQL
Austin	OS/2EE

Table 3 is a graphical representation of IBM's distributed database evolution strategy. The three critical phases of IBM's plan are:

Phase II	Remote Unit of Work
Phase III	Distributed Unit of Work
Phase IV	Distributed Request

In Phase II an application may send discrete committable units of work to different remote databases. However, each committable unit must go to only one physically remote database.

This requirement is loosened somewhat in Phase III, where each committable unit of work may consist of a number of discrete SQL statements, each of which is constrained to a single physical site.

In Phase IV, Distributed Request, the restraints of physical locations are removed and individual SQL statements may execute over data that is located at diverse sites. It is in Phase IV where a true multisite Join and support for replicates are available. IBM's plans call for delivery of Phase IV capabilities in 1994.

To date, IBM's distributed DBMS products are for "like" environments. For example, this means that the DB2 Remote Unit of Work capability that is to be delivered in late 1989 will work amongst diverse DB2 partners exclusively.

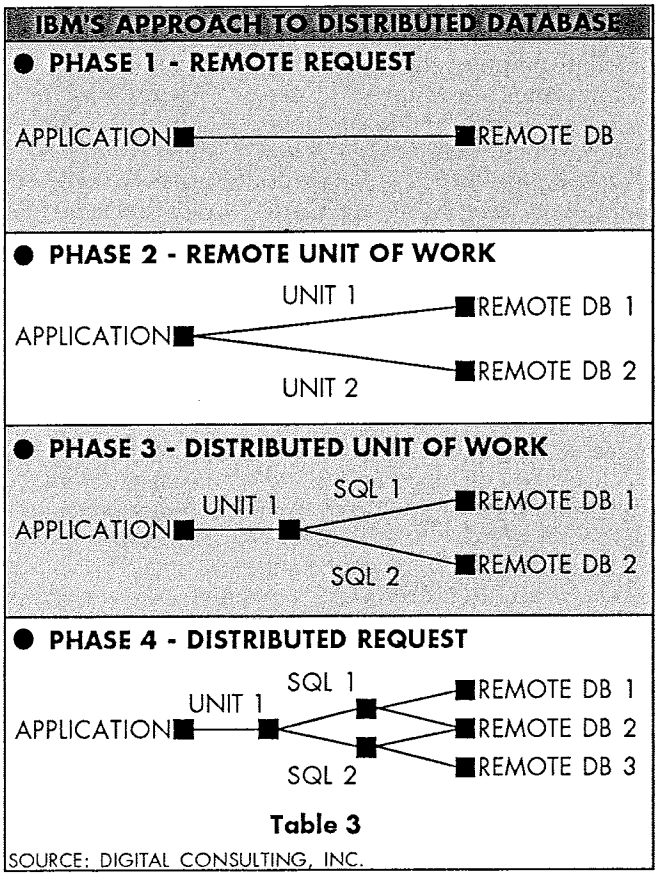

Table 3

SOURCE: DIGITAL CONSULTING, INC.

In comparable fashion, initial releases of SQL/DS and OS/2EE will only distribute over their own equal partners initially. The difficulty in this arena is caused because IBM is pursuing a distributed database strategy with different relational database engines. Distributions over "unlike" partners is expected to be delivered in 1992.

Although IBM developers share research and product development plans amongst different IBM groups, they do not share source code for the DBMS engines. The result of this policy has been that even for the two different SQL DBMS engines that run on the 370 (SQL/DS, DB2) there are significant differences (different return error codes and different handling of nulls).

So, why doesn't IBM take one of its SQL engines, such as DB2, and port it to the diverse operating systems of SAA (like competing products Oracle, Ingres and Datacom are doing) in order to make its job easier?

Santa Teresa staffers argue that differences amongst the SAA operating systems mean each must have its own physical implementation of SQL in order to operate efficiently. This certainly is true for implementation-specific (physical) functions, such as cross-memory services, memory management and I/O management.

It has been estimated, however, that no more than 25% of DB2's code is required to handle these functions. This leaves the large majority of DB2's source code to be taken up with software performing logical SQL functions, such as query management, catalog interaction,

data administration, security, authorization and integrity. The sharing of the code for these functions across the different IBM DBMSs would appear to make the implementation of DBMS easier.

Another reason for the separate product/separate operating system policy is IBM's management and accounting policies.

IBM products have to stand on their own for profitability analysis. That means if two groups are building a DBMS for two different environments and one sends its query management source code to the other, then there has to be a cross-subsidy agreement between the two.

IBM does not seem to want to tackle this problem.

For whatever reasons, IBM is committed to an SAA-SQL solution that involves four different products, while almost all of IBM's competitors are taking a single DBMS product and distributing it over diverse operating systems. Some of IBM's competitors in the distributed DBMS market will also offer support for non-SAA operating systems, like DEC's VMS and Unix.

The ultimate success of IBM's distributed database strategy is tough to forecast. A number of mainframe competitors, such as Computer Associates, Oracle and Relational Technology, appear to be ready to offer advanced distributed database capabilities to their customers years before IBM's products are delivered.

On PC platforms, companies like Lotus, Microsoft and Ashton-Tate, Torrance, Calif., have the market presence to be able to compete on an even footing with IBM.

PERSONAL SYSTEM ENVIRONMENT

In order to succeed in cooperative processing, IBM has to have a robust, successful operating system environment on its PS/2 product line.

While the OS/2 base operating system and Presentation Manager are sourced from Microsoft, and are openly available to IBM's competitors, the database and data communications extensions of OS/2 are IBM proprietary developments. These products are based upon underlying SQL and LU6.2/APPC product concepts.

Initial reviews are in on the first release of OS/2 Extended Edition and they are not positive. A number of complaints have centered on temporary problems that could be expected to be associated with new software.

For example, initial deliveries did not support the Presentation Manager interface. Also, because APPC is needed for data communications there has been no support for client DOS machines.

On the other hand, early users of OS/2EE report some problems which could affect its marketability for a long time:

1. OS/2EE is complex. It requires a systems programmer for installation and maintenance. It needs lots of support, and the training for users is extensive.

2. There is no support for third-party products, such as Irma boards.

3. The minimum useful configuration requires 6Mb of RAM.

4. For communications, a full physical unit 2.0 must be defined in each workstation. This is a problem because it demands much more mainframe resources than a dumb terminal environment in which a single 3274 cluster controller (defined to the mainframe as a single physical unit 2.0) can handle a large number of terminals.

5. For adequate performance, OS/2EE and Presentation Manager require a 80386-based computer. This means that most of the IBM PS/2s sold to date (models 30, 50 & 60) will not be satisfactory for full-blown OS/2.

In spite of the competitive battle that will be waged for market share in database and data communication products for PS/2, it is likely that OS/2EE will evolve as a standard in those markets where IBM mainframe computers are established. This includes Fortune 500 companies in the banking, insurance and financial services industries.

In smaller companies, and where IBM mainframes are less strongly entrenched, OS/2EE is likely to receive a lukewarm reception.

OS/2's acceptance will begin to improve once the market understands that OS/2 is a replacement for minicomputers, not for MS-DOS. As a platform for personal productivy, DOS will survive well into the '90s. Products like OS/2EE will take market share from midrange computers.

IBM's future is most likely to be governed by acceptance of its software products more than ever before. SAA and its ADE will be the major determinants of IBM's fortunes in the mainframe software market.

SAA is likely to be a major success, although it will take until the mid-1990s to have a significant impact on IBM users. But it will be a major factor in marketing decisions well before then.

MAJOR ROLE FOR NEXT 25 YEARS

The Application Development Environment, Repository Manager and SQL will play major roles in applications development for at least the next 25 years. ADE and the Repository Manager are likely to be used as the basis for not only IBM's development customers, but as the basis for add-on value by many application development tools vendors.

It seems that IBM's ultimate market penetration with its distributed database and PC software products will be modest. In the PC arena, IBM is constrained by depending on outside suppliers, like NeXT, Inc., Palo Alto, Calif., and Microsoft, for much of its technology. These and other competitors are able to directly deliver products into the market.

In distributed database, the technical hurdles that IBM has in choosing to implement over different DBMS products means that its products will be late to market compared to competition. Still, IBM's biggest customers are likely to wait for its products. Others, however, will choose from a wide assortment of alternatives. ■