DATA MANAGEMENT:

IBM Strategies and Solutions

Version 5

White Paper

Data
Management
Solutions

IBM®

Fifth Edition (May 1996)

Abstract

This overview of the data management market describes IBM's participation in this segment of the Information Technology industry. The paper provides information about current trends and describes how IBM is addressing the market conditions with strategies and technologies to provide solutions to our customers.

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of the intellectual property rights of IBM may be used instead of the IBM product, program, or service. The evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the responsibility of the user.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Corporation
IBM Director of Licensing
208 Harbor Drive
Stamford, Connecticut 06904
U.S.A.

Trademarks

Terms denoted by an asterisk (*) in this publication, are trademarks or registered trademarks of the IBM Corporation in the United States or other countries or both:

Terms denoted by a double asterisk (**) in this publication, are product or company names that are trademarks or registered trademarks of other companies:

Contents

Introduction

This overview of the data management market describes IBM′s* participation in this segment of the Information Technology industry. The paper provides information about current trends and describes how IBM is addressing the market conditions with strategies and technologies to provide solutions to our customers.

Executive Summary

In today′s global, highly-competitive economy, turning organizational data into an informational asset is key to gaining competitive advantage. To create and manage this informational asset, customers need to store, retrieve, access, analyze, and distribute their data. The core componentry IBM offers customers to do this includes database engines together with middleware for distributed databases, Internet enablement tools, access enabling tools for data replication and decision support, data analysis tools for data warehousing and data mining, tools for systems management and administration, and tools for application development. These data management offerings integrate with complementary components for transaction management, networking, and workgroup solutions. In addition, these products combine with appropriate service, support, and application offerings to provide complete and open data management solutions to customer business problems.

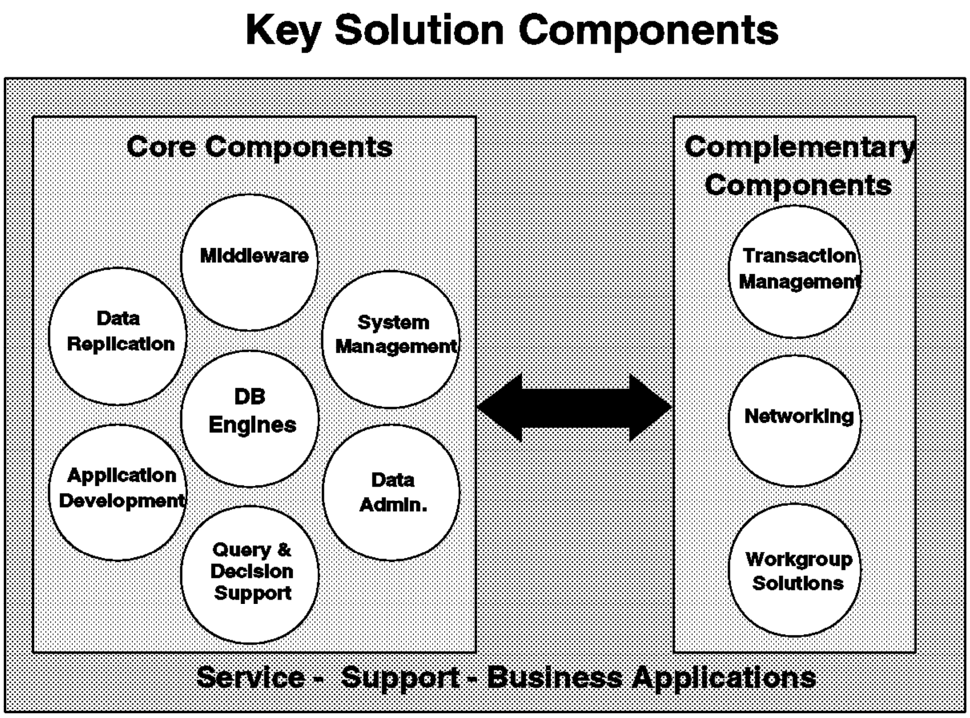

To satisfy customer needs for uniform and open software across diverse environments, IBM develops these solutions for a variety of IBM and non-IBM platforms, including PCs, workstations, LANs, mid-range, and large systems. IBM offerings are designed to run on uni-processor and parallel systems, in simple and complex environments, and are focused on making distribution of data and client/server solutions easy to use while maximizing customer value.

The Importance of Data Management to Customers

A Computerworld article [1], Grow with the (information) flow, provides a good description of today′s information explosion:

″Information once flowed through fairly predictable channels: databases, department reports and accountants′ spreadsheets. Such structured data could be easily indexed, retrieved or queried.

Now a variety of alternate data pipelines - including electronic mail, mobile computers, groupware, voice mail and on-line informational services - are making information systems′ managers jobs more difficult by delivering torrents of unstructured data to computers.

And some observers predict the surging data stream will only widen: Heavier traffic from portable office devices and high-bandwidth channels such as the Internet, Integrated Services Digital Network and multimedia are transmitting more real-time communications and video.″

The ability to digest this exponentially growing stream of data and the ability to turn it into an informational asset is strategically important to most organizations. This asset offers a competitive advantage in the delivery of quality products and services.

Two business examples illustrate the value of data management[2]

- Timely analysis of information is key to success in the women′s fashion industry. Fashions are seasonal (3 months), the goods have a short shelf life, a high turn-over, and are very profitable. It makes a big difference to catch onto trends quickly. So if you can predict exactly what clothing people will buy during a certain week at a certain store, then you can reduce unwanted store inventory and lower capital outlay for the same profit. Analyzing store sales data by product on a daily basis can make a 2% to 3% difference in margin -- and in a business where the margin is typically 4%, that is a significant business factor.

  Technologies that facilitate the above scenario are immediate access to point of sale information, database mining for analysis to detect trends immediately, high performance, multi-vendor database connectivity, and cooperation among heterogeneous clients and servers.

- Success in certain heavy machinery industries depends on mechanical CAD design automation done with world-wide teams wherever the experts with certain skills reside. So world-wide groupware is needed, with real-time

1 Jean S. Bozman, Computerworld, May 3, 1993, 105

2 Patricia G. Selinger, ″Database Management in the Year 2000: Projections, Predictions and Star Gazing″, Proceedings of the 18th International Conference on Very Large Databases, August 1992

changes of parts -- not three month integration intervals. Concurrent re-engineering with one-day turn around for updating parts is what′s required.

Technologies that facilitate this scenario include real-time remote data access, high bandwidth interactions to remote databases, lots of processing power to do the engineering analysis, and support of a variety of cooperating heterogeneous databases.

The IBM Database Heritage

IBM has a long and distinguished record as a pioneer in database technology and as an industry leader in delivering ″industrial-strength″ data management solutions to customers.

- IBM′s highly successful IMS/ESA* Database Manager was first introduced in 1968. IMS DB today supports mission-critical applications around the world.
- Relational database systems, along with the SQL language, are an outgrowth of research done in IBM′s research division. The success of this technology has been phenomenal; in 1993, for the first time, worldwide revenues (all vendors) from relational products were higher than revenues from non-relational products. Today, IBM′s relational DB2* family of products is offered on a diverse set of IBM and non-IBM platforms from PCs to super computers.

Current Trends in the Data Management Market

Like the Information Technology industry as a whole, the data management market is experiencing rapid change. A once stable, mostly host-centric market, has given way to a volatile, dynamic, highly competitive market. Essential to success in the 1990s is the need to provide complete solutions addressing key customer interests and technology trends for open client/server environments, system platforms, data management, communications and access, data analysis, as well as tools, applications, and services. Following are some key trends in more detail:

Open Client/Server Environments

A global economy is fostering the need to enable new ways of doing business. For example, a large retailer engages in inter-enterprise data access. Their inventory and point-of-sale data are open to suppliers in order to optimize deliveries.

- A fundamental trend is towards networked client/server environments. Availability of adequate price/performance, capacity, and communication allows computing to be done much closer to the end user, even on the end-user′s desktop system. When the management or organization of a businesses is decentralized, data is often kept on distributed servers. In many cases, decision makers use personal-data-server products on their desktops. They also connect to LAN servers for locally relevant data, as well as work with enterprise data servers or even outside information service providers for data affecting large groups of users.
- Customers want freedom to choose products which best fit these needs and to be able to connect them using open system interfaces. Most organizations use hardware and software from a variety of vendors. End users need to communicate between unlike environments and they need their applications and data to work together regardless of supplier.

System Platforms

Customers are looking for an optimal mix of heterogeneous platforms, operating systems, and technologies as they try to reduce costs and improve efficiencies of business processes.

- For many organizations, the computing environment contains a wide variety of clients, including platforms such as: DOS, Windows**, multiple UNIX** variants, various terminals, OS/2*, and Macintosh**.
- Investments in UNIX are still growing; customers want software to operate on a variety of UNIX platforms. Also popular are NetWare** NLM database products as well as database products for the Windows NT** platform and Intel**-based symmetrical multi-processors (SMP) systems.
- The mainframe and high-end mid-range systems are evolving to high capacity, highly available WAN superservers and repositories of enterprise data, as applications and informational data are distributed closer to end users. These systems often connect to LANs and support massive storage, large scale computing, and high bandwidth communications.

- Parallel technology in both hardware and software is attracting strong customer interest. Symmetrical multi-processors (SMP) are popular today; massively parallel systems (MPP) are gaining acceptance in commercial environments. This trend will allow customers to exploit their data in ways that were not possible or affordable before. For example, retail customers can use database mining to derive information about buying patterns in order to better manage store inventories.

Data Management

The complexity of managing data increases customer demand for cost-effective scalable solutions that provide good performance and high availability. In addition, customers want the option to decentralize data using multiple vendor products adhering to international standards.

- There will be a richer set of data types supported by the core database engines along with compatibility for user-defined data-type extensions. Data-type extensions will include data structures and methods to operate on those new data types.
- Relational databases are expected to have richer support for structured data as SQL is enhanced with object-oriented extensions. Object-oriented databases will see more set-oriented query extensions to complement support for user-defined data structures.
- To support richer function and data types, the optimization of database requests will become more sophisticated and be a performance differentiator.
- Databases will employ parallel processing techniques in addition to supporting parallel functions in operating systems and hardware.
- Support for multi-dimensional analysis (MDA) or On-line Analytical Processing (OLAP) is evidenced by extensions to relational engines as well as specialized multi-dimensional database engines.
- Typical SQL support today includes SQL92 entry level as well as extensions like stored procedures, triggers, and event alerters.
- X/Open** interfaces become the standard for interoperability.
- Demands for round-the-clock availability are addressed through more granular backup and recovery features along with software-based data mirroring and remote/central management.
- Trends in the distributed database area include data replication, two-phase commits, distributed queries, heterogeneous data access and join capability.

Data Analysis

The ability to store and process large volumes of historic data and the need to use this data to gain competitive advantage is leading to investments in data warehousing and data mining.

- Tools are proliferating to populate and manage complex enterprise data warehouses as well as subject matter data marts.
- Additionally, new and enhanced tools are available for generalized queries, multidimensional analysis (also know as OLAP), and data mining.

- Tools for data mining encompass ″verification″ tools and newly emerging ″discovery″ tools. Verification tools help validate hypotheses while discovery tools help detect valuable knowledge that might otherwise go undetected.

Communications and Access

Technology advances make it cost effective for organizations to enhance the communication of information to a broader base of users.

- Larger communication bandwidths make it feasible to transmit a broader set of data types and larger volumes of information.
- The implementation of a telecommunications superhighway, linking worldwide information networks, broadens the base of users who want interactive access to remote information.
- Remote devices enable more users to access information. A proliferation of hand-held devices for users is evidence of the exploitation of remote capabilities.
- Access to information by mobile field employees is growing rapidly due to the use of wireless radio frequencies.
- Interactive television will increase the base of users accessing product and service information.
- Infrastructures are being developed for ″on demand″ use of multimedia information.
- End Users expect to be able to use portable computers, regardless of where they are, and expect applications and data to be available to them whenever they need to use these resources, regardless of where these applications execute or where the data resides.
- Corporations and Organizations increasingly focus on their core competencies; that makes many business processes available for outsourcing. A first step gaining popularity is to outsource the IT operation. The second wave expected is the outsourcing of the execution of the business process. Advances in communication and access technologies allow the physical location of the IT resource that executes the business process to become irrelevant.

Tools and Applications

Customers want the cost reduction and productivity benefits of easy-to-use and integrated tools for systems administration, systems management, and application development while choosing from different tool providers. Optionally, customers want to purchase application packages that store data using popular data management offerings.

- Tool and application builders play an increasing role as organizations of all sizes increase purchases of applications and tools rather than building and maintaining them internally.
- Tools often have graphical interfaces, are often object-based, and integrate with a variety of other tools and underlying functions.
- Development tools for distributed applications are emerging.

- Suites of integrated tools and applications are gaining popularity at all levels - desktops through enterprise systems.

IBM Strategies and Solutions

IBM is investing significantly in the fast-growing and highly competitive LAN and workstation market for data management products. IBM is dedicated to being an open, client/server vendor capable of providing industry-leading client/server solutions to customers. For PC and UNIX servers, DB2 is available on OS/2, Windows NT, AIX*, Sun Solaris**, HP-UX, and Sinix.

IBM data management products such as IMS DB, DB2 for MVS, DB2 for VSE and VM (SQL/DS*), and DB2 for OS/400* enjoy considerable success on mainframe and midrange platforms. To support customer needs, IBM intends to continually enhance these products in performance, availability, and functionality to address customer requirements in decentralized, open, client/server environments.

The following figure illustrates the breadth of IBM's DB2 offerings across various platforms.

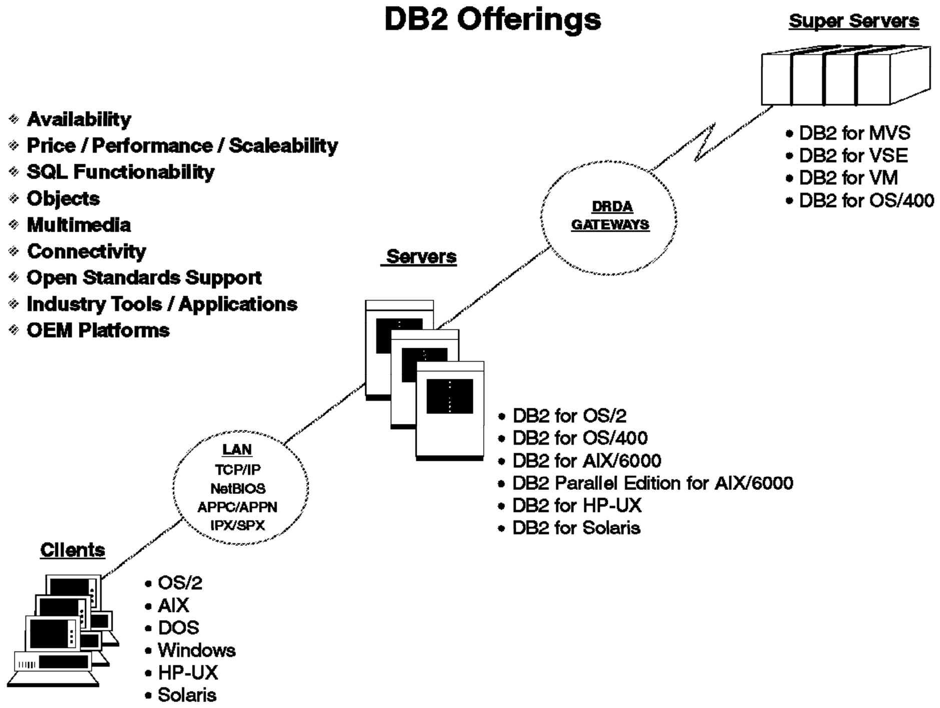

IBM is expanding core database technologies to deal with the explosion of information types such as text, sound, image, and video. We are incorporating new technologies into our data management solutions to improve price/performance and capacity, to enhance ease-of-use and productivity, and to enable new application types.

Our support for open systems, distributed data, and heterogeneous systems, including non-IBM platforms has been greatly expanded. We have increased our investment in tools for middleware, Internet enablement, data administration, systems management, data replication, data warehousing, data mining, data analysis, and design. And, we are continuing to fund research that will maintain our industry-leading position for new and emerging data management technologies as well as accelerating the evolution of our current data management offerings.

Customers want open products that can be mixed and matched regardless of vendor. IBM strategies allow relational database engines to work on IBM and non-IBM hardware platforms, and allow tools to work with key database engines and networks. We work with application builders to enable application execution on the same platforms as our database engines. Also, IBM products are designed to be open and complementary to business partner offerings and to interoperate with other vendors′ data management components. To facilitate this open concept of mixing and matching solution components, IBM technology uses pervasive industry standards; this applies to programming interfaces, user interfaces, communication and networking interfaces, as well as interfaces for systems and network management, security, and auditing. In addition, to support our world-wide customer base, IBM technologies are designed to support multiple languages.

Key components for building industry-leading IBM solutions are discussed in follow-on sections. These components include industrial-strength database engines, middleware to enable access, tools for efficient systems management and data administration, data replication tools to place the data optimally and provide added value, query and decision support offerings to leverage maximum value from the enterprise data assets, and tools for productive application development.

Database Engine Technologies

Customer requirements demand that IBM database engines be state-of-the-art to support mission-critical applications.

IBM Relational Database Technology

IBM Relational Database Technology is designed to incorporate the best optimizing technology in the market, object extensions to combine the power of relational technology with the emerging object and multimedia technologies, as well as industrial strength relational database management functions.

IBM DATABASE 2* (DB2) Databases

IBM′s DB2 family of database products provides a high level of availability, integrity, performance, and security that is demanded by the mission-critical applications businesses rely on today. IBM supplies DB2 on a variety of IBM operating system platforms including, DB2 for MVS/ESA, DB2 for VSE and VM, DB2 for OS/400, DB2 for AIX/6000*, and DB2 for OS/2. IBM also provides DB2 on non-IBM operating system platforms including DB2 for Windows NT, DB2 for HP-UX** ,and DB2 for the Solaris** Operating Environment, and DB2 for Sinix.

The DB2 family of products gives you the ability to grow from a small environment, where database is limited to a single user, to an environment where the database applications may be shared among hundreds of users.

IBM DB2 Relational Extenders

Customers with existing investments in SQL want to be able to manage new complex data types for both existing and new applications. IBM offers a suite of database ″extenders″ (UDFs, UDTs, and triggers written and supported by IBM) for Version 2 of DB2 for OS/2 and DB2 for AIX/6000. These DB2 Relational Database Extenders have the power, flexibility, and extensibility that enable database application developers to access, retrieve, and manipulate text, image, audio, video, and fingerprint information, using SQL.

Relational Extenders encapsulate the attributes, structure, and behavior of new data types and store them in a column of a DB2 table, such that they can be processed through the SQL language as natural additions to the standard set of DB2 data types. Relational Extenders are separate from the database system itself, in the sense that they can be developed, installed, and used independently of full database product releases, but they are part of the database in the sense that they appear to application developers as seamless extensions to SQL and the DB2 product.

IBM Hierarchical Database Technology

IBM Hierarchical Database Manager technology is designed to deliver data integrity, throughput, availability, and proven large-scale applicability.

IBM IMS/ESA Database Manager (IMS DB)

IMS DB, the recognized leader in hierarchical database management, is the product of choice for critical on-line operational applications where support for high availability, high performance and capacity, and low cost per transaction are key factors. Major corporations around the world have a day-to-day dependency on IMS′s capabilities.

IMS DB which can be accessed from IMS/ESA Transaction Manager and CICS*, consists of:

- Full function databases, where the majority of customer data is stored today, allowing complex data relationships and predefined structures to be handled with very efficient I/O operations.
- Fast Path databases, where a growing amount of customer data is stored, allowing high performance, high-volume databases, high data availability, and continuous operability through on-line utilities.

IBM Parallel Database Technology

IBM Parallel Database technology is designed to provide industry-leading price/performance as well as raw throughput and response time. This increases the range of feasible applications that can be deployed to the competitive advantage of our customers.

Products within the DB2 family take advantage of parallelism in various ways using methods specific to the operating systems and hardware with which they run:

- DB2 for OS/2 with OS/2 SMP is designed to provide improved performance in a multi-user environment. Combined with improvements in other areas, DB2/2 for OS/2 will give users in high-contention situations improved transaction throughput.

- DB2 for OS/400 which is integrated with OS/400 takes advantage of an SMP AS/400* system from start-up tasks right through to production environments involving multi-user transaction processing. DB2 for OS/400 uses parallel prefetch processing to shorten the processing time required for long running I/O bound queries.
- DB2 for AIX/6000 is designed to take advantage of I/O parallelism in both uniprocessor and multiprocessor environments. Also DB2 for AIX/6000 is optimized for a shared-memory multiprocessor environment.
- DB2 for VM uses VM Data Spaces Support to perform I/O in parallel.
- DB2 for MVS provides parallel processing for transactions, queries, and for utilities.

IMS DB is architected for parallel processing. Many of the utilities can perform their operations in parallel with on-line applications. Also, IMS DB does parallel I/O processing.

In addition to the parallel database functions described above, there are specialized implementations of IBM database products.

IBM DB2 Parallel Edition for AIX/6000

DB2 PE is designed for users who want an AIX high performance parallel database solution to address requirements beyond the capabilities of DB2 for AIX. DB2 PE is appropriate for informational (decision support) and smaller operational (OLTP) applications. The former allows for implementation of high performance data warehouse and data mining applications. DB2 PE executes on the IBM Scalable POWERparallel* Systems (SP) and provides a flexible, integrated growth path capable of running up to 128 processors (nodes) - all in parallel. Because DB2 PE and SP have the same architecture, this combined solution can offer near linear scalability. This means very little additional overhead when new processors are added to the total solution.

DB2 PE features that aid the implementation of high performance query and decision support applications include:

- All database functions are performed in parallel, e.g., data manipulation, data definition, and utility operations.
- A parallel, cost-based optimizer providing optimal use of parallel resources that results in maximum query speed-up with minimum data administration. All access plans are created for parallel operation.
- Flexible configuration options that exploit the advantages of the SP hardware platform. This allows for incremental growth to increase capacity and accommodate large data volumes.
- Large table support. Maximum table size of 128 nodes is 8 terabytes.
- The ability to rebalance data across nodes when new nodes are added to the complex and analysis and correlation of data distribution by a rebalancing function.

IBM System/390* Parallel Transaction Sysplex with System/390 Parallel
Enterprise Server

This solution evolved from the need for high performance access to transactional data, where detailed business information is processed with a high speed front-end teleprocessing system such as IMS TM or CICS/ESA, with access to database management systems, such as IMS DB or DB2 for MVS/ESA. These OLTP applications have been high performance oriented in the past and are the basis of most company′s operational applications. However, customer demands are challenging even the largest S/390* mainframe processors and their capacities. OLTP workloads have the characteristics of higher transaction volumes, large databases, multiple concurrent users, massive update activity, and big terminal networks - all with the need to maintain rapid response times. To address these emerging very large transactional requirements, IBM has a strategic S/390 platform solution - the S/390 Parallel Sysplex. This system can also be used to handle decision support and batch workloads. IBM services are available to help install and operate a S/390 Parallel Sysplex system.

This multiple processor environment is based in the MVS/ESA operating system and performs like a single system. It is not one product, but a combination of S/390 Parallel Enterprise Servers, along with IBM software, all cooperating to process work in parallel and without interruption.

The S/390 Parallel Sysplex was designed for continuous availability. Through a combination of coupling (hardware) and data sharing (software) technologies, a S/390 Parallel Sysplex can manage up to 32 IMS DB or DB2 for MVS/ESA subsystems. With the themes of high performance and continuous availability, the S/390 Parallel Sysplex is the logical platform for critical OLTP workloads.

IBM POWERquery for the SP2*
POWERquery is an IBM packaged solution offering of highly scalable and flexible ″building blocks″. The components are hardware, software, and IBM services specifically tailored to an installation′s needs. This solution addresses the implementation of large-scale, cost-effective decision support applications. By using POWERquery you can achieve competitive and strategic benefits like successful marketing campaigns, efficient resource allocation, and a way to better analyze your customer′s buying patterns. Industries that can benefit from this solution are: retail, consumer goods manufacturing, banking, insurance, airlines, and telecommunications.

A POWERquery solution includes:

- The two AIX parallel components - the Scalable POWERparallel Systems (SP) and DB2 Parallel Edition for AIX.
- IBM′s most advanced storage offering for RISC-based systems- the IBM 7135 RAIDiant Array.
- IBM services such as planning, implementation, and training assistance.

Optional POWERquery solution components are:

- POWERquery Data Replication Tool: Software tools to populate and maintain decision support databases from operational data sources
- POWERquery Database Tools: Performance monitoring, database administration, and capacity planning tools to manage AIX operations.

DB2 Symmetric Multiprocessing for OS/400

With DB2 Symmetric Multiprocessing for OS/400, a single database operation can run on multiple processors at the same time or, in other words, in parallel. These database operations are typically queries, which are run through Query for OS/400, the DB2 Query Manager, or a PC-based query or report writing tool. Both SQL and native database interfaces are supported.

To achieve this parallelism involving multiple processors, an individual query is split into many smaller subtasks. Each subtask can run independently on a separate processor. Once the subtasks are complete, the results of each subtask are then combined to form the complete query result. Because of the advanced, single-level store architecture of OS/400, these subtasks can efficiently process information on the behalf of a user query request.

This process of splitting queries to run in parallel allows for significant performance increases. These performance increases become more pronounced with the addition of more processors. For example, if a query is running in 20 seconds on an AS/400 with a single dedicated processor, adding a second dedicated processor along with the DB2 Symmetric Multiprocessing for OS/400 feature may allow this query to run in approximately 10 seconds. Adding two additional processors may allow this query to run in approximately five seconds.

Those queries, which include table scans, index scans, joins or group-by operations, will realize the greatest performance benefit from SMP database parallelism.

DB2 Symmetric Multiprocessing for OS/400 can be configured differently for each user of the system. This allows a system administrator to have the greatest control over how parallelism is used on a system and thus greater control over how system resources are used. Part of this enablement process allows the selection of just how much parallelism is used or, in other words, how many subtasks are used for each query. Using fewer subtasks than processors available will allow a greater amount of the total system resources to be used by other users. Using more subtasks than processors available, allows an individual user to use more of the total system resources. This flexibility is needed to allow administrators to balance the needs of all system users with the available resources.

IBM Object-Oriented Database Technology

Object Support and the DB2 family

With Version 2 of DB2 for OS/2, AIX/6000, Windows NT, Solaris, HP-UX, and Sinix , IBM is introducing the first stage of its integration of object support into the DB2 family. This support brings many of the key benefits of object technology to the relational DBMS world, enabling customers to build more flexible business applications and to capture more of the semantics they associate with their data. This is possible because IBM supports new data types (for handling multimedia data as well as user-defined data types), programmable SQL functions, new integrity mechanisms, and object-oriented programming languages. By incorporating object extensions into its DB2 family, IBM is offering customers a way of capitalizing on many of the benefits of object technology without sacrificing the investment they've made in relational technology.

IMS Client Server Object Manager (IMS CSobject)

The IMS Client Server Object Manger, packaged with the IMS Transaction Manager (IMS TM), provides IMS support for object distribution. IMS CSobject manages client access to host data through a local server, using DB2 for OS/2 as a cache, and works with IMS TM applications to manage and protect host data. Multiple client applications can access data locally, on the server, without repeated host requests. IMS CSobject supports a network of multiple client and server workstations, running OS/2, and supports TCP/IP, NetBIOS, and Netware IPX/SPX for communication between client and server. Users can define data objects that map to IMS DB and DB2 for MVS data structures. IMS CSobject uses the Object Management Group (OMG) Persistent Object interface to manipulate server data objects and is compatible with other OMG persistent service providers. And, IMS CSobject provides an object-oriented interface for client applications written in languages that support IBM′s System Object Model (SOM)

ObjectStore**

ObjectStore from Object Design, Inc. in partnership with IBM, enables development of new solutions using complex inter-related data . ObjectStore is a distributed, heterogeneous, object-oriented database. It is available for the AIX/6000, OS/2, Windows NT and other popular workstation platforms. ObjectStore supports a flexible client/server architecture executing TCP/IP, NetBios, Anynet/SNA, and other networks. ObjectStore allows users to store, retrieve and manage complex data such as engineering, business, and multimedia objects in their original form as seen by applications which manipulate them.

IBM Multimedia Database Technology

IBM Multimedia Database Technology extends the capability for multimedia to the DB2 family of products. IBM application development software enables ad hoc multimedia query and the building of multimedia operational applications. These capabilities for the exploitation of multimedia in the business environment include:

- Extensions to Relational Database Management: Enabling seamless integration of multimedia data with traditional data, preserving existing investment while providing for the new applications customers need. (See the previous topic on ″IBM DB2 Relational Extenders″ and also the topic on ″Object Support and the DB2 Family″).
- New Search Engines: Continuing IBM leadership in search technology with image content search and contextual text search capabilities. (See the next section on ″IBM Ultimedia Manager*″ for more information on image content search).
- Rapid Object-Oriented Application Development: Delivering a new object-oriented visual application builder for developing graphical client/server and multimedia database applications, enabling customers to get the benefits of OO programming and multimedia. (See the section ″IBM VisualAge*″).

IBM Ultimedia Manager

Ultimedia Manager helps you to find images based on their content. In addition to storing and manipulating image files you can sort through thousands of pictures, analyze them, and get feedback about their content.

Ultimedia Manager uses patented technology developed in association with IBM′s Almaden Research Center. This state-of-the-art technology has broad application, for example, in retail systems, cataloging, medical imaging, and is available for OEM licensing.

Ultimedia Manager provides tools to create Catalogs of images. Catalogs can be for your own use or for distribution to others. A sample catalog is provided to help you get up to speed quickly. Images may be classified by color, shape, and texture features. You may also create and assign your own text categories to your images.

You can build your own collection of predefined colors, shapes, and textures to use in creating a virtually unlimited number of possible image queries.

Middleware Technologies

Organizations want to connect hardware and software from different vendors to make data accessible across the enterprise.

IBM Middleware Technology for Distributed Databases

Data access middleware technology shields applications and users from the underlying diversity and complexity of distributed environments through application programming interfaces that provide consistent and transparent access to data.

IBM DataJoiner

DataJoiner is designed to enable transparent, consistent, single-command access to heterogeneous multi-vendor data sources. It enables users to join data from disparate databases through a single SQL statement and a single interface, hiding the various database interfaces from the user and the user′s application. The technology is designed to support distributed queries like multi-site joins across all IBM relational databases, key competitive relational databases (such as Oracle** and Sybase**), IMS DB, VSAM file systems, and other data sources through standardized interfaces.

DataJoiner supports a broad range of standards and applications, including ODBC**, X/Open CLI and Client Application Enabler. DataJoiner supports a variety of client operating systems including AIX, OS/2, DOS, Windows, HP-UX, and the Solaris Operating Environment. Host languages for DataJoiner′s ANSI SQL single-database API include popular programming languages such as C, FORTRAN and COBOL. Supporting these interfaces makes it possible to connect multi-vendor application development, database, and query tools to DataJoiner.

DataJoiner incorporates IBM′s most advanced optimization techniques for highly efficient distributed join processing that ″turbo charges″ SQL performance for complex as well as simple SQL queries in a heterogeneous environment.

IBM Distributed Relational Database Architecture* (DRDA*)

IBM′s middleware technology is complemented by the Distributed Relational Database Architecture, a set of formats and protocols that define how client and server systems can interoperate with each other to provide access to remote data. It defines single-site access, multi-site updates with enforced integrity, and the use of stored procedures and DCE security. IBM is committed to open client/server database systems for production-level, enterprise-wide network computing. DRDA is an open architecture in support of this commitment and remains the industry leading common protocol, available and deployed today.

DRDA enabled client systems can access and manipulate data on any DRDA enabled server, whether the clients and servers are provided by the same or different vendors. With DRDA, client systems do not require proprietary software at the server, resulting in more flexible, better performing, less complex, and less expensive solutions. Leading client/server database vendors - IBM, Attachmate**, Informix Software, Inc.**, Oracle Corporation**, Rocket Software, Inc.**, StarWare, Inc.**, Sybase/MicroDecisionware, Inc.**, Wall Data, Inc.**, and XDB Systems, Inc.** - have delivered DRDA client support. FileTek, Inc.** and XDB Systems, Inc.** have delivered DRDA server support.

DRDA clients can be used directly by applications or they can be used together with middleware solutions for added transparency and function. DRDA provides the enabling technology for data access byusers, applications, data replication products, and systems management tools. It is the backbone of IBM integrated data management solutions.

DRDA and the DB2 Family of Products

IBM provides DRDA clients and servers throughout the DB2 family of products. DB2 for MVS, DB2 for VM, DB2 for AIX/6000, DB2 For OS/400, and DB2 for OS/2 can function both as DRDA clients and servers. DB2/VSE is a DRDA server. And, DB2 for HP-UX**, DB2 for the Solaris Operating Environment, DB2 for Sinix, as well as DB2 for Windows NT are DRDA clients. The DRDA client implementation for OS/2, AIX/6000, HP-UX, the Solaris Operating Environment, Sinix, and Windows NT is called Distributed Database Connection Services (DDCS). It allows DOS, Windows, OS/2, and UNIX applications to work with data stored in any DRDA enabled server, regardless of its location.

The DRDA protocol and IBM's implementation of clients and servers throughout the DB2 family of products yields an any-to-any solution that is powerful, flexible, and scalable.

IBM Internet Enablement Technologies for Database Access

Organizations want to develop Internet applications that access information in databases and display that information on World Wide Web pages.

IBM DB2 for the World Wide Web (DB2 WWW)

DB2 World Wide Web (WWW) Connection delivers a Web gateway and application development tool that can make a company's information available to your customers and employees on the Web.

Web applications for DB2 built with DB2 WWW Connection allow any Internet user with an HTML-conforming Web browser to access DB2 data without having to make any changes to the existing data structure. Using standard hypertext markup language (HTML) and structured query language (SQL), you can build applications that query DB2 data with standard SQL statements.

DB2 WWW Connection delivers a flexible way to develop DB2 applications that will look and feel familiar to end users while shielding them from the complexities of Internet traffic and relational database transactions. Application developers can use the templates provided to create query applications tailored to your customer set, database design, and market activity needs. Your customers and staff interact with a DB2 WWW Connection application that has the same look and feel as a typical Web page with fields that allow input and selection and buttons to start transactions. The DB2 WWW Connection run-time engine receives the user input, directs the query to DB2, and returns the results to the user in a Web page report.

In a two-tier client/server environment (local Internet server with one or more Web browser clients), DB2 WWW Connection lets you build an application that can access DB2 for OS/2, AIX, NT, the Solaris Operating Environment, and HP-UX data on an Internet server. In a three-tier environment (local Internet server with one or more Web browser clients, plus access to remote servers), an application can access DB2 data on the Internet server and data on servers connected to the Internet server with Distributed Database Connection Services (DDCS), Client Application Enabler (CAE), or DataJoiner. If the three-tier environment includes DataJoiner, the DB2 WWW Connection application can also access ORACLE, Sybase, Microsoft SQL Server, and other relational and non-relational data sources.

IMS and the Internet

IMS customers are able today to make their IMS applications and data available to users across the Internet. World Wide Web access is provided to IMS applications and data through currently available communications protocols (eg. MQSeries, DCE/RPC, APPC, 3270 emulation), from commonly available Web Browsers and servers (eg. Netscape, MVS ICS). And, IMS users have a broad range of options today for Web application development. These options include ones using, for example, VisualAge Smalltalk or Java, as well as commercial packages providing 3270 emulation.

Database Systems Management Technologies

Customers can′t productively implement applications without tools to manage complex, heterogeneous environments in a cost effective manner. Database management systems are increasingly becoming a critical building block of a distributed environment and must be managed to ensure reliability, availability, and serviceability when mission critical applications are deployed.

IBM Database Systems Management Technologies

IBM recognizes that businesses are searching for an optimal mix of heterogeneous platforms, operating systems, and complex environments, and that these challenges require a broad range of skills. IBM′s Database Systems Management solution is a comprehensive tool suite that offers an open, central point of control for managing small as well as complex heterogeneous database systems with minimal skill and resource requirements.

The IBM Database Systems Management solution addresses client/server and distributed database systems that can be complex and difficult to manage. The solution is comprised of:

- Database Administration (DBA) Tools
- Advisor Tools
- Data Replication Products
- Utilities

Database Administration (DBA) Tools

DBA tools simplify administration and operational tasks. You can use DBA tools to manage both IBM and non-IBM databases, facilitate expansion into a client/server environment, serve as a launching platform for management tools, and leverage your database administration skills across many databases, including databases from different vendors.

IBM DataHub* for OS/2

DataHub for OS/2 allows you to manage complex client/server relational database environments from a single OS/2 control point (the DataHub/2 workstation). DataHub for OS/2 supports DB2 for MVS, DB2 for VM, DB2 for OS/400, DB2 for AIX/6000, and DB2 for OS/2.

With DataHub for OS/2 you can perform the following systems management tasks from the DataHub for OS/2 workstation control point:

- Display catalogued database objects and relationships between those objects.
- CREATE/DROP/ALTER database objects.
- Copy objects, such as tables and their associated authorizations, between databases.
- Manage database authorizations.
- Display the status of relational units of work.
- Invoke and schedule database utilities and commands, such as BACKUP and RECOVER, for remote or local databases.
- Automate database management tasks, either via the DataHub for OS/2 command scheduler or via interoperation with IBM′s workflow management product, FlowMark*.

DataHub for Unix

DataHub for Unix Operating Systems provides powerful functions DBAs can use to manage a variety of databases, all from a single control point. As your business expands and diversifies, your DBAs are faced with the challenge of managing geographically dispersed, multiple vendor database environments on different hardware and software platforms. DataHub for UNIX OS provides three control points from which DBAs can effectively manage your enterprise′s data:

- AIX
- HP-UX
- Solaris Operating System

From these control points, DataHub for UNIX OS can manage DB2 for AIX, DB2 for HP-UX, DB2 for Solaris, Oracle, and Sybase databases. DataHub for UNIX OS offers the following database systems management functions:

- Display catalogued database objects and relationships between those objects.
- Create, change, delete, rename, and analyze database objects.
- Copy tables, including contents and column definitions, indexes, primary keys and manage data type conversions between databases.

- Manage database authorizations, including user authorizations and privileges, roles and their authorizations and privileges and profiles.
- Auditing and accounting (monitoring database objects, SQL statements, and user sessions).
- Invoke and schedule database utilities and commands, such as BACKUP for remote or local databases. Automate management tasks via continuous monitoring of system and database conditions. Automated tasks can be initiated or scheduled based on the correlation of different events.
- Monitor system and database conditions on request, such as disk layout, CPU and disk utilization, locks, user and other statistics.

Note: DataHub for OS/2 and DataHub for Unix are two different products with a somewhat different scope of functionality. Over time, their functionality will converge. Currently, DB2 for AIX/6000 can be managed by either control point.

Advisor Tools

Advisor tools provide functions for performance and capacity prediction, as well as event monitoring, alert management, and network management integration. Event monitoring tools can be used to manage your production applications. And, these tools can be used for designing key transactions during stress testing of performance critical applications.

DB2 Visual Explain for OS/2 and AIX/6000

DB2Visual Explain is and easy-to-use, graphical Explain tool that provides detailed information about the access plan chosen by the DBMS optimizer to access data. DB2 Visual Explain will support Version 2 of DB2 for AIX/6000 and DB2 for OS/2.The product provides functions for:

- Graphical presentation of Explain output
- Detailed optimizer information
- What-if modeling capability for SQL statements

IBM Database 2 Performance Monitor for OS/2 and for AIX/6000

DB2 Performance Monitor for OS/2 and AIX/6000 provides comprehensive performance data collection, viewing, reporting, analysis, and alerting capabilities.

DB2 Performance Monitor supports over 200 performance attributes, including bufferpool, lock and deadlock, sorting, communication, agent, and logging information. DB2 Performance Monitor also supports performance measurements derived by performing calculations on existing measurements. Additionally, data can be analyzed to filter out only those records applicable to certain database objects or to a certain period of time.

For any performance measurement, you can define exception conditions by specifying a threshold value. When the threshold value is reached, you can specify any or all of the following actions:

- Notification through a window or audible alarm.
- Logging of a record in a log file.
- Execution of a command or program.

DB2 Performance Monitor for MVS (DB2 PM)

DB2 PM provides full performance monitoring and problem determination support for all functions for DB2 for MVS. You can use DB2 PM to:

- Measure an application′s performance and resource utilization.
- Evaluate an application′s impact on other applications and the system.
- Determine DB2 subsystem performance and efficiency.
- Determine tuning requirements for DB2.
- Diagnose the cause of performance problems and provide suggestions for resolution.
- Prevent performance problems by optimizing DB2 applications (with extensive EXPLAIN capability).

DB2 PM also provides you interactive report generation, background reporting, and on-line exception processing. You can customize reports to meet your requirements, and control the volume and content of accounting and statistics reports.

DB2 Estimator for Windows

DB2 Estimator is a low-cost, easy-to-use, stand-alone tool for estimating the performance of DB2 applications, transactions, database utilities, and the performance of SQL statements. DB2 Estimator is also optionally invokable from DataHub. During the initial design phase, DB2 Estimator can be used to easily model an application to help evaluate alternative designs. By comparing the cost and performance of these designs, you can determine which design provides the optimum cost/performance ratio -- before valuable time and resources are invested coding an application or creating a real database.

While an application is in production, you can easily evaluate the alternative SQL designs and evaluate the resulting effects on system performance before having to modify any production database objects. DB2 Estimator also helps determine the impact resulting from hardware or workload changes.

Data Replication Products

IBM Data Replication products address the operational, informational, and administrative requirements of data replication. The following replication products support a range of operating environments from mainframes to LANs to workstations, allowing client/server applications to work seamlessly with legacy and relational database management systems.

DataPropagator* Relational

DataPropagator Relational provides replication support between the IBM DB2 family of databases. Update propagation and full refresh propagation will be provided between DB2 for MVS, DB2 for OS/2, DB2 for OS/400, and DB2 for AIX/6000 and DB2 for Windows NT. In addition, other vendor databases, such as Sybase and Oracle, will be supported as targets of full refresh and update propagation through interoperation with DataJoiner.

Note: DataJoiner is a multi-vendor database server providing heterogeneous joins of data between diverse relational database management systems.

Copies are automatically maintained at user-specified intervals, or can be initiated through application events (for example, to coincide with end-of-business reconciliation).

DataPropagator Relational provides powerful data enhancement capabilities for tailoring copies to specific applications or end-user requirements. Enhancements are specified through standard SQL and support joins, subsetting, summarization, new data calculations, and other transformations as part of the replication process. DataPropagator Relational′s components are fully independent, providing for continued operation in the event of network or system failures. DataPropagator Relational recovers automatically from such failures, identifying where it left off and how best to continue processing. Additionally, DataPropagator Relational′s component design allows for adding, modifying, and deleting of copy requests while the rest of the system continues to run.

DataPropagator Relational′s data staging area enables sophisticated customizing and tool interoperability among IBM and multi-vendor products and platforms -- it defines the architecture for product interoperability. DataPropagator NonRelational and DataRefresher* interface with DataPropagator Relational through this interface to extend the solution to other legacy data sources. Customers or other third party vendors can also use the data staging area to extend the IBM Data Replication solution to multi-vendor sources and targets.

You can invoke DataPropagator Relational via a graphical user interface, enabled to DataHub for OS/2, for specifying sources, targets, data enhancements, and copy frequency. Initialization functions are automated, and failure scenarios are automatically handled.

DataPropagator NonRelational

DataPropagator NonRelational helps you leverage your IMS asset for greater business advantage by allowing IMS data to be replicated across the IBM DB2 family of databases. DataPropagator NonRelational′s synchronous copying helps you execute business plans to migrate applications from IMS to DB2, providing a structured, low-risk way to achieve this migration over time.

With synchronous replication, you can copy from IMS to DB2, from DB2 to IMS, and in both directions simultaneously. Changes are applied immediately, in the same unit-of-work as the updating application, providing fully synchronized copies between IMS and DB2.

By propagating changes from IMS databases to copies of those databases in DB2, you can implement new decision-support applications that use relational technology, without affecting any existing IMS/DB programs. Additionally, DataPropagator NonRelational provides support for point-in-time copies or complete histories.

Using DataPropagator NonRelational′s synchronous two-way propagation capabilities, you can support two master copies of the same data, one in IMS and one in DB2. Full synchronization allows for coexistence of IMS and DB2 applications or for an incremental, controlled migration from IMS to DB2. Shared data across IMS and DB2 is synchronized, without requiring changes to the applications using the data. Both databases must be on the same MVS image and must be used with IMS Transaction Manager.

DataRefresher

DataRefresher provides full refresh copying from any MVS source (IMS, VSAM, DB2, flat files, and others) to targets in MVS, VM, VSE, AS/400, AIX/6000, and OS/2, including the DataPropagator Relational staging area. It can perform complex refresh processing of data from several sources, directing it to several relational targets. DataRefresher can initially synchronize IMS and DB2 databases for later update propagation using IBM′s DataPropagator NonRelational, and can also feed legacy data to IBM′s DataPropagator Relational for further propagation.

You can use DataRefresher to refresh your target database daily. Once an extract has been built, it is saved on your workstation and can be sent to the host for processing at any time.

DataRefresher supports heterogeneous joins across up to 16 different sources and can create up to 99 different targets in a single pass of the source data. Generic interfaces allow you to provide exits that support extracting data from or writing data to multi-vendor data stores. Building on the proven technology of its predecessor, Data Extract (DXT) V2.5, DataRefresher provides full upward compatibility for existing customers.

DataRefresher′s graphical user interface (GUI), enabled to DataHub for OS/2, is structured to match tasks, frequency of execution, and skill level of the various users. This GUI expands the traditional user set of DBAs to application programmers, LAN administrators, and other end users, who can be shielded from the underlying complexities of the data being extracted.

Utilities

IBM Database Engine Utilities

Utilities address routine operational utility requirements, including automated and remote utility invocation and scheduling. Utilities include LOAD, REORG, UNLOAD, RUNSTATS, BACKUP and RECOVER. Utilities are provided with the database engines.

IBM ADSTAR* Distributed Storage Manager

DB2/6000 databases can be backed up using ADSTAR Distributed Storage Manager (ADSM). ADSM provides support for remote backup to MVS, VM, AIX/6000 and OS/2 2.1 based systems. DB2/6000 includes the ADSM API Version 1.1.0, which is used during the backup and restoration of DB2/6000 databases.

IBM Decision Support Technologies

IBM Decision Support Technologies are designed so end users, with various skill levels, can easily access and analyze appropriate data. Relevant pieces of information are brought together in a logical fashion in order to assist making decisions. IBM tailors offerings to address needs of users as diverse as power users, professionals, and end users unskilled in data retrieval or information technology. Utilizing graphic and multimedia interfaces, users can browse or request data without previous education about data structures or interfaces; disparate data is transformed into useful information for making better business decisions.

Locating Information for Analysis

Decision makers faced with a business problem need to locate and analyze data relevant to the problem.

IBM DataGuide*

DataGuide provides a powerful business-oriented solution to help end users locate, understand, and access enterprise data. In client/server information catalogs, business metadata (data about data) can be described in business terms, organized into subject areas, and customized to your workgroup's or enterprise's needs. DataGuide is IBM's Information Warehouse* facility for integrating and managing end-user business metadata.

DataGuide User enables you to:

- Locate data in the information catalog by keyword text search or navigation ("drill down") through a subject area and business groupings
- Understand your data by browsing metadata descriptions in business terms
- Access the data by launching applications directly from DataGuide
  - DataGuide for OS/2 User can launch OS/2, Windows, DOS, and client/server applications. Using DDCS for OS/2 as a gateway, DataGuide can launch client/server applications for access to remote databases such as DB2 for MVS.
  - DataGuide for Windows User can launch popular Windows and DOS applications.
- Communicate comments and corrections on metadata descriptions to the DataGuide administrator via a comments object
- Share administration of the DataGuide information catalog (when authorized)

- Be more productive when looking for data by using an easy-to-use graphic interface in either the powerful OS/2 or the familiar Microsoft** Windows environment
- Provide a workgroup solution for information sharing of objects including spreadsheets, charts, reports, tables, queries, images, videos, and more
- Share the same information catalogs with other DataGuide users across a LAN or enterprise
- Identify the owners for data via the contact object

DataGuide Administrator enables you to:

- Establish your information system: set up the information catalogs and identify administrator and user update authority
- Classify, organize, and annotate metadata objects from data sources (such as databases, tables, columns, and warehouse) and informational sources (such as queries, reports, spreadsheets, charts, documents, and internet documents)
- Populate your information catalog from a variety of metadata sources using provided extractors or your own customized extractors
- Associate application programs to metadata objects for launching in OS/2, Windows, or DOS environments to deliver timely data and information to your end users

- Define virtually any metadata object, or object relationship, your end users need within DataGuide's extendible, flexible data model
- Integrate and manage your business metadata through the externally published tag language and Application Programming Interface (API)
- Maintain the information catalog through an easy-to-use graphic interface in the powerful OS/2 environment

DataGuide Products:

- Provide a client/server information catalog
- Help you share and reuse company information
- Store information catalogs on DB2 for OS/2 on a LAN, DB2 for MVS, DB2 for AIX, DB2 Parallel Edition, or DB2 for OS/400

Analyzing the Data

Once data has been found that can help resolve business problems, tools may be needed to analyze the data. Several classes of tools are available for different types of analysis:

- General purpose Data Inquiry tools such as IBM's Query Management Facility,and Lotus Approach
- Data Interpretation tools such as IBM's Application System
- Multidimensional Analysis tools such as IBM's Intelligent Decision Server
- Information Discovery tools such as IBM's Intelligent Miner

Following are highlights of IBM's broad selection of decision support tools.

Lotus Approach

Designed for anyone who can use a spreadsheet, Lotus Approach is the award-winning relational database that guarantees everyone on your team can track, manage and analyze business information within two hours of opening the box. And with more than 100 new and enhanced features, Lotus Approach 96 is easier and more powerful than ever before. The intuitive interface with familiar terms and a notetab layout enables you to become productive quickly and puts the commands you use most frequently right at your fingertips. With the enhanced PowerClick reportwriter, you can generate professionally designed reports with unprecedented simplicity. And new and improved charting, crosstab and worksheet analysis capabilities let you review and present your data like a pro.

Lotus Approach 96 is not just easy to use, it's fast. Lotus Approach 96's X10 speed technology is more powerful than ever, allowing you to get real results, real fast. What's more, the enhanced PowerKey technology delivers the best compatibility and connectivity in the business, letting you work seamlessly with dBASE, Paradox, Lotus Notes and other popular databases. That means no wasted time importing, converting, or "attaching" your data. If you're a 1-2-3 user, you'll love Approach. Approach's superior integration with 1-2-3 means you can access Approach's data management capabilities and use them to manipulate live 1-2-3 data -- all transparently from within 1-2-3 menus.

And, if it's custom solutions you're looking for, the new LotusScript** 3.0 programming language can help you automate tasks and create sophisticated, cross-product business solutions.

So, if you need to track customers or personnel, report on sales information, manage projects, analyze your corporate data and more, check out Lotus Approach 96. You′ll find the usability you′ve dreamed of plus the speed and power you need to tackle tough business problems.

Become productive quickly and get the most out of your information

- Enhanced! Intuitive interface with familiar terms like ″report″, ″form″, and ″find″ and an elegant notetab design helps you get started fast. In fact, studies have shown that the average Approach user becomes productive with the software within 2 hours of opening the box!
- New! SmartMaster applications enable you to create complete database applications from scratch in minutes. Input your own data and use the professionally designed default reports, forms, and analyses, or customize your new application to meet your needs.
- Enhanced! PowerClick reporting with point-and-click report design lets you see exactly how your reports will appear in print as you work. You can sort, group, and calculate data by simply clicking on SmartIcons.
- Enhanced! Chart, crosstab, and worksheet views help you analyze your information and make sound business decisions.
- Relational power enables you to tackle complex, information management problems. You can build applications which combine information from up to 50 different tables in one file.

Easily share information with your team

- Team data management with plug and play networking and multi-user support enables your entire team to seamlessly and securely share common pools of information.
- Enhanced! PowerKey data access provides fast, direct access to virtually all of your business information, regardless of file format. Update, manipulate and report on information stored in dBase, Paradox, Notes, and more while enjoying native database speed and functionality.
- New! TeamSecurity maintains data security with ease and flexibility by providing customized security privileges for individual team members or entire groups.

Create custom business solutions

- Lotus 1-2-3 integration leverages your investment in 1-2-3 by using Approach to create forms, reports, mailing labels, and crosstabs on your live 1-2-3 data. Transparently obtain Approach functionality from menus within 1-2-3. Or, choose any 1-2-3 named range from within Approach and create the forms, reports, etc. you desire.
- New! LotusScript 3.0 programming language provides a development environment for robust, cross-product applications.
- Enhanced! Full OLE 2.0 support enables you to seamlessly share information with other OLE 2.0 applications. Approach will serve as either an OLE 2.0 server or container.

IBM Query Management Facility (QMF*)

QMF integrates a query tool, report writer, table editor and application support. Three query methods are provided: Prompted Query, Structured Query Language and Query-by-Example. Reports can be displayed in a wide variety of formats including tabular, matrix, free format and graphical. QMF offers all this in an environment that is easy for a novice to use but powerful enough for an application programmer. With QMF, you can easily access relational data stored in DB/2 on MVS and SQL/DS VM or VSE, DB2 for AIX Version 2, and DB2 for OS/2 Version 2 and DB2/NT.

Applications written in many languages can use QMF query and report writing facilities through the callable interface. QMF lets you write procedures with logic, using the full power of the REXX programming language to process your instructions. A QMF Governor is supplied to provide resource control capability. Access to data at remote locations allows QMF users and application programs to connect to any of the DB2 Family of databases within a distributed network.

QMF also provides a set of integrated tools to enhance performance management and to facilitate the administration of QMF:

- The QMF HPO/Manager consists of a group of functions that improves governing capabilities, including a preemptive governor to analyze QMF queries and establish controls to protect production applications and satisfy the business need for ″on demand″ information. The QMF HPO/Manager allow you to display critical information on active batch and online QMF session as well as provide powerful object tracking and management capabilities, including a broad range of statistics on usage and administration of QMF objects. .
- The QMF HPO/Compiler provides the ability to automatically convert queries and reports into efficient programs in OS/VS COBOL or VS COBOL II.
- QMF HPO/Manager and QMF HPO/Compiler now include a Windows client component called Shuttle. Shuttle allows the user to execute a QMF host query and bring the results to the workstation for embedding into other workstation products such as Lotus 1-2-3 or Microsoft Excel spreadsheets.

QMF HPO/Complete combines the benefits of the QMF HPO/Compiler and the QMF HPO/Manager into one easy-to-order and economical feature.

IBM Application System (AS)

Application System is IBM′s proven, integrated decision support system for the VM and MVS environments. Application System provides business control tools for customers and enables customers to make decisions based on facts rather than guesswork. User interfaces suit the novice to the expert, with general data processing facilities and specialist business tools. It can be used for simple and ad hoc queries through the development of complex applications, using data from many sources to provide information to all organization levels.

- AS Server is a mandatory base product which provides the basic structure for other optional host components. On its own it provides standard host facilities for access via client products such as AS Client .
- AS Client provides full host terminal support for non-programmable terminals. All the standard features of AS are available apart from the facilities of the Business Tool feature.
- AS Business Tools provide facilities for project management, business planning, business statistics and optimization (linear programming).

AS provides users access to its valuable information store and its analytical power on System/390 systems. It enables seamless client/server applications to be designed which exploit the different strengths of the mainframe and workstation environments.

IBM Intelligent Decision Server (IDS)

The IBM Intelligent Decision Server (IDS) provides an analytic engine that functions in a three-tier architecture. This engine provides the capability to perform querying and reporting, OLAP-type analysis and to rapidly develop decision support applications that manipulate relational data sources. The results of these data manipulations are available to a wide variety of clients, making Intelligent Decision Server analytical applications available throughout the enterprise.

There are two major categories of OLAP products: those which store data in a proprietary, hierarchical, multi-dimensional database (MDDB) architecture or data cube, and those which map a multi-dimensional view to data stored in an RDBMS, sometimes referred to as ROLAP products. Intelligent Decision Server falls into the second, ROLAP, category, which is known for its flexibility in handling multiple, complex views of the data without a need to repeatedly extract separate, multi-dimensional data structures for analysis.

IDS Analytical applications, known as Capsules, residing on shared servers can be accessed from a broad set of clients, passing parameters and receiving results. Clients including popular Lotus Notes and Web browsers such as IBM WebExplorer and Netscape Navigator. Together, a wide variety of platforms, ranging from Windows 3.1**, Windows 95**, Windows NT, to OS/2 Warp*, Macintosh, and UNIX systems such as IBM AIX, HP-UX, Sun Solaris, Silicon Graphics, can all be used to access IDS Capsules.

In addition, Capsules can also be accessed from customized applications running on Windows 3.1, Windows 95, and Windows NT as well as Sun, OS/2 Warp, and AIX.

IDS WebKit supports Web servers to be set up on SunOS V4.1.3, Sun Solaris V2.4, OS/2 Warp, and AIX 3.2.5.

IDS provides a gateway for IBM DataJoiner, the multi-database server, which is capable of transparently accessing and joining a number of IBM and non-IBM relational and non-relational data sources, and can extend access to additional data sources via ODBC data drivers.

In addition, a number of other relational database gateways are available from IDS, including DB2 for MVS, DB2 for OS/2, DB2 for AIX, Sybase SQL Server and Oracle (on RISC System/6000*, HP 9000, and Sun SPARCstations), and Red Brick Warehouse (on RISC System/6000 and HP 9000). In addition, access to data stored on SQL/DS and DB2 for OS/400 is available through DRDA client support with DB2 for OS/2.

The Intelligent Decision Server provides multi-dimensional data analysis capability against relational data. This support is provided via IDS′s ability to manipulate data cubes stored in a relational star schema. (A relational star schema separates data into fact and dimension tables.) The tools then allow the developer to graphically construct complex analyses with multi-dimensional views.

The DBA tool enables the management of metadata, which maps the physical attributes of the data sources to a business dimension. This metadata is stored in relational tables along with the star schema for easy manipulation.

For the developer, IDS provides a graphical, object-based development environment that allows for sophisticated analysis of relational data for presentation and distribution to end users in the form of reports, spreadsheets and graphics. There is a set of powerful objects called transformers which are used by developers to create complex applications. Transformers cover a wide range of functions, from data cleansing, data manipulation, to statistical analysis. No programming language is required to create applications.

The Capsule development environment currently runs on OS/2 and IBM intends to extend the environment to Windows 95 and Windows NT.

Users of the IDS will have access to functions provided by the IBM Intellegent Miner. Functions within the Toolkit will be accessible via additional IDS transformers. These data mining transformers can be used with existing transformers to allow for data mining applications to be developed. Like any other Capsule application, these data mining applications can be accessed from a wide variety of clients. Results from data mining can also be analyzed by IDS.

IDS provides ODBC access from Windows and OS/2 clients. This allows ODBC-enabled data retrieval, analysis, and reporting tool, including Lotus Approach and other vendor tools, to execute IDS Capsule applications with a restricted set of SQL statements, and retrieve results in tabular form.

IBM Intelligent Miner

The IBM Intelligent Miner is an integrated solution for large scale, sophisticated analysis of data. It allows data analysts to harvest valuable information from databases and present it to business users for decision making.

Data mining is the process of extracting valid, previously unknown, and ultimately comprehensible information from large databases and using it to make crucial business decisions. The extracted information can be used to form a prediction or classification model, identify relations between database records, or provide a summary of the database(s) being mined.

The objective of the Intelligent Miner is to expedite the information discovery operation while maintaining the quality of the extracted information and enabling the decision-maker to take advantage of niche or emerging opportunities.

The Intelligent Miner is a suite of tool functions which support data mining operations and deploy a variety of techniques to

- Create classification and prediction models,
- Discover associations and sequential patterns in large databases,
- Automatically segment databases into groups of related records,
- Find similar patterns of behavior within special time sequences.

and provide functions to:

- Prepare the data for mining, and
- Visualize the discovered information using advanced graphical techniques.

These functions are implemented in a consistent manner and accessible interactively by the data analyst or by decision support applications via a published API. The data mining process is iterative and functions can be used independently or in combination to produce the desired result. Analysis of mined results may result in the repetitive execution of the same or altered mining operation.

Four operations are associated with discovery-driven data mining:

1. Creation of prediction and classification models to use the contents of the database, which reflect historical data, i.e., data about the past, to automatically generate a model that can predict a future behavior;
2. Link analysis to identify explicit relations between the records in a database;
3. Database segmentation to partition databases into collections of related records either as a means of obtaining a summary of each database, or before performing a data mining operation such as model creation, or link analysis;
4. Deviation detection to identify outlying points in a particular data set, and explain whether they are due to noise or other impurities being present in the data, or due to causal reasons.

The techniques supported by the Intelligent Miner are:

- Clustering - for example, group a set of individuals based on demographics.
- Classification - for example, "Buyers of expensive sport cars are typically young urban professionals whereas luxury sedans are bought by elderly wealthy persons." transaction.
- Association Discovery - for example, of the shoppers who purchased milk, 55% also purchased some other type of dairy product, and 42% also purchase bread.
- Sequential Pattern Discovery For example, "42% of new checking account customers who apply for an ATM card in the same transaction, will also apply for our charge account within 90 days".
- Similar Time Sequence Discovery - for example, given NorthSouth Airlines daily closing price for 1995, find all other stocks with similar behavior.
- Prediction - for example, predict likelihood of fraud for a credit card.

In establishing a Data Warehouse users address data preparation issues such as cleansing, consistency, and reconciliation to create a well-defined, integrated data store for analytics and operational applications. However, for a given data mining operation, additional steps may be required to perform field selection, sampling, and transformations such as data conversions or creation of derived attributes, as dictated by the type of data mining operation performed and the data mining technique used.

Some of the most typical of these data processing steps have been incorporated in the Intelligent Miner to improve the analyst's productivity. These tasks have been automated and can be invoked dynamically in the iterative process of preparation, mining and verification -programming is not required.

The processing functions for data preparation work with data in DB2 tables. Examples of processing functions incorporated include:

- Data collection functions to extract data from other DBMSs and/or systems and load in the mining base. For example, extract data from an Oracle database and load into a DB2 database for mining.
- Data selection functions to select, sample, aggregate, project, group, summarize. For example, merge pharmaceutical transaction data with purchased demographic data regarding physicians.
- Data transformation functions to perform symbol mapping, encoding, conversion, or to discretize values based on ranges or quantities. For example, categorize vehicle speed as high, medium or low.
- Data cleansing functions to filter or discard invalid data. For example, discarding customer demographic data with invalid state code.
- Data derivation functions to derive mathematical or statistical values. For example, calculate frequency count of transactions by customer.

Visualization allows users to explore the relationships in data interactively using advanced graphical techniques. The extracted information is analyzed and the best information is identified and presented to the decision-maker. Therefore, the purpose of result interpretation is not only to visualize (graphically or logically) the output of the data mining operation, but also to filter the information that will be presented to the decision-maker through the decision support system. The analyst may iterate through additional data mining and preparation steps to refine, filter or extend the data mining result.

Data mining results can be further analyzed, queried, or presented to business analysts through the use of popular decision support products including IBM's Intelligent Decision Server and Query Management Facility, using standard interfaces which support DB2 access. Over time IBM and other Independent Software Vendor (ISV) DSS products may be extended to directly invoke Intelligent Miner functions by use of its published API.

Furthermore, results of data mining operations may be taken directly to the business user or be used to drive core business systems using customized data mining applications.

IBM has developed three customizable, cross-industry applications.

1. Customer Segmentation -- segments and scores customers data from marketing databases, including private and public data sources, in an effort to better understand customer behavior. Results are used for target marketing, cross-selling, customers retention campaigns, propensity to purchase and consumer vulnerability analysis campaigns.
2. Item Set Analysis -- also known as market basket analysis, this application aims to understand customer buying behavior, and to predict their future behavior by identifying affinities among their choice of products and services.
3. Fraud Detection -- this application identifies deviations from established usage norms in order to flag suspicious transactions which may be indicative of fraudulent activity. Results are used to identify potentially fraudulent claims practices, target suspect providers, and discourage further abusive practices. Customers and other Independent Software Vendors will also be able to build customized applications using the Intelligent Miner's published API.

Your Data Warehouse is an obvious data source for mining operations, but sources also include extracts from operational databases and purchased data such as demographics, lifestyle, etc. data. The Intelligent Miner is designed to support data mining against DB2 databases or flat files.

The flat file support not only facilitates handling of purchased and transaction data, but may handle extracts from operational data and other DBMSs as well. In addition, a high speed extract is provided to import into DB2 data from Oracle and Sybase databases.

IBM DB2/400* Query Manager and SQL Development Kit

The DB2/400 Query Manager and SQL Development Kit* provides an interactive query and report writing interface similar to QMF, as well as precompilers and tools to assist in writing Structured Query Language (SQL) application programs in high-level programming languages. DB2/400 Query Manager and SQL Development Kit is the new name for the enhanced version of its predecessor product, SQL/400*.

Information Warehouse* Technologies

The previous section discussed tools for analyzing data. Often an infrastructure is put in place to ″house″ the data needed by decision support applications - hence the term ″data warehouse″. Data for the warehouse is typically extracted from operational data and then subjected to transformations, aggregations and cleansing to make it more suitable for business analysis.

A data warehouse is a store of data that is optimized for applications that support analytical processing and decision making, called informational applications. IBM offers a comprehensive capability for implementing a data warehouse called the IBM Information Warehouse Solution. This solution provides data warehousing capabilities by means of an open architecture, a comprehensive set of products, consulting services, and extended capabilities not readily available from other vendors. These extended capabilities are the processes, tools, and facilities to create, manage and maintain the data warehouse.

IBM also offers flexible alternatives for implementing data warehouse solutions. These alternatives vary from starting with a ″tops down″ approach and resulting in the creation of more customized ″global″ data warehouse, to starting from the ″bottoms up″ approach and resulting in the creation of a smaller data warehouse that services a department or workgroup and is sometimes called a ″datamart″. IBM has the products and services to help customers with either approach.

IBM Visual Warehouse Solution

The Visual Warehouse is a LAN-based product designed for creating a data warehouse (or datamart) in a departmental or workgroup environment. It is ideal for this type of environment because of the low relative start-up costs, fast implementation capability, and easy administration support for building and managing the datamart.

The Visual Warehouse can natively access data from a variety of sources, transform it, store it, and deliver it to end users through a variety of client tools and applications. It is also highly scalable and can support databases that are quite large. And, it can be integrated with other Information Warehouse products to be part of a larger enterprise data warehouse.

The IBM Visual Warehouse Solution includes an end-user information catalog (DataGuide) and an end-user decision support tool (Lotus Approach), along with the Visual Warehouse product, to provide a complete datamart implementation. Consulting services are also available to provide any implementation help the customer needs.

The Visual Warehouse can access source data from DB2, VSAM, IMS, Oracle, Sybase, Informix, Microsoft SQL Server, flat files, and binary large objects (BLOBs) and populate the Visual Warehouse data store on either DB2 for OS/2, for AIX/6000, for OS/400, for HP-UX, for SUN Solaris, and for Windows NT**. This is a powerful datamart that can be located close to the end users to reduce network traffic, provide high availability, and enable them fast and easy access to the data. For example, if one branch office keeps sales history data on an AS/400 and another keeps sales history data on a Sybase system, Visual Warehouse can access the data from both sources, transform it as necessary, and store it in the warehouse data store. The data can then be made available to the end users through a variety of popular end-user tools, either using ODBC or the IBM Client Application Enabler (CAE).

IBM Information Warehouse Solution

In addition to Visual Warehouse, IBM offers a comprehensive set of products and services that can be used to build a more customized ″global″ data warehouse solution. These products and services are described elsewhere in this paper but are summarized here to show how the offerings fit into an overall data warehouse solution:

- The IBM DB2 Family of relational database management systems provide the flexibility, security, availability, and scalability to support the requirements of any data warehouse implementation.
- The IBM Data Replication products provide the power to transform, filter, and distribute data to meet the requirements of a specific datamart or an enterprise-wide data warehouse.
- IBM DataJoiner can be used to access and join data from a variety of heterogeneous data sources. The results can then be used to populate the data warehouse.
- IBM DataHub offerings can be used to automate many of the data administration functions required to maintain a data warehouse in a heterogeneous environment.
- IBM FlowMark can be used to define, manage, test and automate data warehouse processes to minimize the technical skills requirements and workload of the I/S operations organization.
- IBM DataGuide provides a common interface, or ″window″, into the warehouse data for end users. It is an information catalog that describes what data is available in the warehouse in business terms that the end users can understand. It provides the management, control, and point of integration for all the metadata in the data warehouse.
- IBM Decision Support tools such as Lotus Approach, the Intelligent Decision Server, Application System, and Query Management Facility can be used to manipulate and analyze data and present the results, as information, to the end user.
- IBM Intelligent Miner provides a suite of tools that expedite information discovery.

Application Development Technologies

Customers can build applications more productively with application development offerings integrated with other data management products

IBM Application Development Technology

IBM is focused on providing a full line of application development (AD) products and service offerings as complete solutions, to solve AD problems across a range of development and execution environments. The IBM AD strategy focuses on tools that address the implementation of open client/server applications, development of applications using object-oriented technology, expansion of the AD family to multiple IBM and non-IBM platforms, and stronger linkages to IBM′s robust data management products.

It is to this vision that IBM is building various AD solutions:

- Enterprise 3GL Client/Server: For application developers with a huge inventory of existing applications who build or maintain applications primarily in COBOL or PL/I for an enterprise execution environment, and who want to take advantage of workstation technology, for both building and deploying applications. (Cobol, PL/I, etc.)
- Enterprise Rapid Application Development: For developers of enterprise applications that leverage the power of productive techniques, such as fourth-generation languages and Computer Aided Software Engineering (CASE). (VisualGen, TeamConnection, etc.)
- Workgroup Rapid Application Development: For developers primarily working on new client/server applications using powertools: technology which includes visual programming, object technologies, and new methodologies for execution primarily in the PC/Workstation local area network (LAN) environments. (VisualAge, IBM Smalltalk, etc.)
- Workgroup C++: For the programming professional whose primary languages are C and C++, developing applications in a LAN-based PC/workstation environment.
- AS/400*: For application developers using traditional technologies, but with a focus on visual, object, or client/server technologies for the AS/400 environment. (ILE Cobol/400, ILE RPG/400, ILE C/400, Application Development Toolset/400, etc.)
- Engineering/Scientific: For the technical professional solving mathematically intensive problems, and who develop applications primarily in FORTRAN on various platforms. C and, for some advanced users, C++ are becoming important additional languages in this domain. (Fortran, etc.)

Highlights for Selected IBM Application Development Solutions

IBM VisualAge for Smalltalk

VisualAge is a client/server powertool, provides the professional developer working individually or within a team, with a capability for creating graphical clients which can access local and remote databases and support a wide variety of LAN-based applications with Windows** and OS/2 Clients and SOM Support.

VisualAge for Smalltalk enables you to access a broad range of databases created by IBM and other major database vendors. The base VisualAge product provides native support for DB2/2* that exploits the 32-bit power of OS/2 (This includes support for DB2/2 2.1 function, including user-defined functions and types (UDF and UDT) and binary large objects (BLOBs)). DB2/2 supports a complete suite of functions, such as server-enforced referential integrity, stored procedures, and industry-standard SQL support.

Teamed with the Distributed Database Connection Services/2* (DDCS/2) product, VisualAge provides host database access for DB2/6000*, DB2 for VSE and VM (SQL/DS), and DB2 for AS/400.

VisualAge′s ODBC support, which is included with VisualAge and IBM Smalltalk, enables you to create applications that access all your data sources, from spreadsheets to relational databases like ORACLE, INFORMIX**, and SYBASE** SQL Server.

VisualAge for Smalltalk, Database for ORACLE, a new feature, is available as an add-on to provide you with native access to specific ORACLE functions for excellent performance. It′s never been easier to access the database of your choice.

VisualAge C++ for OS/2

VisualAge C++ for OS/2 is a powerful development package, allowing AD professionals to exploit the full potential of OS/2 32-bit capabilities, and includes a comprehensive set of class libraries.

The Data Access Builder allows you to create new object-oriented database applications quickly and reliably by generating the source code and embedded SQL for you. Add, update, delete, and retrieve methods are generated for each class.

These database parts can be used directly in your programs, or you can import them into Visual Builder. By using Visual Builder to connect them to the GUI, or other parts, you can create high quality applications quickly.

Object-Oriented COBOL Across Multiple Platforms

COBOL developers can now take advantage of object-oriented technology with IBM′s deliveryof VisualAge for COBOL for OS/2*, COBOL Set for AIX, and COBOL for MVS*. The availability of COBOL with object-oriented extensions allows programmers to continue to use a powerful, popular programming language, create advanced client/server applications on the desktop, and deploy objects across various platforms and applications.

VisualGen

VisualGen provides a new generation of visual tools for development and generation of both client and server application components. It is fully resident on LAN-based workstations. Its high level definition facilities target both mainframe and workstation environments, and transparently supports various client/server communications protocols.

With VisualGen V2.0, you can rapidly develop applications for the DB2 family: VisualGen V2.0 fully exploits Distributed Relational Database Architecture at definition, test and execution. VisualGen supports non-IBM relational databases via DataJoiner for AIX. Future releases will add additional non-IBM database support capabilities.

IBM DB2 Software Developer′s Kits (SDK)

DB2 Software Developer′s Kits provide programmers with application development support for building application programs that manipulate data stored in DB2 family databases. These products provide documentation, tools, and code samples for developing database application for execution in DOS, Windows, OS/2, AIX/6000, HP-UX, and the Solaris Operating Environments. All SDKs provide facilities for developing applications using either embedded structured query language (SQL) or a call level interface (CLI) that is also provided.

IBM DataBasic

DataBasic is a powerful suite of rapid application development tools built around an advanced implementation of the BASIC* programming language. It makes creating, debugging, and testing DB2 stored procedures and user-defined functions point-and-click simple.

DataBasic lets application programmers work in an OS/2, AIX or Windows 3.1 development environment to easily and quickly develop BASIC* client applications as well as BASIC stored procedures and user-defined functions which run with DB2 for OS/2 and DB2 for AIX. From this development environment, the programmer can automatically distribute the procedures and functions to the target server, and build, debug, and execute the stored procedures and user-defined functions on the server. If the programmer is developing the database procedure or function on OS/2 or AIX and has DB2 for OS/2 or DB2 for AIX installed locally, the procedure or function can be unit tested and debugged locally before being distributed to the target server for final test and execution.

IBM Clients and Client Application Enablers (CAE)

Client access enabling products are available for DOS, Windows, OS/2, AIX/6000, HP-UX, and the Solaris Operating Environment. The interfaces provided support embedded SQL, CLI, and ODBC. Also included is runtime support for any client application running against DB2. Access is enabled to DB2 for OS/2, DB2 for AIX/6000, DB2 for HP-UX, DB2 for the Solaris Operating Environment and any remote IBM database server or any other DRDA-enabled database server.

IBM Client Access/400

The Client Access/400 Family contains all the necessary functions for connecting PC desktops to AS/400. Distinct products are provided for DOS, Microsoft Windows, and OS/2 operating systems. Functions are provided for communications, 5250 emulation, full network printing capability, file transfer to DB2/400, PC file serving, application enablers such as SQL and ODBC, and office integration.

Complementary Technologies for Complete Data Management Solutions

Key data management technologies must be complemented by related technologies for building a complete system for Information Technology business solutions. These key complementary technologies can be grouped as follows:

Transaction Management

Transaction management products provide such functions as security, integrity, and control for on-line and general interactive applications. Because many transaction management applications access databases, it is imperative that data management and transaction management products work efficiently together.

IBM IMS/ESA Transaction Manager (IMS TM)

IMS/ESA Transaction Manager is a high-performance, high-capacity, low-cost enterprise transaction server supporting database management systems, such as IMS DB or DB2. IMS TM:

- Provides Advanced Program-to-Program Communications (APPC) application enabling support with applications on a programmable workstation and other products that support APPC/MVS.
- Enhances system management.
- Supports large networks.
- Simplifies network administration.
- Manages and secures the IMS TM terminal network.
- Routes messages from terminal to terminal, from application to application, and between application programs and terminals.
- Queues input and output messages.
- Provides message integrity.
- Schedules application programs.

CICS

CICS is a transaction processor that provides high-performance client/server support on IBM and many non-IBM platforms.

CICS schedules and monitors execution of business application programs on behalf of tens to thousands of users. To handle the needs of these users in a real-time environment, it provides the facilities needed to protect against failure either of an individual transaction or failure of the system while transactions are busy.

An important feature of CICS is the common application programming interface (API) which allows applications to be ported easily from one platform to another to enable rightsizing and migration of systems. Function can be distributed across different elements in a network, allowing best use to be made of hardware investment.

CICS also provides:

- Support for multiple network protocols, including SNA and TCP/IP.
- Dependable, high integrity platform interoperability.
- Transparent use of applications across LANs and WANs.
- Support for broad range of industry-specific terminals and devices.
- Consistent, dependable performance.
- Support for DB2 for MVS, DB2 for VSE, DB2 for AIX/6000, DB2 for OS/400, DB2 for OS/2, DB2 for HP-UX and other RDBMSs and data systems.
- Concurrent sharing of data with complete integrity.

IBM IMS Client Server/2* (IMS CS/2*)

IMS CS/2 allows you to execute IMS TM transactions from a PC environment. It increases the productivity of application developers, enabling them to generate business solutions that fully utilize the capabilities of the PC. IMS CS/2 does this while protecting current investments in IMS TM 3270 applications.

IMS CS/2 connects the application programming power of PCs with the current IMS TM application base. Through PC functionality, IMS CS/2 enables application developers to build PC applications that use existing host data in new and different ways.

Tuxedo, and Encina

These transaction managers can be used with DB2 for AIX/6000 through the X/Open XA interface.

Networking

Networking provides the basis for management of data distribution, data access, and operation of a distributed database. Data Management solutions must be able to operate on all pervasive networking architectures and protocols.

Communication Protocols supported by IBM Database Products

A variety of protocols are supported for client-server communication; some examples are TCP/IP, NetBIOS, SNA (APPC/APPN*), and IPX/SPX. When connecting IBM and other vendors′ database products through DRDA, i.e. outside a LAN, the communication protocol supported is SNA (APPC/LU6.2).

Workgroup Solutions

Workgroup solutions consist of a collection of software offerings to facilitate the sharing of information among a group of workers. The goal is to align people with the business processes that they perform in their daily work activities by bringing together personal productivity tools, data (from wherever it resides), line of business applications, and groupware together in a single presentation at the user′s desktop.

Groupware Applications

Lotus Notes**

Lotus Notes is the world′s leading messaging and groupware product. Millions of people in thousands of companies are using Notes today to break down traditional barriers and create alliances that extend across departments and beyond the walls of the organization. So teams of employees, customers, suppliers and business partners have simultaneous access to the information they need, no matter where they′re located or what computer platform they are using.

Notes is a product of proven technologies, that works with your current computing environment. Lotus Notes is the only product family available today that combines three essential technologies: Notes Mail, the world′s most reliable and innovative client/server messaging system; a powerful distributed document database; and a rich application development environment. Together they have an impact on the way you do business that is far greater than the sum of their parts.

And Lotus Notes works with what you already have. It runs on all the most popular platforms and operating systems, including IBM OS/2 Warp, Microsoft Windows 3.1, Windows 95, Windows NT, Apple Macintosh, Novell** NetWare and UNIX platforms including IBM AXI, Sun Solaris, and HP-UX. Notes applications developed on any platform can be immediately deployed across all Notes platforms.

- Integrated client/server messaging.

  Notes integrated client/server messaging provides a single reliable infrastructure for all of your messaging and groupware applications. Using Notes Mail every Notes user can send and receive e-mail, and every Notes application is automatically mail-enabled. Users can easily exchange mail across a variety of mail systems, including cc:Mail** and X.400. Notes Public Key/Private Key Security, digital signatures, encryption, and directory synchronization ensure safe and reliable delivery. With Notes, you can build messaging and routing applications to automate the flow of work and information throughout your organization. Least cost routing, centralized SNMP management, an easy-to-use Administrative Control Panel, and Notes single copy object store (only one copy of any given message is stored on a single mail server) makes Notes less expensive to operate.

- A powerful document database.

  Once you build a Notes application, the Notes database becomes a container for your organization′s business knowledge - a central access point from which you can easily access, track, store and organize the shared knowledge and expertise of your people. Notes′ powerful database technology gives you the ability to assemble, manage and share compound documents created in Notes or integrated from popular desktop applications you already use, such as Lotus SmartSuite**, Microsoft Office, or Novell Perfect Office. You can also integrate information from World Wide Web pages, relational databases and legacy systems into Notes documents. Or link and embed objects, sound, video and other critical information directly into Notes documents.

  What′s more, Notes also includes facilities to make information easy to find. The Full Text Search feature lets you scan all or part of a database to find information. Users can even create their own custom views and folders to see the information in a way that′s most meaningful to them. Most important, all this information is available to those who need it in a secure environment. Lotus Notes offers multiple layers of security, so you can control who can access a particular Notes server, what level of access users should have to specific databases and even limit the access to individual documents. Additionally Notes′ digital signatures and encryption ensure that no one has falsified his or her Notes identity or gained unauthorized access to information. Importantly, Notes′ robust security model doesn′t upset the flow of work it′s virtually transparent to users.

In a system unique to Lotus, Notes databases on distributed servers are automatically copied and synchronized. All team members have access to the same information. Mobile users have the unique ability to maintain local replicas of Notes databases and work off-line. Once they reconnect to the server, Notes replication synchronizes the changes across the enterprise. Only those fields which have changed need to be replicated, minimizing communications costs. Notes applications are also distributed via replication, making new applications and enhancements easy to distribute.

- An unsurpassed environment for rapid application development and deployment.

  With Notes, customers can quickly develop and deploy custom business applications that are optimized for managing a rich variety of document-oriented information, but can also incorporate structured database information. Notes so accelerates the pace at which users can model, deploy, and validate applications that line-of- business professionals now work side-by-side with MIS staff to develop and prototype applications at many Notes sites. And now with LotusScript(TM), a BASIC-compatible structured programming language, Notes designers have all the cross-platform tools they need to build the next generation of business process applications.

Lotus SmartSuite

With SmartSuite, you get five of the hardest-working, easiest-to-use desktop applications in business today.

SmartSuite is packed with productivity tools to help you work smarter. Like Lotus SmartMasters, a series of easy-to-use, pre-designed templates that give you a head start on all kinds of everyday business tasks.

What′s more, SmartSuite is tightly integrated to help you perform even the trickiest cross-application tasks in no time. And with features like TeamConsolidate**, TeamReview** and TeamMail**, SmartSuite gives you some exciting new ways to collaborate with your colleagues, co-workers and associates whether you′re using it by itself or in tandem with Lotus Notes.

With SmartSuite, Team Computing truly becomes a reality. What′s Team Computing? It′s nothing more than what you′re already doing every day -- working together with others to get a job done. SmartSuite just makes it easier, in lots of different ways.

For example, with TeamReview, you can quickly distribute documents to team members for their comments -- and limit their editing rights if you wish. Then, with TeamConsolidate, you can easily put everyone′s thoughts together in one file. And with TeamMail, you can quickly communicate the results to everyone involved. For the ultimate in teamwork, SmartSuite works together seamlessly with Lotus Notes, the Groupware Standard**. Through Notes/FX**, you can turn individual SmartSuite applications into powerful workgroup systems that dramatically improve productivity.

Get The Numbers: Lotus 1-2-3 - Lotus 1-2-3 has everything you and your team need to make all your everyday spreedsheet tasks easier. It works harder so you don't have to.

- One-step charting. Create charts with titles, labels and legends -- or color-coded geographical maps -- in a single step.
- Range Routing. Sequentially route portions of spreadsheets to others for comments, then roll everyone's ideas back into your file.
- Version Manager**. The easy way to track worksheet changes. Save multiple sets of the same or various ranges of data in one file.

Get The Data: Lotus Approach - Lotus Approach is the high-powered database the whole team can use. It's the easy way to access, manage, analyze, and report on all your business information.

- Unprecedented ease of use. Create reports, forms, mailings and analyses fast.
- Revolutionary PowerClick** reporting. See how your report will appear in print as you create it. Sort, group or calculate data by clicking on SmartIcons**.
- Blazing speed and power. Enjoy unparalleled speed with Approach's exclusive X10 Technology. Join up to 50 database tables on one or more fields.

Get The Words: Lotus Word Pro - Lotus Word Pro is the first team word processor, offering revolutionary new ways to create, review and edit documents, making you and everyone else on your team more productive.

- Easy document collaboration. Review and consolidate edits in one file.
- Improved productivity. Navigate quickly through large documents. Automatically correct common typing mistakes.
- Superior compatibility. Import documents from other word processors, including WordPerfect**, without losing formatting. Built-in Internet support.

Get Visual: Lotus Freelance Graphics - Lotus Freelance Graphics is the fastest, easiest way to create and collaborate on plans, reports and proposals that double as high-impact presentations.

- Best start-to-finish process. Breakthrough technology for everything from generating new ideas, to consolidating input from co-workers, to publishing your finished product, on paper or online.
- Customizable content SmartMasters. Expert advice on structure and flow for 30 common business documents, from launch plans to sales proposals.
- Team Computing tools. Use TeamReview to collaborate with co-workers for the best finished product.

Get Organized: Lotus Organizer - Lotus Organizer is the personal and group scheduling standard. There's simply no better way to schedule your appointments and manage our workload.

- Stay organized in the office. Integrates Calendar, To Do list, Planner, Address book, Call manager, Notepad and Anniversary reminder.

- Schedule meetings fast. Set up meetings in record time with anyone via Notes, cc:Mail or with a single LAN server.
- And on the road. Lets mobile users merge laptop file updates with files on the office LAN.

Information Management

Information Management addresses the need for people to create, access, and manage information with applications such as document management, image management, and search and retrieval.

Visual Document Library*

Visual Document Library is an electronic document library and retrieval system for LANs. It can store, organize, and locate virtually any kind of electronic data files.

- takes full advantage of the OS/2 WorkPlace Shell and System Object Model.
- provides sophisticated search capabilities such as keyword, parametric, and contextual searches.
- offers leading edge function: version control, check in/out, access control and automatic archiving.
- includes major elements of DB2/2 and SearchManager/2*.

ImagePlus* VisualInfo*

ImagePlus VisualInfo is a multiple-object enterprise document management system. Multiple objects include images of scanned paper, text documents, spreadsheets, voice, video, etc. VisualInfo′s major components are image processing, basic work management, and library management.

- Scalable Library Server, from LAN to RISC to Host, through the use of the DB2 family.
- industrial strength by aligning with systems and network management facilities such as First Failure Support Technology* (FFST), Configuration, Installation, and Distribution (CID), and Systems Managed Storage (SMS).
- Line of Business Application integration via a rich set of APIs.

The other ImagePlus family members: ImagePlus/2, ImagePlus/400, and ImagePlus/ESA also utilize the DB2 product family.

Work Management

Work Management addresses the need for people to manage their work, whether structured or unstructured, through a series of recurrent or non-recurrent (ad hoc) events. Workflow applications automate business processes and ensure that tasks are routed and completed by the appropriate individuals. For example, on the local server, a series of aggregation reports could be automatically executed at the close of business every day. These reports could be automatically transmitted back to regional and headquarters servers.

FlowMark for OS/2

FlowMark for OS/2 provides a set of workflow management functions to help users precisely and consistently document the business processes, control the execution of business processes, and progressively improve the business processes.

- provides building blocks for integration with VisualInfo, Lotus Notes, VIM-enabled e-mail systems, FormTalk, and TaP/2.
- uses state of the art technologies: Object Oriented Database, C++, and IBM Class Library.
- improves processes through the use of modeling and animation.
- automatically invokes applications on the workstation or on the host.

Providing Full Solutions

Customers need comprehensive solutions to data management problems. Not only is IBM addressing the technical issues just described, we are addressing customer needs for full solutions, IBM is:

- Complementing data management software offerings with support and service offerings that provide a complete solution for an organization′s business problems.
- Fostering the development of cross-industry applications, vertical applications, and tools based on IBM data management products.
- Supporting the specialization of our field sales force and linking them to our development and research laboratories so as to build stronger customer partnerships.
- Increasing investment in both direct and indirect channels to facilitate learning about and acquiring IBM products.

Summary

Data Management solutions are important to gaining competitive advantage from using data as an informational asset. IBM investments in research, development, marketing, sales and support focus on providing customer-oriented data management solutions. IBM is pursuing strategies that position IBM as the leader in the platform-independent data management market and make IBM the vendor of choice for data management solutions.

Additional Information

Related White Papers

This white paper is one of a suite of Data Management papers available from IBM. The complete set consists of:

- *Client Server Solutions for IMS*
- *Data Management: IBM Strategies and Solutions*
- *Data Replication: the IBM Story*
- *Database Systems Management: IBM's Solution*
- *DataJoiner - A Multidatabase Server*
- *DB2 Relational Extenders*
- *Decision Support Solutions: IBM's Strategy*
- *High Performance Parallel Database Solutions*
- *IBM Information Warehouse Solution: A Data Warehouse - Plus!*
- *IBM's Data Mining Technology: ehp1.*
- *IMS and the Internet*
- *IMS in an Object Oriented World*
- *Into the Future with IMS*
- *Multi-dimensional Analysis: Extending the Information Warehouse Framework*
- *Object Support and the DB2 family*
- *Open IMS*

Solutions Directory

*For a list of complementary tools and applications available to work with IBM database products, please refer to the DB2 Client/Server Solutions Directory; this directory describes over 1300 applications and is available on CompuServe** in the IBMDB2 forum - please look for the DB2VEND file in the DB2/2 section.*

Also available is the AS/400 Client/Server Application Directory which includes over 1500 DB2/400 client/server applications organized by industry/market segment. The IBM publication number is G325-6202

Technical Report

In addition the following is available as IBM Santa Teresa Laboratory Technical Report: Object and Relational Databases: IBM's Strategy.(03.457)

Ordering Information

To obtain free copies of the DB2 Client/Server Solutions Directory, the AS/400 Client/Server Application Directory, the IBM white papers listed above, or the technical report on Object and Relational Databases: IBM's Strategy, readers in the U.S. and Canada may call 1-800-IBM-3333 and ask for extension STAR 706.

To send a fax requesting these materials, please fax your name, mailing address, and the titles of specific materials that you want to 1-408-463-2750. Please address your fax to "Data Management Information".

IBM employees may obtain these papers online through the following conference disks:

- *DB2INFO: using the command request* **name** *from db2info at torolab2*
- *MKTTOOLS: using the command tools sendto usdist mkttools mkttools get* **name**

Document	**name** on DB2INFO	**name** on MKTTOOLS
Client Server Solutions for IMS	imsspcs	imsspcs package
Data Management: IBM Strategies and Solutions	datasolu	datasolu *
Data Replication: The IBM Story	drs_wp	drs_wp *
Database Systems Management: IBM's Solution	dsm_wp	dsm_wp *
DataJoiner - A Multidatabase Server	dj1wp	dj1wp *
DB2 Relational Extenders	db2relex	db2relex *
Decision Support Solutions: IBM's Strategy	dssstrat	dssstrat *
High Performance Parallel Dabase Solutions	hpdbsolu	hpdbsolu *
IBM Information Warehouse Solution: A Data Warehouse - Plus!	iwsolu	iwsolu *
IBM's Data Mining Technology	datamine	datamine *
IMS in an Object Orinted World	imsoo	imsoo package
IMS and the Internet	imsin	imsin package
Into the Future with IMS	imsitf	imsitf package
Multi-dimensional Analysis: Extending the Information Warehouse Framework	mdaandiw	mdaandiw *
Object Support and the DB2 Family	db2obj	db2obj *
Open IMS	imsop	imsop package
Object and Relational Databases: IBM's Strategy		oostrat *
DB2 Client/Server Solutions Directory		db2vend package

Printed in U.S.A.