

The Big Picture: IBM DB2® Information Management Software and DB2® Universal Database™

(Reprinted courtesy of developerWorks DB2, at www.ibm.com/developerworks/db2)

[Jeff Jones](#)

IBM Data Management Solutions, Silicon Valley Laboratory
January 2003 (Updated February 2004)

Table of contents

Table of contents..... 2
Introduction 3
A brief history of DB2 3
DB2 database servers today 4
DB2 software and e-business on demand 4
A closer look at the technology 5
 Proven performance and scalability 5
 Administration (made easier) 6
 Application development and deployment for your chosen environment 6
DB2 technology is at the core of a wide variety of solutions 7
 Business intelligence..... 7
 Content and records management..... 8
 Federation and information integration 9
What a value! 10
Strong and varied partnerships 10
Summary 10
Resources 10
About the author..... 11

Introduction

In 1970, E. F. Codd of IBM Research published a paper that led to a new way for computers to manage information. His paper, "A Relational Model of Data for Large Shared Data Banks," proposed a new architecture for storing, managing and interacting with digital data. This new relational model freed application developers from having to know details about the data being managed.

Four years later, IBMers [Don Chamberlin](#) and Ray Boyce published "SEQUEL: A Structured English Query Language," which became the basis for the SQL language standard. Questions written in the new SQL language became more important than how the data was stored and organized on disk. New, more powerful questions could be asked and answered. Applications could be built much more quickly. The relational database system itself took on more of the burden of managing the data, leaving applications more freedom to focus on business logic.

Since 1970, IBM has developed a complete family of relational database management system (RDBMS) software now called DB2® Universal Database™ (DB2 UDB). In addition, IBM has built additional information management software with DB2 as its "engine" for purposes that include data warehousing, data analysis, data mining, media asset management, enterprise content management and information integration. DB2 UDB and the IBM DB2 Information Management Software portfolio represent one element of what IBM calls middleware, software that serves as the glue uniting systems and software applications. DB2 is one of five IBM Software brands: DB2, WebSphere®, Lotus®, Tivoli® and Rational. This article introduces DB2 UDB and the IBM DB2 Information Management portfolio.

A brief history of DB2

A series of research projects have been a steady source of technology for the DB2 family since the beginning:

- The System R project resulted in the first IBM implementation of the relational model. A project called ARIES delivered row-level locking technology used throughout the database industry today.
- Cost-based query optimization has been an area of intense effort and innovation ever since the System R days. The R Star project extended the relational model to distributed system environments.
- The Starburst project focused on making the relational model extensible to handle new forms of information and new kinds of optimization strategies.
- The Garlic project brought an emphasis on data federation, allowing data in diverse systems, not just DB2 systems, to be managed together.
- Most recently, a technical preview based on DB2 has demonstrated the integration of information from Web services and the use of XQuery as an additional and powerful query language for managing XML content.

The first implementation of relational technologies from the initial System R project was the database integrated into the System/38 server in 1980. In 1982, the SQL/DS™ product was delivered on the mainframe operating systems VM and VSE, also based on System R. DB2, formally called DATABASE 2, was born in 1983 on MVS™. The database manager in OS/2® Extended Edition in 1987 was the first relational database on distributed systems. SQL/400® for the new AS/400® server emerged in 1988. New DB2 editions were delivered on AIX® (1993), HP-UX and Solaris (1994), Windows® (1995), and Linux (1999).

DB2 database servers today

Today, the DB2 family spans a wide variety of UNIX®, Linux and Windows platforms and the IBM iSeries™ (OS/400® operating system) and zSeries™ (OS/390®, z/OS®, VM, VSE, and Linux) server lines. DB2 Everyplace™ supports handheld devices and embedded Linux environments and provides data synchronization with larger systems. Common tools have been delivered for application development and database administration across the family. Innovations from all family members, and the Informix database line acquired in 2001, feed the growth of the entire family.

DB2 technologies of today address emerging customer requirements in several new areas:

- Autonomic computing requires that servers, operating systems and middleware including DB2, diagnose and correct problems without human intervention. Database self-management and automation for the database administrator are areas of particular emphasis in the most recent edition of DB2.
- Standards-based [Web services](#) have emerged as a new style of application processing with full support from DB2.
- [Grid computing](#), or the idea of large-scale computing resources used as a utility or service, including database services, takes advantage of the vast clustered scalability of DB2 to support large databases and large numbers of simultaneous users in a highly available manner. Standards-based Web services are another key component of grid computing supported by DB2.
- The "[e-business on demand](#)" business model requires an operating environment built on open standards to allow quick and cost-effective innovation and reconfiguration. The infrastructure to support e-business on demand must be reliable, scalable and secure. DB2 is part of that infrastructure.

In addition to strong and innovative technology, DB2 provides high value to customers of all sizes. DB2 pricing on UNIX, Linux and Windows systems is recognized by industry analysts as being roughly half that of its main competitor. DBA automation and self-management enhancements combine with low price to provide remarkable value to DB2 customers.

The purpose of this article is to give you a broad brush overview of particular technology areas addressed by capabilities in DB2 and of the role of DB2 in the IBM e-business on demand strategy. Just as SQL took more of the burden of data management away from application developers in the beginning, the DB2 technologies described here work together to enable today's and tomorrow's application developers and database administrators to focus more and more on solving business problems. For them, this means increasing freedom from the mechanics of managing information. For businesses, this means increasing responsiveness to customer demands and market opportunities.

DB2 software and e-business on demand

The past 40 years of IT evolution have left most companies with an enterprise computing infrastructure that is heterogeneous, widely distributed and increasingly complex. With more digitized information collected in recent years than in the history of the civilized world, companies are being challenged to extend traditional data management techniques to new forms of content from many different sources.

The next phase of e-business -- e-business on demand -- will drive a new kind of business transformation to address these challenges. Companies want a flow of transactions, work, ideas and opportunities of all kinds to ripple immediately through the entire enterprise and beyond,

across the extended value chain on which every business depends. The underlying information infrastructure is central to making this transformation possible.

IBM DB2 software plays a critical role in this infrastructure -- the on demand operating environment. All elements of the portfolio (database servers, business intelligence software, enterprise content management software, data management tools, and information integration software) are developed with four key e-business on demand attributes in mind: integrated, open, virtualized, autonomic. Here are a number of DB2 capabilities supporting these attributes.

- **Integrated** -- Built-in support for both Microsoft and Java™-based operating environments; integration into WebSphere, Tivoli, Lotus and Rational products and plans; cross-platform DB2 family capabilities; integration with Web services and message queuing technologies; heterogeneous data source support via DB2 Information Integrator; support for both structured and unstructured information.
- **Open** -- Deep commitment to and support for Linux and standards for Java, XML, Web services, grid computing, and distributed database interoperability; multi-vendor multi-platform exploitation.
- **Virtualized** -- Federation and integration technologies in DB2 Universal Database and DB2 Information Integrator that provide a pragmatic alternative to data centralization; clustered scalability to support expansion of a virtualized information environment.
- **Autonomic** -- Self-tuning capabilities of DB2 Universal Database; rapid DB2 deployment via optimized configuration tooling; dynamic adjustment and tuning; simple and silent installation processes; integration with Tivoli® for system security and management.

A closer look at the technology

DB2 is designed to be powerful for those who need that power. However, there has been increased focus on usability and on ease of development. Let's take a look at the technologies that together make DB2 excel:

- Proven performance and scalability
- Administration (made easier)
- Application development and deployment for your chosen environment

Proven performance and scalability

To simultaneously meet the requirements for broad operating system support and for high performance and scalability, DB2 has been developed in ways specific to each environment.

- On [OS/390 and z/OS](#), DB2 is developed in concert with enhancements to the operating system and the server hardware. This tight integration led to the delivery of DB2 "data sharing," the shared-resource clustering architecture that exploits the IBM System/390® and zSeries Parallel Sysplex® hardware architecture. Some of the largest databases in the world are built on DB2 in this environment, as is noted in the periodic study of large databases done by [Winter Corporation](#).
- On [OS/400](#), the operating system for the IBM iSeries server line, formerly the AS/400, DB2 is implemented as part of the operating system itself, which supports single-server and multi-server parallel processing and clustering.
- On [UNIX, Linux and Windows platforms](#), DB2 has a "shared-nothing" architecture that enables a common code base to be used across these environments. Servers in a DB2 shared-nothing cluster work independently and in parallel on a subset of the overall data and on a subset of the SQL requests received by the cluster. On Linux, DB2 Integrated Cluster Environment (ICE) provides a building block approach for rapid deployment of

clustered DB2 systems, from 2 servers to 1,000. [Benchmark results](#) for both transactional (e.g., TPC-C) and decision support (e.g., TPC-H) workloads demonstrate the enormous scalability of DB2 with this portable architecture, in both clustered and non-clustered configurations.

Clustering for high availability and scalability, and support for the newest processors and interconnect technologies, are aspects of DB2 that ensure a smooth growth path for customers. DB2 failover and standby support provide the high availability required today. DB2 support for the latest 64-bit processors (Intel Itanium 2 and AMD Opteron) means larger databases can be built and faster performance can be achieved. DB2 on Linux supports all IBM server lines, including 64-bit POWER-based pSeries and iSeries servers. Simpler and faster clustering and connectivity technologies such as [InfiniBand](#) allow DB2 to scale easily. Smooth growth paths exist for customers and developers of all sizes.

Administration (made easier)

Economic conditions and the desire to improve the bottom line mean that many DBAs are becoming more and more overloaded as the amount and variety of information to be managed increases without the resources to hire additional administration support.

DB2 eases the burden of database administration in many ways:

- Its Control Center provides a central place for DBAs to perform their work across networks of DB2 systems.
- An array of advisor tools provides expert resource monitoring, problem diagnosis and corrective action. A recent example of this is the Configuration Advisor used to rapidly achieve peak DB2 performance for new installations on UNIX, Linux and Windows. Another is the Health Center, which serves as a centerpiece for much of the recent DB2 work on self-management. Its rules-based problem diagnosis and corrective action capabilities complement the [DB2 Performance Expert](#) and [DB2 Recovery Expert tools](#), an emerging class of [IBM database tools](#) providing more expert guidance and automatic action than previously possible.
- DB2 UDB Express, introduced in 2003, provides automatic installation and configuration for smaller businesses, on both Linux and Windows platforms with one or two processors.
- Continued advancements in cost-based optimization and automatic query rewrite technologies, there since the beginning of DB2, continue to remove the burden of DB2 performance management from the database administrator. The goal for each new version of DB2 is to require fewer and fewer database administration resources. DB2 benefits from the overall IBM focus on and investment in [autonomic computing](#).

Application development and deployment for your chosen environment

If you are an application developer, you have a wide variety of options for developing applications that use DB2 as the database server. The DB2 team has worked hard to make it easy to develop applications. Work efforts with the [IBM WebSphere Studio](#) product team, with the IBM Rational XDE team, with the Eclipse open source community, and with the [Microsoft Visual Studio](#) group have yielded plug-ins for DB2 application development and visual data modeling. DB2 also supports Microsoft Common Language Runtime (CLR). Through it, application developers can develop application logic in Visual Basic .NET or C# for both client- and server- sides. [Borland](#) development tools (Kylix, Delphi, C++Builder) support DB2 UDB as well. Finally, DB2 UDB itself provides a [Development Center](#) for building server-side pieces of the application, such as stored procedures and user-defined functions, in Java and SQL procedural language.

DB2 has deep Java roots. In late 1996, Java support was first provided in DB2. Stored procedures and user-defined functions could be built then in Java. JDBC, a communication mechanism between Java applications and database systems, was supported at that time. Since then, DB2 Java support has grown to include SQLJ for static SQL communication with DB2, and IBM has participated in the creation of JOLAP, a Java-based data analysis standard. Management tooling in Java makes Web-based database administration possible with DB2. Finally, DB2 supports J2EE as an application processing environment.

At the same time, DB2 developers work closely with Microsoft's Windows and .NET teams to ensure DB2 is a strong citizen of this application environment. IBM is committed to supporting both J2EE and .NET with DB2. The DB2 commitment to Windows is strong. Evidence of this includes DB2 support for Windows 2000 on the first day of its availability to customers, and membership in the Microsoft Gold Certified Partner Program for Software Products. To achieve Gold Certified status, DB2 was certified on three Windows 2000 server packages: Server, Advanced Server and Datacenter Server. In addition, DB2 provides a high-speed native interface to Microsoft OLE DB data sources. In 2003, IBM demonstrated DB2 integration with the Windows Server 2003 .NET platform, including Common Language Runtime and .NET Data provider support.

DB2 technology is at the core of a wide variety of solutions

A strong commitment to research and development means that DB2 is the core of a wide variety of data management products and solutions in the areas of:

- [Business intelligence](#)
- [Enterprise content and records management](#)
- [Federation and information integration](#)

Business intelligence

Business intelligence (BI) capabilities are built into the DB2 engine, and BI applications have DB2 at their core. Business intelligence tools span the areas of data warehousing, data analysis and data mining. The DB2 Data Warehouse Center provides an interface for defining, building and maintaining data warehouses. The [DB2 Warehouse Manager](#) also provides an Information Catalog for managing warehouse meta data, and tools for reporting and governing complex query execution.

Online analytic processing (OLAP) is possible with DB2 in two ways:

- Provided with DB2 are built-in functions for CUBE and ROLLUP, popular OLAP operations for exploring information in a database. DB2 also contains a library of statistical analysis functions and aggregation functions like [rolling sum and rolling average](#).
- IBM and Hyperion have worked together to create [DB2 OLAP Server](#)[™], a complete OLAP solution built on Hyperion Essbase analytics. The current version of DB2 OLAP Server, built on DB2 UDB, provides both multidimensional and relational data storage. Hybrid analysis, combining the speed of multidimensional storage and the scalability of relational storage, and automatic deviation detection (data mining) of data in an OLAP cube, are capabilities of the latest version of DB2 OLAP Server. IBM also partners with a variety of data analysis software vendors who have enabled their tools to work with DB2 databases.

- [DB2 Cube Views](#), delivered with [DB2 UDB Data Warehouse Editions](#), improves OLAP query performance and accelerates OLAP application deployment through the use of materialized query tables and dimensional metadata.

Another analysis tool called the [DB2 Spatial Extender](#), jointly developed with partner ESRI, extends the DB2 SQL language to understand concepts like distance between points on a map or relationships such as "inside" or "outside" a defined area.

Data mining represents the frontier of business intelligence capabilities. Data mining is the process of discovering patterns in data that cannot be found by other means. Techniques for clustering information based on important attributes and for predicting behaviors of customers based on prior behavior patterns are two examples of data mining. Since 1996, IBM has provided [DB2 Intelligent Miner](#). Its algorithms prepare and transform data in preparation for mining, perform mining operations, and visualize mining results. In 2001, these capabilities were implemented as extensions to DB2 in the form of "[scoring services](#)." Scoring services enable data mining to be performed in real time on small segments of data using SQL. Today, DB2 Intelligent Miner Modeling, Visualization and Scoring are optional features of DB2.

Several new features of DB2 UDB are designed to benefit customers using DB2 for business intelligence purposes. [Multi-dimensional clustering](#) keeps related information physically together on disk for faster retrieval. Materialized query tables provide dramatically faster performance for complex queries requiring information from a variety of diverse data sources at the same time. Null and default data compression technology reduces the disk storage requirements for data warehouse, as well as other forms of databases.

Content and records management

The future of data management involves managing and uniting all kinds of information, structured and unstructured, to solve business problems. DB2 has traditionally focused on managing structured data, meaning rows and columns of numbers and letters. Another part of the DB2 software portfolio focuses on managing "content," or unstructured information such as images and other multimedia information, word processing documents, and computer-generated reports. The products [DB2 Content Manager](#) and [DB2 Information Integrator for Content \(formerly Enterprise Information Portal\)](#) address customer requirements for content management solutions. They are built on DB2 UDB.

DB2 Content Manager provides support for two kinds of content management: media asset management, and enterprise content manager. Media asset management is the business of storing and managing large collections of large multimedia objects. Customers include art museums, university music libraries, and television broadcasters. Enterprise content management is the second kind of content management involving the management of large collections of smaller objects such as scanned check images, in the case of a bank, as well as bank statements, invoices, and reports.

DB2 Information Integrator for Content provides a programming layer above DB2 Content Manager and other data sources, structured or unstructured, for the purpose of accessing and searching across the sources using a common interface. For example, all information pertaining to a particular client can be retrieved, regardless of data type or document type. Also provided are Web crawling, workflow management and information mining services.

Development work with a partner Tarian Software, then the acquisition of Tarian, led to the creation of [DB2 Records Manager](#). This product adds electronic records retention and lifecycle management to the IBM content management portfolio.

The acquisition of Green Pasture Software led to the creation of [DB2 Document Manager](#) which provides extended document management services for DB2 Content Manager and Lotus Domino.doc.

Federation and information integration

At the center of the IBM DB2 Information Management software philosophy is the belief, backed by customer requirements, that integrating information in heterogeneous data environments is more important and provides quicker return on IT investments for many application than data centralization in single, large database systems. DB2 and related information management software reflects this belief in integration and federation.

DB2 software supports a wide variety of methods for accessing information remote to itself. These include ODBC and JDBC, SQLJ and OLE DB. Both .NET (Microsoft) and J2EE (Java) application environments are supported. Since 1995 with the delivery of DB2 DataJoiner®, IBM has provided optimized SQL access to information in non-DB2 databases such as those from Oracle, Microsoft and Sybase. DB2 applications can query information in both DB2 and non-DB2 databases using DB2 SQL. This ability to federate diverse relational databases has evolved from DB2 DataJoiner and DB2 Relational Connect into the new [DB2 Information Integrator](#).

Proper data placement and management of copies is another important aspect of information integration. To accomplish this, [data replication](#) technology is also provided with DB2. Both log-based change capture and refresh styles of replication are supported across the DB2 family. Combining the heterogeneous capabilities of DB2 and DB2 Information Integrator with DB2 data replication means non-DB2 databases can also be the replication target and/or source.

Information of all types can be managed by DB2 today. [DB2 Extenders](#)™ have been built to manage text, XML, image, audio, video and spatial information. These Extenders are the result of the evolution of DB2 from a purely relational system to an object-relational system. The universe of data sources available to DB2 applications has expanded in recent years to include WebSphere MQ [message queues](#) and [standards-based Web services](#). Data in file systems can be managed by DB2 as if it were stored in a DB2 table by way of [DB2 Data Links Manager](#), an optional feature of DB2. DB2 applications use SQL to manipulate data both inside and outside DB2 tables.

Increasing use of XML as a method to describe, organize, and interchange information has led to a variety of XML support enhancements in DB2. Today, more than 100 extensions to the SQL language have been implemented in DB2 to support the management of XML data. The [DB2 XML Extender](#), first given to customers in 1999, provides the foundation for native XML data management. Recent enhancements include automatic schema validation for XML documents composed from data in DB2 and automatic style transformation using XSLT. DB2 also supports the SQLX (or SQL XML) publishing functions and XPath expressions, and has demonstrated support for XQuery through a [public prototype](#) in early 2002. DB2 is on course to become a truly bilingual database, supporting both SQL and XQuery.

Many of these federation and integration technologies, and new software wrapping techniques, have come together in an IBM DB2 Information Management solution called [DiscoveryLink](#)® for the life sciences industry. DiscoveryLink enables life sciences applications to join information using SQL across dramatically diverse sources unique to the industry (e.g., genomics file data, toxicology spreadsheets, clinical trial and regulatory text, and assay results databases).

The breadth of data types, data sources and means of connectivity supported by DB2 and the SQL language uniquely qualify DB2 as an information integration engine. Combine this with the content management capabilities described earlier, and the magnitude of the IBM commitment to

helping customers integrate their information, no matter the type, the quantity or the location, becomes clear.

What a value!

As important as the technological strengths illustrated above is the value that DB2 provides customers. DB2 is priced to challenge the competition at all levels. DB2 pricing is based largely on a per-processor model for simplicity and clarity. In high-availability settings, DB2 alone is priced for only one processor on idle standby servers performing no active DB2 work. Combine price considerations with increased self-management capabilities and rich function, and a strong total cost of ownership (TCO) story emerges. The five-year TCO advantages of DB2 over its competitors are documented in several industry analyst reports.

In February 2003, IBM announced [DB2 Express](#) (on Windows and Linux servers with one or two processors) for business partners and for customers with 100 to 1,000 employees. Its transparent installation capability, enhanced self-management features, and lower total cost of ownership further emphasize the TCO advantages of DB2.

Strong and varied partnerships

Partners are recognizing the technological strength and the high value of DB2. Business application partners including SAP, Siebel, PeopleSoft, i2 and J. D. Edwards have in various ways chosen to standardize their applications on DB2. In many cases they now lead in sales situations with their applications on DB2. Because IBM focuses solely on providing middleware and not application software, our partners are not threatened with competition from IBM as they partner with IBM. IBM DB2 Information Management software runs comprehensive partnership programs to attract and retain tool, application and business partners of all kinds.

Our partners also include IBM Systems Group and the other IBM Software brands (WebSphere, Tivoli, Lotus and Rational). DB2 developers work with Systems Group teams during DB2 development and benchmark testing. WebSphere integrates DB2 to manage information on applications, databases, users and other resources under its control. An edition of WebSphere Commerce Analyzer contains DB2 Intelligent Miner technology. Together DB2 and WebSphere Application Server deliver support for standards-based Web services. Lotus plans to extend the scalability of Notes® and Domino™ through the integration of DB2. DB2 is a Tivoli-managed system resource. These are examples of the increasing integration across the IBM middleware and server portfolios resulting in faster deployment of high-performance e-business applications.

Summary

From structured data to unstructured content, from handheld devices to clustered server configurations, and from transaction processing workloads to data mining, DB2 and the IBM DB2 Information Management software portfolio supports customers as they grow and succeed as on demand businesses. With IBM DB2 Information Management software, customers can take advantage of e-business on demand advancements and strategies such as Web services, autonomic computing and grid computing. The continued growth of DB2 software in the database market is a strong indicator of the continuing success of IBM DB2 Information Management investments and strategies on behalf of IBM customers around the world.

Resources

A good summary of the functions and features of DB2 for Linux, UNIX, and Windows is available at:

- <http://ibm.com/software/data/db2/udb/features.html>

The following book is a little outdated, but is still the best-written introduction to DB2:

- Chamberlin, Don. *A Complete Guide to DB2 Universal Database*. Morgan Kaufmann Publishers August, 1998.

The latest news on Total Cost of Ownership:

- DH Brown Associates. *IBM DB2 Universal Database v8.1 vs. Oracle 9iR2: Total Cost of Ownership*. November, 2002. At <http://www-3.ibm.com/software/data/pubs/pdfs/dhbrown.pdf>
- Market Magic Research. *Database Comparative Cost of Ownership*. January, 2003. At <http://www-3.ibm.com/software/data/highlights/mmr-db-cco.pdf>

Interesting historical perspective of the beginning of SQL:

- McJones, Paul, ed. *The 1995 SQL Reunion: People, Projects, and Politics*. SRC Technical note 1997-018, August, 1997. At http://www.mcjones.org/System_R/SQL_Reunion_95/sqlr95.html
- *Bibliography of the System R project* at http://www.mcjones.org/System_R/bib.html
- From DB2 Magazine, a 20-year [timeline of DB2](#)

Covers almost everything in DB2 7.2 on Linux, UNIX, and Windows:

- Melnyk, Roman B, and Paul C. Zikopoulos. *DB2: The Complete Reference*. McGraw-Hill Osborne Media, October, 2001.

A good overview of the mainframe DB2:

- *An Introduction to DB2 for OS/390 V7* at <http://publib.boulder.ibm.com/cgi-bin/bookmgr/BOOKS/dsnith12/CCONTENTS>

DB2 Developer Domain, a Web site of technical resources for DB2 software, at:

- <http://www.ibm.com/software/data/developer>

Journals:

- *DB2 Magazine* at <http://www.db2mag.com>
- *IDUG Solutions Journal* at <http://www.idug.org/journal/index.cfm>

Home base for developing skills:

- <http://www-3.ibm.com/software/data/education.html>

About the author



Jeff Jones is senior program manager in IBM DB2 Information Management Software, the home of DB2 Universal Database. He serves as a spokesperson for IBM's information management software portfolio to the customer, press, and industry analyst communities in addition to serving on the Editorial Board of Directors for DB2 Magazine. He has 23 years of technical, staff and managerial experience with IBM. Prior to his current position, Jeff was a brand manager for IBM's data replication and database systems management middleware. He has also served in the DB2 for OS/390 product group, in IBM Research Division on experimental engineering graphics and analysis projects, and on engineering and logistics application projects for the former IBM Storage Systems division. Jeff is located at IBM Silicon Valley Laboratory in San Jose, California.

AIX, AS/400, DataJoiner, DB2, DB2 Universal Database, DB2 Extenders, DiscoveryLink, IBM, iSeries, Lotus, Lotus Notes, Lotus Domino, MVS, Tivoli, WebSphere, OS/2, OS/390, SQL/DS, SQL/400, zSeries, z/OS, and OS/390 are trademarks or registered trademarks of IBM Corporation in the United States, other countries, or both.

Microsoft and Windows are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, and service names may be trademarks or service marks of others.