


# DB2 History

A short history of DB2 for z/OS  
in buttons and images for the  
DB2Night Show  
Monday 25 April, 2011

This presentation shows much of the history of DB2, with milestones of new versions and new inventions. It also shows some of the people who contributed key inventions and innovations in DB2. It uses many images and mementoes to show the process and progress in DBMS.

- Understand some of the key inventions in DB2
- Note the teams and some key inventors for the technology
- See the cross-product design and implementation
- Provide insight into the past and direction for DB2
- Show some of the reasons for the success of DB2




This presentation shows some of the history of DB2, with milestones of new versions and new inventions. It also shows some of the people who contributed key inventions and innovations in DB2.

Note inventions across products: data sharing, locking, compression, design, standards, consistency across family, convergence

Need to add: C. Mohan, Yun Wang, Josephine Cheng, Jim Teng, Jeff Josten, Tim Vincent, Matt Huras, Curt Cotner, Berni Schiefer, Mark Anderson, Bruce Lindsay

Also get picture of larger groups of developers.



IBM's history of innovation to meet the needs of human endeavors.

- 1) To support sending man to the moon, invented hierarchical database to run on System 360
- 2) As business found more ways to make their businesses better with computers, IBM invented the relational db and introduced DB2 to make computers more than just a technical experiment, but a true business tool
- 3) The world wide web make the world smaller and information more accessible. IBM introduced the first DB with SOAP based web services support to help business use the web for a real business advantage
- 4) As IT grew in importance and usage, it also grew in complexity. To help customers reduce the manpower cost of managing their DB's. IBM was a leader in advanced autonomic capabilities.
- 5) And finally, as the world become yet smaller with globalization, where customers and supplier on the other side of the world are no different from the customers in your back yard, IBM introduced the first hybrid XML and relational database... leverage the XML standard that's quickly becoming the language of business.

All of these innovations were meant to help our customers do what they needed to do better... better for their customers, and better against their competitors. Looking forward, businesses are facing a ever changing set of new pressures... changes in the marketplace require businesses to run faster, leaner and more flexibly to survive.

Globalization of trade, production and skills is leading to an explosion of new emerging opportunities but also competition from every corner of the world. Natural resources are becoming scarce. There are new threats to national security. There are also great advancements in technology that business will need to adopt to compete. Scientific management is adding precision and rigor to how business are automated and optimized. The next wave of the web is truly delivering on the promises that seems like science fiction just a decade ago.

With all these economic, competitive, and technological pressures... business need to focus on establishing a firm competitive advantage to survive and flourish. But how will they find that competitive advantage?

# The Evolution of Relational Database Technology

## IBM DB2 Timeline

1970	1980	1990-1996	1997-2000
<ul style="list-style-type: none"> <li>IBM's Don Chamberlin and Morton Armstrong's paper "Implementation of a Structured English Query Language," describes the first implementation of SQL as part of the System R work.</li> <li>IBM's Jim Gray publishes "Terminology of Locks and Degree of Consistency in a Shared Data Base," describing the formal definition of a database transaction and degree of consistency.</li> <li>The IBM System R team publishes "System R: Relational Approach to Database Management," a paper describing their relational database prototype.</li> <li>System R prototype is installed at three IBM customer locations.</li> <li>IBM's Pat Selinger writes about industry's first relational query optimizer in her paper, "Access Path Selection in a Relational Database Management System."</li> </ul>	<ul style="list-style-type: none"> <li>In 1968, IBM ships IMS V1 – industry's first hierarchical database management system.</li> <li>IBM's E. F. Codd publishes the industry's first paper on relational database technology entitled "A Relational Model of Data for Large Shared Data Banks."</li> <li>System R project is established at IBM Research to build a relational database management system.</li> <li>IBM's Don Chamberlin and Ray Boyce publish the paper, "RELQUE: A Structured English Query Language," which later becomes the basis for the SQL standard definition.</li> <li>IBM delivers S/36 with an integrated database server. The database technology is based on IBM's System R research prototype. The API is consistent with the SQL S/32 API to support application migration.</li> </ul>	<ul style="list-style-type: none"> <li>IBM extends DB2 V1 to non-IBM UNIX platforms and delivers DB2 Common Server Technology on HP-UX and Sun-Solaris platforms.</li> <li>IBM ships DB2 Common Server Technology V2, the industry's first object relational web enabled database, for OS/2, AIX, HP, Sun, SCO, S/390, and Windows NT platforms.</li> <li>IBM announces support for multimedia data types, including Text, Image, Video and Audio.</li> <li>IBM ships DB2 Parallel Edition V1 for the RS/6000 SP2 platform. This is the birth of IBM's clustered, highly-scalable architecture, with focus on large data warehouses on AIX for complex query work.</li> </ul>	<ul style="list-style-type: none"> <li>IBM ships its Web-ready DB2 for OS/390 V5, the industry's only database that supports up to 64,000 concurrent users and hundreds of terabytes of information.</li> <li>IBM ships DB2 Universal Database – the database for e-Business.</li> <li>At the Winter Olympic Games in Nagano, the DB2-powered Web site handles 100,000 hits per minute and processes 8.2 terbytes of data – more information than is stored in the U.S. Library of Congress. The site is recognized by the Guinness Book of World Records for setting two world records in Internet traffic.</li> <li>IBM announces DB2 Universal Database V5.2 with Data Link and built-in OLAP technology (CUBE and ROLLUP functions). Data Link enables DB2 Universal Database users to manage data in files outside of IBM databases.</li> <li>One year after its launch, DB2 Universal Database has been adopted by nearly two million users worldwide. Some 2,000 industry software vendors are porting applications to DB2 Universal Database, adding to the more than 4,300 applications already developed.</li> <li>IBM announces beta program for DB2 Universal Database for the Linux operating system, industry's first database to support Linux.</li> <li>IBM shatters the Windows NT scalability barrier with the industry's first-ever one terabyte TPC-D benchmark on Windows NT.</li> <li>IBM announces a host of additions and enhancements to the award-winning DB2 family: DB2 Universal Database Version 6.1 with improved support for e-Business and business intelligence, DB2 Universal Database for OS/390 and AIX/600, DB2 Universal Database Satellite Edition for mobile computing, and DB2 Everywhere for pervasive computing.</li> <li>IBM announces beta program for DB2 Universal Database for Linux – industry's first database to support the new operating system.</li> <li>IBM announces DB2 Universal Database Version 7, world's first fully integrated e-business database. It also introduced a first-of-its-kind pricing structure, allowing application service providers (ASPs) to purchase DB2 Universal Database at a minimal investment.</li> <li>DB2 Universal Database delivered a gold medal performance at the 2000 Summer Olympics, it handled 8.3 billion hits for 17 days with the highest traffic reaching 1.2 million hits per minute – more hits than the Nagano Olympics.</li> <li>IBM announces the general availability of IMS version 7.</li> <li>IBM announces Extended Software Leadership on Linux including the availability of DB2 Universal Database for Linux on Intel-based clusters – the only commercial Database software available for Linux on the mainframe.</li> </ul>

This is a relational database timeline used in IBM.

**20 YEARS OF INNOVATION DB2 1983-2003**

**1983** IBM announces the first release of DB2, a relational database management system for the System/370 mainframe.

**1984** IBM introduces DB2/VS, a version of DB2 for the VS (Virtual Storage) architecture.

**1985** IBM releases DB2/PC, a version of DB2 for the PC (Personal Computer) platform.

**1986** IBM introduces DB2/AS/400, a version of DB2 for the AS/400 (Advanced System/390) platform.

**1987** IBM releases DB2/ESA, a version of DB2 for the ESA (Enterprise System Architecture) platform.

**1988** IBM introduces DB2/OS/2, a version of DB2 for the OS/2 (Operating System/2) platform.

**1989** IBM releases DB2/NT, a version of DB2 for the NT (New Technology) platform.

**1990** IBM releases DB2/DB2, a version of DB2 for the DB2 (Database/2) platform.

**1991** IBM releases DB2/DB2, a version of DB2 for the DB2 (Database/2) platform.

**1992** IBM releases DB2/DB2, a version of DB2 for the DB2 (Database/2) platform.

**1993** IBM releases DB2/DB2, a version of DB2 for the DB2 (Database/2) platform.

**1994** IBM releases DB2/DB2, a version of DB2 for the DB2 (Database/2) platform.

**1995** IBM releases DB2/DB2, a version of DB2 for the DB2 (Database/2) platform.

**1996** IBM releases DB2/DB2, a version of DB2 for the DB2 (Database/2) platform.

**1997** IBM releases DB2/DB2, a version of DB2 for the DB2 (Database/2) platform.

**1998** IBM releases DB2/DB2, a version of DB2 for the DB2 (Database/2) platform.

**1999** IBM releases DB2/DB2, a version of DB2 for the DB2 (Database/2) platform.

**2000** IBM releases DB2/DB2, a version of DB2 for the DB2 (Database/2) platform.

**2001** IBM releases DB2/DB2, a version of DB2 for the DB2 (Database/2) platform.

**2002** IBM releases DB2/DB2, a version of DB2 for the DB2 (Database/2) platform.

**2003** IBM releases DB2/DB2, a version of DB2 for the DB2 (Database/2) platform.

**DB2 From the User's View... Online, Anytime!**


**www.iQugora.com**

**International Group of Users (IGU)**

**IBM**

This is a poster commissioned by IDUG and provided in DB2 Magazine, as well as a pdf on the web.


**Dr. Ted Codd  
invented the  
Relational Model in 1969  
published in CACM 1970**



# Dr. Don Chamberlin invented SQL



# Dr. Jim Gray



- Leader in System R and DB2 locking and transaction processing
- Missing at sea Jan. 28, 2007

# Query Optimization

## Dr. Pat Selinger



1975



2005

# Query Optimization

## Dr. Pat Selinger



1975



2005



[http://www.mcjones.org/System\\_R/SQL\\_Reunion\\_95/SRC-1997-018.pdf](http://www.mcjones.org/System_R/SQL_Reunion_95/SRC-1997-018.pdf)

<http://it.toolbox.com/blogs/db2luw/what-is-the-halloween-problem-in-databases-12618>

# A short history of DB2 for z/OS

## In buttons and images



**Don Haderle**



Here we are at 25 years, looking forward and building for the next 25. Don, the mother of DB2 and our fearless leader for formative years. The old story notes that success has many fathers, while failure is an orphan. Many people have claimed to be the father of DB2. We could take one or two more steps back. It would be useful to start with IMS and the heritage it brings, then Dr. E.F. Codd, and the relational database journey. Include Mike Stonebraker, the Informix, Illustra history, and RedBrick too.

## Many more

- C. Mohan, Yun Wang, Josephine Cheng, Jim Teng, Jeff Josten, Tim Vincent, Matt Huras, Curt Cotner, Berni Schiefer, Mark Anderson, Bruce Lindsay, Roger Reinsch, Bob Jackson, Akira Shibamiya, ...


C. Mohan, Yun Wang, Josephine Cheng, Jim Teng, Jeff Josten, Tim Vincent, Matt Huras, Curt Cotner, Berni Schiefer, Mark Anderson, Bruce Lindsay



There have been dozens of major changes and probably hundred of minor changes in the database world over the last forty years.

Having moved from a basic file storage system (albeit somewhat sophisticated ones), database systems are now called on to do more and more. Simply storing the data provides no value to a company – it’s only when data is used to provide information to give a company a competitive edge that the true value of a data management system is realised.

This chart outlines some of the massive changes that have occurred in the Data Management sphere over the past 40 years. One can only imagine what can be expected over the next 10, 20 or 30 years (who had ever heard of ‘Autonomic Computing’ 40 years ago, or even only 5 years ago?). We can justifiably look forward to the future with some anticipation and excitement.




Over the years we have seen greater emphasis and responsibility being placed on the database management system, and consequently more responsibility on those database professionals charged with supporting them! The early databases were effectively file systems (albeit very sophisticated ones) whose prime role was to store the data and make it accessible to the applications that were doing most of the work.

With the advent of the Relational Database system and DB2, it became possible for more and more of the processing that was traditionally done in the application to be 'pushed' down to the database, and we have seen this trend continue to expand as more functionality has continued to be added to DB2 over the years. While increasing the productivity of application developers, it has meant that the database professionals who support the database systems now need to be more cognisant of the application and the business requirements it is seeking to satisfy. An example of such functionality was database maintained referential integrity, introduced in DB2 on the mainframe in Version 2 in 1988.

As the database systems have continued to evolve and become even more sophisticated, we have seen them take even more of the onus of keeping a company's IT systems running and operational. Certain 'SMARTs' are built into databases like DB2, and there is a high degree of integration with other database systems. Database professionals are finding that they need to know more and more about previously esoteric (at least to them) topics like .....

But we're getting ahead of ourselves.....

# Let's Take a Step Back...



## System R: An architectural overview

by M. W. Blasgen, M. M. Astrahan, D. D. Chamberlin, J. N. Gray, W. F. King, B. G. Lindsay, R. A. Lorie, J. W. Mehl, T. G. Price, G. R. Putzolu, M. Schkolnick, P. G. Selinger, D. R. Slutz, H. R. Strong, I. L. Traiger, B. W. Wade, R. A. Yost

## An Access Specification Language for a Relational DataBase System

R Lorie  
J Nilsson

In the early days of mainframe computing, file systems of varying flavours were generally acceptable to store a company's information. These were useful for storing data and making it available to applications that would process it. However, as data volumes and complexity started to increase, it became apparent that basic file systems would not be sufficient to provide the appropriate level of support.

As it was for so many industries, the Manned Space Program proved to be the push to develop a more rigid data management system. Put simply, a better tool needed to be found that was able to store and manage the inventory of 3.6 million parts that made up the Apollo Saturn V rocket. Thus was born the IMS Database system, whose hierarchical structure lent itself to the massive Bill of Materials application that storing the inventory for the Apollo spacecraft had become. Over time, this same structure had applicability to hundreds of commercial business requirements, such as banking systems and airline reservation systems and continues to do so today.

However, the hierarchical structure of the IMS system was also its drawback, as far as application developers were concerned. In order to develop applications, developers needed to understand the structure of the hierarchy, and needed to know where in the hierarchy the data resided. It has served the business community for more than 30 years, and was borne out of a specific need. By contrast, the relational model has been built on the foundation of the experience gained from IMS, as well as an academic mathematical model. Extracts from the IBM Systems Journals from the time outline the basic tenet of the relational model – that data is logically stored in tables comprised of rows and columns, and the user does not need to know the location of the data. The user only needs to worry about **what** data is needed, the system needs to worry about **where** the data is and **how** to access it. This was a much simpler model than the hierarchical model, as most people can understand the concept of data being stored in tables (at least anyone who can read a bus timetable or read a phone book). The era of the relational database was nigh.....

## Let's Take a Step Back...



### System R: An architectural overview

by M. W. Blasgen, M. M. Astrahan, D. D. Chamberlin, J. N. Gray, W. F. King, B. G. Lindsay, R. A. Lorie, J. W. Mehl, T. G. Price, G. R. Putzolu, M. Schkolnick, P. G. Selinger, D. R. Slutz, H. R. Strong, I. L. Traiger, B. W. Wade, R. A. Yost



### An Access Specification Language for a Relational DataBase System

R Lorie  
J Nilsson

In the early days of mainframe computing, file systems of varying flavours were generally acceptable to store a company's information. These were useful for storing data and making it available to applications that would process it. However, as data volumes and complexity started to increase, it became apparent that basic file systems would not be sufficient to provide the appropriate level of support.

As it was for so many industries, the Manned Space Program proved to be the push to develop a more rigid data management system. Put simply, a better tool needed to be found that was able to store and manage the inventory of 3.6 million parts that made up the Apollo Saturn V rocket. Thus was born the IMS Database system, whose hierarchical structure lent itself to the massive Bill of Materials application that storing the inventory for the Apollo spacecraft had become. Over time, this same structure had applicability to hundreds of commercial business requirements, such as banking systems and airline reservation systems and continues to do so today.

However, the hierarchical structure of the IMS system was also its drawback, as far as application developers were concerned. In order to develop applications, developers needed to understand the structure of the hierarchy, and needed to know where in the hierarchy the data resided. It has served the business community for more than 30 years, and was borne out of a specific need. By contrast, the relational model has been built on the foundation of the experience gained from IMS, as well as an academic mathematical model. Extracts from the IBM Systems Journals from the time outline the basic tenet of the relational model – that data is logically stored in tables comprised of rows and columns, and the user does not need to know the location of the data. The user only needs to worry about **what** data is needed, the system needs to worry about **where** the data is and **how** to access it. This was a much simpler model than the hierarchical model, as most people can understand the concept of data being stored in tables (at least anyone who can read a bus timetable or read a phone book). The era of the relational database was nigh.....

# Announcement of DB2 for MVS, Version 1.1

- o Relational Data Model
- o Structured Query Language (SQL)
- o Continuous Operations
- o Concurrent Access to Data
- o Application Programming
- o DB2 Interactive (DB2I)
- o Large Processor Support

**Announced: June 7, 1983**

**General Availability:  
April 2, 1985**

**V1R2 March 7, 1986**

**V1R3 June 26, 1987**



Following the academic work and mathematical modelling that led to the relational model, IBM announced its first Relational DataBase Management System (RDBMS) in 1983 to great fanfare. As well as introducing the new-fangled Relational Model, it provided a whole new database language in Structured Query Language (SQL - which, in spite of its name, provided much more than just the ability to **query** the data) and features pertinent to uninterrupted access to data.

There was an entire volume of the Systems Journal dedicated to DB2 in 1984, from which many of these images are taken. The names on this particular article (i.e. Haderle) will be familiar to many of you.

Some of you (either those with greying hair, or those – like myself - that started their IT careers when they were about 7 years old) may recall some of the mementoes of the time. The brass #1 was distributed to customers and IBM Support staff who took part in the Early Support Program for DB2 V1.

# DB2 Version 2

- o Relational Enhancements
  - System Enforcement of Referential Integrity
- o Performance Enhancements
  - Extended Use of MVS/XA and MVS/ESA Facilities
  - Significant Improvements in Transaction Processing
  - Advances in Query Optimization & Access Path Selection
  - Faster Sorting of Large Data Volumes
- o Operational Enhancements
  - Flexibility of System control and Authorization
  - DB2 Governor to Limit CPU Resource Consumed by a Dynamic Query
  - Selective Audit Trail Capability
  - New and More Flexible Data Storage Organization Technique
  - Enhancement to Utilities Across the Entire Recovery Scenario
- o Host Language Support
  - Support of IBM C Language
  - Support of VS COBOL II Release 3

V2 R1 Sept 23, 1988    R2 Sept 22, 1989    R3 Oct 25, 1992

Over the next few slides, I'll just touch on the main new features and functions that have come out with each new version of DB2 (in the interests of time, I'll skip the point releases, although I'm sure many of you recall those features. After all, some of those point releases contained a **lot** of new functions. (Who can forget the functionality included in DB2 Version 2.3?)

During this time, IBM introduced a number of peculiarly database terms, some of which have remained over the years, and some of which have been changed to reflect community standards. Has anyone ever used the word **SPUFI** in a game of Scrabble and lived to tell the tale? And who ever used the term **SARGABLE** in general conversation?

## DB2 Version 2



V2 R1 Sept 23, 1988    R2 Sept 22, 1989    R3 Oct 25, 1992

Over the next few slides, I'll just touch on the main new features and functions that have come out with each new version of DB2 (in the interests of time, I'll skip the point releases, although I'm sure many of you recall those features. After all, some of those point releases contained a **lot** of new functions. (Who can forget the functionality included in DB2 Version 2.3?)

During this time, IBM introduced a number of peculiarly database terms, some of which have remained over the years, and some of which have been changed to reflect community standards. Has anyone ever used the word **SPUFI** in a game of Scrabble and lived to tell the tale? And who ever used the term **SARGABLE** in general conversation?

# Version 3

- o Availability Improvements
  - o Partition Independence
  - o Utility Improvements
  - o DISPLAY and ALTER BUFFERPOOL Commands
  - o Data Compression
  - o CONCURRENT COPY Function of DFSMS
- o Connectivity Enhancements
  - o Multi-site Update of Distributed Data
  - o Increased Maximum Concurrent Threads to 10000 (2000 active)
  - o DPROP Release 2 Support
- o Performance Improvements
  - o Query I/O Parallelism
  - o Hiperpool Support
  - o CREATE INDEX With DEFER Option
- o Systems Integration
  - o Optical Storage Devices

**V3 December 17, 1993**

And on to DB2 Version 3. As new functions were added, DB2 professionals found themselves more and more having to become familiar with concepts and systems that they had previously not needed to know about. Over the course of several releases in DB2 V2, DB2 professionals not only had to get their heads around system controlled referential integrity, but with the advent of distributed databases, they also had to become familiar with communications protocols, DRDA, SNA etc. With DB2 Version 3, closer integration to storage subsystems was introduced, and DB2 professionals found themselves having to become more familiar with I/O systems, DFSMS, optical storage, hardware-assisted data compression and the like. The upside was that they were able to use this new-found knowledge to impress members of the opposite sex at cocktail parties which DB2 professionals regularly attend!

Fortunately, at about this time, IBM also realised that clever, esoteric terms were actually turning people off, and sought to redress this issue.

## Version 3



V3 December 17, 1993

And on to DB2 Version 3. As new functions were added, DB2 professionals found themselves more and more having to become familiar with concepts and systems that they had previously not needed to know about. Over the course of several releases in DB2 V2, DB2 professionals not only had to get their heads around system controlled referential integrity, but with the advent of distributed databases, they also had to become familiar with communications protocols, DRDA, SNA etc. With DB2 Version 3, closer integration to storage subsystems was introduced, and DB2 professionals found themselves having to become more familiar with I/O systems, DFSMS, optical storage, hardware-assisted data compression and the like. The upside was that they were able to use this new-found knowledge to impress members of the opposite sex at cocktail parties which DB2 professionals regularly attend!

Fortunately, at about this time, IBM also realised that clever, esoteric terms were actually turning people off, and sought to redress this issue.

# Version 4

- Availability
  - DB2 data sharing for incremental processing growth and concurrent access to the same data at all times.
  - Locking enhancements and new indexes.
- Client/Server Support
  - Stored procedures for improved client/server performance and application design.
  - Support for up to 25,000 concurrent client threads per DB2 server and up to 800,000 concurrent threads in a data sharing group.
- Performance Improvements
  - Query parallelism for processing-intensive queries, noticeably improving complex queries on very large tables.
  - Faster utilities
- Usability Improvements
  - Outer join to retrieve unmatched rows in a JOIN operation.
  - Table check constraints for domain definition by users to enforce business rules (part of the ANSI/ISO SQL92 standard).
  - Thread cancellation for any DB2 allied thread.

**V4 November 30, 1995**

And on to DB2 Version 4. This was the version that introduced Type 2 indexes (clearly IBM's Chief Imaginative Feature Namer had departed by this time), stored procedures, row-level locking, mind-blowing scalability, significant performance improvements and – the biggie – DB2 Data Sharing across a Parallel Sysplex. This version was a quantum leap in DB2's evolution, and while it contained many features that DB2 professionals had been wanting for a long time, it also increased the onus on what these professionals had to know.

## Version 4



And on to DB2 Version 4. This was the version that introduced Type 2 indexes (clearly IBM's Chief Imaginative Feature Namer had departed by this time), stored procedures, row-level locking, mind-blowing scalability, significant performance improvements and – the biggie – DB2 Data Sharing across a Parallel Sysplex. This version was a quantum leap in DB2's evolution, and while it contained many features that DB2 professionals had been wanting for a long time, it also increased the onus on what these professionals had to know.

# Version 5

- Availability for Data Warehouse and Operational Applications
  - o Enhances complex queries with reduced processing time
  - o Delivers data very large tables, up to a terabyte in size
  - o Delivers on-line reorganization of indexes and table spaces
  - o Enhances utilities with reduced processing time
- Client/Server Applications
  - o Enhances client/server applications with reduced communication and processing times
  - o Increases connectivity options for client applications
  - o Delivers support for direct TCP/IP connectivity
  - o Delivers Call Level Interface for ODBC and X/Open's CLI
  - o Enhanced stored procedures support
  - o Increases application portability
- User Productivity for All Applications
  - o Global temporary tables
  - o Enables use of object-oriented languages for application development

**Version 5    June 27, 1997**

And to DB2 Version 5. So much new function, so much for the DB2 professional to absorb. This version of DB2 emphasized the common approach that IBM was taking for its DB2 systems across all platforms, mainframe and otherwise. As a result, DB2 professionals who had been supporting DB2 on the mainframe some years found themselves having to converse with their colleagues who supported DB2 on other platforms.

In fact, this release had so much new function, that IBM turned it into a whole new version. If anyone still has one of these frisbees, keep it, because it is truly a collector's item. The story behind this frisbee needs 2 hours and a box of tissues to tell, and is probably not appropriate to tell here.

## Version 5 1997

- Availability for D
- o Enhances com
  - o Delivers data v
  - o Delivers on-line
  - o Enhances utilit
- Client/Server Ap
- o Enhances client
- times
- o Increases conn
  - o Delivers suppo
  - o Delivers Call L
  - o Enhanced stor
  - o Increases appl
- User Productivit
- o Global tempora
  - o Enables use of



processing

Version 5 June 27, 1997

And to DB2 Version 5. SO much new function, so much for the DB2 professional to absorb. This version of DB2 emphasized the common approach that IBM was taking for its DB2 systems across all platforms, mainframe and otherwise. As a result, DB2 professionals who had been supporting DB2 on the mainframe some years found themselves having to converse with their colleagues who supported DB2 on other platforms.

In fact, this release had so much new function, that IBM turned it into a whole new version. If anyone still has one of these frisbees, keep it, because it is truly a collector's item. The story behind this frisbee needs 2 hours and a box of tissues to tell, and is probably not appropriate to tell here.

# DB2 Versions 6 & 7

## DB2 UDB for OS/390 V6

- Capacity improvements
  - 16 Terabyte tables(!)
  - More open datasets
- User Productivity improvements
- Network computing improvements
- Object-Relational Extensions
- Performance and Availability improvements
  - Rebalancing partitions
  - Alter varying length columns
    - Faster recovery
    - Data Sharing enhancements

**V6 June 25, 1999**

**V7 March 30, 2001**

## DB2 UDB for OS/390 and zOS V7

- Application enhancements
  - Scrollable cursors
  - SQL enhancements
  - Row expressions
- Scalability enhancements
  - Parallel load
  - UNLOAD utility
  - Improved parallelism
- Availability enhancements
  - Dynamic alter ZPARMS
  - Cancel thread without rollback
- Manageability enhancements
  - Statistics history
  - More DBADM authority
- Accessibility enhancements
  - Improved DB2 Connect
  - JDBC and ODBC
- Extensibility enhancements
  - XML extender
  - Unicode

DB2 V6 brought the concept of a **Universal** database to the mainframe, by supporting many functions and data types previously only supported on the distributed platform. It was this version that brought the Universal moniker to the mainframe, effectively putting the **U** in UDB. At this point, our DB2 professionals found themselves having to learn a whole new set of data types. You mean pictures, text and movies can be stored in DB2 tables?! Ya gotta be kidding!

These concepts were expanded with the release of DB2 UDB V7 in 2001. As you can see, there is ever more for the DB2 professional to learn.

# DB2 Versions 6 & 7

DB2 UDB for OS/390 V6

DB2 UDB for OS/390 and zOS V7



**V6** June 25, 1999

**V7** March 30, 2001

- Improved DB2 Connect
- JDBC and ODBC
- Extensibility enhancements
  - XML extender
  - Unicode

DB2 V6 brought the concept of a **Universal** database to the mainframe, by supporting many functions and data types previously only supported on the distributed platform. It was this version that brought the Universal moniker to the mainframe, effectively putting the **U** in UDB. At this point, our DB2 professionals found themselves having to learn a whole new set of data types. You mean pictures, text and movies can be stored in DB2 tables?! Ya gotta be kidding!

These concepts were expanded with the release of DB2 UDB V7 in 2001. As you can see, there is ever more for the DB2 professional to learn.

## Greatest Hits: DB2 for z/OS V8

- ✓ Continuous availability
- ✓ Scalability and very large database
- ✓ Java and the web
- ✓ Queries and data warehouses
- ✓ Migrating and porting applications
- ✓ Application packages



V8 March 26, 2004

**Greatest Hit 3: Java and the web:** Improvements in SQLJ and JDBC support, a new Java Universal Driver, enhanced Unicode support, integration with WebSphere and new XML functions make Java and web applications more robust and more productive.

**Greatest Hit 4: Queries and data warehouses:** Optimization changes provide the best performance improvement opportunities in V8. Faster response and reduced processing time come from improved optimization and better information for the optimizer. New database design options for indexes, clustering and materialized query tables provide more gains. Warehouses often need to have the new rotate partition capability.

**Greatest Hit 5: Migrating or porting applications from other platforms:** Many SQL enhancements provide better compatibility with the DB2 family and with the industry. If customers develop on Windows, Unix or Linux, and then move to z/OS, the process is much easier. Early customers reported success at porting applications.

**Greatest Hit 6: Application packages: SAP, PeopleSoft, Siebel, etc. ...** About 50 improvements, including everything mentioned in the "Greatest Hits" section, are provided for most of the key vendor packages. SAP R/3 4.6 and PeopleSoft PeopleTools 8.45 certified for V8 just two months after general availability.

## DB2 9 for z/OS

- Business needs
  - Reduce CPU time
  - Improve business agility
  - Service Oriented Architecture
- Application developers need
  - Powerful new SQL enhancements
  - Portability with SQL and data definition compatibility
  - PureXML for a powerful SQL and XML interface to XML data
- Database Administrators need
  - Improve availability and performance
  - More flexible security and easier regulatory compliance
  - Better web application & data warehouse function and performance
  - LOB function, performance, usability

June 16, 2007




DB2 9 has a lot for everyone. Here are just a few of the highlights.

The business needs include CPU cycle reductions that deliver in most utilities, improved query optimization, improved business agility via faster implementation cycles, and new pureXML™ that builds a strong foundation for SOA and XML initiatives. Kevin Campbell, an Application Architect at Univar USA said it better than I can, “This is not a bolt-on or band-aid approach, DB2 9 for z/OS is XML without compromise.”


Database Administrators (DBAs) need improved database availability and performance including LOBs, reorganization, backup and recovery, and partitioning enhancements. DBAs also get more flexible trusted network context and role-based security to help with regulatory compliance. A wide range of enhancements improve ERP application and data warehouse functionality and performance. Large object (LOB) function is added with file reference variables and REORG, while performance is improved.

Application developers are most excited by PureXML, which adds a powerful SQL and XML interface to access XML data stored in a native format. Application developers need powerful new SQL enhancements including MERGE and TRUNCATE statements, INTERSECT and EXCEPT set operations, and spatial support for geographical data. Text handling is improved with the XML changes, many new built-in functions, and an upcoming text server. Improved SQL and data definition compatibility with other DB2 platforms makes porting much easier.




These next two charts illustrate the “big picture” of the versions of DB2 announced on the mainframe, and year these releases were made available. I’ve restricted myself to just the version releases in this time, although I’m sure many of you will recall the huge functional enhancements shipped in some of IBM’s point releases. (Who could forget the amount of new function shipped in in DB2 **V2.3**, for example?)

As you can see, DB2 has been vibrant since it was first shipped, with new functions and features, and enhanced performance being shipped so often. Of course, these functions are there for a purpose, and that is to support the ever-increasing complexities and requirements of the applications. At the end of the day, it’s the DB2 professionals charged with implementing and maintaining these systems that need to maintain their DB2 skills.



Part 2 of previous chart. This might bring back some memories for some of you.



DB2 for z/OS V7 became generally available (GA) March 2001, and V8 delivered three years later. DB2 9 became generally available in March 2007, three more years. We expect the next version will be 2.5 to 3 years from V9 GA to DB2 10 or DB2 X or whatever the name becomes.

The themes for future versions will continue to focus on core platform strengths of performance, scalability, reliability, stability, availability, resilience, and security. PureXML and Schema evolution or data definition on demand will be ongoing for a long time. In contrast, most of the 64 bit evolution should be completed in DB2 X.

The key interfaces for customers and vendors expand for both XML and for SQL. Information is a key leg of the SOA platform, and DB2 for z/OS provides many advantages for data management in SOA.

Standards, interoperability, portability and security along with secure access using the latest technologies are key touch points. Productivity improvements for application developers and for database administrators are very important as data grows in scale and complexity.

## IBM Data Servers – DB2 Linux, UNIX, Windows, z/OS, IDS Application Development & Administration Client


- API support
  - Java programming JLinQ
  - Open Source (PHP, Python, Ruby, ...)
  - .NET, EGL, & other technologies
- SOA / Web Services / Web Access
  - Data Server Web Services
- AD tooling
  - .NET, DWB, RAD
- Administration tooling
  - Rational Data Architect (RDA)
  - IBM Data Servers Administration Console



This slide shows the scope and mission for IBM's Common Application Development and Administration organization. A wide range of administration and application development function delivers for IBM relational database: Informix Dynamic Server, DB2 for Linux, UNIX and Windows, DB2 Connect and DB2 for z/OS. The new IBM Data Servers organization is changing the face of IBM relational database as it provides the client for Informix and DB2 for Linux, UNIX, Windows, i5/OS and z/OS.

The new IBM Data Servers offering consolidates database administration and application development (APIs, developer tools) across IBM's relational databases. You will see names of some components changing to ones which include IBM Data Server, for instance IBM Data Server Administration Console and IBM Data Server Developer Workbench. Other Data Server deliveries include the Client, Runtime Client, and Drivers for ODBC, CLI, .NET, JDBC, SQLJ, Ruby, PHP, Perl, and Python. Application Development APIs strategy includes new support for PHP and Ruby on Rails, dramatic improvements in Java with JLinQ and keeping Microsoft support very current. The next generation of SOA web services comes from Data Server.

Recent improvements in application development tooling for both DB2 and Rational make the application life cycle more productive with better quality. The new Web-based administration console project provides a much improved, more productive way to administer IBM relational databases.




Slide objective: set up for onion peel. First we give you a comprehensive look at the total Consul portfolio

Points:

1. Consul provides the depth of security management on the mainframe, right side of the diagram and the breadth across the enterprise with comprehensive compliance management from the distributed environment on the left side of the diagram
2. Distributed portfolio is InSight
  1. Focus is on access monitoring, log management and compliance reporting
3. Mainframe portfolio is zSecure
  1. Focus is on integrated mainframe audit, monitoring, compliance and administration
  2. z/OS
  3. RACF
4. The intersection is the overlap and illustrates how Consul integrates mainframe data into InSight to provide the comprehensive compliance management view across the enterprise

# But Don't Panic

"So I started off having to understand tables and rows, and now you tell me I need to know how much ??"



So over the last 20 years, the typical DB2 professional has had to learn and keep up to date with so much more than just the standard 'run of the mill' database technology. We've seen how his/her job has extended a variety of areas, and how, at the very least, the DB2 professional's new best friends have become the network coordinator, the business analyst, the application developer, the storage supervisor, the Unix and Windows operators, and especially the OS/390 systems programmer.

Over the last 10 years or so, open systems (aka. Distributed Platforms, or Multiplatforms) have become more prominent in running applications to run a company's business. Along with these, DB2 has also developed on these platforms, and is used by more customers and manages more data than ever.

In many cases, companies that have traditionally run applications on DB2 on mainframes have expanded their use to DB2 on distributed platforms, and in some cases, have even migrated completely off the mainframe to the distributed platform. In other cases, companies run a mixture of mainframe and distributed system applications. The common point here is that many of the DB2 professionals who have maintained DB2 on the mainframe for years are often now finding themselves in the position where they need to maintain DB2 on a Windows or Unix (and increasingly Linux) platform. Even if the two platforms are maintained by different groups, both groups of support staff (i.e. those on the mainframe side and those on the distributed side) are finding that they need to have an understanding of how the other side works.


# DB2 on OS/2, UNIX, Windows, Linux, ...

DB2 has been available on desktop and servers in various guises since 1989, starting with OS/2 Data Base Manager (DBM) .....




DB2 has been available on the desktop in various guises since 1989, starting with Data Base Manager (DBM).....

# DB2 on Distributed Platforms



Contrary to popular belief, DB2 is not just available on the mainframe. Since the late 1980s, DB2 has made its presence felt on the open system platforms. This chart gives a brief overview of the history of DB2 on distributed platforms (or open systems). This culminated in the announcement of DB2 **Universal** Database on the next slide.



DB2 Universal Database (or DB2 UDB) was a culmination of years of development and experiences gained with earlier versions of DB2. It was, in effect, a ‘merging’ of DB2 Common Server (rich in function) and DB2 Parallel Edition (designed for scalability and high performance), as well as the addition of significant new function. It set the direction for which IBM would be applying its development resources and established the “6 Universals” that future RDBMSs from IBM would adhere to, namely:

- Universal Applicability (being able to run a mixture of different types of applications)
- Universal Availability
- Universal Scalability
- Universal Manageability (being able to manage the system easily)
- Universal Accessibility (being able to access the data easily)
- Universal Extensibility (support for non-traditional data types such as pictures, video, audio, etc.)

# DB2 UDB on UWOL

(or a Lesson in Nested Acronyms)

## DB2 UDB Version 6

- Capacity and performance
- Usability
- DB2 Satellite Edition
- DB2 for Linux
- Application Development

## DB2 UDB Version 7

- Integrated Warehousing
- Integrated OLAP
- Federated databases
- Native OLE DB client
- XML extender
- Spatial extender
- Abstract Data Types
- SQL stored procedures

This chart just outlines some of the features that came with DB2 UDB for Unix/Windows/OS/2/Linux versions 6 and 7.

# DB2 UDB on UWOL

(or a Lesson in Nested Acronyms)

DB2 UDB Version 6

DB2 UDB Version 7



This chart just outlines some of the features that came with DB2 UDB for Unix/Windows/OS/2/Linux versions 6 and 7.

# DB2 UDB for Linux, UNIX & Windows V8 "Stinger"

## **Increase Programmer Productivity**

- Microsoft.NET Integration
- Java Enhancements
- SQL Support Extensions



## **Reduce Deployment & Management Costs**

- Autonomic Computing
- Design Advisor
- Self Tuning Backup & Restore

## **Provide a Robust Infrastructure**

- Performance, Scale, Availability, Security
  - High Availability Disaster Recovery
  - Automated Client Re-route
  - Linux Scalability

## **Extend the Value of Information**

- Replication Enhancements
- Geodetic Extender
- Mobility on Demand

*and more...*

By coming to the conference you have access to the Stinger technical preview. In order to take advantage of this make sure you visit the demonstration pavilion. There will be a special technology preview education track. Visit to start on your journey, register and pick up your CDs at the booth.

# Key Business Value DB2 9 Iuw

- *Greater business insight faster*
- *Manage risk and streamline compliance*
- *Resilient, low cost operation*
- *Lower cost, optimized infrastructure*
- *Faster, lower cost development*
- *Reduce information complexity*

## ■ SOA / XML

- Faster, Simpler Access to your XML data
- Eliminate external XML processes
- Put DB2 Viper at the core of your SOA strategy

## ■ SAP

- Optimized BI application performance
- Single step deployment
- Automated management

## ■ Storage

- Compression significantly lowers cost
- Storage groups reduce DBA effort
- Auto-resize tablespaces simplifies storage management
- Range Partitioning increases performance, reduces maintenance
- DB2 can handle your largest databases

## ■ Simplicity

- Significant Ease-of-Use Enhancements
- Faster development & Open standards
- Autonomic Object Maintenance
- Automatic Statistics Collection
- Self-Tuning Resource Management

## ■ Security

- LBAC Security / LDAP Authentication
- New security administrator role

## ■ Performance and Savings

- Continued commitment to leading performance
- Continued commitment to reduce cost
- Unbundled features = buy only what you need

**28 Years...**



**DB2 Means  
More Than  
Ever  
Before...**

	DB2 Information Integrator
	DB2 Cube Views
	DB2 Express
	IDS V9.4, Red Brick V6.2
03	DB2 V8 for z/OS
	DB2 e-Records Manager
	DB2 Data Grid Prototype
02	DB2 V8 Autonomic Computing
	Informix Acquisition
	DB2 Web Services
	DB2 Content Manager
	IMS & DB2 Tools
01	DiscoveryLink
	Integrated OLAP
	Integrated Mining
00	DB2 Web Integration
	DB2 for Linux
	Digital Library
	DB2 Universal Database
	DB2 Parallel Edition
90s	DB2 for Unix, Windows
	ImagePlus
80s	DB2 for Mainframes
70s	Invented Relational
60s	IMS

**1983**

Over the last hour, I've just touched on the changes that we've seen in the DB2 world over the last 20 years, and hopefully given you a taste of the myriad of aspects that comprise the life of a DB2 Professional in the 21<sup>st</sup> century. Life is a lot more complex these days than just worrying about tables, columns and rows, but consequently, it's a lot more interesting and exciting too. The

# Beginnings ...



**1983:** IBM gives DATABASE 2 (DB2) for MVS (once known internally as Eagle) its wings, shipping V1



**1988:** IBM announces and ships SQL/400 on the new AS/400 server with integrated RDBMS. The International DB2 Users Group (IDUG) is formed.



**1993:** IBM announces and ships DB2 for OS/2 V1 and DB2 for AIX V1

For a full, comprehensive DB2 timeline, visit: [http://wiki.ibmdatasemag.com/index.php/DB2\\_History\\_-\\_A\\_Timeline](http://wiki.ibmdatasemag.com/index.php/DB2_History_-_A_Timeline)

For the full timeline: [http://wiki.ibmdatasemag.com/index.php/DB2\\_History\\_-\\_A\\_Timeline](http://wiki.ibmdatasemag.com/index.php/DB2_History_-_A_Timeline)

## A little bit of press...



### Happy Birthday DB2! Happy Birthday Business Intelligence!

On June 7, 2008, IBM celebrated the 25th anniversary of Database 2 (DB2). Released to the market in 1983, DB2 has become one of the leading relational databases for data warehousing. DB2 was an outgrowth of research into relational data modeling by E.F. Codd of IBM Research — work that led to several related commercial innovations that have had a significant influence on the data warehousing market, among them the SQL language.

-Dan Vessett, IDC Link

**IDC LINK**

[http://www-306.ibm.com/software/data/db2/25th-birthday/?wm=7115001f1779&cm\\_sp=ZZ999-\\_-SWB00-\\_-1779&weight=30.0](http://www-306.ibm.com/software/data/db2/25th-birthday/?wm=7115001f1779&cm_sp=ZZ999-_-SWB00-_-1779&weight=30.0)

## A little bit more...

### IBM DB2 Turns 25!

On June 7, 1983, IBM issued a press release announcing a "relational data base management system [RDBMS] for large enterprises", and so began the life of Database 2, now simply called DB2. Tomorrow we celebrate the 25th anniversary of DB2. Here are some of the key happenings that led not only to the birth of DB2 but to the formation of the RDBMS industry as a whole.

*-from the 25<sup>th</sup> Anniversary release*



[View the original DB2 press release from 1983!](#)

[View the 25<sup>th</sup> Anniversary release](#)

1983 Press Release: [http://download.boulder.ibm.com/ibmdl/pub/software/data/sw-library/db2/DB2\\_Press\\_Release\\_6-7-83.pdf](http://download.boulder.ibm.com/ibmdl/pub/software/data/sw-library/db2/DB2_Press_Release_6-7-83.pdf)

25<sup>th</sup> release: [http://www-306.ibm.com/software/data/db2/25th-birthday/?wm=7115001f1779&cm\\_sp=ZZ999-\\_-SWB00-\\_-1779&weight=30.0](http://www-306.ibm.com/software/data/db2/25th-birthday/?wm=7115001f1779&cm_sp=ZZ999-_-SWB00-_-1779&weight=30.0)

## Did you know?



- **IBM DB2 is at the core of business applications in...**
  - 25 of the Top 25 Worldwide Banks
  - 9 of the Top 10 Global Life/Health Insurance Providers
  - 23 of the Top 25 US Retailers
- **In 20 years, DB2 for z/OS has had fewer than 5 critical security patches**
  - Oracle had to ship 89 in 1 qtr (October 19th, 2005), ... 51 more October 17, 2007, 26 more January 15, 2008, 41 more April 15, 2008, 45 more July 15, 2008, ...

[http://www.computerworld.com/action/article.do?command=viewArticleBasic&articleId=9042942&source=rss\\_topic17](http://www.computerworld.com/action/article.do?command=viewArticleBasic&articleId=9042942&source=rss_topic17)

### eWeek Oracle Issues 41 Security Fixes in Latest CPU

By Brian Prince 2008-04-15

January's CPU featured 26 security fixes for Oracle products. The next CPU is slated to be released July 15.

<http://www.eweek.com/c/a/Database/Oracle-Issues-41-Security-Fixes-in-Latest-CPU/>

<http://www.eweek.com/c/a/Security/Oracle-Ships-45-Security-Fixes-in-Latest-Update/?kc=EWKNLDATE071708FEA1>

### Oracle Patches 45 Security Holes in Database, App Server and More

July 15, 2008

# DB2 Support and Social Networking

DB2 proves to be a tool vital to the day-to-day work of businesses around the world. Support groups and social networks are abundant in support of DB2.



facebook

developerWorks.

channelDB2

DB2 videos, podcasts, blogs, photos, discussions, etc.



Facebook: <http://www.facebook.com/pages/IBM-DB2/10442975871>

IDUG: <http://www.idug.org/wps/portal/idug>

Channel DB2: <http://channeldb2.ning.com/>

DB2.community: [http://www.db2-community.org/DB2\\_Certification\\_-\\_General\\_Informa.db2-cert-gen.0.html](http://www.db2-community.org/DB2_Certification_-_General_Informa.db2-cert-gen.0.html)

IBM Database Magazine:

<http://www.ibmdatasemag.com/newsletter/070129.shtml>

Developer Works: <http://www.ibm.com/developerworks/wikis/display/db2xml/Home>

# DB2 hits the Blogosphere



"IBM puts the shizzel into the Data Warehouse!"  
- Vincent McBurney's Blog,  
Consultant, Solution Architect

## Bloggers from inside and outside of IBM dish on DB2

- [An Expert's Guide to DB2 by Chris Eaton](#)
- [DB2 Portal Blog by Craig Mullins](#)
- [Getting the Most Out of DB2 for z/OS and System z by Willie Favero](#)
- FreeDB2 by Leon Katsnelson  
<http://FreeDB2.com>
- [DB2usa: Blog About DB2 for z/OS](#)

An Expert's Guide to DB2: <http://blogs.ittoolbox.com/database/technology>

DB2 Portal Blog: <http://www.db2portal.com/blog.html>

Getting the Most Out of DB2 for z/OS and System z:  
<http://blogs.ittoolbox.com/database/db2zos>

FreeDB2 by Leon Katsnelson: <http://FreeDB2.com>

DB2usa: Blog About DB2 for z/OS: <http://db2usa.blogspot.com/>

## Come and Join Our Party!

IBM DB2 celebrates 25 years as a provider of high performance, reliable and secure data servers.



### Happy Birthday DB2 from Dan Vasset at IDC

"In 1983, the launch of DB2 from IBM marked the beginning of data warehousing and business intelligence as we know it today"



"DB2 was the first database that enabled businesses to understand their business"

"Ecommerce would be impossible without the birth of DB2 "



"the traces of Web 2.0 can be attributed to DB2"

"DB2 established credibility for real-world enterprise data management software"



### Dan Vasset at IDC

Program Vice President of IDC's Business Analytics Research

## Some DB2 Product Managers

### z/OS

Ed Altman  
Bob Jolls  
Don Sallan  
Eric Miles  
Sam Kahn  
Marilyn Bohl  
Bob Jackson  
Roger Reinsch  
Janet Perna  
Gary Ferdinand  
Diane Jablonski  
Al Zollar  
Carl Chamberlin  
Dan Wardman  
Beth Smith  
Fiona Gleeson  
Greg Lotko  
Rick Bowers



*Curt Cothner  
IBM Fellow, VP  
CTO for  
Database Servers*



### Linux, UNIX, Windows

Janet Perna



Herschel Harris

Bob Picciano

Peg Williams



Sal Vella



Here is my list of DB2 for MVS, DB2 for OS/390 and DB2 for z/OS product managers, sometimes for all database, sometimes more than one in DB2 across the history so far: Ed Altman, Bob Jolls, Don Sallan, Eric Miles

Sam Kahn, Marilyn Bohl, Bob Jackson, Roger Reinsch, Janet Perna, Gary Ferdinand, Diane Jablonski, Al Zollar, Carl Chamberlin, Dan Wardman, Beth Smith, Fiona Gleeson, Greg Lotko, and TBD, with pictures of some of the more recent ones.

# System R Web Pointers

[http://www.mcjones.org/System\\_R/](http://www.mcjones.org/System_R/)  
SQL Reunion conference paper  
[http://www.mcjones.org/System\\_R/SQL\\_Reunion\\_95/](http://www.mcjones.org/System_R/SQL_Reunion_95/)  
<http://infolab.stanford.edu/~widom/cs346/selinger.pdf>  
<http://www.cs.washington.edu/homes/lazowska/cra/database.html>  
<http://homes.cerias.purdue.edu/~bhargav/cs526/p242-griffiths.pdf>  
<http://redbook.cs.berkeley.edu/redbook3/oldeditions.htm>  
<http://www.cs.cmu.edu/~natassa/courses/15-823/F02/papers/p378-mohan.pdf>  
<http://user.it.uu.se/~torer/publ/Mimer-PG.pdf>  
<http://www.scs.stanford.edu/nyu/01fa/sched/recovery.pdf>  
<http://portal.acm.org/citation.cfm?doid=320455.320457>  
<http://www.informatik.uni-trier.de/~ley/db/journals/tods/AstrahanBCEGGKLMPTWW76.html>  
[http://en.wikipedia.org/wiki/System\\_R](http://en.wikipedia.org/wiki/System_R)  
<http://www.cs.berkeley.edu/~brewer/cs262/SystemR.pdf>

These are a few pointers that I found looking quickly.

# DB2 History Web Pointers

[http://i.cmpnet.com/db2mag/db\\_area/archives/2003/q1/AnvPostrTimelineSide.pdf](http://i.cmpnet.com/db2mag/db_area/archives/2003/q1/AnvPostrTimelineSide.pdf)  
[ftp://ftp.software.ibm.com/software/data/db2zos/IDUGDBfortomorrow\\_2007dist.pdf](ftp://ftp.software.ibm.com/software/data/db2zos/IDUGDBfortomorrow_2007dist.pdf)  
<ftp://ftp.software.ibm.com/software/data/db2zos/IDUGathenskeynote2007Nov.pdf>  
[http://www.mcjones.org/System\\_R/](http://www.mcjones.org/System_R/)  
early relational, SQL Reunion conference paper  
[http://www.mcjones.org/System\\_R/SQL\\_Reunion\\_95/](http://www.mcjones.org/System_R/SQL_Reunion_95/)  
DB2 for LUW history <http://www.informit.com/articles/article.aspx?p=375536>  
[http://en.wikipedia.org/wiki/IBM\\_DB2](http://en.wikipedia.org/wiki/IBM_DB2)  
<http://www.ibm.com/developerworks/db2/library/techarticle/0301jones/0301jones.html>  
<http://mainframe-watch-belgium.blogspot.com/2008/01/history-of-db2.html>  
<http://blogs.ittoolbox.com/database/db2zos/archives/db2-history-101-version-1-22046>  
<http://blogs.ittoolbox.com/database/db2zos/archives/db2-history-101-version-12-22083>  
<http://blogs.ittoolbox.com/database/db2zos/archives/db2-history-101-version-13-22097>  
<http://idealprogrammer.com/databases/db2/what-is-db2-history-editions-and-video-illustrations/>  
<http://www.databaseblog.net/?p=24>  
<http://ezinearticles.com/?What-is-DB2?-History,-Editions,-and-Video-Illustrations&id=704386>  
<http://www.ibmssystemsmag.com/mainframe/septemberoctober05/administrator/10021p1.aspx>

These are a few pointers that I found looking quickly. With the 25<sup>th</sup> anniversary coming up in June 2008, there will be more to come.

# DB2 History Web Pointers

The business value of DB2 chapter 2 “The evolution of DB2 for z//OS”

<http://www.redbooks.ibm.com/abstracts/sq246763.html>

## Presentations on the web

- <http://www.eweek.com/c/a/Database/IBM-DB2-Turns-25-and-Prepares-for-New-Life/>
- <ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/overview/database-keynote-idug-eu-2007-selinger.pdf>
- [database-keynote-idug-na-2007-selinger-haderle.pdf](ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/overview/database-keynote-idug-na-2007-selinger-haderle.pdf)
- [db2-history-2008-haderle.pdf](ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/overview/db2-history-2008-haderle.pdf)

These are a few pointers that I found looking quickly. With the 25<sup>th</sup> anniversary coming up in June 2008, there will be more to come.

# DB2 History Web Pointers

E.F. Codd's paper on relational database, "A Relational Model of Data for Large Shared Data Banks"

<http://www.seas.upenn.edu/~zives/03f/cis550/codd.pdf>

Original Chamberlin and Boyce paper on SQL publish "SEQUEL: A Structured English Query Language,"

<http://www.almaden.ibm.com/cs/people/chamberlin/sequel-1974.pdf>

DB2 history wiki

[http://wiki.ibmdatasemag.com/index.php/DB2\\_History\\_-\\_A\\_Timeline](http://wiki.ibmdatasemag.com/index.php/DB2_History_-_A_Timeline)



These are a few pointers that I found looking quickly. With the 25<sup>th</sup> anniversary coming up in June 2008, there will be more to come.

# Bibliography

- 1. Codd, E. F. 1970. A relational model of data for large shared data banks. Communications of the ACM 13(6): 377-387.
- 2. Astrahan, M. M., et al. 1976. System R: Relational approach to database management. ACM Transactions on Database Systems 1(2): 97-137.
- 3. Stonebraker, M. 1976. The design and implementation of Ingres. ACM Transactions on Database Systems 1(3): 189-222.
- 4. Bernstein, P., et al. 1998. The Asilomar report on database research. ACM SIGMOD Record 27(4).  
<http://www.sigmod.org/record/issues/9812/asilomar.html>.
- 5. Chaudhuri, S., and Weikum, G. 2000. Rethinking database system architecture: Towards a self-tuning RISC-style database system. The VLDB Journal: 1-10. <http://www.vldb.org/conf/2000/P001.pdf>.
- 6. Stonebraker, M., and Cetintemel, U. 2005. One size fits all: An idea whose time has come and gone. Proceedings of the 2005 International Conference on Data Engineering (April). [http://www.cs.brown.edu/~ugur/fits\\_all.pdf](http://www.cs.brown.edu/~ugur/fits_all.pdf).
- 7. Broussard, F. 2004. Worldwide IT asset management software forecast and analysis, 2002-2007. IDC Doc. #30277.  
<http://www.idc.com/getdoc.jsp?containerId=30277&pid=35178981>.
- 8. Gray, J., and Reuter, A. 1993. Transaction Processing: Concepts and Technologies, 397-402. San Mateo, CA: Morgan Kaufman Publishers.
- 9. Computer history museum RDBMS through 1983  
<http://www.computerhistory.org/collections/accession/102658267>