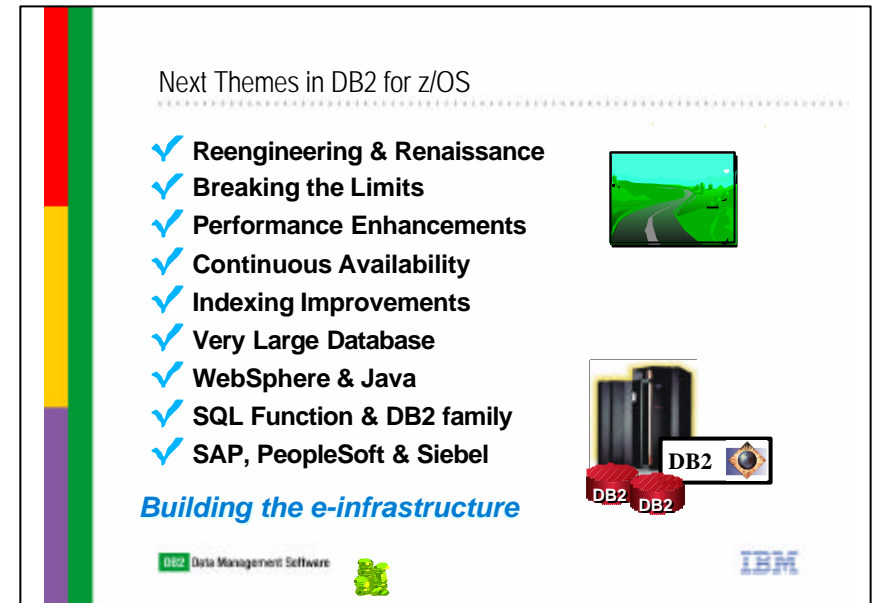




- ▶ This talk will focus on recent and future improvements in DB2 for z/OS. The key focus points will be on performance, synergy with zSeries and the DB2 family, continuous availability, applications and productivity. This session will highlight some of the recent changes delivered in DB2 Version 7, as well as several other key changes coming later.
- ▶ The organization for this presentation is in themes, with the V7 changes first within a theme, and then later enhancements. The foils generally have a version indication or the work is beyond version 7.



- ▶ We expect our next delivery to make a fundamental change in many areas, reengineering much of DB2 and helping deliver more growth. There are key improvements that help improve scalability, improved ability to port applications, much better ability to deliver continuous availability. The ability to handle very large databases has fundamental changes. Support for key vendor applications is compelling. The enhancements for 64 bit virtual storage will make management simpler, improve scalability and availability.


Reengineering & Renaissance: History & Future



- Indexes
- Data Sharing
- Stored Procedures
- Object Relational
- Storage & Space
- SQL & Optimization
- Unicode
- DRDA & Common Client

Building the e-infrastructure

- ✓ Continuous Availability
- ✓ Systems Management, Security
- ✓ Performance & Scalability
- ✓ Connectivity
- ✓ Productivity
- ✓ Synergy: DB2 family & zSeries


= Total Cost of Ownership





DB2 Data Management Software  

Limits: DB2 for z/OS

Breaking the limits



– Virtual Storage 2 GB	2 ³¹ to 2 ⁶⁴
– Table name sizes	18 to 128
– Column name sizes	18 to 30
– Partitions	254 to 4096
– SQL statement length	32K to 2 MB
– Index key size	255 to 2000
– Literals	255 to 32704
– Short strings	255 to 32704
– Predicates	255 to 32704

DB2 Data Management Software  

- ▶ The renaissance of DB2 has continued as an ongoing process, refreshing some components, updating the work during each release. Some of the architecture has made more dramatic changes, such as changing to type 2 indexes, data sharing, stored procedures, object relational and moving to larger address spaces. These are examples of major changes. Unicode changes our definition and handling of characters, changing the foundation. A common client across the DB2 family provides better family consistency and a stronger DRDA.
- ▶ Building the e-infrastructure shows the value proposition of DB2 for z/OS.
- ▶ Rebuilding, rearchitecting and renewing continues.

- ▶ One of the keys to reengineering is breaking through the limits of the current architecture. Increasing some of the limits improves the ability to scale DB2. Increasing other limits improves productivity, portability and family consistency.
- ▶ Increasing the amount of virtual storage we can address directly can help with the ability to scale and simplify management for main storage. It will require more storage, but permit increased scalability and availability.
- ▶ Increasing name sizes and SQL statements makes porting from other DBMS much easier and improves DB2 family compatibility.
- ▶ Increasing the maximum number of partitions helps DB2 scale farther and makes management much easier for cases that need to have one partition per day for a number of years.

Limits: DB2 for z/OS

Breaking the limits ...



– Tables in a join	15 to 225
– Active logs	31 to 93
– Archive logs	1000 to 10,000
– Current optimization	8 to 128
– CURRENT PACKAGESET	18 to 128
– CURRENT PATH	254 to 2048
– CURRENT SQLID	8 to 128
– USER	8 to 128
– VIEW	18 to 128
– ALIAS	18 to 128

DB2 Data Management Software




IBM

- ▶ The larger number of tables in a join helps with porting applications and improves our ability to handle vendor applications.
- ▶ The additional active and archive logs provides larger volumes and better flexibility for the amount of log data. Customers could keep up to 372 GB of active log data and 40 terabytes of archive log data.
- ▶ Increasing the special register lengths allows more flexibility for applications.

64 bit evolution

- zSeries, z/OS, z/Architecture & large real storage
 - ▶ Improve all versions
 - ▶ V6 data space advantages
 - zSeries & (z/OS or OS/390 V2R10 & 64 bit addressing)
- 64 bit virtual
 - ▶ See z/OS 64-bit Virtual Storage Roadmap
- Many incremental changes ...



DB2 Data Management Software



IBM

- ▶ The biggest impact of the zSeries architecture on DB2 is the ability to have large real memory support. Prior to the zSeries, customers were limited to 2 GB real storage due to the 31-bit addressing of the S/390 architecture. The real storage limit of 2 GB is a leading performance inhibitor for many high end customers. Another performance inhibitor is the 2 GB virtual storage limit for the main DB2 (DBM1) address space. Moving virtual pool buffers to hiperpools offers some relief, but many customers need more. If you have zSeries & OS/390 V2R10 64-bit mode or z/OS, use V6 buffer pools in data spaces, but not otherwise. See V7 Performance Topics red book and the web.
- ▶ There will be many more steps as real and virtual memory sizes increase, moving more above the line and above the bar. See the Roadmap, GM13-0076 ibm.com/servers/eserver/zseries/library/




64 bit virtual storage

24 31 64
Line 2 Bar 2 2

- Above the bar
 - ▶ Buffer pools, EDM pool
 - ▶ Sort Pool, Rid Pool
 - ▶ Parse trees
 - ▶ Compression dictionaries
- Most large storage above the bar
- Real storage requirements will increase
- Management of virtual storage is simpler
- Improve scalability and availability

DBM1 16 EB





2 GB

Performance: DB2 for z/OS

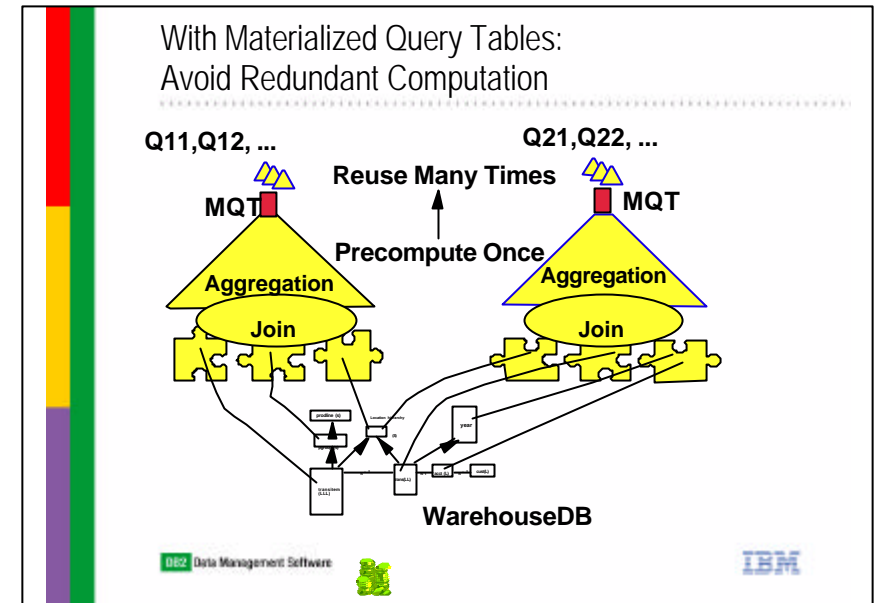
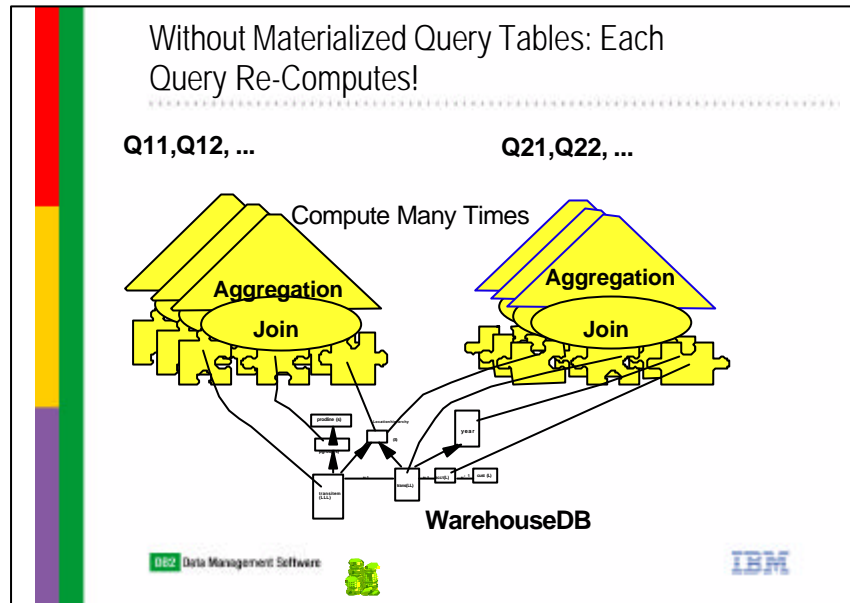
Performance Enhancements

- Ability to use indexes more often
 - More matching in predicates
 - Indexes bi-directional, varying length
 - Distribution statistics on non-index columns
- Multi-row fetch & insert
- Materialized Query Tables
- More Parallel Sorting


- ▶ A statement of direction was included in the September 11, 2001 announcement, IBM z/OS Version 1 Release 2: Enabling and Protecting Your e-business and Preview: z/OS Version 1 Release 3
- ▶ IBM plans to deliver 64-bit virtual storage addressing for the DB2® for z/OS product in a future release. The future release of DB2 for z/OS, with 64-bit virtual address support, can only execute on IBM (elogo)server zSeries 900 (z900), or equivalent, running z/OS V1R2, or later.
- ▶ Instead of hiperspaces or data spaces, the single large address space can allow easier management of storage. We expect real storage needs to increase as the scalability and availability is addressed.

- ▶ Many of the performance enhancements come from the enhanced index capabilities. Comparing a decimal column to a floating point number could not be done in the past. Being able to compare these values can mean choosing a better index or even an index access instead of a table space scan. An index can be used for a backward scan, so some indexes may not be required, reducing the overhead for inserts & deletes. Support for varying length indexes can save space and can use index-only access with a varchar. Having distribution statistics for columns which are not part of the index can give the optimizer better information so it can perform better optimization. Gathering these statistics was a separate program, rather than part of RUNSTATS.
- ▶ Materialized query tables can provide a one or two order of magnitude performance improvement by rewriting queries to use the precalculated information.





- ▶ A materialized query table (MQT) can avoid redundant work of scanning, aggregation and joins. Multiple levels of summary tables have been used in warehouses and complex applications for years. One of the major issues is communicating the summaries to the users. In some cases, the users want to query the base data. With MQTs, the query users do not have to be aware of the MQT.

- ▶ Even though the query is submitted for the base table, the optimizer can rewrite the query to use the MQT. Using the precalculated information can improve subsequent queries by as much as two or three orders of magnitude. Materialization or precalculation and parallelism resolve the long response times.
- ▶ A database administrator can use an MQT much as she or he would use an index for optimization. Controls for usage, initial loading and refresh are part of the definition.




Materialized Query Tables


- Sometimes called
 - ▶ Automatic Summary Tables
 - ▶ Automatic Materialized Query Tables
 - ▶ Materialized Views
- Pre-computed information
- Managed by user or system
- Automatic rewrite or manual

DB2 Data Management Software  

- ▶ Many names are used for the MQTs in various implementations, including automatic summary tables, aggregate summary tables, automatic materialized query tables and materialized views. They all have a common objective, to precompute information for later use by other statements.
- ▶ The objective is to make the use as transparent as possible. Customers may choose to have the MQTs managed and refreshed by the system or manually. Most customers will have the optimizer rewrite queries, but MQTs can be names in SQL statements.





Availability: DB2 for z/OS




Continuous Availability

- Schema Evolution
 - ALTER instead of DROP / CREATE
- Data Partitioned Secondary Indexes (DPSI)
- Add a column to an index
- Improved LPL Recovery
- Additional online zparms



DB2 Data Management Software  

- ▶ The most important change for many customers is the ability to use ALTER in many places instead of needing to drop and redefine. We call this schema evolution, and it can reduce outages by hours or days for a major structure change on an application.
- ▶ The ability to have secondary indexes that are partitioned with the data can improve recovery times by an order of magnitude. It can also eliminate the outage for online reorganizing a single partition or BUILD2 phase.
- ▶ We have some additional cases where subsystem parameters can be changed while the subsystem is running.




Online Schema Evolution

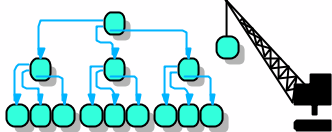
- Table or column Changes
 - ▶ Increase column size within datatype integer, float
 - ▶ Change to expand datatype: char, numeric
 - ▶ Change varchar to / from char
- Dynamic Partitions
 - ▶ Add partition
 - ▶ Rotate partition
- Unbundling Partitioned Table Attributes
 - ▶ Partition without an index
 - ▶ Cluster on any index
 - ▶ May be able to have one less index

DB2 Data Management Software  

- ▶ Three types of changes are very high on our priority list.
 - ▶ We need to be able to change the data type for columns. In V5 we could increase the size of varchar columns, but the change allows us to extend numeric and character columns and to change between char and varchar.
 - ▶ Adding a new partition to an existing partitioned tables space is very important. Rotating the partitions, such as keeping a rolling 36 months of data is also key.
 - ▶ Partitioning and clustering are bundled in current DB2. Some of the time we are required to make a difficult choice. We also want to partition without an index and be able to cluster on any index. These changes will allow us to have one less index and less random IO in some cases.





Index: DB2 for z/OS



Index Improvements

- Variable length index keys
- Index-only access for varchar data
- Maximum index key 2000 bytes
- Predicates indexable for unlike types
- Backward Index Scan
- Partitioning separate from clustering
- Data-partitioned secondary indexes


DB2 Data Management Software  


- ▶ DB2 provides many new opportunities for improving index processing.
- ▶ We are able to use indexes more effectively, reducing the space in variable-length indexes, being able to have index-only access with variable-length data and being able to use the index when the predicates do not match.
- ▶ In some cases, such as backward index scans or partitioning, we will be able to work as efficiently with one less index. This will improve the insert, delete, LOAD, REORG and update processing.
- ▶ We have more flexibility in indexes, with longer index keys, the ability to partition secondary indexes and the ability to have more effective clustering.

VLDB: DB2 for z/OS

Very Large Database Improvements

- Separate partitioning from clustering
- Partition without an index
- Data-partitioned secondary indexes
- 4096 Partitions
- Add partitions
- Rotate partitions



DB2 Data Management Software  IBM

- ▶ Very large databases face the combined challenge of very high performance needs, continuous availability and complexity. Improvements in scale and flexibility are more important in this area. Being able to have more partitions and to add them with ALTER are a big improvement.
- ▶ Often it is useful to partition by date, so that we can archive or delete an entire partition, but processing will be much more efficient with another clustering order, such as by customer. Before this change, the clustering order was the same as the partitioning. This flexibility offers many opportunities for improved performance and availability.
- ▶ Some customers have an index that is used only for partitioning the data or have extra columns at the beginning of the index. Being able to avoid the extra index or columns can improve our efficiency a lot.
- ▶ For these very large tables, the ability to have more partitions, to add new partitions and to be able to rotate partitions is crucial.

Coming soon to a web near you
DB2 Family

IBM DB2 Universal Database

**SQL Reference
for Cross Platform Development**


A new SQL Reference book for the DB2 UDB family, not just one platform.



DB2 Data Management Software  IBM

- ▶ We are expecting to have a new SQL Reference book in the next month or two. This is not an SQL Reference for a single product, but is designed to help you design and build applications which will deploy on the DB2 UDB platform of your choice easily.
- ▶ I could not resist the hummingbird family picture, taken out the window at Silicon Valley Lab.

SQL: DB2 for z/OS



SQL function & DB2 family


- Scalar Fullselect
- Materialized Query Tables
- Multiple DISTINCT
- INSERT within SELECT Statement
- Sequences
- Identity column enhancements
- Long names
- Long statements (2 MB)

DB2 Data Management Software



IBM

SQL: DB2 for z/OS



SQL function & DB2 family ...

- Long strings
- Long predicates
- Multi-row FETCH & INSERT
- Additional Built-in Functions
- Unicode
- Stored Procedure & UDF enhancements
- Dynamic cursor scrolling

DB2 Data Management Software



IBM



- ▶ Standards & DB2 family consistency drive the SQL and application work. As in other DB2 family members, longer names for tables, columns & longer SQL statements will help application portability. We have more work on cursor scrolling & Unicode. Summary tables or materialized query tables provide both improved family consistency and a big performance enhancement.
- ▶ The ability to use INSERT within a SELECT statement provides a technique to return values. Sequences provide a counter that is not a column in a table. Identity columns are improved with better flexibility.

- ▶ The lengths of literal strings and predicates are increased from 255 to 32704, so that SQL will be more portable. The ability to fetch and insert multiple rows in an SQL statement makes some applications easier and more efficient.
- ▶ The ability to use Unicode more flexibly will expand the range of applications using this option.

Family: DB2 for z/OS

Other DB2 Family Improvements

- Control Center
- Development Center
- Warehouse Manager
- Replication & Replication Center
- DB2 Connect
- Common Client / Java Combined Client
- XML



DB2 Data Management Software



IBM

ERM: DB2 for z/OS



SAP, PeopleSoft & Siebel

- 64 bit virtual storage
- Unicode
- 64 bit DB2 Connect for Linux zSeries
- Multi-row fetch & insert
- Multiple DISTINCT Clauses
- Lock Contention on volatile Tables
- Fast Retrieval of Most Recent Value
- Control Center Enhancements

DB2 Data Management Software




IBM

- ▶ A wide range of other improvements are being made in components that work across the DB2 family. These graphical interfaces, wizards and clients can be developed once. The improvements are consistency across the family and usability.


- ▶ Enterprise Resource Management (ERM) includes Enterprise Resource Planning (ERP) and Customer Relationship Management (CRM). Most of the key items in this version help our key enterprise application partners: SAP, PeopleSoft and Siebel.
- ▶ The 2 GB storage limit is a major problem for customers who have large dynamic statement cache, many concurrently open data sets and long running units of work.
- ▶ Unicode helps vendors support multinational companies and avoid character conversions which lose in the translation.
- ▶ These applications often insert and fetch many rows, using a technique that is sometimes called array fetch & insert.
- ▶ The SQL flexibility improvements allow DB2 to be efficient in performance and in productivity for our partners.

ERM: DB2 for z/OS




SAP, PeopleSoft & Siebel ...

- Longer Table Names & Column Names
- Additional statistics
- Convert Column Type
- Altering CLUSTER Option
- Adding Columns to Index
- Index-only Access Path for VARCHAR
- Adding New Partitions
- Separate Clustering from Partitioning


DB2 Data Management Software  IBM

- ▶ The longer table names and column names mean that these applications are restricted on our platform. The additional statistics can improved optimization, reducing the needed tuning. The ability to change data type with ALTER, to ALTER CLUSTERING and to add new partitions improves the time to convert from one application version to another.
- ▶ Index-only access path with VARCHAR data removes a choice with a potential incompatibility.


UNICODE Evolution



- UNICODE encoding scheme can represent codepoints of characters of many different geographies & languages
 - ▶ More than one byte to represent many characters
 - ▶ Different languages in the same table
 - ▶ CCSID UNICODE
- V5 ASCII data and Unicode for Java
- V7 UNICODE data storage & manipulation
- OS/390, z/OS, z/Architecture and zSeries


DB2 Data Management Software  IBM

- ▶ Multinational companies that engage in international trade often store data from more than one country. Some countries use different coded character set identifiers.
- ▶ Previous releases of DB2 have offered support for numerous code sets of data in either ASCII or EBCDIC format. However, there was a limitation of one code set per system.
- ▶ Version 7 of DB2 for OS/390 & z/OS delivers support for Unicode encoded data. The encoding scheme can represent the code points of many different geographies and languages. You can easily store multilingual data within the same table or on the same DB2 subsystem.
- ▶ These changes are supported by new function in OS/390, z/OS, the z/Architecture and the zSeries machines. See www.unicode.org and session G5.





Unicode Enhancements


- Able to join Unicode with EBCDIC or ASCII
- SQL in Unicode, statements, literals
- Most DB2 character catalog columns in Unicode
- Collating sequence change
- Lengths and maximum lengths can change




www.unicode.org

- ▶ The key improvements for Unicode allow much more flexibility, with the ability to join a Unicode table to one that is ASCII or EBCDIC. SQL statements and literals can be Unicode or EBCDIC. Many of the DB2 catalog character columns will be converted to Unicode, so Unicode is for everyone.
- ▶ The collating sequence for Unicode is similar to ASCII, with numbers sorting lower than letters. The lengths and maximum lengths of strings can change as they are converted to Unicode, so treat lengths as variable.



DB2 for z/OS Vnext





Migration Considerations:

- DB2 and IRLM will be 64-bit exclusive
- Base Pre-requisites:
 - zSeries z800 z900 or later
 - z/OS R3 or above
- Major catalog changes:
 - Long fields
 - Unicode
- Migration Process:
 - From Version 7 only
 - Multi-step process
 - V7 and Vnext data sharing coexistence supported in compatibility mode only


Building the e-infrastructure

- ✓ Continuous Availability
- ✓ Systems Management, Security
- ✓ Performance/Capacity
- ✓ Connectivity
- ✓ Productivity
- ✓ Synergy: DB2 family & zSeries
- = Total Cost of Ownership

Ready for e-business






- ▶ Here is some early information for your planning, while the information is not announced yet, what you will want to do to get ready for the next version of DB2 is to migrate to current versions of hardware architecture, operating system and DB2.




DB2 Catalog Vnext

- Longer names (128 byte varchar Unicode)
 - ▶ table, view & alias
 - ▶ column (30)
 - ▶ schema or authorization id
 - ▶ ...
- Some page sizes greater than 4K
- Indexes longer than 255 bytes
- Contains Unicode data
- Migration only from V7
- Prerequisite zSeries, z/OS R3 or later



DB2 Data Management Software  



- ▶ This is a discussion about possible changes in the DB2 catalog beyond Version 7. The changes are more substantial for the catalog than for any prior release, with the possible exception of V1R2.
- ▶ Note that you will migrate to Version 7 before you can use the standard process to migrate. The migration process from Version 6 is to Version 7, then another migration.




Four step process for Migration

- Test with New Install & Migrate
- Apply the Fallback SPE to all members
 - ▶ Start all members at SPE level
- Migrate to new release without new function (COMPAT mode)
- Enable New Function Mode
 - ▶ Then in New Function Mode


DB2 Data Management Software  

- ▶ We are thinking that the process of migration will take four steps.
- ▶ Use a new install on a separate subsystem and practice migration on the separate subsystem. This will be used to ensure that you are ready to move forward.
- ▶ Apply the fallback SPE on every member of each data sharing group and on all subsystems to migrate. Applying the SPE was required only for data sharing, but is now required for all subsystems.
- ▶ Migrate to the new release in compatibility mode. You are using the new code but almost no new function.
- ▶ Finally, after all of the members of a data sharing group are running the new version, you can switch to new function mode.




Tighter Migration Customer Value

- More robust migration process
 - ▶ Fewer migration & fallback errors
 - ▶ All customers apply the Fallback SPE
- Customer control for migration process & timing for using new function
- Reduced risk of problems in Fallback SPE
 - ▶ Reduced number of modules, fewer changes
- Improved ability to deliver new function, while still having fallback & coexistence


IBM Data Management Software 

- ▶ We think that the new process is formalizing what customers have done and adding some new controls.
- ▶ The process should make the process more robust by avoiding a technique that has caused many problems.
- ▶ There is a way to control the use of new function, so that falling back to a prior release is easier from COMPAT mode.
- ▶ There will be fewer modules to change in the old release to tolerate changes from the later version.
- ▶ We think the new technique provides a method that will make it easier to deliver new function and still support fallback and coexistence, even with many fundamental changes.



Incompatible Changes

- ▶ Use valid host variable declaration for PREPARE
- ▶ Valid VCAT required no x.y (V5 and up)
- ▶ No COMPJAVA, Use Java JIT
- ▶ CREATE, ALTER PROCEDURE WLM ONLY
- ▶ Unicode used in
 - DB2 catalog, changes order
 - SQL statements, translation, lengths may change
 - Instrumentation

IBM Data Management Software 

- ▶ The new version will probably have some new incompatible changes:
- ▶ Host variable declarations for PREPARE and EXECUTE IMMEDIATE must declare the correct length. V7 made this change, removed in PQ50494.
- ▶ Note that a valid VCAT (one level) is enforced in V5, V6 and V7 with APAR PQ53145.
- ▶ Only the Java JIT compiler will be supported. COMPJAVA or HPJ depend upon JDK 1.1.8, which is out of support.
- ▶ While stored procedures created earlier will run, creating or altering a stored procedure will require WLM environment.
- ▶ Unicode changes will require migration work.
- ▶ Migration will be permitted only from Version 7

Any Questions?



DB2 Data Management Software





IBM

Disclaimer and Trademarks

Information contained in this material has not been submitted to any formal IBM review and is distributed on "as is" basis without any warranty either expressed or implied. Measurements data have been obtained in laboratory environment. Information in this presentation about IBM's future plans reflect current thinking and is subject to change at IBM's business discretion. You should not rely on such information to make business plans. The use of this information is a customer responsibility.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not imply giving license to these patents.

TRADEMARKS: THE FOLLOWING TERMS ARE TRADEMARKS OR ® REGISTERED TRADEMARKS OF THE IBM CORPORATION IN THE UNITED STATES AND/OR OTHER COUNTRIES: AIX, AS/400, DATABASE 2, DB2, e-business logo, Enterprise Storage Server, ESCON, FICON, OS/390, OS/400, ES/9000, MVS/ESA, Netfinity, RISC, RISC SYSTEM/6000, iSeries, pSeries, xSeries, SYSTEM/390, IBM, Lotus, NOTES, WebSphere, z/Architecture, z/OS, zSeries,  

The following are TRADEMARKS OR REGISTERED TRADEMARKS OF THE MICROSOFT CORPORATION IN THE UNITED STATES AND/OR OTHER COUNTRIES: MICROSOFT, WINDOWS, WINDOWS NT, ODBC, WINDOWS 95

For additional information see ibm.com/legal/copytrade.phtml

DB2 Data Management Software



IBM

- ▶ The most important change for many customers is the ability to use ALTER in many places instead of needing to drop and redefine. We call this schema evolution. Allowing more partitions is important for customers who need to have a partition for each day and keep the data for years.
- ▶ Standards and DB2 family consistency drive the application items. As in other DB2 family members, longer names for tables, columns and longer SQL statements will help application portability. We will continue with more work on cursor scrolling and Unicode. Materialized query tables are crucial to complex SQL performance & consistency across the DB2 family. Multi-row fetch & insert reduce cpu time. Improved optimization & parallel processing are expected in every release.

- ▶ Some parts of this presentation are more like looking into a crystal ball than at measurements. This crystal ball is cloudy, and gets fuzzier the farther we look into the future. The only near certainty is that there will be changes. My best guess is that fewer than 10% of the items will change their delivery time. I would expect some new items to come in, some to come early, and others to deliver in stages. More will have major changes in their design.
- ▶ Do you need Answers to Frequently Asked Questions?
 - ▶ Hints and Tips? Product Information?
 - ▶ Technical Presentations? Redbooks? White Papers?
- ▶ Support is on the web. Click Support on DB2 web pages.

Session Title: DB2 Technology Highlights

Roger Miller
IBM Silicon Valley Lab
millerrl@us.ibm.com

ibm.com/software/db2zos

DB2 Data Management Software



IBM

Bibliography: DB2 for z/OS & OS/390

- Squeezing the Most Out of Dynamic SQL, SG24-6418
- Large Objects with DB2 for z/OS and OS/390, SG24-6571
- Version 7 Selected Performance Topics, redpaper
- Version 7 Performance Topics, SG24-6129
- Version 7 Using the Utilities Suite, SG24-6289
- Powering the World's e-business Solutions, SG24-6257
- Version 7 Presentation Guide, SG24-6121
- Performance Management Tools, SG24-6508
- Application Programming Topics, SG24-6300

DB2 Data Management Software



IBM

- Are you looking for additional support? Have you had problems finding answers to your questions? Would you like to see hints and tips for some of the techniques you use? Do you want to get recent technical presentations? Do you need to be able to search for DB2 Redbooks more effectively? Do you need white papers?
- Improved Support is on the web now. Click on the Support entry on your DB2 web page. We have added more than one thousand answers to frequently asked questions and hints and tips documents. Click on the Frequently asked questions (FAQs) line to see answers to the most frequently asked questions. Click on Hints & Tips to get brief information on installation, configuration, troubleshooting, and usage. The DB2 library has been on the web, and it is also accessible from this page. Many recent DB2 technical presentations can be downloaded by clicking on Technical Presentations. Some customers need to have a more effective search for Redbooks. The Support page lists recent red books, but you can also narrow down the search and filter results by adding keywords. You can also find a selection of recent white papers.

- If you want to look across the DB2 family, start at
 - ibm.com/software/db2
 - Click on Support or start at
 - <http://www.ibm.com/software/data/support/>
 - Then choose a product.
- If you want to look at DB2 for z/OS & OS/390, start at
 - ibm.com/software/db2zos
 - Click on Support or start at
 - <http://www.ibm.com/software/data/db2/os390/support.html>
- If you want to look at DB2 for UNIX & Windows, start at
 - ibm.com/software/db2/udb
 - Click on Support or start at
 - <http://www.ibm.com/software/data/db2/udb/support.html>