DB2 10 for z/OS technical preview

This session will recap the key features of DB2 9 for z/OS, followed by a look at DB2 10, the next release of DB2. If your organization needs to lower operating costs through CPU cycle reductions while still building a strong foundation for SOA and XML initiatives, you'll learn how much DB2 10 for z/OS will have to offer. Database administrators will find improved database performance, scalability and availability, while memory management is dramatically reduced, so growth is much simpler, as are many other tasks. You'll also have more flexible security to help with regulatory compliance. A wide range of enhancements improve the functionality and performance of ERP applications and data warehousing. Application developers will learn about the second release of pureXML®, which improves the SQL and XML interface to access XML data stored in a native format. We'll discuss powerful new SQL enhancements, including improved SQL and data definition capability with other DB2 platforms, making porting much easier.

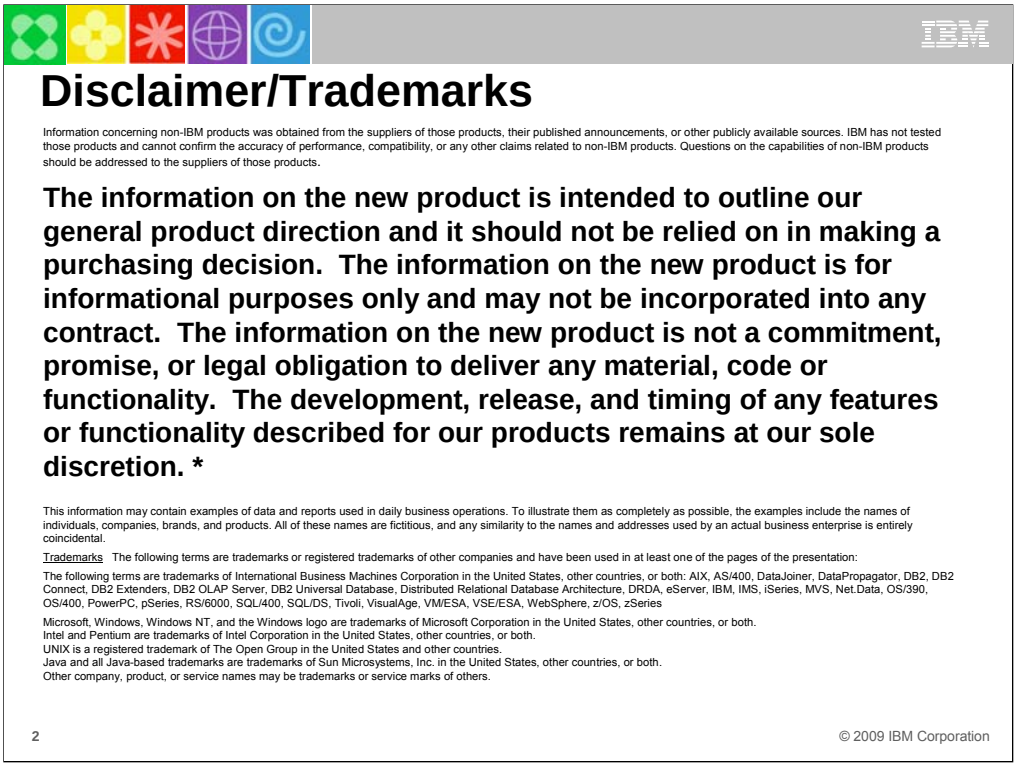

Disclaimer/Trademarks

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements, or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

The information on the new product is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information on the new product is for informational purposes only and may not be incorporated into any contract. The information on the new product is not a commitment, promise, or legal obligation to deliver any material, code or functionality. The development, release, and timing of any features or functionality described for our products remains at our sole discretion. *

This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious, and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Trademarks The following terms are trademarks or registered trademarks of other companies and have been used in at least one of the pages of the presentation:

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both: AIX, AS/400, DataJoiner, DataPropagator, DB2, DB2 Connect, DB2 Extenders, DB2 OLAP Server, DB2 Universal Database, Distributed Relational Database Architecture, DRDA, eServer, IBM, IMS, iSeries, MVS, Net.Data, OS/390, OS/400, PowerPC, pSeries, RS/6000, SQL/400, SQL/DS, Tivoli, VisualAge, VM/ESA, VSE/ESA, WebSphere, z/OS, zSeries

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Other company, product, or service names may be trademarks or service marks of others.

2 © 2009 IBM Corporation

At this point, the beta announcement for DB2 10 was made February 9, 2010. The general availability announcement has not been made, so this is not an available product. We have used the code name DB2 X for z/OS up to this point.

See the web for current information, the announcement, …

http://www.ibm.com/software/data/db2/zos/db2-10/

http://www.ibm.com/common/ssi/rep_ca/5/877/ENUSZP10-0015/ENUSZP10-0015.PDF

http://www.ibm.com/support/docview.wss?uid=swg27017960

http://it.toolbox.com/blogs/db2zos/db2-10-for-zos-beta-announced-today-36790

http://www.ibm.com/developerworks/spaces/db2zos

http://davebeulke.com/?p=625

http://community.solutionscenter.techweb.com/community/mainframe/blog/2010/02/09/db2-10-for-zos-beta-starts-today

http://www.triton.co.uk/blog/?p=415

DB2 10 for z/OS provides the best reduction in CPU for transactions and batch for 22 years, since V2R1. We expect most customers to reduce CPU times between 5% and 10% initially, with opportunity for more. Applications which can take advantage of additional benefits, such as hash access, can have larger CPU and memory reductions. Scalability is the second major benefit, with the ability to run five to ten times as many threads in a single subsystem by moving 80% to 90% of the virtual storage above the bar. Schema evolution or data definition on demand enhancements improves availability. SQL, pureXML, and web services extend usability and application portability for this platform. Productivity improvements for application developers and for database administrators are very important as data grows in scale and complexity. Warehousing continues to evolve, with key trends matching System z and DB2 for z/OS strengths of performance, scalability, reliability, stability, availability, resilience, and security.

DB2 10 has a lot for everyone. Your organization needs to lower operating costs thru CPU cycle reductions while still building a strong foundation for SOA and XML initiatives. Database Administrators (DBAs) will find improved database performance, scalability, and availability. Memory management is dramatically reduced, so growth is much simpler. DBAs also get more flexible security to help with regulatory compliance. A wide range of enhancements improve ERP application and data warehouse functionality and performance. Large object (LOB) performance and flexibility are improved. Many DBA tasks are simplified.

Application developers will be most excited by the second release of pureXML, which improves the SQL and XML interface to access XML data stored in a native format. Application developers need powerful new SQL enhancements. Improved SQL and data definition compatibility with other DB2 platforms makes porting much easier.

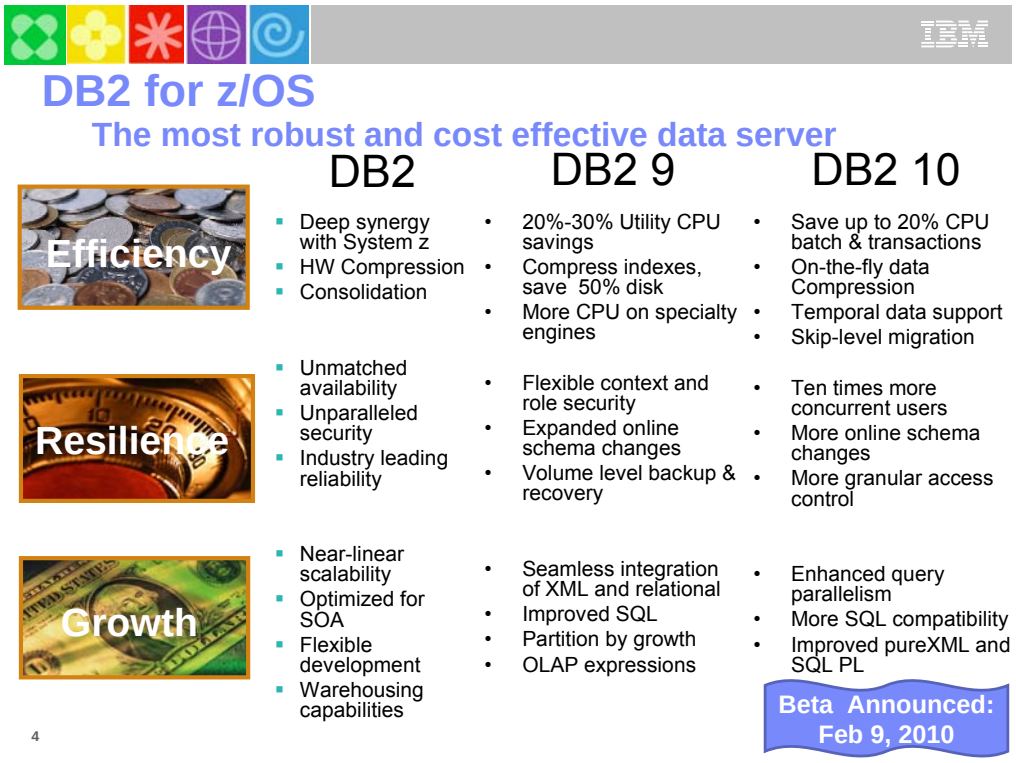

Customers have come to know DB2 as the most robust and cost effective data server. With every version of DB2, we are focused on the needs of our customers to operate efficiently, to be up and running 24x7, and to grow with their business. With V9, customers get CPU and disk savings as well as a boost in application productivity with the new pureXML technology. For the next DB2 version, a lot of customers are getting excited. We are putting a lot of focus on out-of-the-box performance improvements and productivity improvements such as online schema, temporal data support, and fine-grain security controls. DB2 continues to be the choice for mission critical business data and we continue to make it easier for customers to keep data on the platform.

DB2 9: One of the key initiatives of V8 was online schema evolution, and that theme is expanding and changing to be data definition on demand. These are key improvements for resilience. One of the important changes is to be able to replace one table quickly with another. Another is to be able to rename a column or an index. A new type of table space combines the attributes of segmented and partitioned, without a partitioning key. Rebuild index can be run with much less disruption. Online table space reorganization for a few partitions is improved a lot, removing the BUILD2 phase for all types of secondary indexes. Table space and index logging can be altered.

Many other improvements help with performance, with scalability and with availability. Index on an expression can be combined with caseless comparisons to improve text search. Improved insert rates can result from improved latching of the log data. Significant reductions in cpu usage are provided with new utilities.

Today's complex applications include both transactions and reporting, so performing both well is imperative. The key improvements for reporting are optimization enhancements to improve query and reporting performance and ease of use. More queries can be expressed in SQL with new SQL enhancements. Improved data is provided for the optimizer, with improved algorithms. Improved cpu and elapsed times can be achieved with the FETCH FIRST clause specified on a subquery. The INTERSECT and EXCEPT clauses make SQL easier to write.

DB2 10: DB2 10 for z/OS provides the best reduction in CPU for transactions, queries, and batch for over 20 years, since V2R1. We expect most customers to reduce CPU times between 5% and 10% as soon as DB2 10 is out of the box. Applications which can take advantage of additional benefits, such as hash access, index include columns, inline large objects, parallel index updates, faster single row retrievals, work file in-memory, index list prefetch, 64 bit memory enhancements, use of the System z10 1 megabyte page size, buffer pools in memory, access path enhancements, member clustering for universal table spaces, efficient caching of dynamic SQL statements with literals, improved large object streaming, and SQL procedure language performance can have additional CPU and memory reductions. As always with performance, individual customer experiences will vary, and individual workloads will vary more.

Scalability is the second major benefit, with the ability to run five to ten times as many threads in a single subsystem by moving 80% to 90% of the virtual storage above the bar. Schema evolution or data definition on demand enhancements improve availability, by using an ALTER where the only prior option was DROP and recreate. Improved concurrency to DB2 catalog access and utilities extends the scaling. Security is enhanced with better granularity for administrative privileges, masking for data, and new audit capabilities.

SQL, pureXML, and web services improvements extend usability and application portability to the System z, z/OS and DB2 for z/OS platform. Temporal or versioned data improves productivity for applications in a wide range of industries. Applications ranging from SAP to warehousing see benefits from every category and item.

The net result is productivity improvements in DB2 10 for application developers, for database administrators, and for systems administrators that are very important as data grows in scale and complexity.

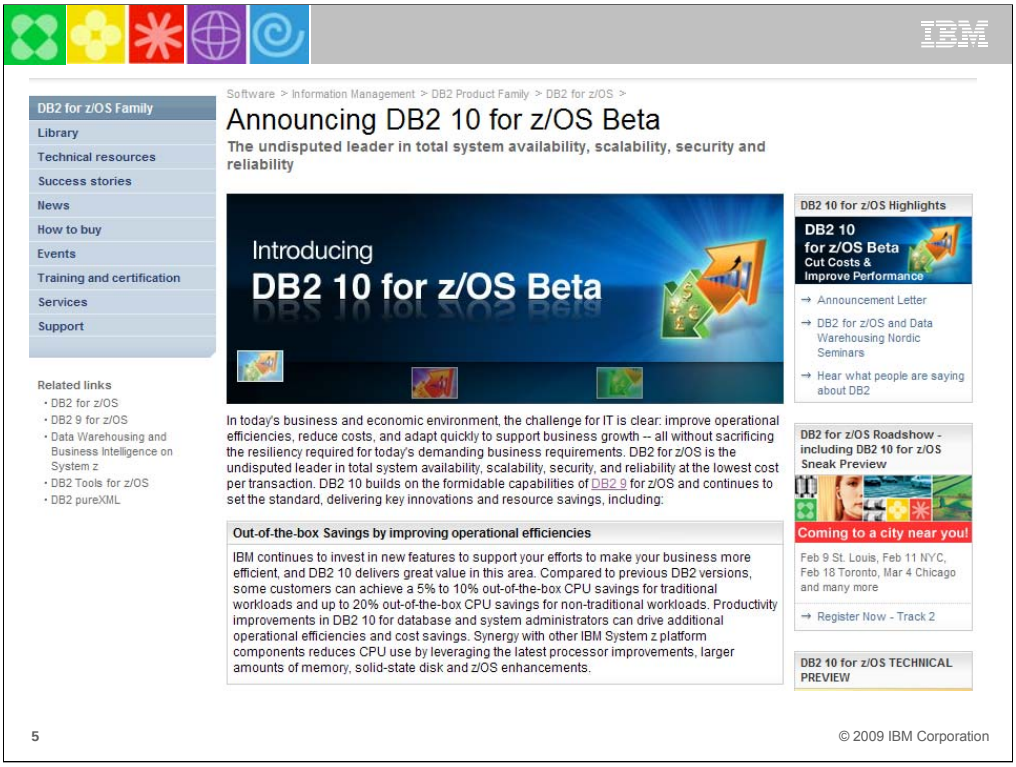

This is the announcement for the beta on the IBM web. See some early pointers and quotes from customers and consultants.

http://www.ibm.com/software/data/db2/zos/db2-10/

The announcement itself is here:

http://www.ibm.com/common/ssi/rep_ca/5/897/ENUS210-015/ENUS210-015.PDF

A version if this presentation is on the web.

http://www.ibm.com/support/docview.wss?uid=swg27017960

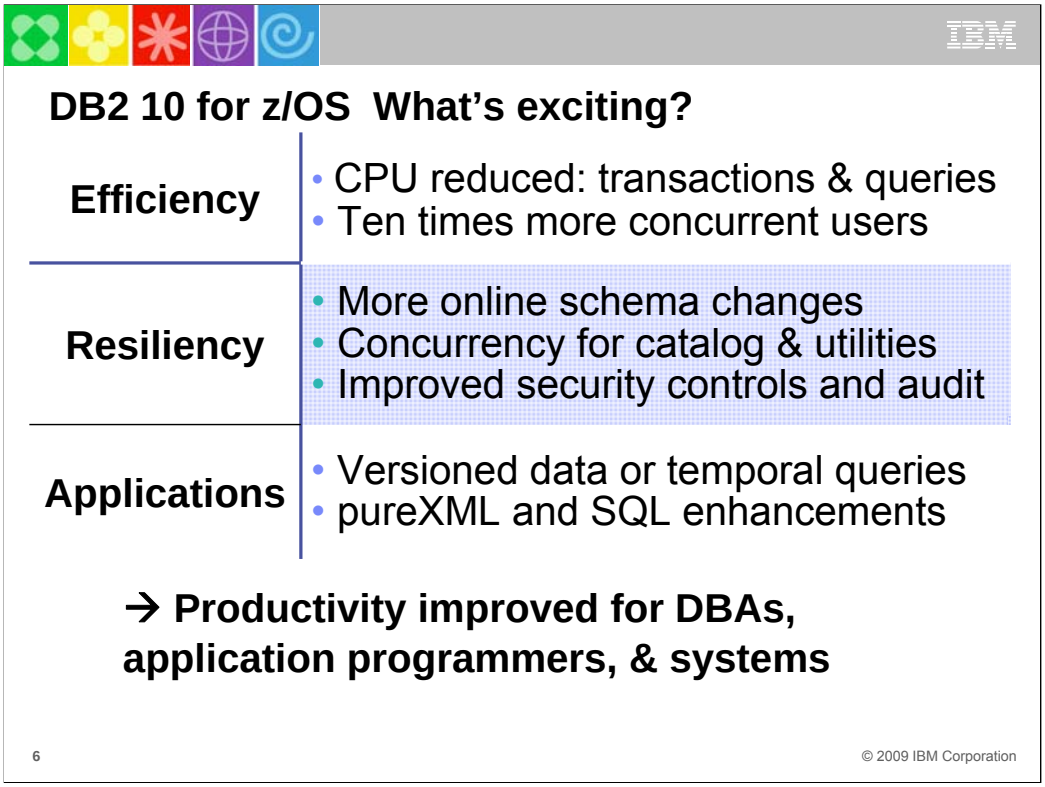

DB2 10 for z/OS provides the best reduction in CPU for transactions, queries, and batch for over 20 years, since V2R1. We expect most customers to reduce CPU times between 5% and 10% as soon as DB2 10 is out of the box. Applications which can take advantage of additional benefits, such as hash access, index include columns, inline large objects, parallel index updates, faster single row retrievals, work file in-memory, index list prefetch, 64 bit memory enhancements, use of the System z10 1 megabyte page size, buffer pools in memory, access path enhancements, member clustering for universal table spaces, efficient caching of dynamic SQL statements with literals, improved large object streaming, and SQL procedure language performance can have additional CPU and memory reductions. As always with performance, individual customer experiences will vary, and individual workloads will vary more.

Scalability is the second major benefit, with the ability to run five to ten times as many threads in a single subsystem by moving 80% to 90% of the virtual storage above the bar. Schema evolution or data definition on demand enhancements improve availability, by using an ALTER where the only prior option was DROP and recreate. Improved concurrency to DB2 catalog access and utilities extends the scaling. Security is enhanced with better granularity for administrative privileges, masking for data, and new audit capabilities.

SQL, pureXML, and web services improvements extend usability and application portability to the System z, z/OS and DB2 for z/OS platform. Temporal or versioned data improves productivity for applications in a wide range of industries. Applications ranging from SAP to warehousing see benefits from every category and item.

The net result is productivity improvements in DB2 10 for application developers, for database administrators, and for systems administrators that are very important as data grows in scale and complexity.

Top 10 in DB2 10 for z/OS

1. CPU reductions for transactions, queries, & batch
2. Ten times more users by avoiding memory constraints
3. More concurrency for catalog, utilities, and SQL
4. More online changes for data definition, utilities and subsystems
5. Improved security with more granularity
6. Temporal or versioned data
7. SQL enhancements improve portability
8. pureXML performance and usability enhancements
9. Hash, index include columns, access path stability, skip migration, … Insert your favorite.
10. Productivity improved for database & systems administrators, and application programmers

© 2009 IBM Corporation

DB2 10 for z/OS provides the best reduction in CPU for transactions and batch for over 20 years, since V2R1. We expect most customers to reduce CPU times between 5% and 10% as soon as DB2 10 is out of the box. Applications which can take advantage of additional benefits, such as hash access, index include columns, inline large objects, parallel index updates, faster single row retrievals, work file in-memory, index list prefetch, 64 bit memory enhancements, use of the System z10 1 megabyte page size, buffer pools in memory, access path enhancements, member clustering for universal table spaces, efficient caching of dynamic SQL statements with literals, improved large object streaming, and SQL procedure language performance can have additional CPU and memory reductions. As always with performance, individual customer experiences will vary, and individual workloads will vary more.

Scalability is the second major benefit, with the ability to run five to ten times as many threads in a single subsystem by moving 80% to 90% of the virtual storage above the bar. Schema evolution or data definition on demand enhancements improve availability, by using an ALTER where the only prior option was DROP and recreate. Improved concurrency to DB2 catalog access and utilities extends the scaling. Security is enhanced with better granularity for administrative privileges, masking for data, and new audit capabilities.

SQL, pureXML, and web services improvements extend usability and application portability to the System z, z/OS and DB2 for z/OS platform. Temporal or versioned data improves productivity for applications in a wide range of industries.

The net result is productivity improvements in DB2 10 for application developers, for database administrators, and for systems administrators that are very important as data grows in scale and complexity.

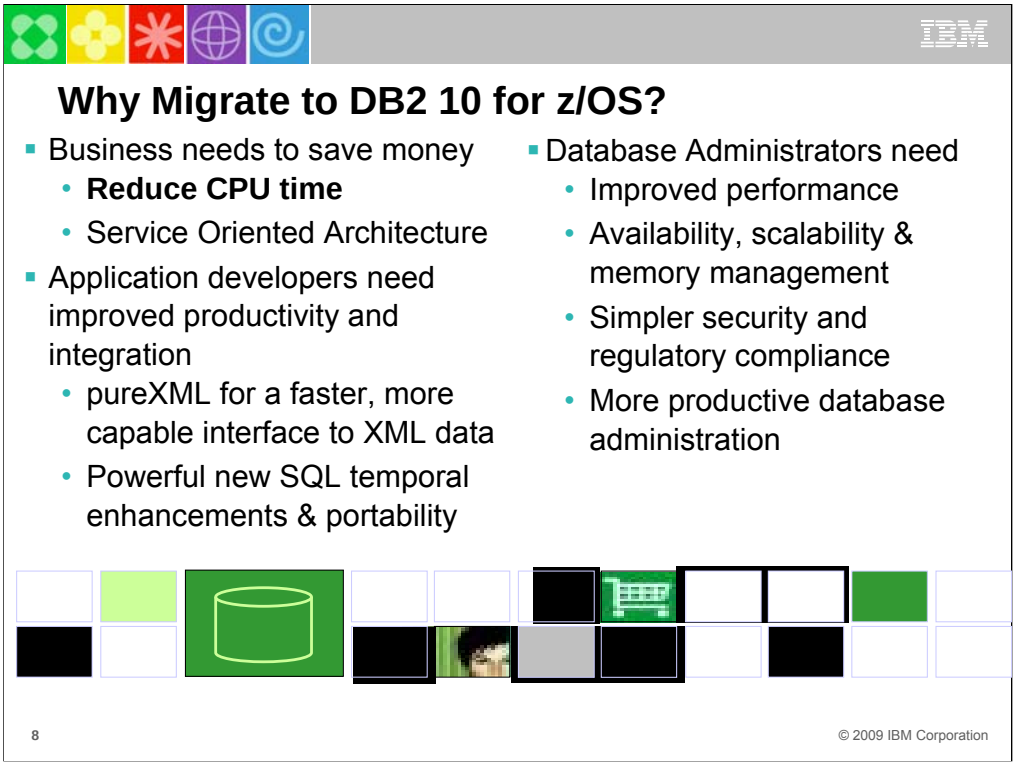

DB2 10 has a lot for everyone. Here are just a few of the strategic highlights.

The business needs include CPU cycle reductions that deliver for batch and transactions, and pureXML™ improvements to build a strong foundation for SOA and XML initiatives.

Database Administrators (DBAs) need improved database performance, scalability, and availability. Memory management is dramatically reduced, so growth is much simpler. DBAs also get more flexible security to help with regulatory compliance. A wide range of enhancements improve ERP application and data warehouse functionality and performance. Large object (LOB) performance and flexibility are improved. Many DBA tasks are simplified.

Application developers are most excited by the second release of pureXML, which improves the SQL and XML interface to access XML data stored in a native format. Application developers need powerful new SQL enhancements. Improved SQL and data definition compatibility with other DB2 platforms makes porting much easier.

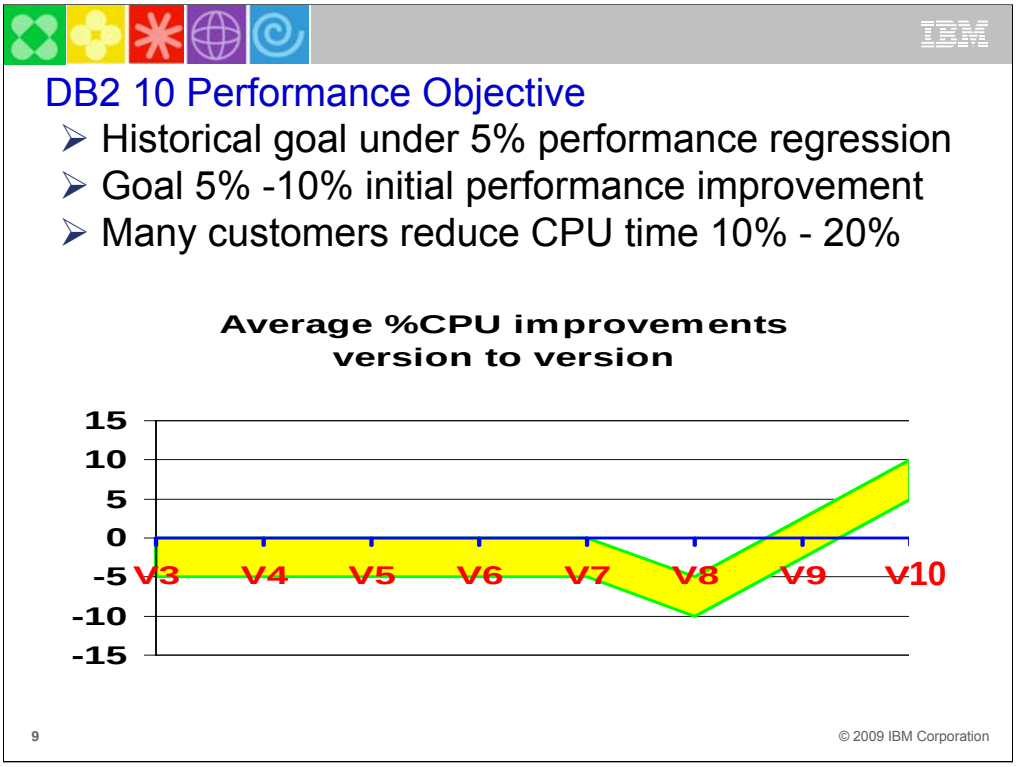

The objective for general transaction and batch performance has been to minimize the regression. Version 2 in 1988 provided a substantial Improvement in transaction and batch work, but the past 21 years have seen the focus on removal of bottlenecks, scalability, query performance, and minimizing performance regression. DB2 for z/OS V8 had more regression, with it's engineering for 64 bit, Unicode, and larger scaling. DB2 9 was better, generally in the +3% to -3% range for transactions and batch. DB2 9 provided much better performance for utilities, often in the range of 20% CPU reduction.

DB2 10 will see many customers with 5% to 10% CPU reduction in transactions and batch just by migrating to DB2 10. As always, customer experiences will vary. Many of the key improvements deliver in Conversion Mode and require no actions from customers. Memory improvements help with scalability. Improvements for CPU efficiency, chaining the open, fetch and close, parallel index IO, index performance, and fewer reorgs occur in CM with no action. The buffer pool enhancements require systems work. Optimization enhancements require rebinding. Some important enhancements, like hash access, index include columns, and inline LOBs require NFM and database administration.

•Reducing CPU from DB2 9 to DB2 10 without significant administration or application changes is the primary thrust of the performance work. Most of the changes are related to CPU caching and path length improvements inside the DB2 engine, so that applications changes aren't needed to benefit from the improvements. DB2 can take advantage of new hardware instructions without needing to have other techniques for older processors which do not have fast implementations of the new instructions.

•This work is preliminary, but the performance plan for DB2 10 is much more aggressive than in any recent version. The last version which contained significant improvements for reducing CPU time in transactions and batch was Version 2 in 1988. Versions 3 to 9 made improvements in queries and in utility CPU time and provided many scalability improvements, but little reduction in transaction CPU time, other than in specific situations.

•As customers move from DB2 V8 to DB2 9 CM, they generally find some CPU improvements, often in the utilities. As customers move to DB2 10 CM, we anticipate a bigger reduction coming from transactions and batch work. REBIND will improve optimization and activate certain internal DB2 performance improvements. The largest improvements are expected for applications that can use the database changes, such as a hash for primary key access, and SQL improvements in DB2 10.

•We expect DB2 10 to run only on z10, z9, z890, z990, and later processors, and to provide CPU reductions from the beginning, with improvements in CM, but more dramatic reductions for applications that can take advantage of the improvements in application design.

•64 bit instructions were more expensive than 31 bit, recovered in DB2 9 by staying flat and now exploiting it in DB2 10 to gain improvement in virtual storage constraint relief.

DB2 10 for z/OS: Out-of-the-Box Savings

Up to 20% CPU reductions for transactions, queries, and batch

- Out-of-the-box CPU reductions of 5-10% for traditional workloads
- Out-of-the box CPU reductions of up to 20% for new workloads
- Up to additional 10% CPU savings using new functions

Scales with less complexity and cost

- 5-10x more concurrent users – up to 20,000 per subsystem
- Significant scale-up capabilities in addition to existing scale-out support
- Consolidate to fewer LPARs and subsystems

Improved operational efficiencies and lower administration cost

- Automatic diagnostics, tuning, and compression

Even better performance

- Elapsed time improvement for small LOBS and Complex Queries

© 2009 IBM Corporation

The most exciting improvements are out-of-the-box savings by reduced CPU usage, improved productivity, and expanded scale-up capabilities for all workloads.

Just getting in DB2 10 in CM with some REBINDs will generally yield 5% to 10% CPU reductions. More work to understand your workloads and make the appropriate changes can have extra savings.

Many more concurrent users for a single subsystem result from moving most memory above the bar. Additional changes help with more concurrency. Scaling improvements will help many customers.

Administrators and programmers will be more productive with a wide range of improvements.

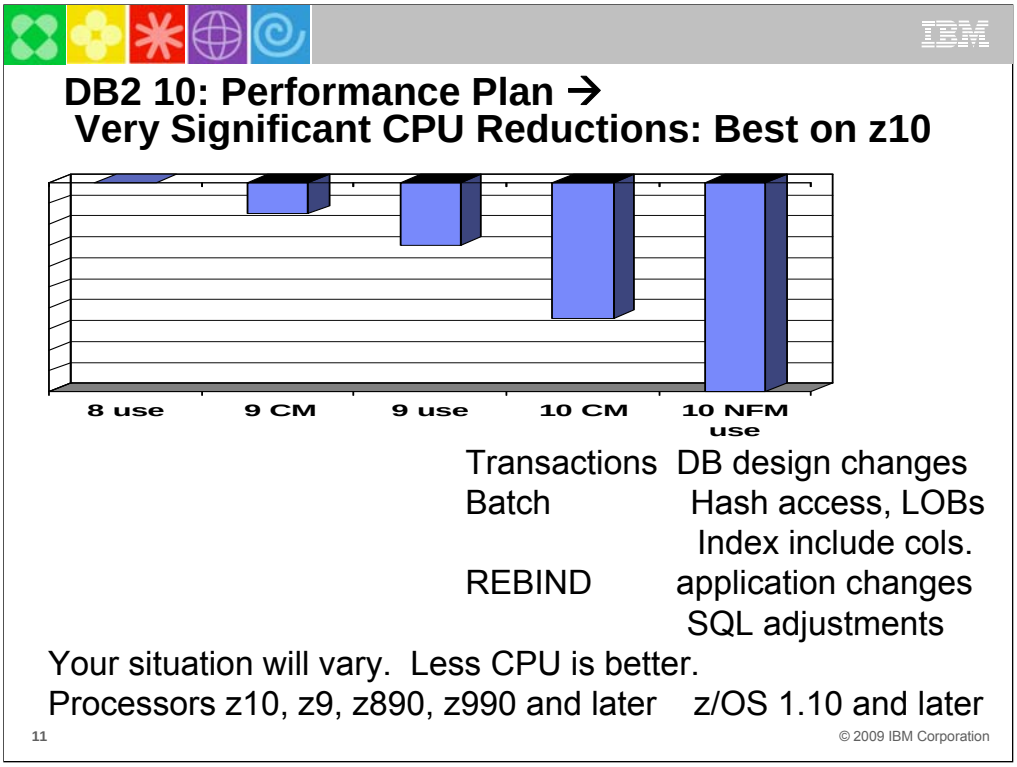

Reducing CPU from DB2 9 to DB2 10 without significant administration or application changes is the primary thrust of the performance work. Most of the changes work with CPU caching and path lengths, so that applications are not changed. We can take advantage of new instructions without needing to have other techniques for older processors which do not have fast implementations of the new instructions.

This work is preliminary, but the performance plan for DB2 10 is much more aggressive than in any recent version. The last version which contained significant improvements for reducing CPU time in transactions and batch was Version 2 in 1988. Versions 3 to 9 made improvements in queries and in utility CPU time and provided many scalability improvements, but little reduction in transaction CPU time, other than in specific situations.

As customers move from DB2 V8 to DB2 9 CM, they generally find some CPU improvements, often in the utilities. As customers move to DB2 10 CM, we anticipate a bigger reduction coming from transactions and batch work. REBIND will improve optimization. The largest improvements are expected for applications that can use the database changes, such as a hash for primary key access, and SQL improvements in DB2 10.

We expect DB2 10 to run only on z10, z9, z890, z990, and later processors, and to provide CPU reductions from the beginning, with improvements in CM, but more dramatic reductions for applications that can take advantage of the improvements in application design.

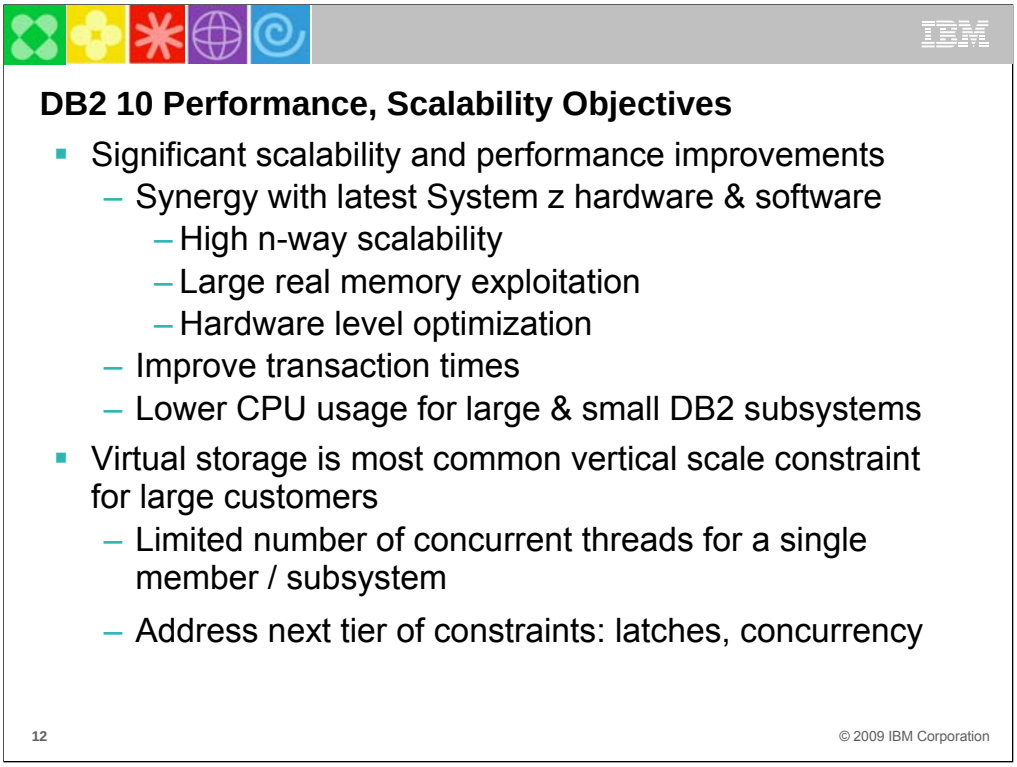

Providing significant scalability and performance improvements is an important DB2 10 objective. Synergy with the latest System z10 processors and follow-on machines provides part of the improvements. Being able to deliver high scalability for increasing numbers of processors is important for growth and costs. Being able to use large real memory effectively is required for scalability. Working with the hardware to improve CPU time by using new instructions and improving memory access and cache access is growing more important.

Synergy with z/OS 1.10 and later helps with managing larger volumes, and can help with memory, such as using 1 MB pages to manage the large amounts of memory.

The results are expected to be improved transaction times, with lower CPU usage for both large and small DB2 subsystems on transaction and batch workloads.

The most significant barrier to vertical scaling is virtual storage below the 2 GB bar. Moving storage above the bar will allow many more threads, five to ten times as many for most customers. Increasing the number of concurrent threads will expose the next tier of constraints, which will also be addressed.

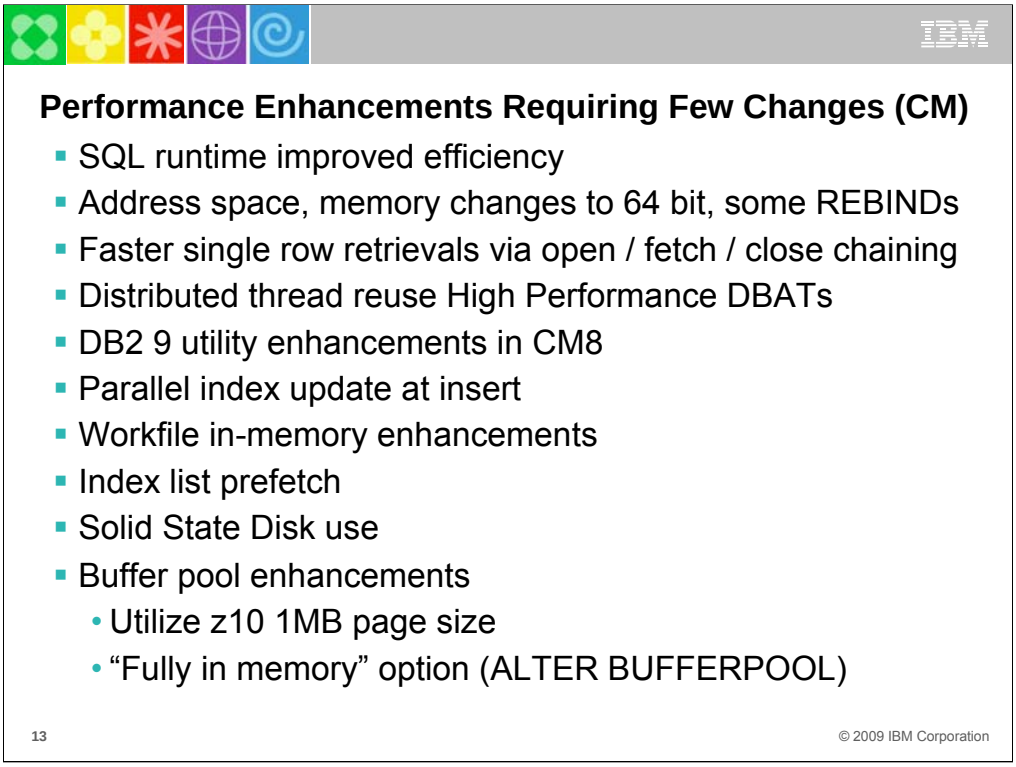

These are the improvements which we expect almost every customer to see as soon as DB2 10 is running, even in conversion mode.

CPU times are reduced for SQL running transactions and batch which are generally the peak customer workload. These techniques take very little change, but the buffer pool enhancements do need an ALTER BUFFERPOOL command.

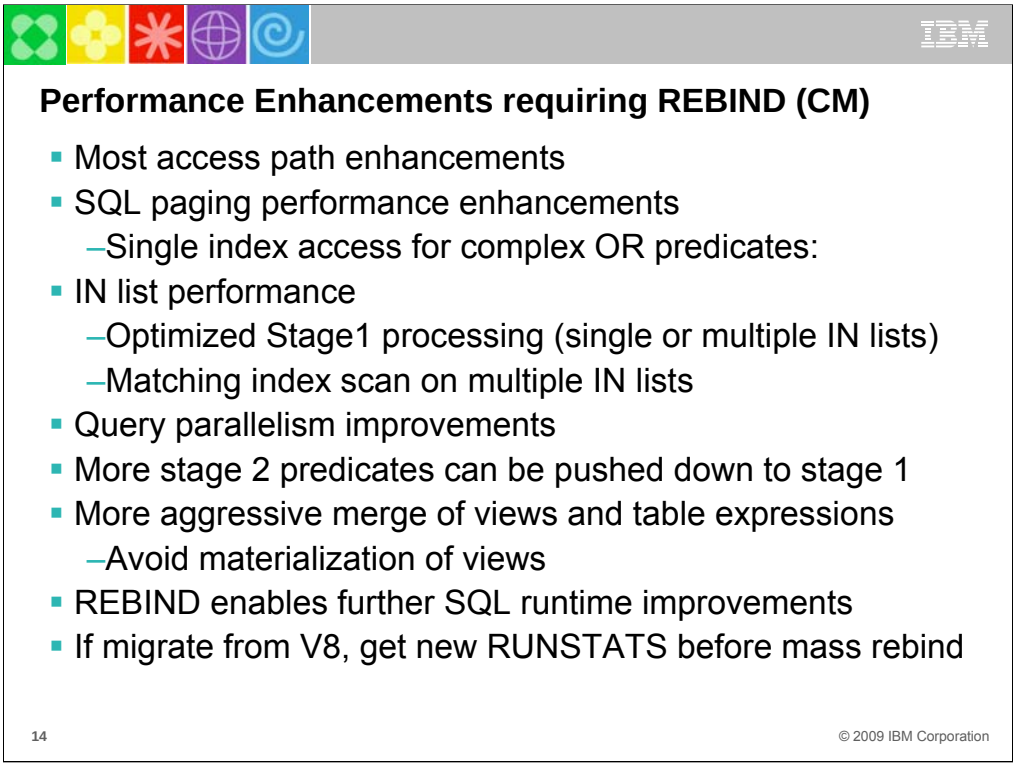

Rebind is required for a long list of improvements in optimization and parallelism. The key improvements from REBIND in CM include SQL inlist improvements, SQL paging enhancements, query parallelism improvements, and more aggressive view and table expression merge.

If you are migrating from DB2 V8, then you will want to get improved statistics for cluster ratio, data repeat factor and high cardinality non-uniform distribution of data by running RUNSTATS before you REBIND.

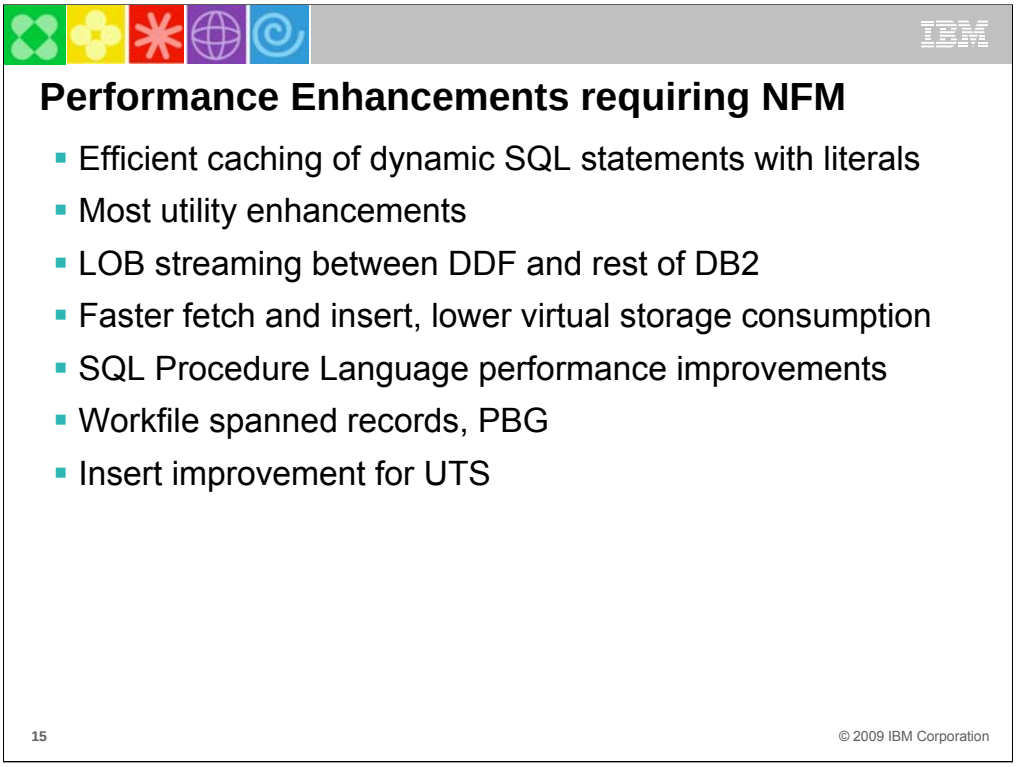

Some of the performance improvements require new function mode and some work by database administrators to tune the database design and often to REBIND. Efficient caching for literals needs a rebind in NFM.

DB2 10 supports partition-by-growth table spaces in the WORKFILE database and provides in-memory work file enhancements in the WORKFILE database.

In the WORKFILE database, DB2 supports simple predicate evaluation for work files. This enhancements reduces the CPU time for workloads that execute queries that require the use of small work files.

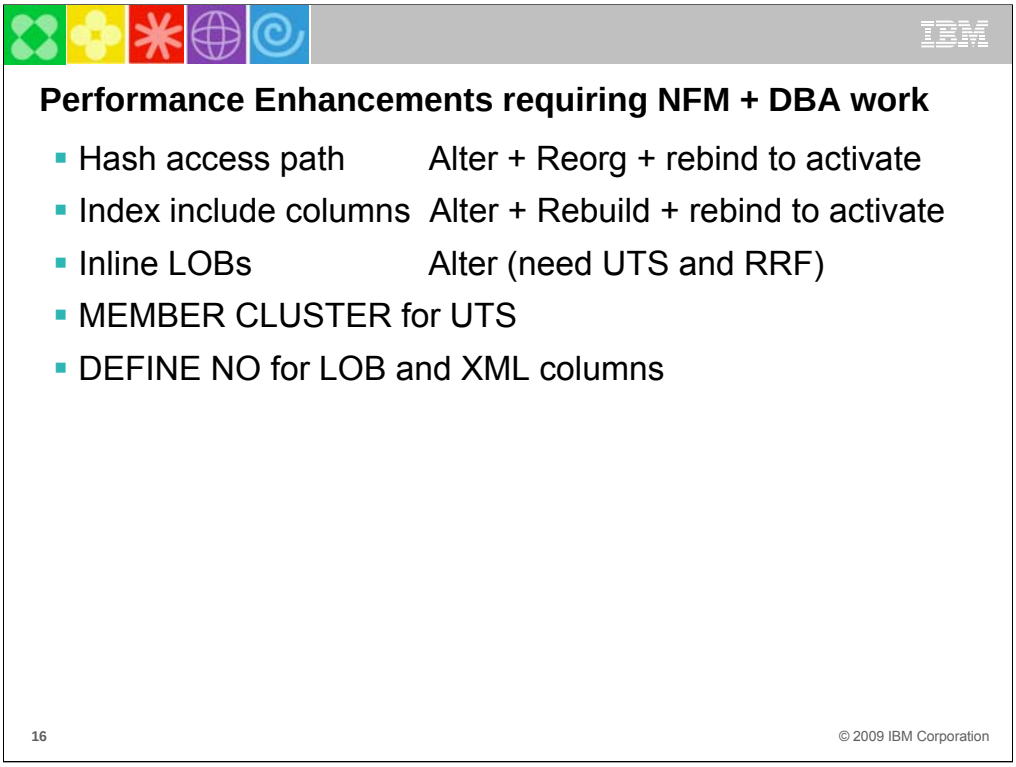

Some of the performance improvements require new function mode and some work by database administrators to tune the database design and often to REBIND. Efficient caching for literals needs a rebind in NFM.

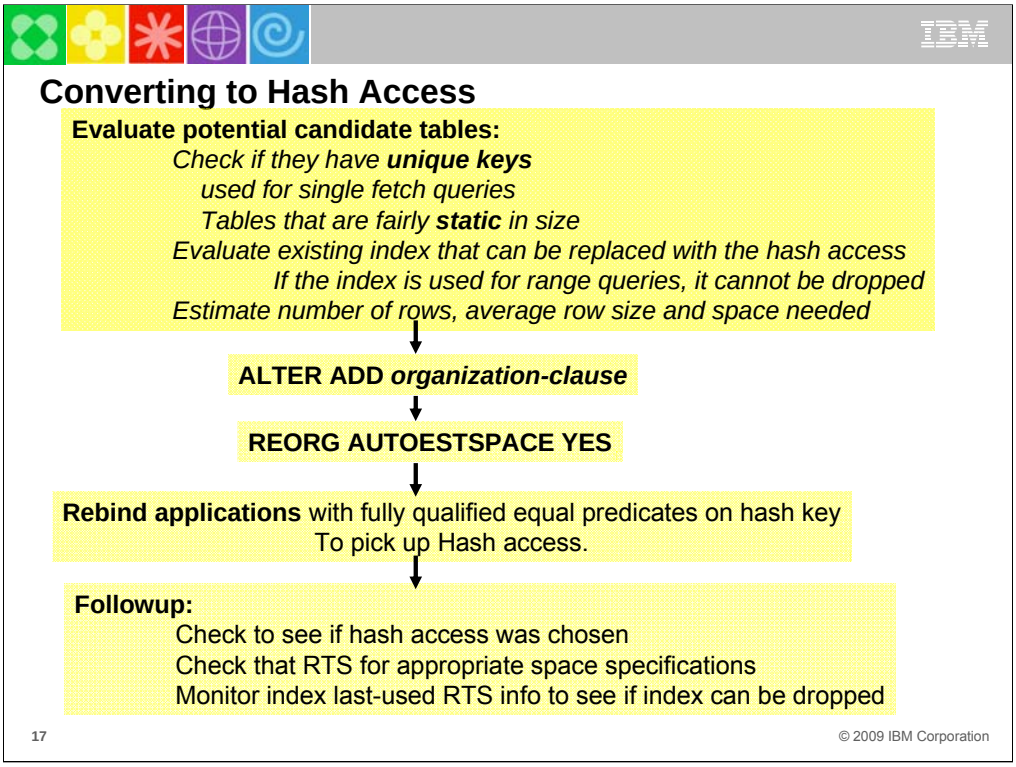

This is an example of the process needed to convert to hash access from index access. The first step is to see if the key conditions are met.

The process for converting includes the ALTER, a REORG and REBIND for the applications.

Then you need to be sure that the benefits are delivered, checking that hash access is used, space is accurate and the performance improvements are provided. At some point in the future, you can see if the index can be dropped or if some processes need the index.

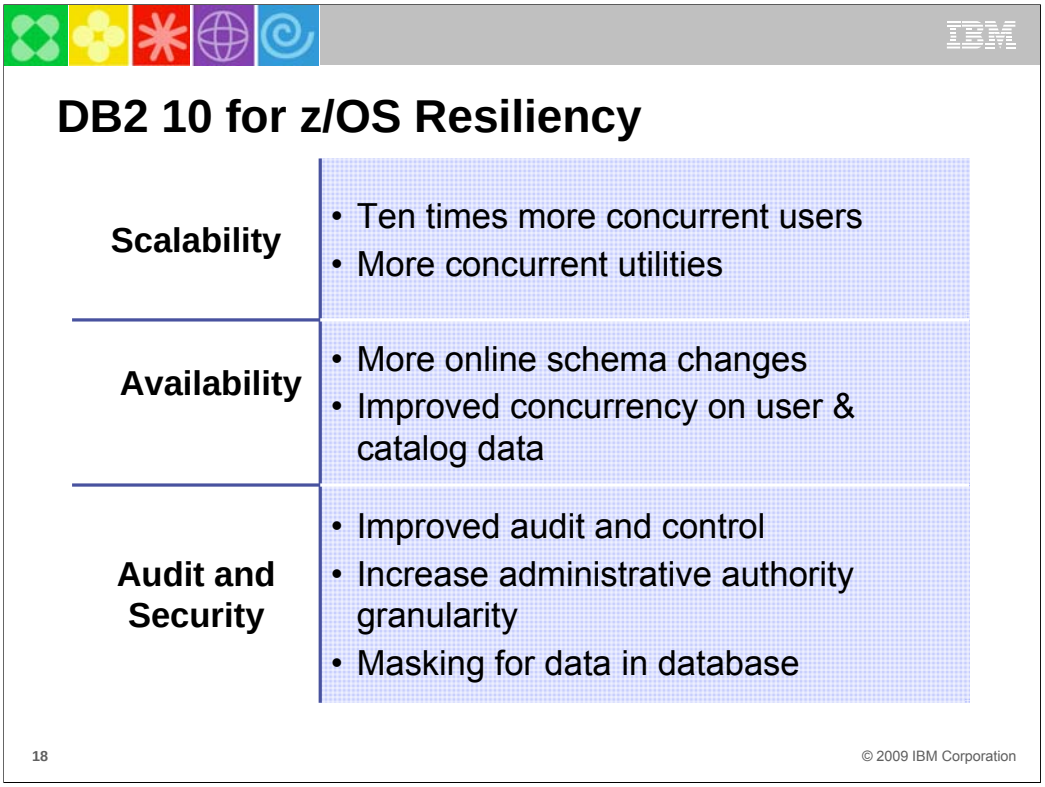

Continuous availability requirements continue to escalate. Large batch and maintenance windows are in the past. Those windows are being closed on the fingers of DBAs. DBAs increasingly need the ability to make all changes and to do all maintenance activities online or around the clock. DB2 10 allows more online schema changes with an ALTER for a PENDING change, then an online REORG to take effect.

Scalability improvements make DB2 simpler to run, simpler to scale and keep the database up when the load is much more than expected.

Improved audit and security controls reduce your risk and make regulatory compliance easier with improved granularity for administrative privileges, the ability to mask data and new audit capabilities.

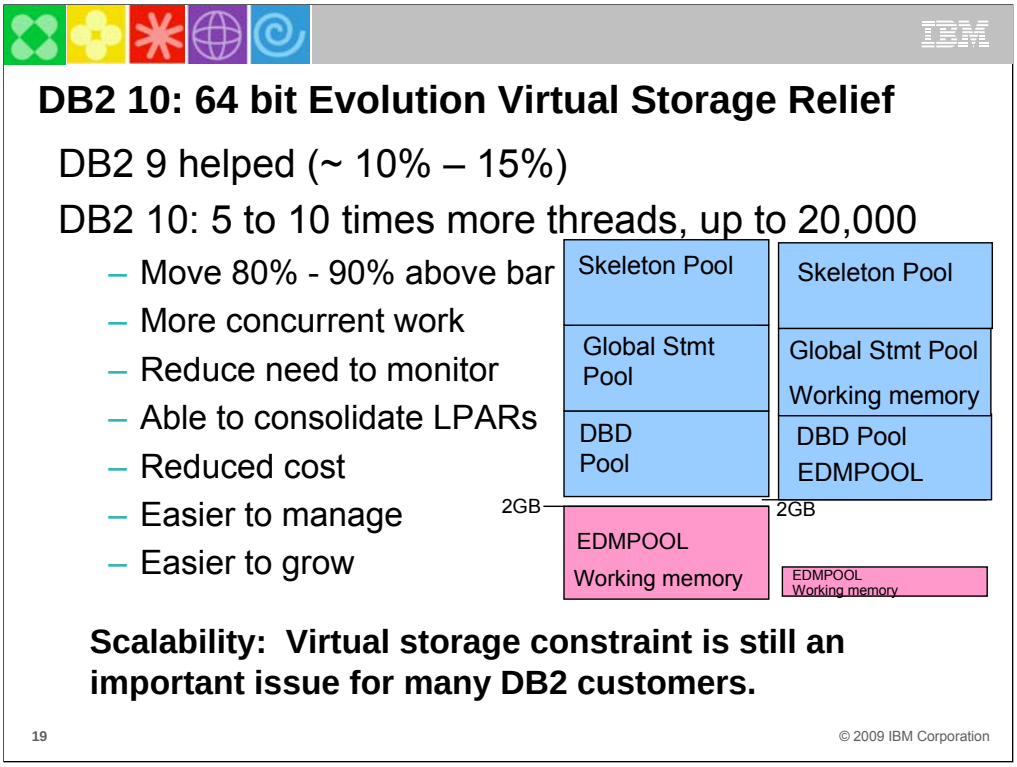

Virtual storage is most common constraint for large customers. Virtual storage can limit the number of concurrent threads for a single member or subsystem.

The DB2 9 virtual storage objective was 10-15% relief. The DB2 10 target is 80% to 90% of the DBM1 address space. We expect the result to be the ability to run much more concurrent work, with an early guess of 3 to 5 times more threads.

Storage monitoring should be drastically reduced. Customers are consolidating LPARs. Sometimes they need to have more than one DB2 subsystem on an LPAR, costing real storage and CPU. With these changes, work can run in one DB2 subsystem, rather than needing more members.

The net for this change is expected to be reduced cost, improved productivity, easier management, and the ability to grow DB2 use much more easily.

Increasing the number of concurrent threads will expose the next tier of constraints. DB2 10 will address a number of the next items, such as utility locking, catalog concurrency.

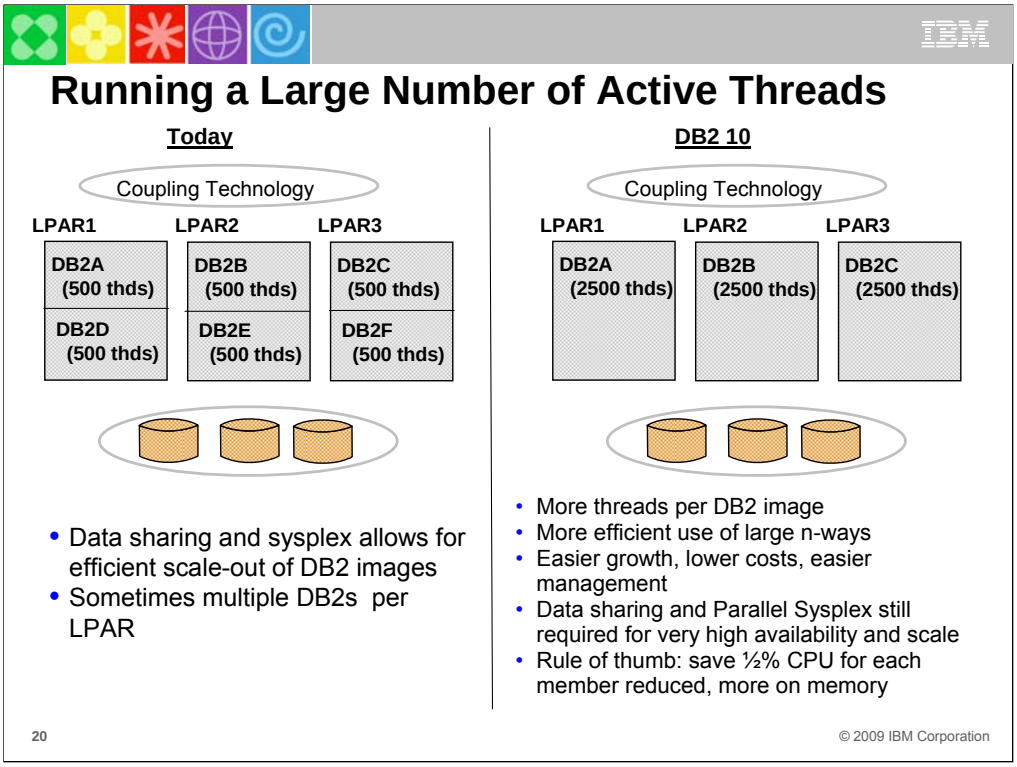

Customers are constrained by virtual memory to various degrees. This slide shows a relatively extreme situation experienced by some customers today. With a maximum of 500 threads (very dependent upon workload) in a DB2 subsystem, this customer is using two DB2 subsystems in the same data sharing group on a single LPAR. This is not efficient for memory of CPU, but avoids the memory constraints with fewer LPARs. Additional relief for virtual storage comes with IMS V11 and other products.

This example allows customers to run 10 times as many threads in a single DB2 subsystem, improving efficiency for storage and CPU. The biggest change is easier management and simpler growth. Most customers use data sharing for high availability, and that need still exists. Extreme scale continues to need data sharing, but fewer data sharing members can mean easier management and reduced resource consumption. In this example, changing from 6 members to 3 can mean a reduction of 1.5% in CPU time, as a rule of thumb.

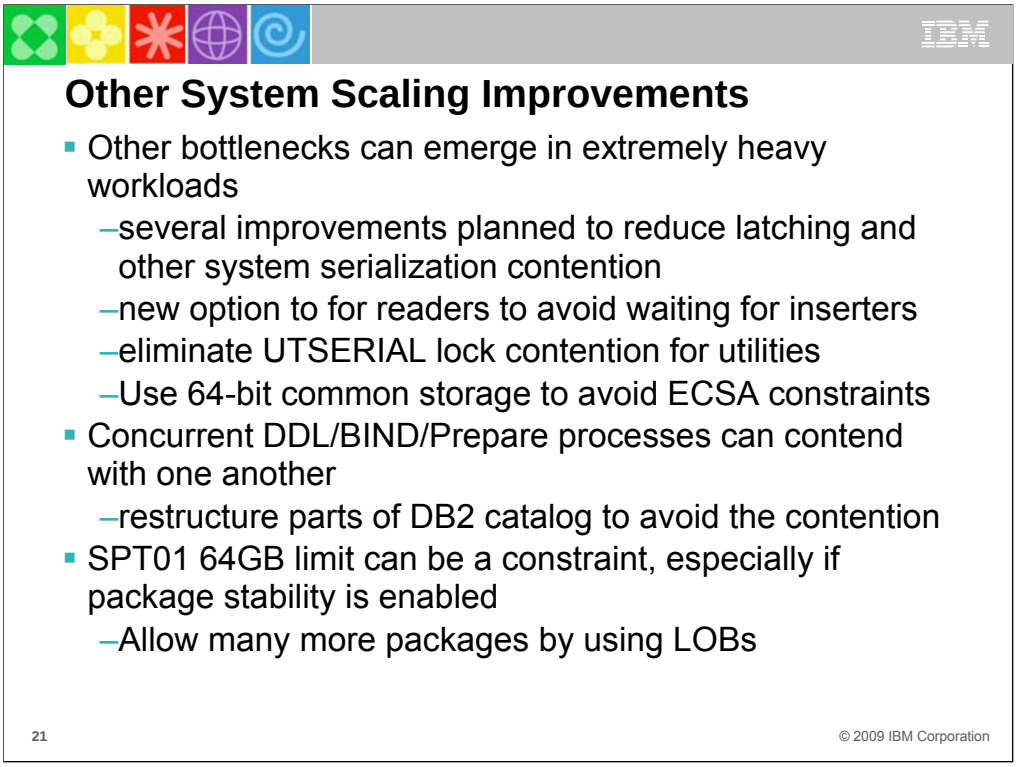

Increasing the number of concurrent threads will expose the next tier of constraints. DB2 10 will address a number of the next items, such as utility locking and catalog concurrency.

The UTSERIAL lock means that scheduling 20 concurrent REORGs for hundreds of partitions in each one will result in deadlocks too often. Reducing the granularity by removing this lock means that the jobs run. DB2 10 eliminates the use of UTSERIAL by DB2 utilities. This enhancement prevents the majority of timeouts on the global UTSERIAL lock resource.

Improving the catalog structure to allow row level locking can improve concurrency substantially.

The DB2 catalog structure is changed to move most of the large fields with repeating rows of data into LOB columns, eliminating the 64 GB limit and making the information more readable by separating character from binary data. The LOB columns are inline for improved performance.

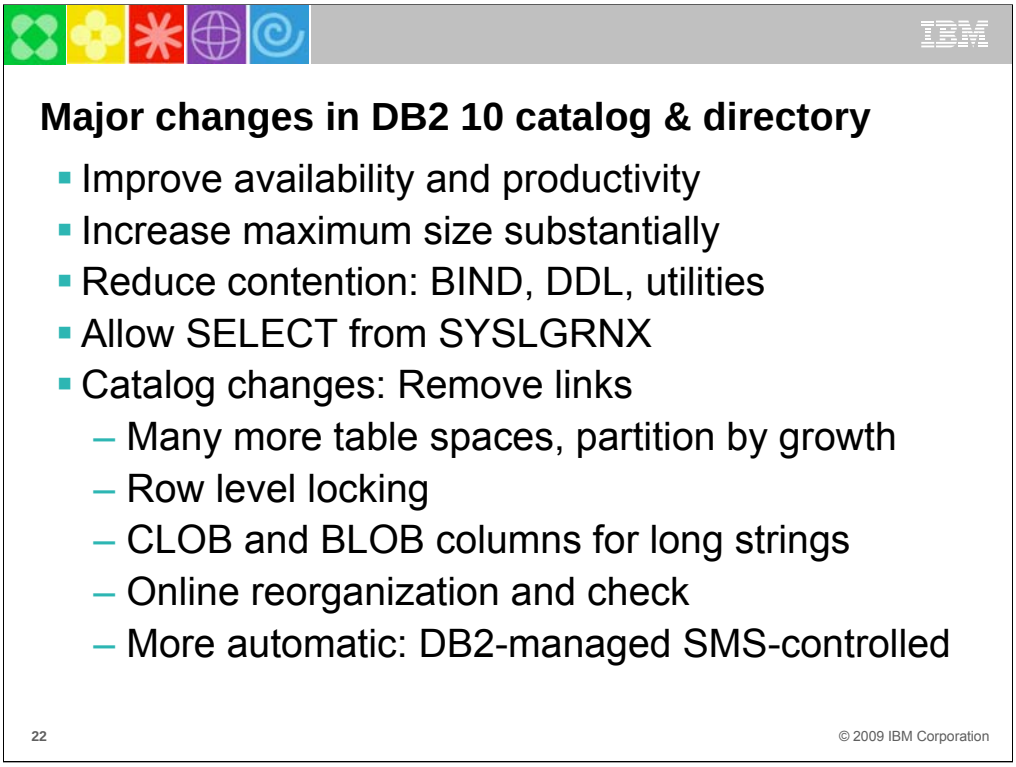

The DB2 catalog and directory are restructured in DB2 10 to improve productivity and availability. The current size limits are increased substantially and contention among process like BIND, dynamic SQL, data definition and utilities is reduced.

The primary techniques are changes in the DB2 catalog to remove links and the special structures for the catalog. These table spaces change from many tables to one table per table space in a partition by growth table space defined as DSSIZE 64 GB and MAXPART 1. Row level locking is used in place of page level locking. The new catalog tables use a partition by growth universal table space structure. Each table space holds a single table, so many more table spaces are needed. Rather than repeating columns with parts of long strings, the catalog will use CLOB and BLOB columns to store the data, expanding maximum sizes. The new structure allows more standard processes, so that all catalog tables can be reorganized and checked online.

The DB2 catalog changes from using manual definition and extension to DB2 managed data sets under SMS control. The changes improve productivity and availability, but take time to set up.

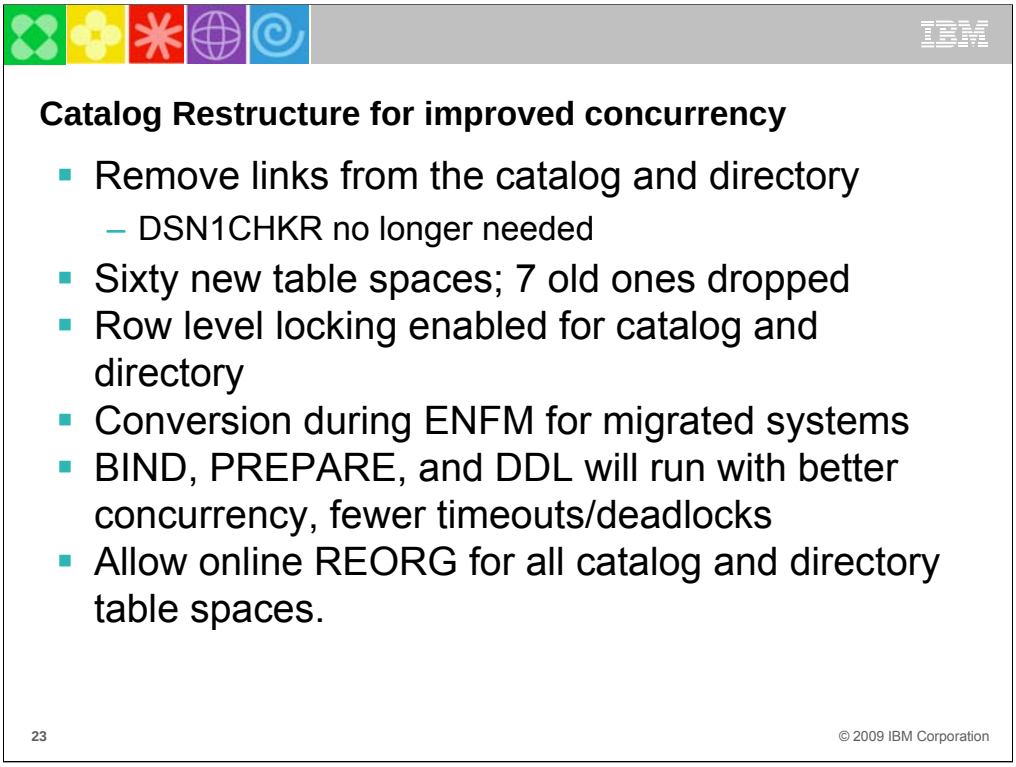

All the links in the catalog and directory will be replaced with referential integrity by running the DSNTIJEN job, which is also known as the ENFM processing. Links in the following table spaces are removed: DSNDB06.SYSDBAUT, DSNDB06.SYSGROUP, DSNDB06.SYSDBASE, DSNDB06.SYSPLAN, DSNDB06.SYSVIEWS, DSNDB01.DBD01.

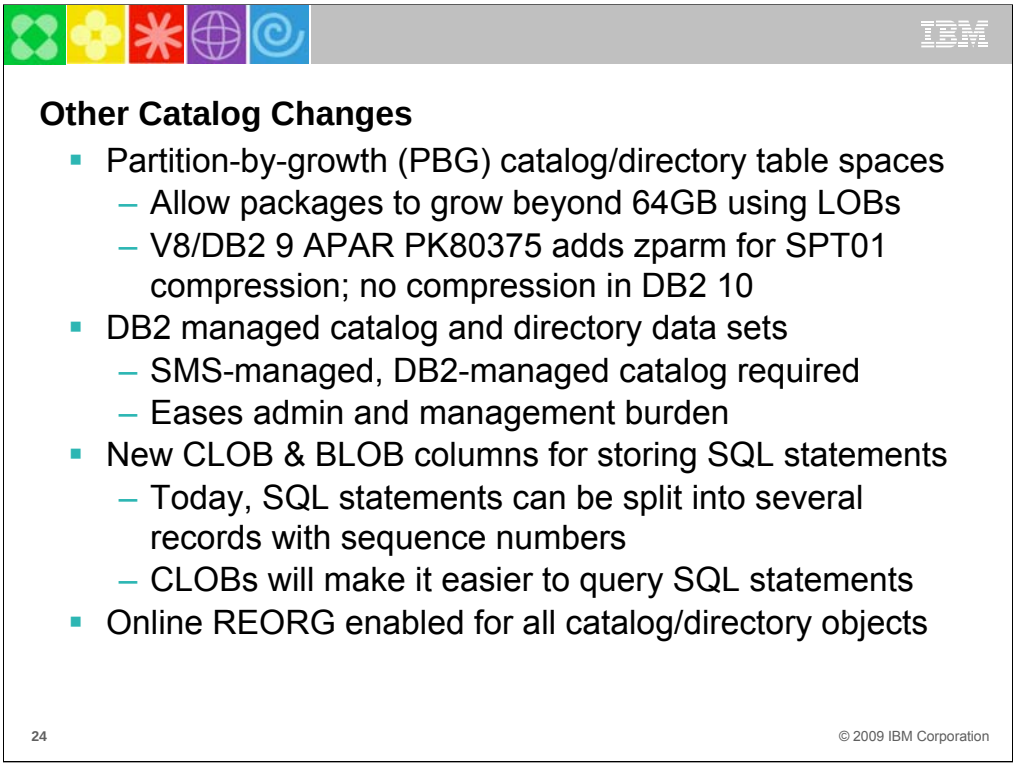

Moving some tables to partition by growth table spaces with DSSIZE 64 GB, MAXPART 1 and LOB columns will make access slower in some cases, but much more standard, more scalable and more available. Inline LOBs are used to help the performance when the amount of data is small. While compression is provided for a few customers with an APAR and a zparm change, the CLOB and BLOB columns do not have compression.

The following table spaces in the catalog and directory are moving to the new PBG table spaces: DSNDB06.SYSOBJ, DSNDB06.SYSVIEWS, DSNDB06.SYSPLAN, DSNDB06.SYSDBASE, DSNDB06.SYSDBAUT, DSNDB06.SYSGROUP, DSNDB01.DBD01, DSNDB01.SPT01. Every table in these table spaces will have its own PBG table space.

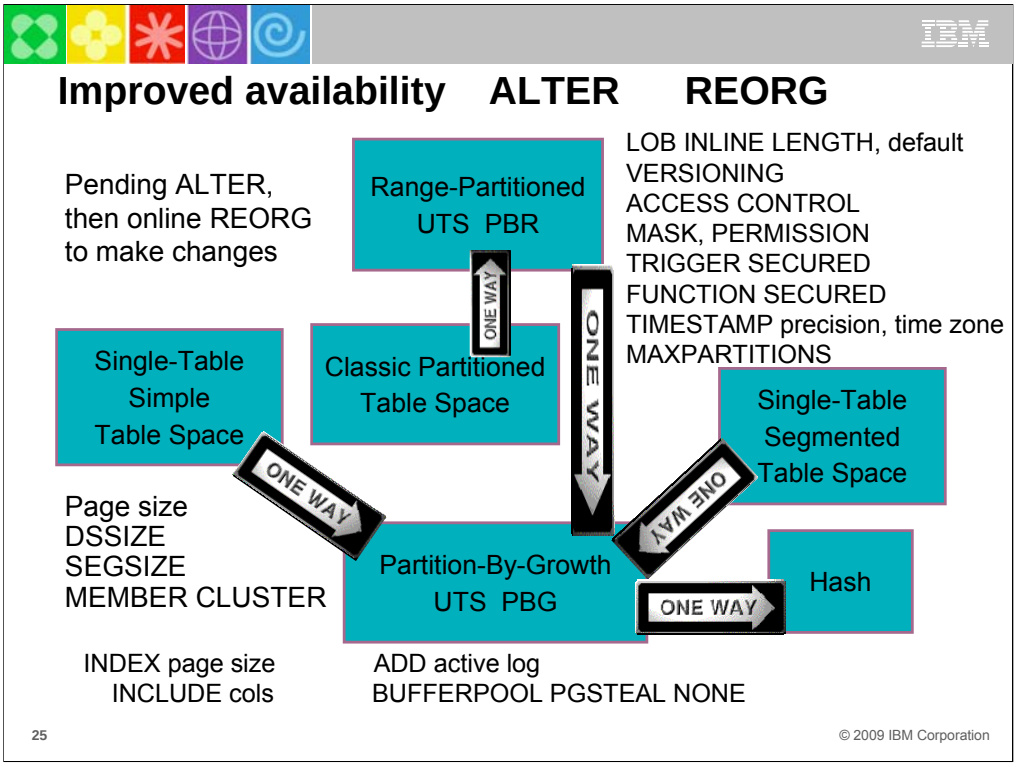

DB2 10 brings many new options for ALTER BUFFERPOOL, TABLE, INDEX, and TABLE SPACE. These are the changes in table space type in diagram form, adding the ability to change from single table segmented, simple or partitioned table spaces to universal table spaces. This release also adds the ability to modify some new attributes, the page size, the dataset size, and the segment size. These attributes are pending changes when the ALTER Is performed, then the changes take place when the online REORG occurs. If a mistake is made before the REORG, then DROP PENDING CHANGES allows you to start again. More alters are provided for universal table spaces, adding the ability to change to MEMBER CLUSTER and the ability to ALTER inline length for LOB columns.

Indexes can now be altered to add INCLUDE columns and index page sizes can be altered, as a pending change. Bufferpools can be altered to PGSTEAL NONE, meaning that they stay resident.

What is not done? Change from multi-table segmented table space. Change back to classic simple, segmented and partitioned. The strategic choice for table space type is the universal table space. Simple table spaces are deprecated, and this version provides a migration path. The ability to add a new active log data set is included. Many online REORG restrictions are removed, to allow more online operations.

If you need more improvements in table spaces, then universal table spaces – either partition by range or partition by growth should be your choice.

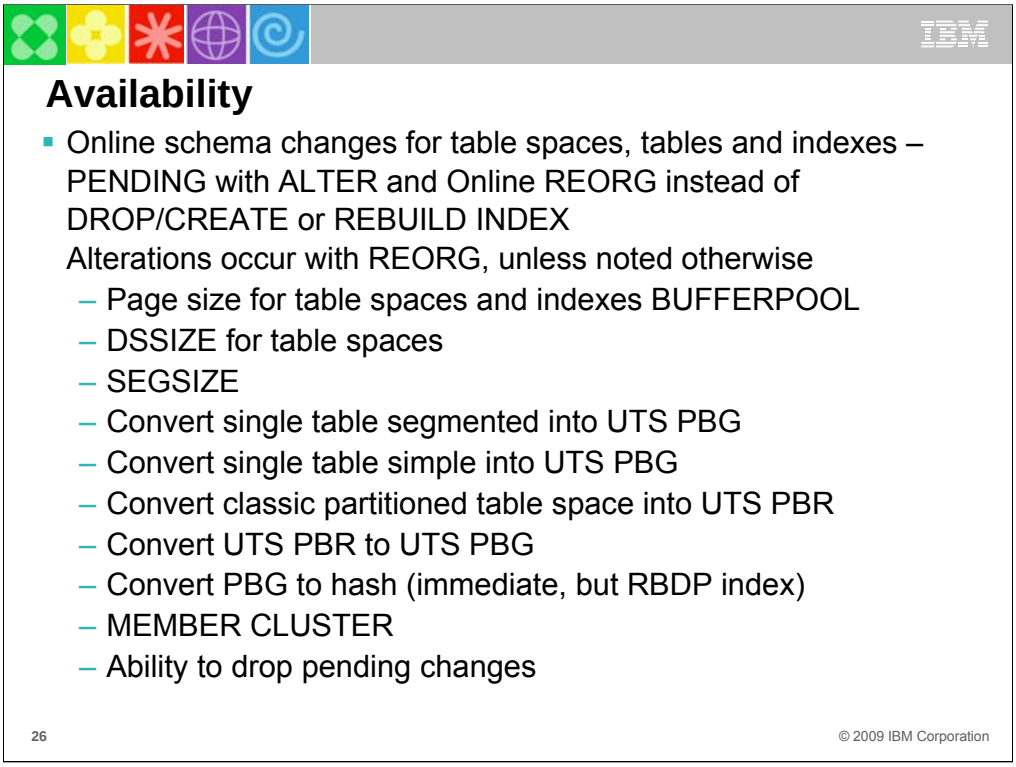

Continuous availability requirements continue to escalate. Large batch and maintenance windows are in the past. Those windows are being closed on the fingers of DBAs. DBAs increasingly need the ability to make all changes and to do all maintenance activities online or around the clock. DB2 10 allows more online schema changes with an ALTER for a PENDING change, then an online REORG to take effect. ALTER a simple or segmented table space containing a single table or a partitioned table space to a universal table space. Page size and member clustering can be altered. Index changes become less disruptive. Pending changes which have not been completed with a REORG can be dropped.

REORG is improved to allow SHRLEVEL(CHANGE) for LOBs.

Consistent image copies can be provided without a quiesce.

Inline copies to allow for dataset-level FlashCopy.

Online REORG usability and performance enhancements are provided.

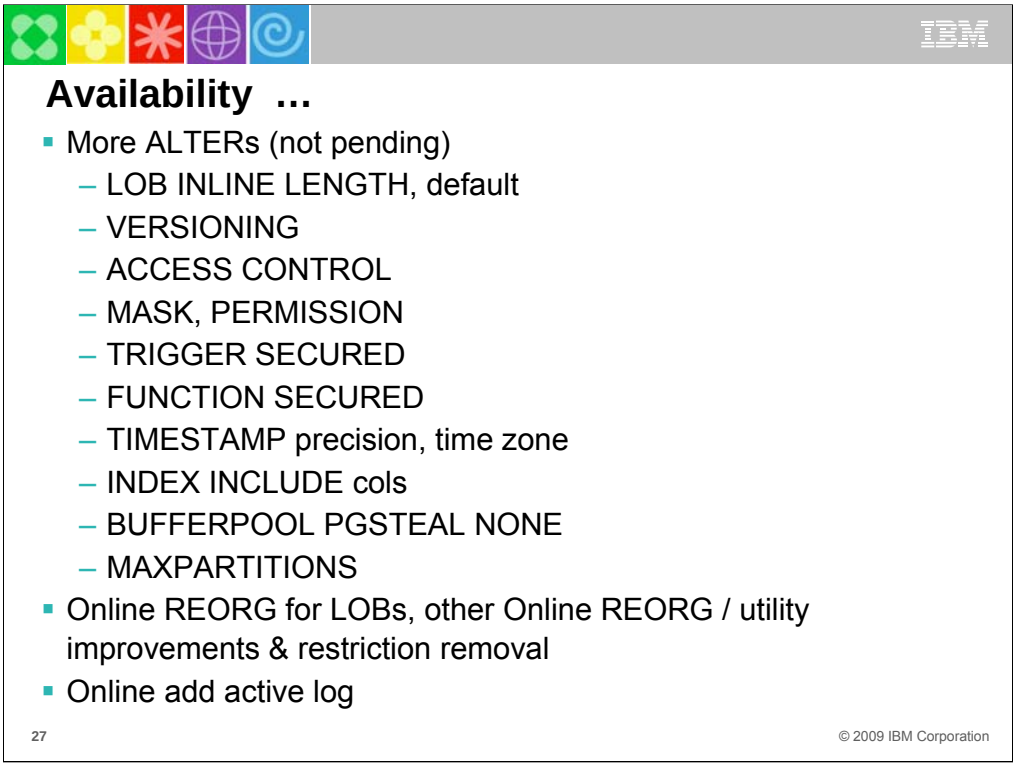

Continuous availability requirements continue to escalate. Large batch and maintenance windows are in the past. Those windows are being closed on the fingers of DBAs. DBAs increasingly need the ability to make all changes and to do all maintenance activities online or around the clock.

DB2 10 allows more online schema changes with an ALTER for a PENDING change, then an online REORG to take effect. ALTER a simple or segmented table space containing a single table or a partitioned table space to a universal table space. Page size and member clustering can be altered. Index changes become less disruptive. Pending changes which have not been completed with a REORG can be dropped.

REORG is improved to allow SHRLEVEL(CHANGE) for LOBs.

Consistent image copies can be provided without a quiesce.

Inline copies to allow for dataset-level FlashCopy.

Online REORG usability and performance enhancements are provided.

DB2 10: Business Security & Compliance

- Protect sensitive data from privileged users & improve productivity
  - SYSADM & DBADM without data access
  - Usability: DBADM for all DB
  - Revoke without cascade
- Separate authorities to perform security related tasks, e.g. security administrator, EXPLAIN, performance monitoring and management
- Audit privileged users
- Row and column access control
  - Allow masking of value
  - Restrict user access to individual cells

Use disk encryption

© 2009 IBM Corporation

Customers are being pressed for a wide range of improved security and compliance. Data retention is a growing need. Protecting sensitive data from the privileged users and administrators is required. Separation of authority for security, access, and some common tasks, like EXPLAIN will help. Auditing for privileged users can also make compliance simpler.

Access control is refined in several ways with better granularity for the administrative privileges and with finer grained access control at the row and column level, including the ability to mask access to some fields. Auditing is also enhanced.

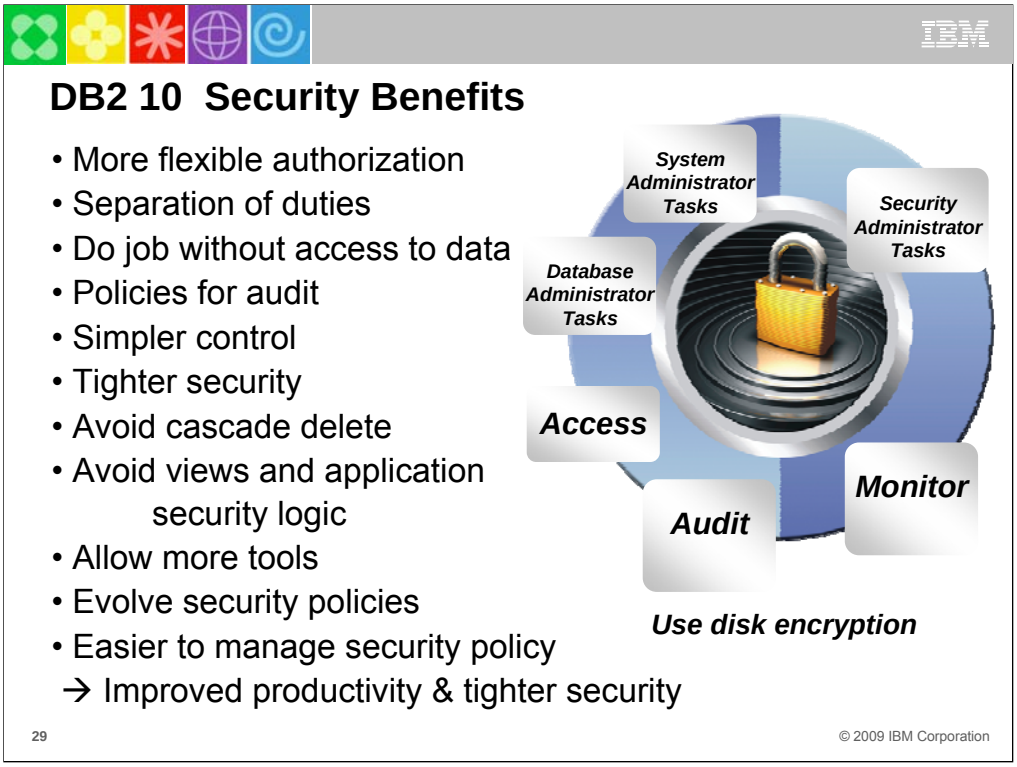

Access control is refined in several ways with better granularity for the administrative privileges and row and column access, including the ability to mask access to some fields. Auditing is also enhanced.

New security administrative authorities allow a flexible solution to allow a company to not use SYSADM by defining separate administrative authorities as required by company security policies

Performance analysts can perform their tasks without having access to user data

New audit capability allows security administrator to audit all use of new administrative authorities as well as define a set of security policies to audit the DB2 security policies

Using SQL to provide tighter security directly on tables with column masking and row filtering.

Eliminating views and security logic in applications

Preventing access from authorities including SYSADM and DBADM

Allowing usage of ad-hoc query tools, report generation tools

Making the evolution of security policies easy to deal with

Relieving the difficulties of managing security policies

DB2 10: Productivity – Doing More with Less!

- Auto statistics collection
- Easier scaling, simpler memory management
- Reduce contention, more online processing
- Access path stability
- Reduced need for REORG
  - Build compression dictionary on the fly
  - Index list prefetch enhancements
- Configure IBM UDFs and stored procedures
- Allow one SDSNEXIT data set for many subsystems
- Monitoring enhanced
  - Timeout / deadlock diagnostics
  - Identify SQL statements

Manual invocation of
•RUNSTATS
•COPY/BACKUP SYSTEM
•QUIESCE
•MODIFY RECOVERY
•REORG

30

© 2009 IBM Corporation

Some of the improvements come with Data Studio for application programming and administration – stronger cross-platform graphical interfaces, better integration with Java, improvements in the ability to develop and debug.

Some of the improvements come within DB2 for z/OS. Improvements in SQL and XML improve productivity for those who develop new applications and for those who are porting from other platforms. Some of the improvements remove complexity from application tasks.

DB2 has a strong focus on making DB2 easier to use by automating tasks and eliminating tasks where possible. Avoiding the manual invocations can also help avoid problems for running the function too often or not often enough. Where the task cannot be eliminated, the frequency and monitoring can be reduced, such as the need to reorganize. The improvements for virtual storage and for availability also help DBA productivity.
Allowing tailored names for DSNHDECP will permit many subsystems to share the SDSNEXIT data set.

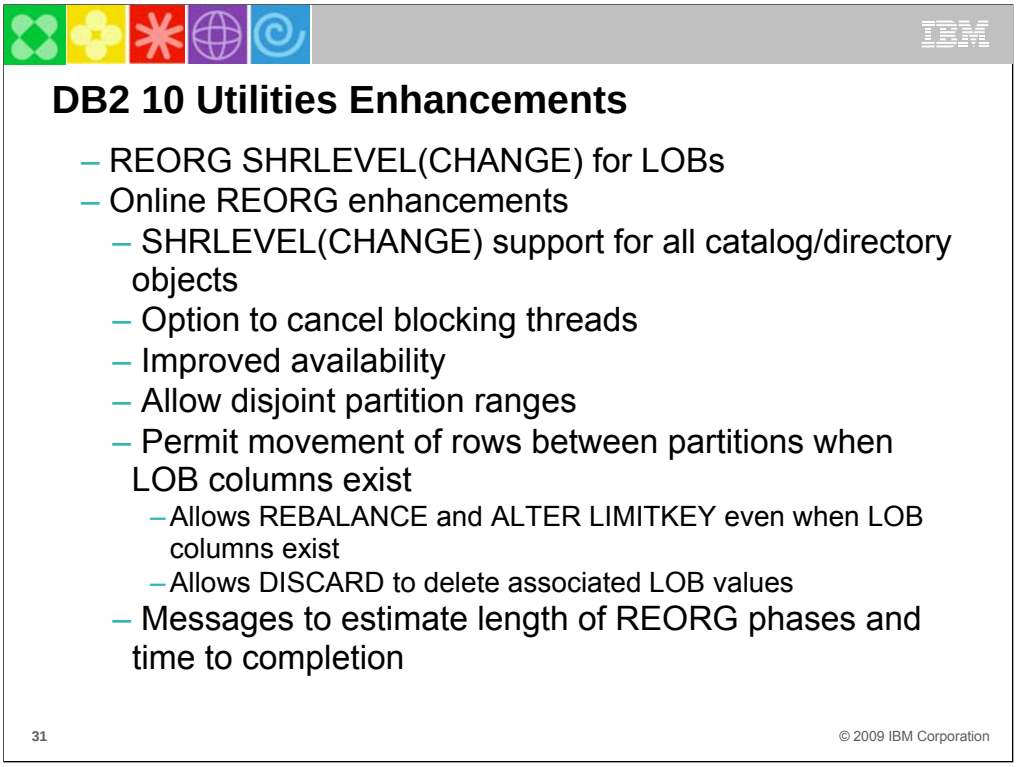

Continuous availability requirements continue to escalate. Large batch and maintenance windows are in the past. Those windows are being closed on the fingers of DBAs. DBAs increasingly need the ability to make all changes and to do all maintenance activities online or around the clock.

DB2 10 allows more online schema changes with an ALTER for a PENDING change, then an online REORG to take effect. ALTER a simple or segmented table space containing a single table or a partitioned table space to a universal table space. Page size and member clustering can be altered. Index changes become less disruptive. Pending changes which have not been completed with a REORG can be dropped.

REORG is improved to allow SHRLEVEL(CHANGE) for LOBs.

Consistent image copies can be provided without a quiesce.

Inline copies to allow for dataset-level FlashCopy.

Online REORG usability and performance enhancements are provided.

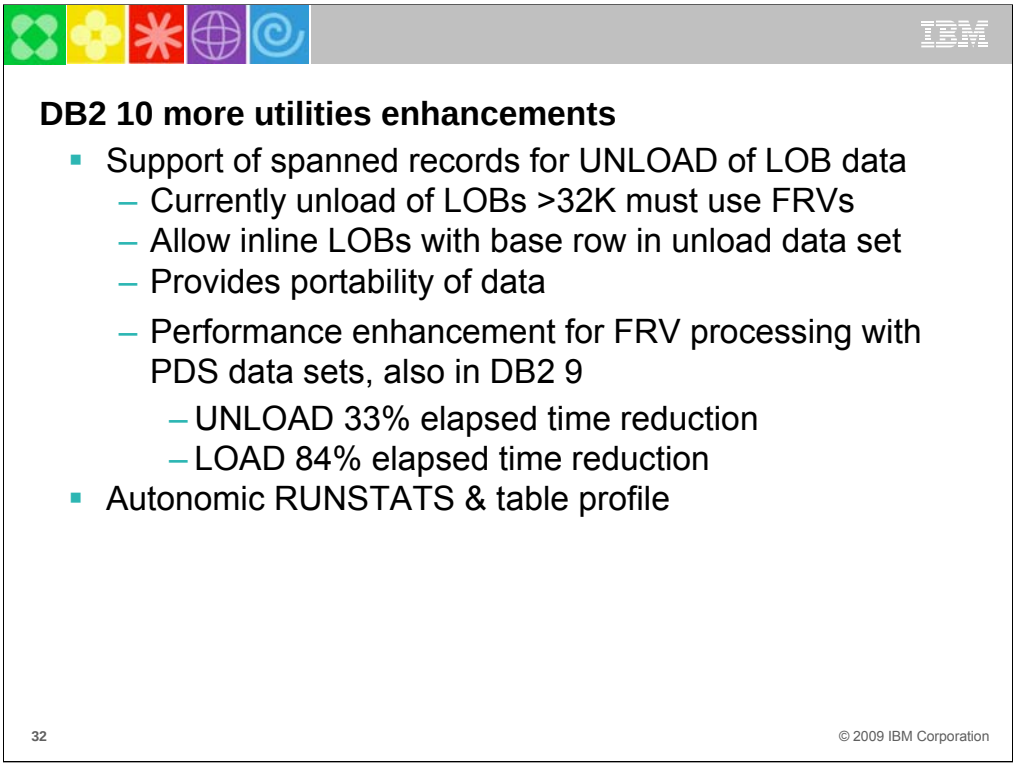

DB2 10 improves the usability and performance of online reorganization in several key ways. This release of DB2 for z/OS supports the reorganization of disjoint partition ranges of a partitioned table space, and improves availability. Also, DB2 10 removes restrictions that are related to the online reorganization of base table spaces that use LOB columns. In new-function mode, the syntax for the REORG TABLESPACE statement is changed. For partitioned table spaces, the PART specification is extended to allow for multiple parts or part ranges, and the SHRLEVEL REFERENCE and CHANGE specifications are extended to add a new keyword, AUX YES/NO. This new keyword allows for the reorganization of LOB table spaces that are associated with the base table.. DB2 10 provides several performance enhancements that reduce the need to reorganize indexes frequently, resulting in a reduction in CPU time and synchronous I/O waits.

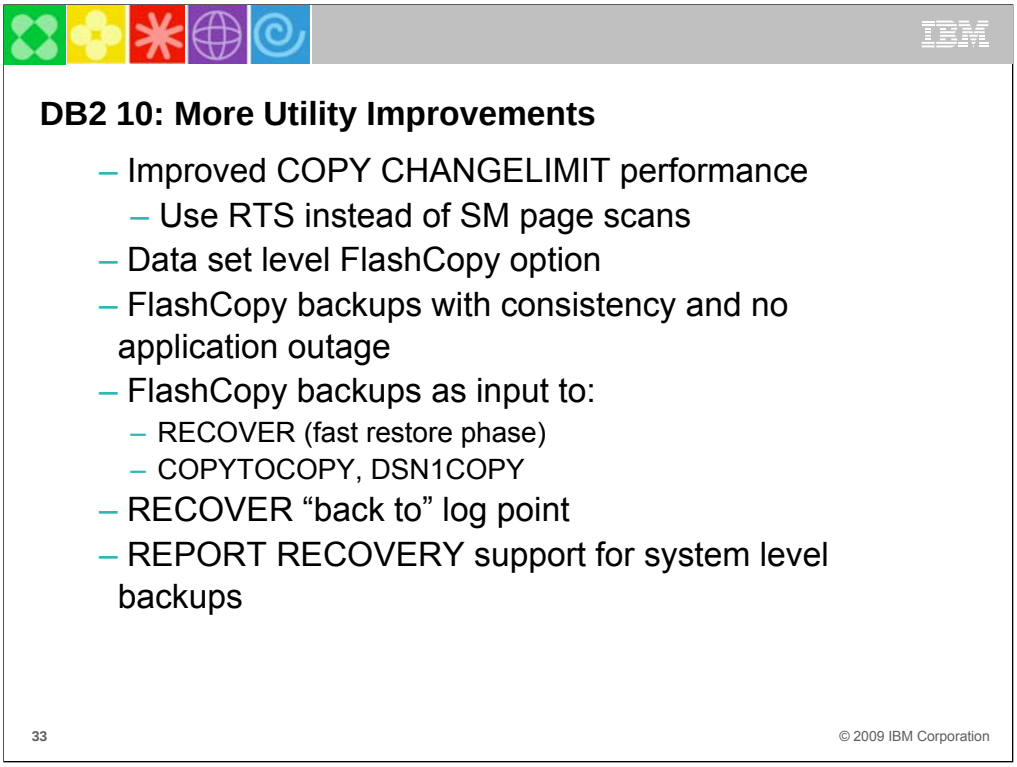

DB2 10 supports enhancements for the use of FlashCopy® technology for both backup and recovery. The goal is to keep data available, but create a consistent image copy, which is accomplished by copying the object and then backing out uncommitted changes. In addition, the RECOVER utility provides an option to backout changes when recovering data to a prior point in time. Data can be available faster than if the most recent recovery base, prior to the point that was recovered to, is restored and logs are then forward applied.

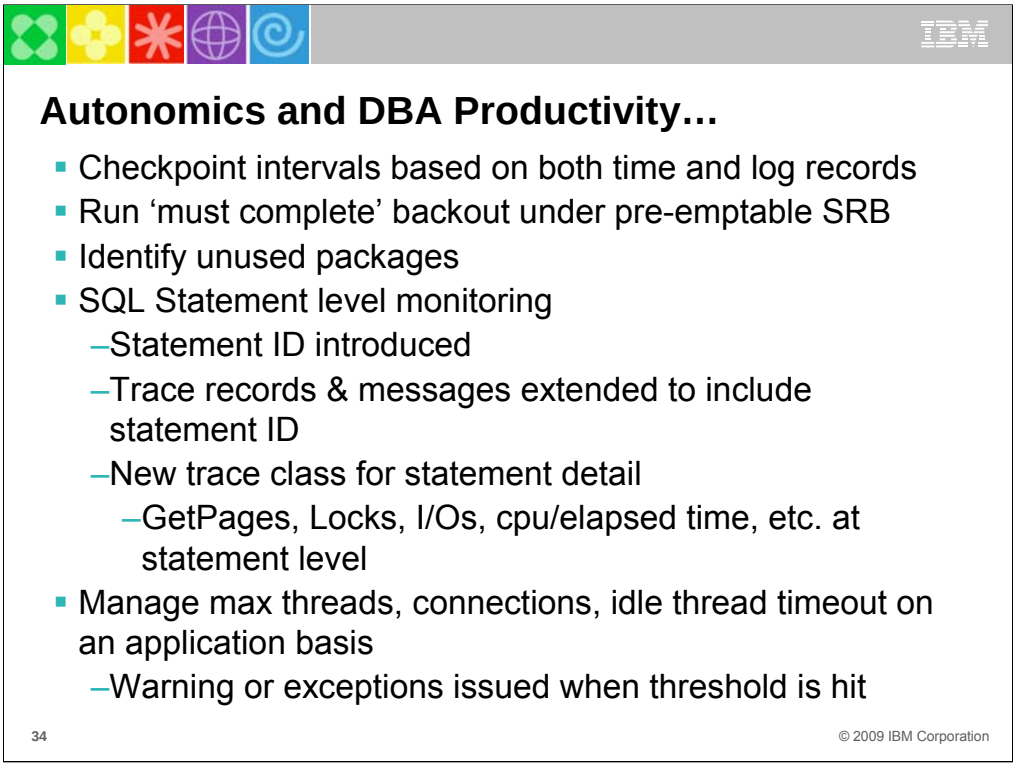

This release improves the support for monitoring within DB2 for z/OS by providing additional performance and diagnostic monitoring capabilities. Version 10 enhances performance monitoring support and monitoring support for problem determination for both static and dynamic SQL. This new support uses the Instrumentation Facility Interface (IFI) to capture and externalize monitoring information for consumption by tooling.

In order to facilitate the collection and correlation of enhanced monitoring data, this release introduces a unique statement execution identifier (STMTID). The statement ID is defined at the DB2 for z/OS server, returned to the DRDA application requester, and captured in IFCID records for both static and dynamic SQL. Through DRDA, the statement ID is returned to the client drivers, along with a compilation source identifier and a compilation time.

To support problem determination, the statement ID is provided in several existing messages, including messages related to deadlocks and timeouts. In these messages, the STMTID is associated with thread information. You can use this thread information to correlate the statement execution on the server with the client application on whose behalf the server is executing the statement.

To support performance monitoring, some existing trace records that deal with statement-level information are modified to capture the new statement ID and new statement-level performance metrics. Also, this release introduces new trace records that provide access to performance monitoring statistics in real time, and allow tooling to retrieve monitoring data without requiring disk access.

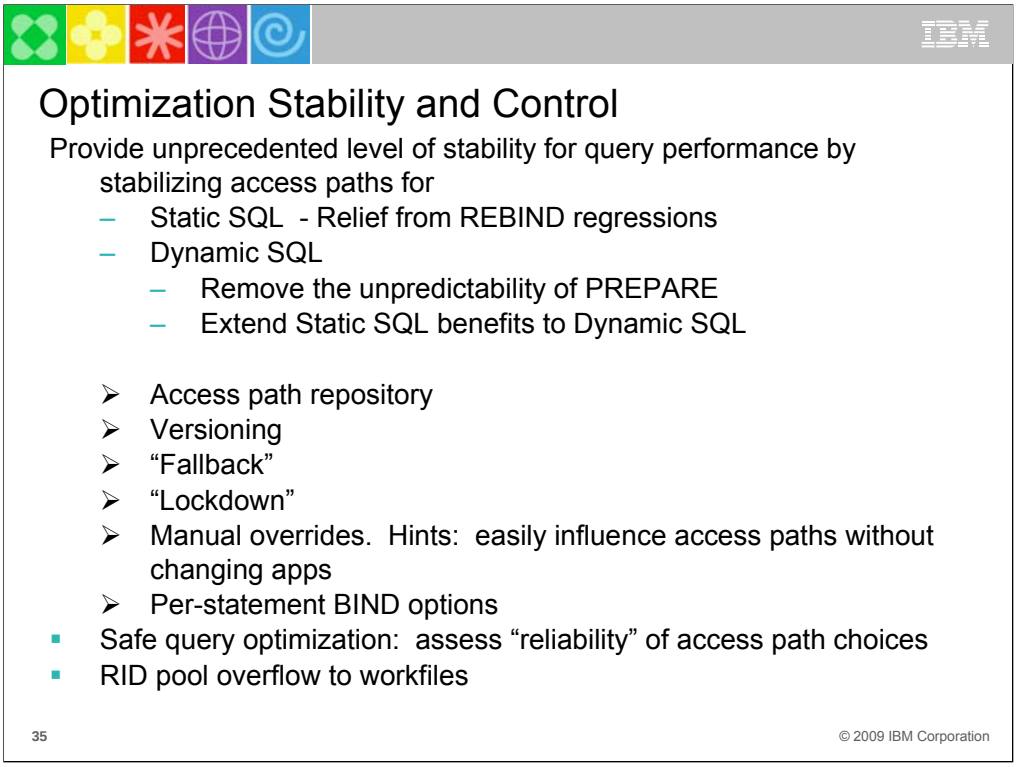

System level plan hints allows you to set a hint that will globally be used for a statement by all users. Before you had to put the hint in to the plan table for all users' schemas. Ability to apply certain optimizer zparms (query //ism, star join, etc) to individual statements.

Easier to specify plan hints for dynamic SQL, no longer need to modify source SQL statements.

CURRENT EXPLAIN MODE special register to allow explain info to be collected at PREPARE time. Dyn SQL programs such as JDBC, CLI, didn't previously have a good way to do this without modifying the source SQL statement.

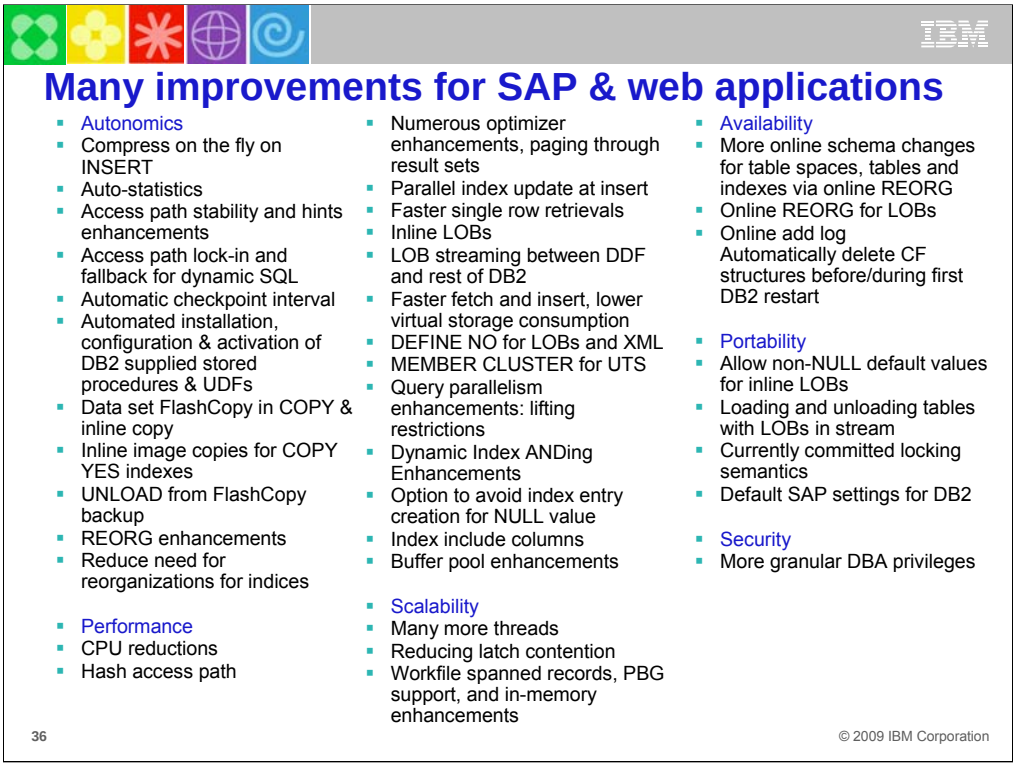

The enhancements cover many aspects of the database technology including new applications support, SQL enhancements, performance and scalability, continuous availability, data warehousing improvements as well as reducing the total cost of ownership.

DB2 10 for z/OS satisfies or partially satisfies many requirements from the worldwide user group communities such as Guide Share Europe, Japan GUIDE/SHARE, and SHARE Incorporated. In addition, this release satisfies many requirements submitted directly to IBM by customers or Business Partners.

As with recent previous releases, Enterprise Applications providers, such as SAP, many other web applications and their customers have been a very important source of the requests for new functions and features.

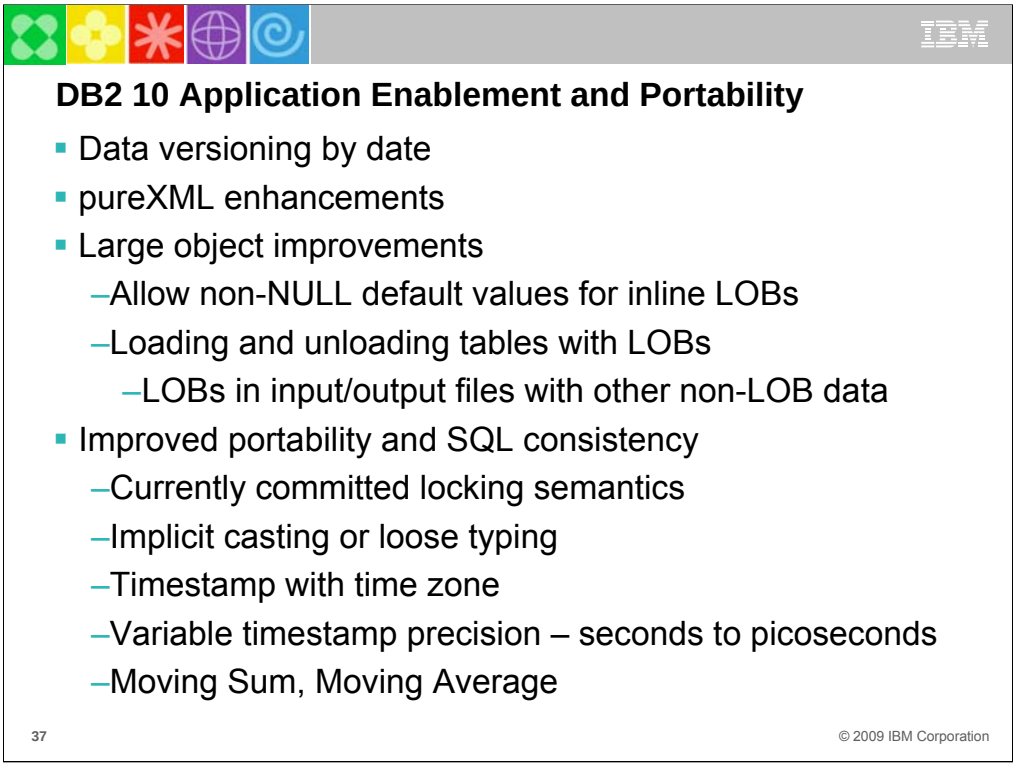

The pureXML improvements are noted two slides later.

Improvements for LOBs include the ability to placed inline in the data page when they are small, and better ability to unload and load to a sequential file, rather than needing a separate file for each LOB.

Instead of needing to explicitly cast each data type, more flexible data typing improves productivity.

Time stamps have the option to include a time zone and to have more precision (nanoseconds, rather than microseconds).

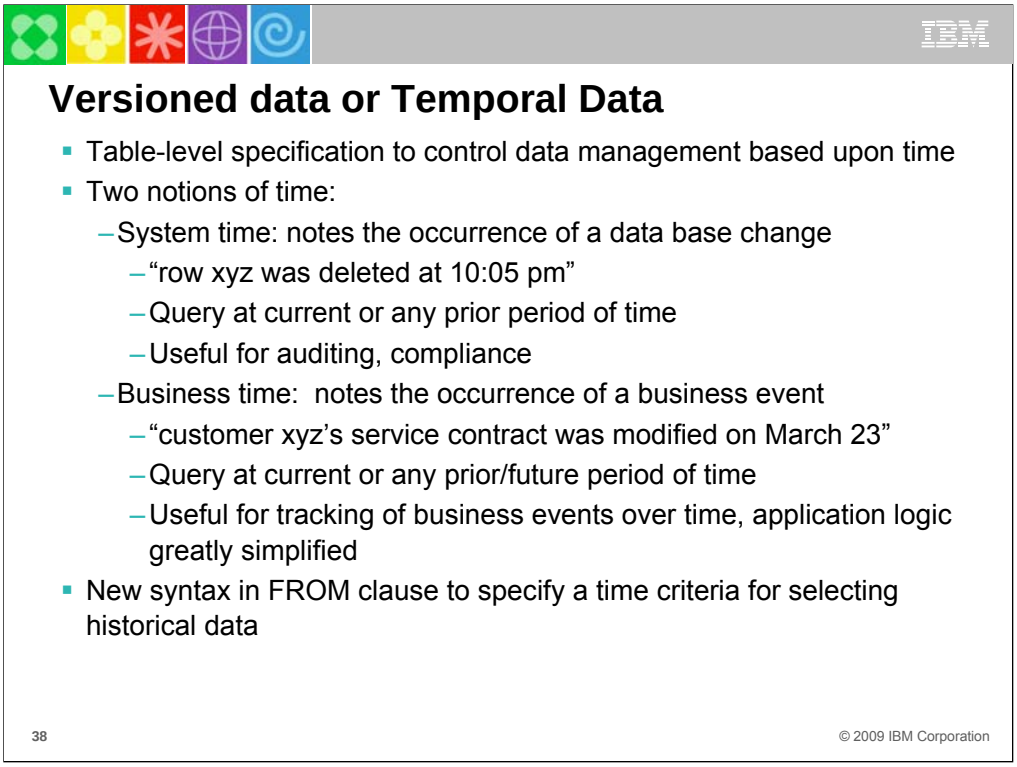

In DB2 10, you can create a temporal table, which is a base table with one or more time periods defined on it. DB2 supports two built-in types of periods, which are the system time period and the business time period. The system time period is a system-maintained period in which DB2 maintains the start and end timestamp values for a row. The business time period is a user-specified period in which you maintain the start and end values for a row.

The SYSTEM_TIME period is meaningful because of versioning. Versioning specifies that old rows are archived into another table. The table that contains the current active rows of a table is called the system-maintained temporal table. The table that contains the archived rows is called the history table. DB2 creates a history table and a table space to hold that table when you define a base table to use versioning, or when you enable versioning on an existing table. You can delete the rows from the history table when those rows are no longer needed.

Using these two built-in periods together in the same table creates a bi-temporal table. You can use a bi-temporal table to keep user-specified period information and system-based historical information. Therefore, you have a lot of flexibility in how you query data based on periods of time.

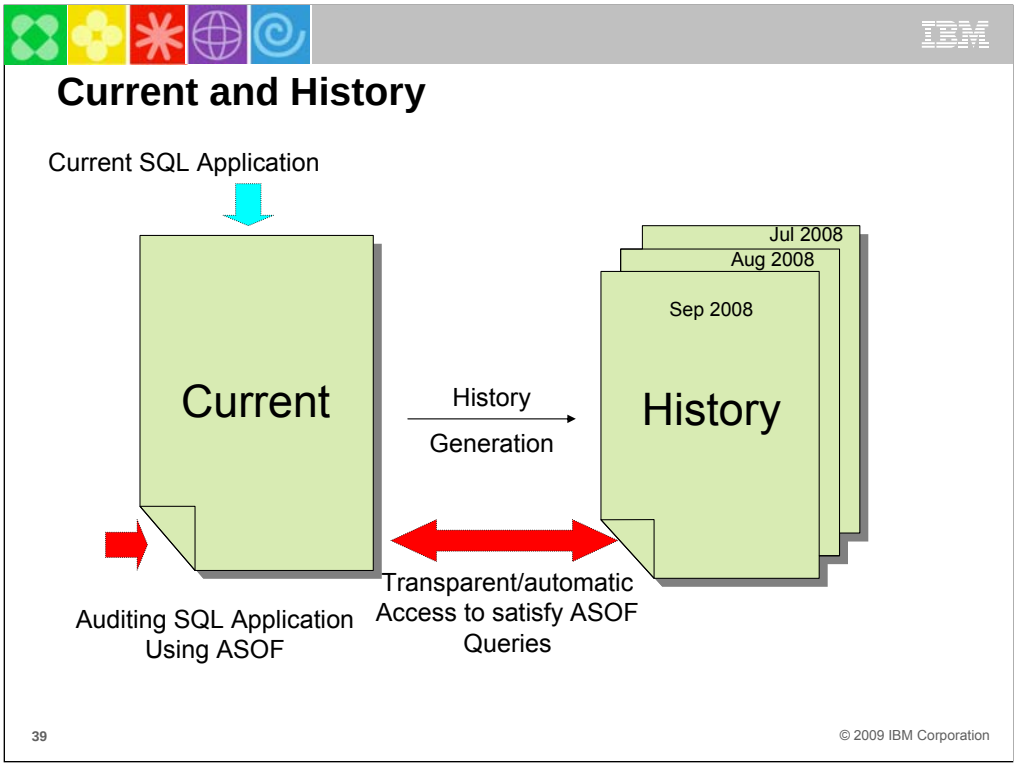

In DB2 10, you can create a temporal table, which is a base table with one or more time periods defined on it. DB2 supports two built-in types of periods, which are the system time period and the business time period. The system time period is a system-maintained period in which DB2 maintains the start and end timestamp values for a row. The business time period is a user-specified period in which you maintain the start and end values for a row.

The SYSTEM_TIME period is meaningful because of versioning. Versioning specifies that old rows are archived into another table. The table that contains the current active rows of a table is called the system-maintained temporal table. The table that contains the archived rows is called the history table. DB2 creates a history table and a table space to hold that table when you define a base table to use versioning, or when you enable versioning on an existing table. You can delete the rows from the history table when those rows are no longer needed.

Using these two built-in periods together in the same table creates a bi-temporal table. You can use a bi-temporal table to keep user-specified period information and system-based historical information. Therefore, you have a lot of flexibility in how you query data based on periods of time.

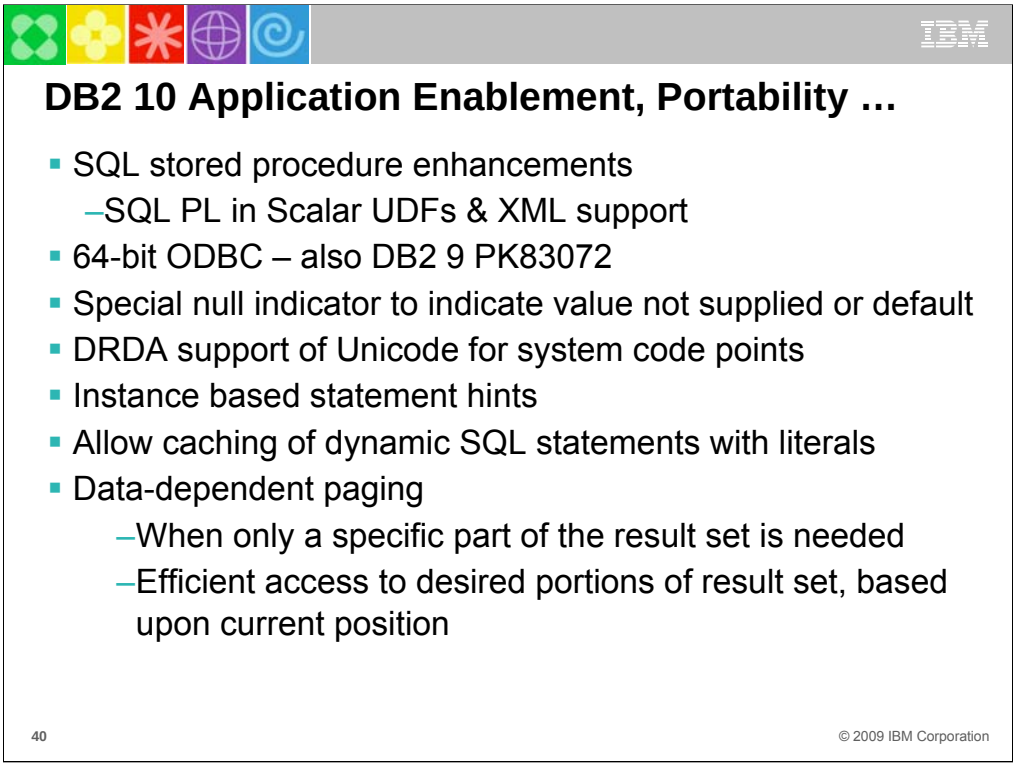

SQL procedure language is now allowed in scalar user-defined functions. Applications can use data above the bar in the new ODBC structures.

Dynamic statement caching has new techniques to use access path hints and to use the dynamic statement cache when literals are used.

Some applications require to access part of a result set based on a certain position. SQL pagination delivers an efficient way for applications to access part of a result set based on a logical key value.

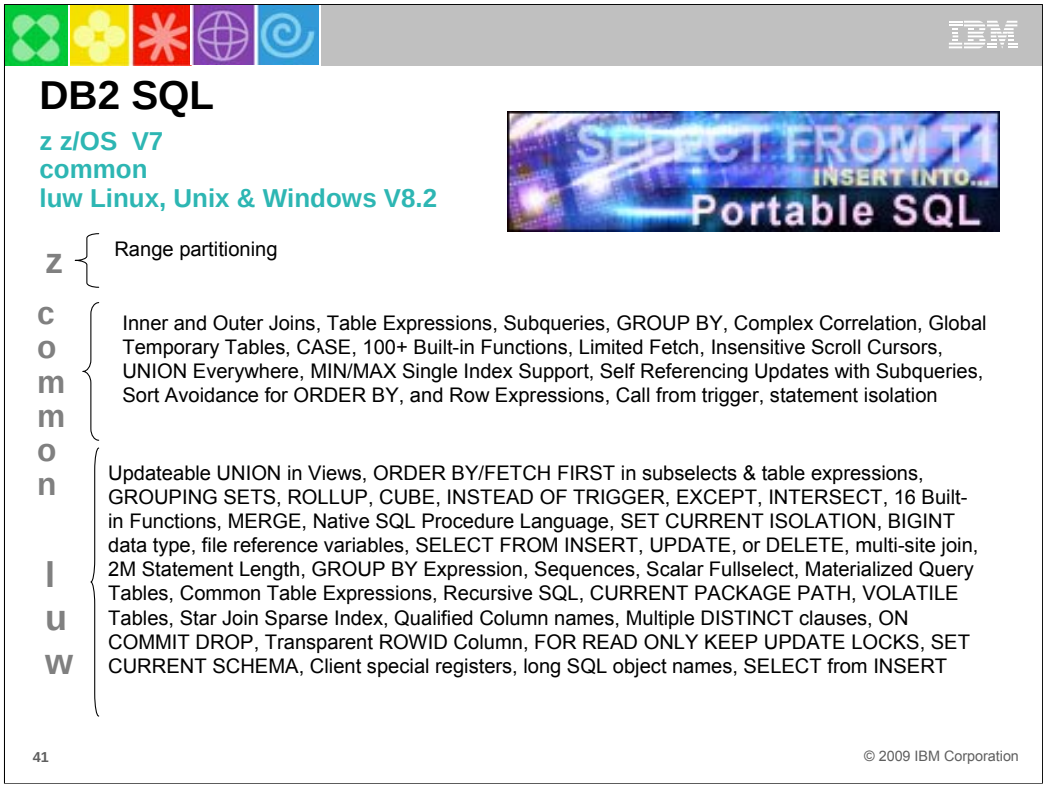

This text just shows the relationship of DB2 for Linux, Unix & Windows with DB2 for z/OS and OS/390 Version 7, comparing a March 2001 z/OS version with an October 2004 LUW version. V7 has almost no unique function, there is a small set of common function, and a larger set of SQL unique to LUW.

The next step in the process is DB2 for z/OS Version 8. There are three sets of SQL noted above, with none that is unique to DB2 for z/OS in the first group, SQL that is common across DB2 for Linux, Unix, Windows and z/OS in the large group in the middle, then SQL that is unique to DB2 for Linux, Unix and Windows in the bottom group.

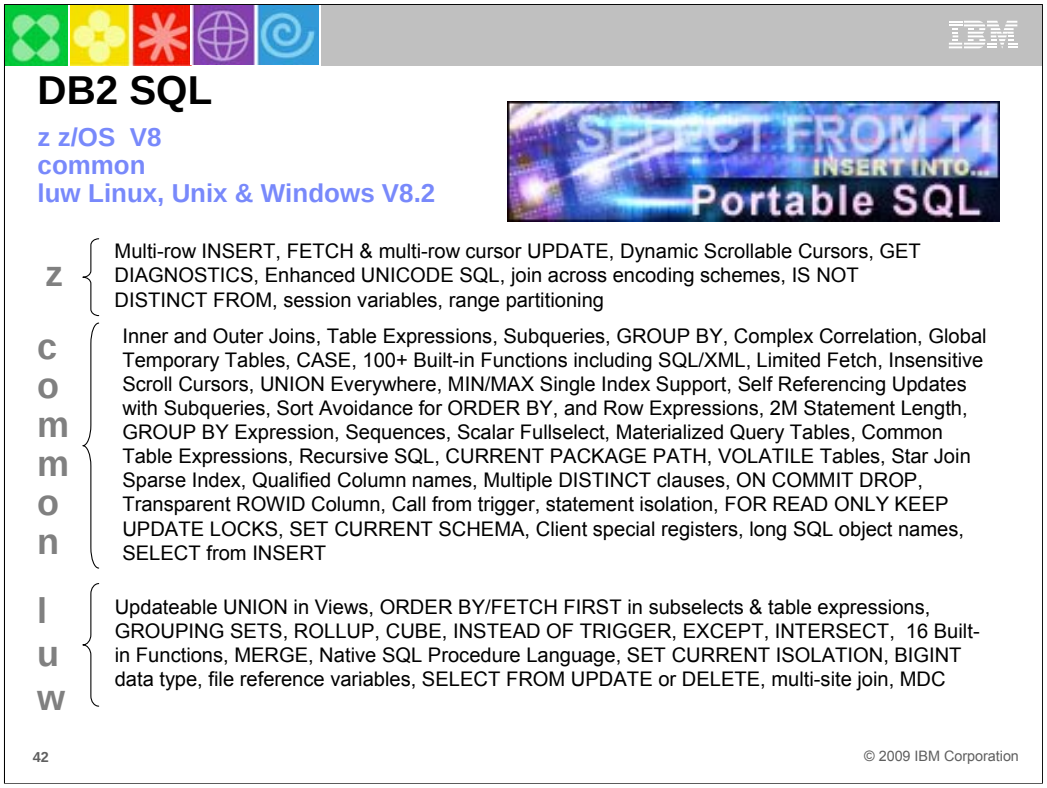

This chart shows the relationship of SQL in the DB2 family comparing DB2 for Linux, Unix & Windows with DB2 for z/OS for key language constructs. This chart compares the z/OS Version 8 from March 2004 with the LUW version from October 2004.

There are three sets of SQL noted above, with some that is unique to DB2 for z/OS in the first group, SQL that is common across DB2 for Linux, Unix, Windows and z/OS in the large group in the middle, then SQL that is unique to DB2 for Linux, Unix and Windows in the bottom group. Sheryl Larsen provided the base for this information, but the mistakes are mine.

If you want to improve DB2 family consistency, then DB2 for z/OS Version 8 is a big step, changing the game from one of catch up to one of leapfrog.

If you want to have a book for SQL across platforms, see the 2004 Cross-Platform SQL Reference.

- Cross-Platform Development, http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html

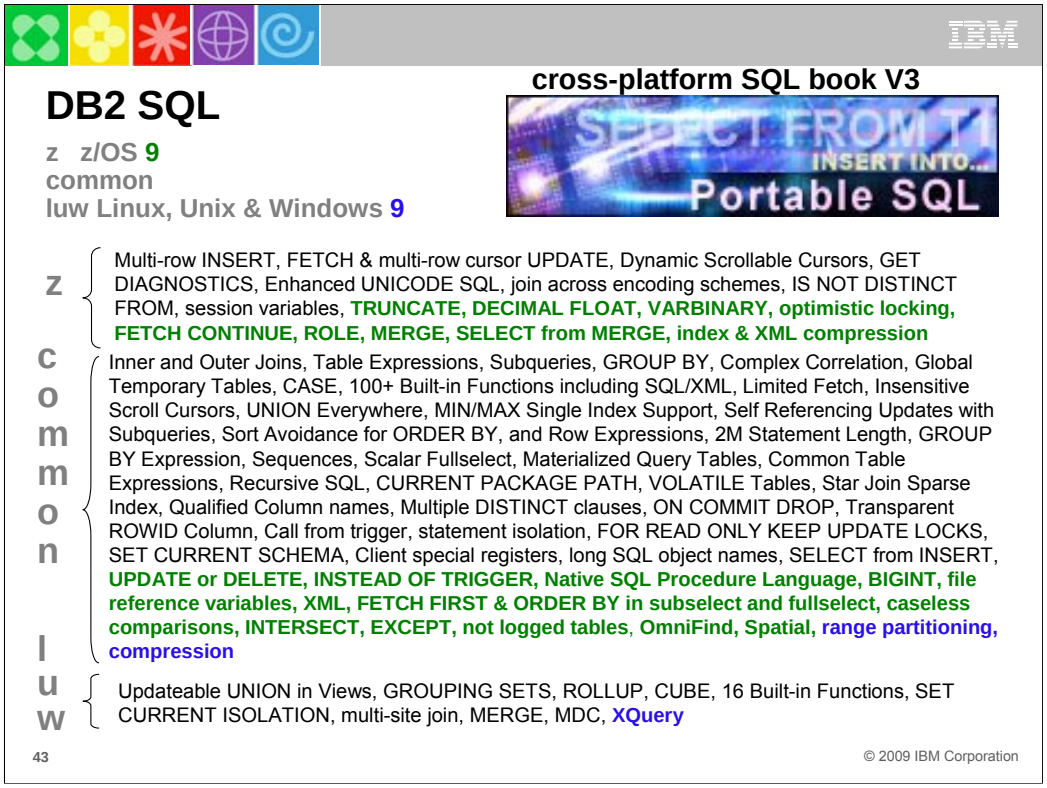

This chart shows the 2007 relationship of DB2 for Linux, Unix & Windows with DB2 for z/OS. This step in the process is DB2 9 for z/OS and DB2 9 for LUW. DB2 9 moves about half of the LUW unique items into the common set and adds a little more that is unique to the z platform. We are able to move more from the z list to the common list with Viper. There are three sets of SQL noted above, with some that is unique to DB2 for z/OS in the first group, SQL that is common across DB2 for Linux, Unix, Windows and z/OS in the large group in the middle, then SQL that is unique to DB2 for Linux, Unix and Windows in the bottom group.

The Cross-Platform SQL Reference Version 3 documents this combination, with DB2 for i5/OS V5R4.

Cross-Platform Development Version 3,
http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html
ftp://ftp.software.ibm.com/ps/products/db2/info/xplatsql/pdf/en_US/cpsqlrv3.pdf

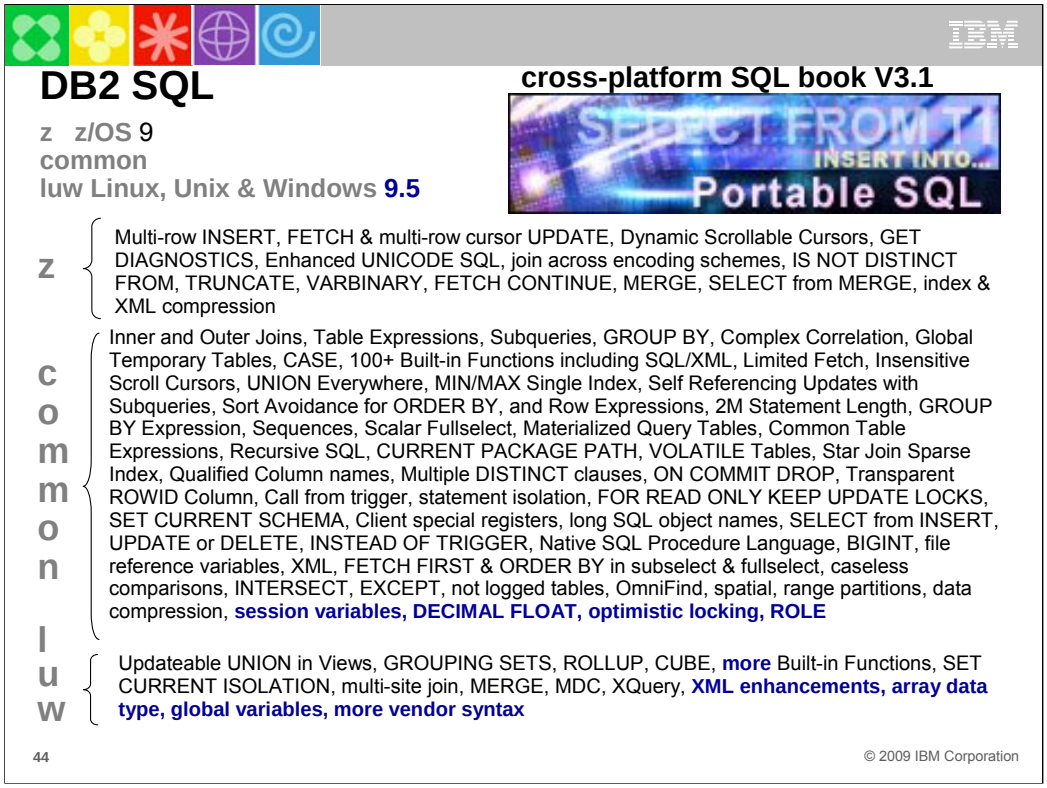

This chart shows the 2008 relationship of DB2 for Linux, Unix & Windows with DB2 for z/OS. This step in the process is DB2 9 for z/OS, (DB2 9). DB2 9 moved about half of the LUW unique items into the common set and adds a little more that is unique to the z platform. DB2 9.5 for LUW, delivered in 2008. We are able to move more from the unique z list to the common list with DB2 9.5 for LUW.

There are three sets of SQL noted above, with some that is unique to DB2 for z/OS in the first group, SQL that is common across DB2 for Linux, Unix, Windows and z/OS in the large group in the middle, then SQL that is unique to DB2 for Linux, Unix and Windows in the bottom group. The changes in a specific version are not consistent. As we introduce new function, sometimes it will be on one platform first, but movement from unique lists into the common list continues to be the strongest trend.

The Cross-Platform SQL Reference Version 3.1 documents this combination, with DB2 for i V6R1.

Cross-Platform Development Version 3.1,
http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html

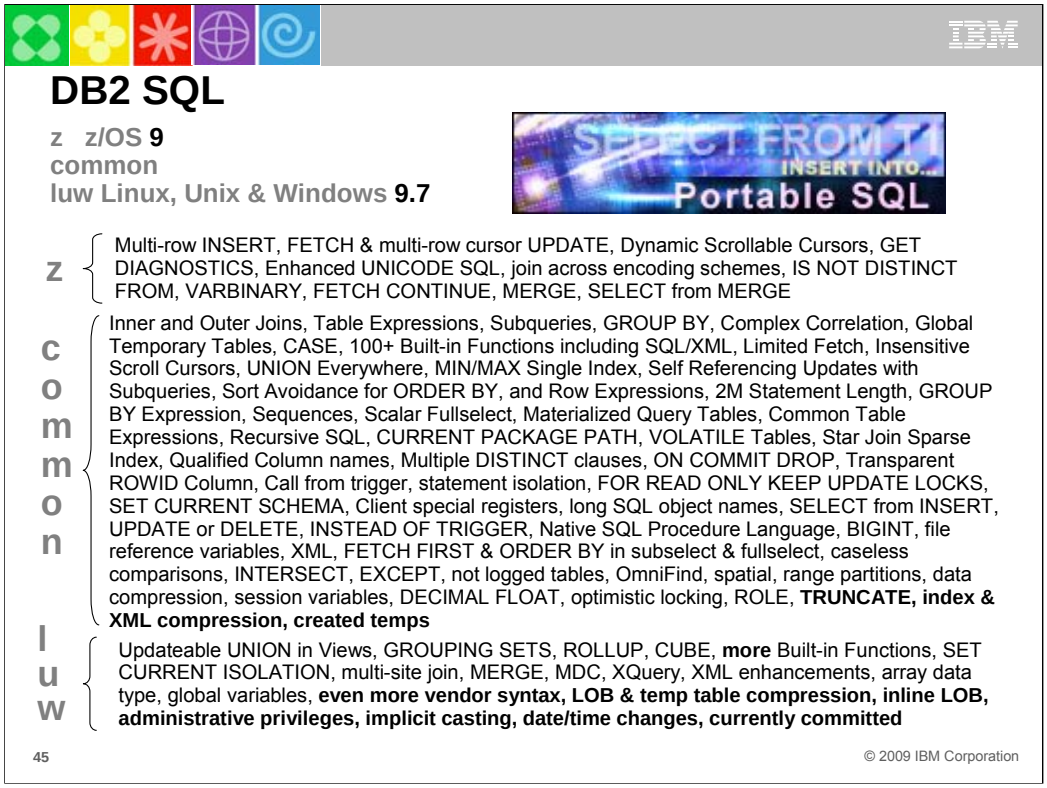

This chart shows the 2009 relationship of DB2 for Linux, Unix & Windows with DB2 for z/OS. This step in the process is DB2 9.7 for LUW. DB2 9 for z/OS moved about half of the LUW unique items into the common set and adds a little more that is unique to the z platform. DB2 9.5 for LUW, delivered in 2008 and 9.7 in 2009. We are able to move more from the unique z list to the common list with DB2 9.5 and 9.7 for LUW, while bringing in some new unique function.

There are three sets of SQL noted above, with some that is unique to DB2 for z/OS in the first group, SQL that is common across DB2 for Linux, Unix, Windows and z/OS in the large group in the middle, then SQL that is unique to DB2 for Linux, Unix and Windows in the bottom group. The changes in a specific version are not consistent. As we introduce new function, sometimes it will be on one platform first, but movement from unique lists into the common list continues to be the strongest trend.

The Cross-Platform SQL Reference Version 3.1 documents the prior combination, with DB2 for i V6R1.

Cross-Platform Development Version 3.1,
http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html

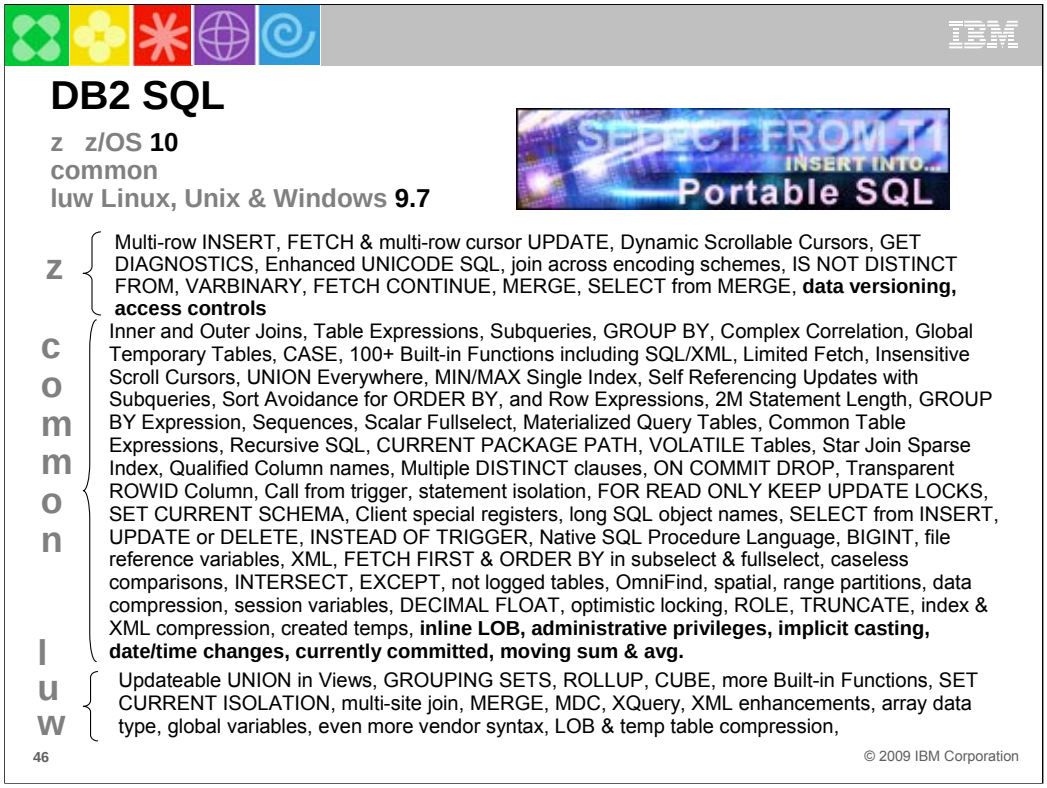

This chart shows the relationship of DB2 for Linux, Unix & Windows with DB2 for z/OS. This step in the process is DB2 10 for z/OS. DB2 10 for z/OS moves more of the LUW unique items into the common set and adds a little more that is unique to the z platform. DB2 9.5 for LUW, delivered in 2008 and 9.7 in 2009. We are able to move more from the unique z list to the common list with DB2 9.5 and 9.7 for LUW, while bringing in some new unique function.

There are three sets of SQL noted above, with some that is unique to DB2 for z/OS in the first group, SQL that is common across DB2 for Linux, Unix, Windows and z/OS in the large group in the middle, then SQL that is unique to DB2 for Linux, Unix and Windows in the bottom group. The changes in a specific version are not consistent. As we introduce new function, sometimes it will be on one platform first, but movement from unique lists into the common list continues to be the strongest trend.

The Cross-Platform SQL Reference Version 3.1 documents the prior combination, with DB2 for i V6R1.

Cross-Platform Development Version 3.1,
http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html

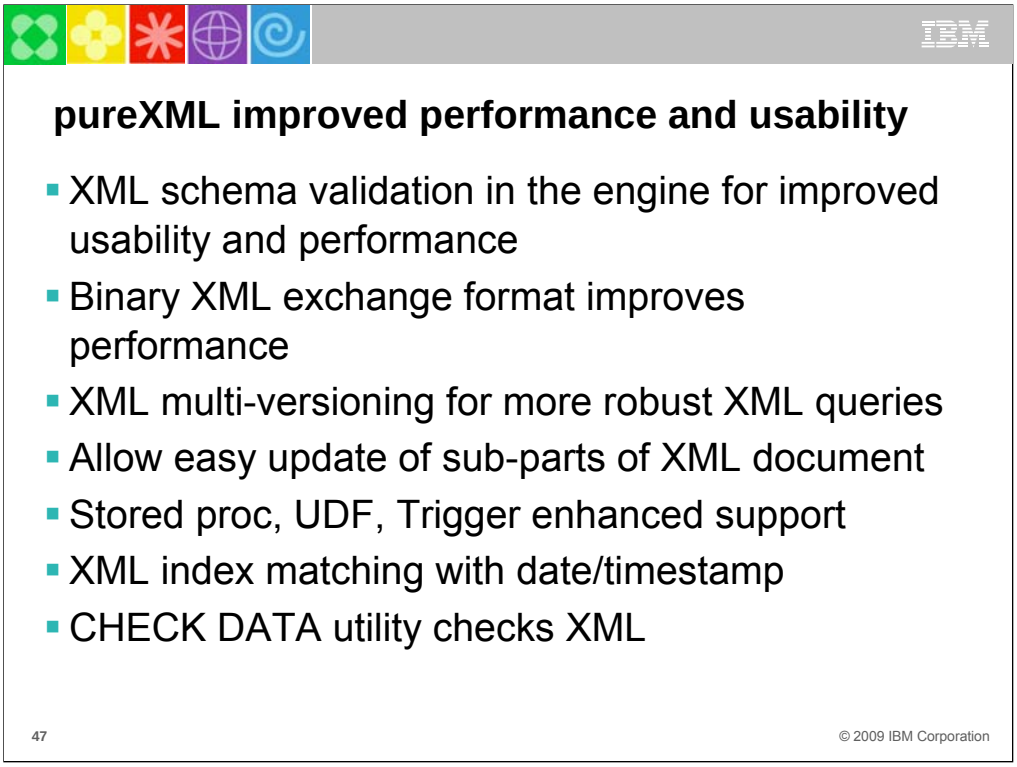

A range of XML improvements delivers a strong release 2 of the pureXML function. Customers use of DB2 9 pureXML shaped this delivery of improved performance and usability.

Multi-versioning: During the execution of a SQL statement, a row with an XML column can be kept in a work file. The row in the work file does not contain the actual XML document. Instead, the information needed for DB2 to retrieve the XML document from the XML table is cached in the work file. The problem occurs if the XML document in the XML table is deleted or updated. When the row in the work file is fetched, DB2 cannot find the expected XML document in the XML table, and the SQL statement fails with an error SQLCODE.

XML UPDATE: Applications which require parts of XML documents to be modified need to break apart the XML document into modifiable pieces, make the modification to a piece, and then construct the pieces back into an XML document.

SP/UDF/Trigger support: XML variables inside SQL PL, XML arguments, transition variables.

The CHECK DATA utility is extended to check XML data.

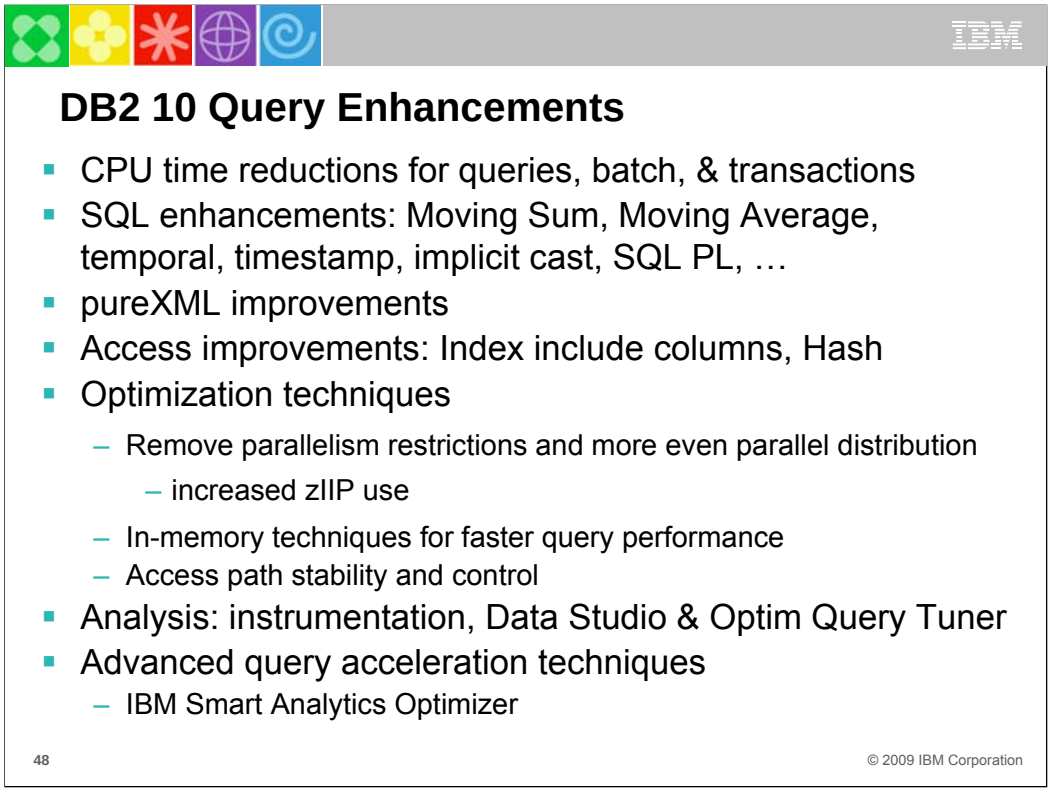

Query enhancements in DB2 build on the improvements in DB2 V8 and 9. The CPU reductions can make a differences to queries. Improved SQL with better ability to query, temporal understanding, and XML improvements make the queries simpler.
The key new access techniques are index include columns and hash access. Improvements in access techniques provide more parallel access by reducing restrictions. More parallel means more ability to redirect the work to zIIP and reduce costs. In memory techniques provide improved performance. Access path stability allows better control for queries. Improvements in the instrumentation help all performance monitors. The Data Studio and Optim Query Tuner have replaced some older function. Advanced query acceleration is being previewed in the IBM Smart Analytics Optimizer.

Technology Preview: IBM Smart Analytics Optimizer

What is it?

- A special purpose, network-attached appliance that is an add-on to a DB2 for z/OS system
- Offloads typical DW/BI queries resulting in predictable and orders-of-magnitude faster query response times while reducing overall TCO

Business Value

- Dramatically lowers the cost for query and reporting on System z
- Advanced in-memory scale-out cluster technologies that keep the complete system centrally managed without having to change any requirements for BI applications
- Complements the many new Data Warehousing features in DB2 9 for z/OS
- Leverages the many new warehousing and business intelligence solutions now available on System z

Targeted Uses for DB2 for z/OS customers:

- Requirements to accelerate a subset of their warehouse or reporting queries
- Looking for more insight and business intelligence from operational data
- Needs to consolidate datamarts or data stores into one enterprise warehouse

© 2009 IBM Corporation

A technology preview for the IBM Smart Analytics Optimizer was delivered at IOD in 2009. This technique provides a new approach to a hybrid structure delivering lower cost query on System z.

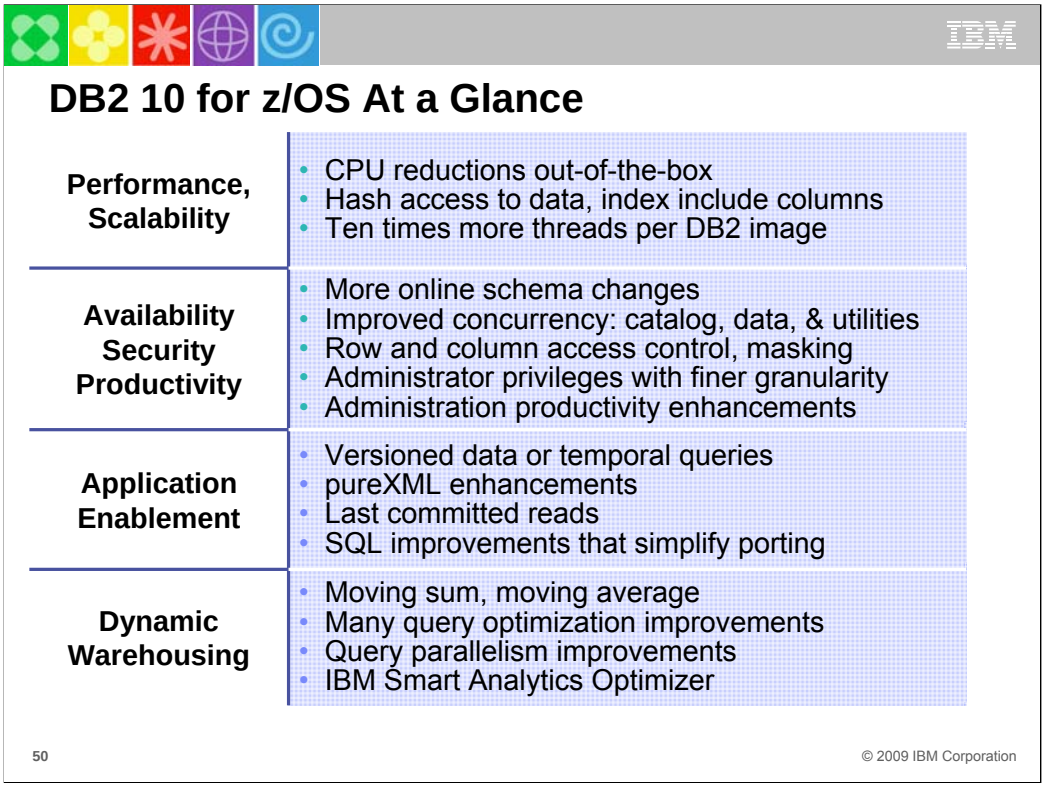

DB2 10 for z/OS provides the best reduction in CPU for transactions and batch for 22 years, since V2R1. We expect most customers to reduce CPU times between 5% and 10%. Applications which can take advantage of additional benefits, such as hash access, can have larger CPU and memory reductions. Scalability is the second major benefit, with the ability to run five to ten times as many threads in a single subsystem by moving 80% to 90% of the virtual storage above the bar. Schema evolution or data definition on demand enhancements improves availability. SQL, pureXML, web services extend usability and application portability for this platform. Productivity improvements for application developers and for database administrators are very important as data grows in scale and complexity. DBAs can avoid running statistics, some REORGs, and benefit from memory, improved query stability, and utilities enhancements. Warehousing continues to evolve, with improvements in SQL and XML, better optimization techniques, increased parallelism and the new IBM Smart Analytics Optimizer.

DB2 10 is an exciting new version for many customers, even though it is still in beta. Many customers will want to know what to do to position themselves for it. Step 1 is wait for general availability and additional information about the new version.
Here are some of the early planning options and work that customers can do to get ready for DB2 10.

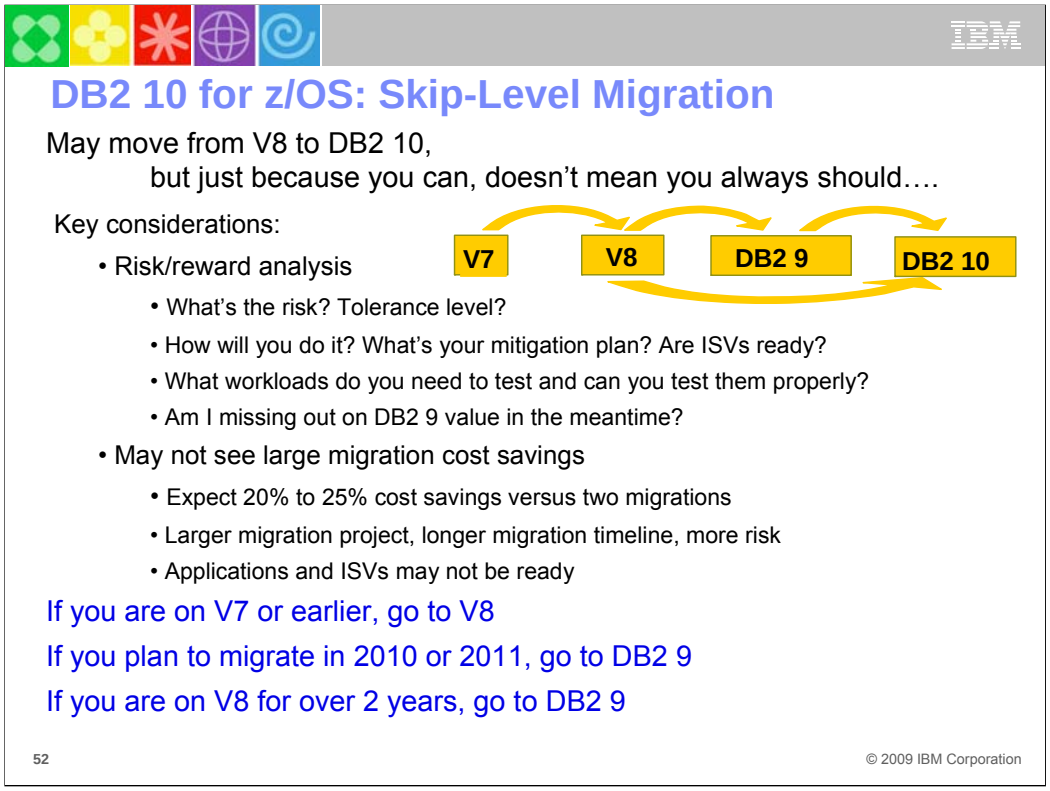

DB2 10 supports migration from DB2 9 NFM or from V8 NFM. Customers not yet running V8 or DB2 9 should plan to migrate to DB2 for z/OS V8 NFM as preparation for a migration to DB2 10. We estimate that about one customer in five migrated using a skip version technique for V5 to V7, and we'll see a similar fraction this time.

Normal migration is moving one version at a time every three years. For customers who have gotten behind, the ability to skip a migration cycle will be attractive, but this ability is not "something for nothing". Customers need to consider the tradeoffs and challenges that we know about in skip version migration. Most customers who migrate to new versions by three years after GA are already on DB2 9.

The project for skipping is larger. While the testing and rollout are only a little larger than a single version migration, the education and remediation work is roughly double the normal size. Most project plans estimate 150%. Consider the timing carefully. Improvements in DB2 9 are delayed for 2 to 4 or more years with a skip plan. You may need to have extended service on V8.

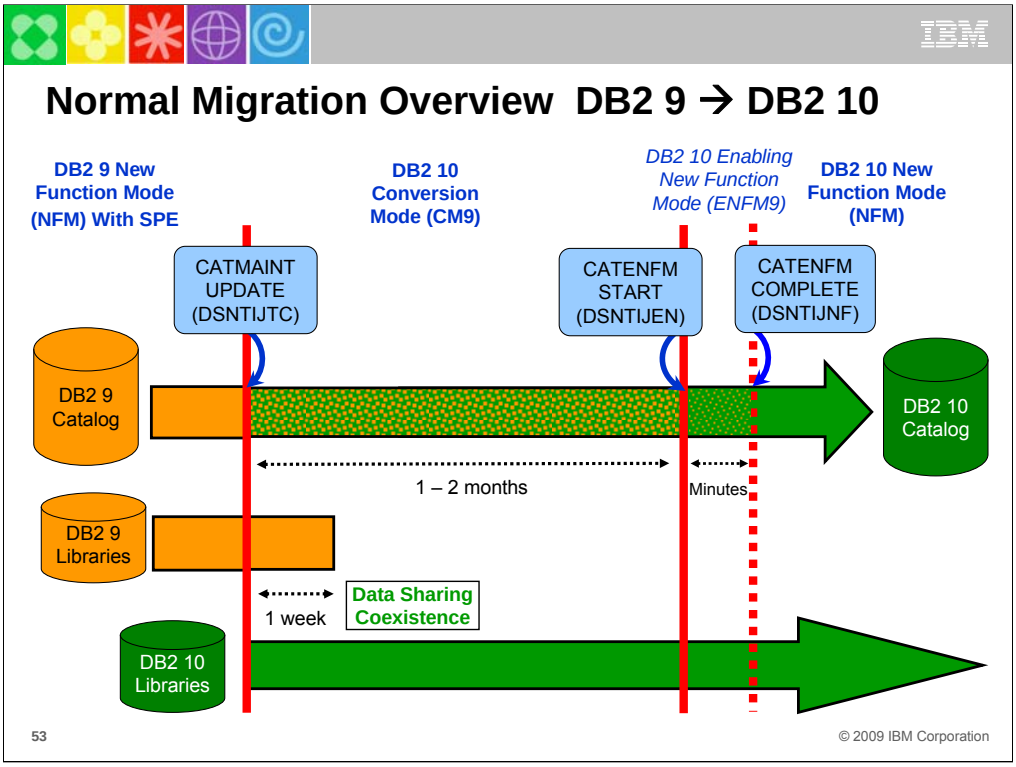

Migration from DB2 9 to DB2 10

The catalog changes will happen in two places. One is the migration from DB2 9 to DB2 10 conversion mode (CM9) using the DSNTIJTC job. The other is the DB2 10 enabling-new-function mode process (ENFM9) using the DSNTIJEN job.

When a system is migrating from DB2 9 NFM to DB2 10 conversion mode the possible DB2 10 modes are:

CM9 Conversion Mode entered when migrating from DB2 9 NFM to DB2 10

ENFM9 Enabling New Function Mode on a system that migrated from DB2 9 NFM to DB2 10. Once this mode has been entered the system can not fallback to DB2 9 and a DB2 9 member can not be started in a data sharing group.

NFM This is the NFM mode when all system changes are made on a system that migrated from DB2 9 NFM to DB2 10 and the system is ready for DB2 10 new function

CM9* The system migrated from DB2 9 NFM to DB2 10 and at one point was in either ENFM9 or NFM on DB2 10.

ENFM9* The system was migrated from DB2 9 NFM to DB2 10 and at one point was in NFM on DB2 10.

Some migration considerations are:

• A DB2 9 system that has started the migration to DB2 10 can only fallback to DB2 9.

• A data sharing group that migrated from DB2 9 NFM to DB2 10 can not have any V8 members.

Note – this is not necessarily to scale!

Note – ONE WAY – fallback to CM9* (covered later) is possible but not to DB2 9

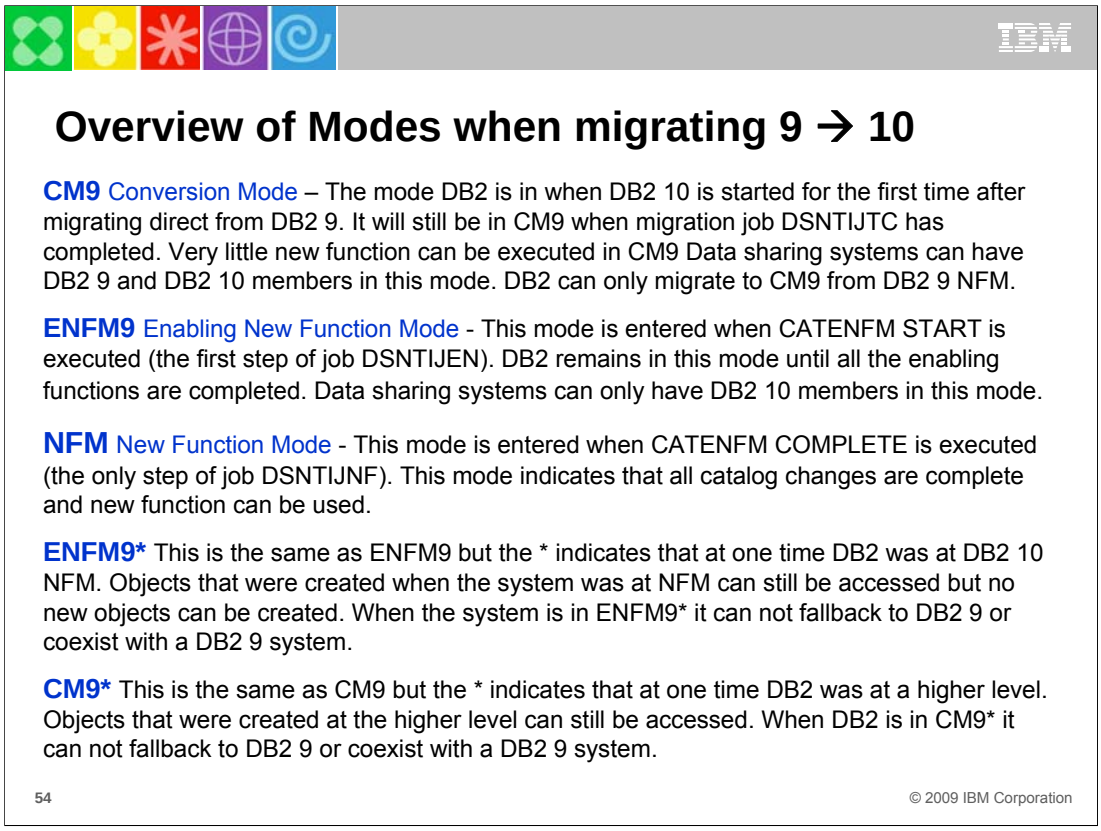

Enabling New Function Mode 9 (ENFM9) is entered when CATENFM START is executed, following a previous migration (CATMAINT) direct from DB2 9 to DB2 10. This is the first step of job DSNTIJEN. DB2 remains in this mode until all the enabling functions are completed.

Data sharing systems can only have DB2 10 members in this mode. The enabling functions are:

•Place the DB2 subsystem in enabling-new-function mode 9 (ENFM9).

•The first time that you run job DSNTIJEN, DB2 saves the RBA or LRSN of the system log in the BSDS.

ENFM9* is the same as ENFM but the * indicates that the at one time the system was at NFM. Objects that were created when the system was at NFM can still be accessed but no new objects can be created. When the system is in ENFM9* it can not fallback to DB2 9 or coexist with a DB2 10 system. When DB2 is in ENFM9* the mode can be changed with the following:

•Migration job DSNTIJCS takes the system to CM9*.

•Migration job DSNTIJNF takes the system to DB2 10 NFM.

CM9* is the same as CM but indicates that at one time DB2 was at a higher level. Objects created at the higher level can still be accessed.

When in CM9*, DB2 cannot fallback to DB2 9, or coexist with a DB2 9 system. CM9* is entered when DSNTIJC is run and DB2 is in ENFM9, ENFM9* or NFM. When DB2 is in CM9*, the mode can be changed as follows:

Migration job DSNTIJEN

•If DB2 went from NFM or ENFM9* to CM9*, DSNTIJEN takes it to ENFM9*.

•If DB2 went from ENFM9 to CM9*, it completes ENFM9 processing and leaves DB2 in ENFM9.

Migration job DSNTIJNF

•If DB2 went from NFM or ENFM9* to CM9*, DSNTIJNF takes it to NFM.

•If DB2 went from ENFM9 to CM9*, DSNTIJNF does one of two things:

- •If ENFM9 processing had completed it takes DB2 to NFM.
- •If not, DSNTIJNF indicates that DSNTIJEN must be run.

Migration job DSNTIJES

•If DB2 went from NFM or ENFM9* to CM9*, DSNTIJES takes it to ENFM9*.

•If DB2 went from ENFM9 to CM9*, DSNTIJES indicates that DSNTIJEN must be run.

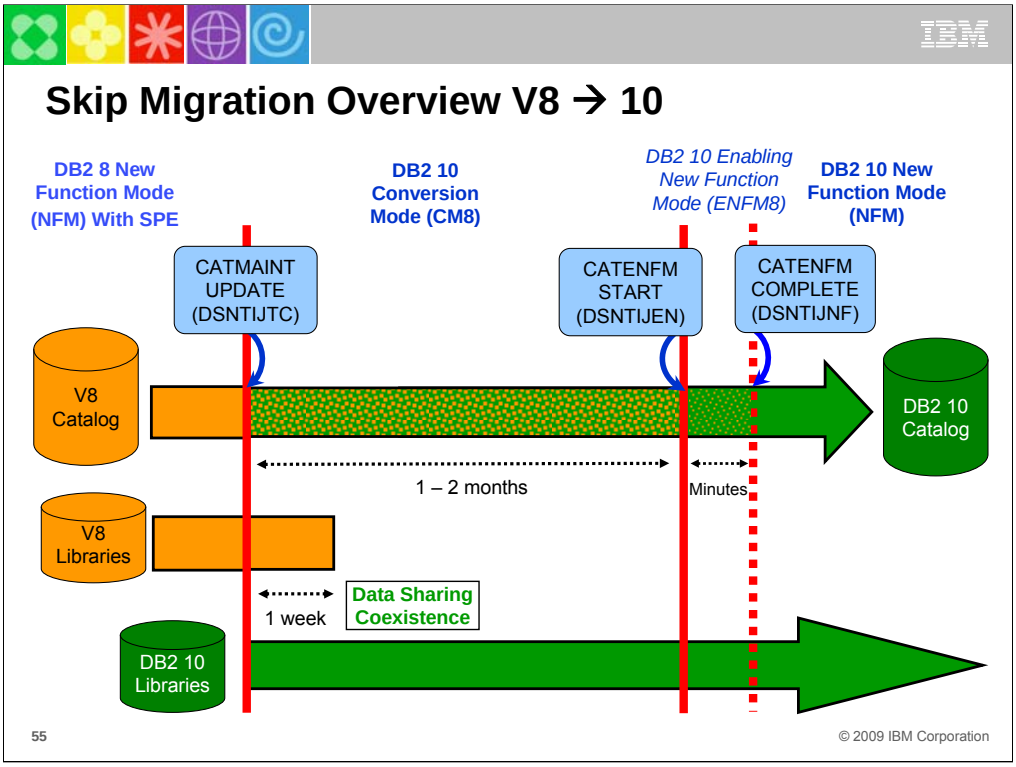

Migration from V8 direct to 10

The catalog changes will happen in two places. One is the migration from DB2 V8 to DB2 10 conversion mode (CM8) using the DSNTIJTC job. The other is the DB2 10 enabling-new-function mode process (ENFM8) using the DSNTIJEN job.

DB2 will support migrating from V8 NFM to DB2 10 without ever starting the system in DB2 9. When a system is migrating from V8 NFM to DB2 10 conversion mode the possible DB2 10 modes are:

CM8 Conversion Mode entered when migrating from V8 NFM to DB2 10

ENFM8 Enabling New Function Mode on a system that migrated from V8 NFM to DB2 10. Once this mode has been entered the system can not fallback to V8 and a V8 member can not be started in a data sharing group.

NFM This is the New Function Mode when all system changes have been made on a system that migrated from V8 NFM to DB2 10 and the system is ready for DB2 10 new function

CM8* The system migrated from V8 NFM to DB2 10 and at one point was in either ENFM8 or NFM on DB2 10.

ENFM8* The system was migrated from V8 NFM to DB2 10 and at one point was in NFM on DB2 10.

Some migration considerations are:

- A V8 system that has started the migration to DB2 10 can only fallback to V8.
- A V8 system that has started the migration to DB2 10 and then performed a fallback to V8 can not migrate to DB2 9.
- A system that migrates from V8 NFM to DB2 10 can not use DB2 9 new function until DB2 10 NFM is reached.
- A data sharing group that migrated from V8 NFM to DB2 10 skipping DB2 9 can not have any DB2 9 members.

Some differences in the C8 and C9 catalog are:

- The RTS is still in a user database.
- The SEQNO column of SYSPACKSTMT is a SMALLINT.
- The SYSOBJ table space uses 4k pages in CM8 and 8K pages in CM9.

Note – this is not necessarily to scale!

Note – ONE WAY – fallback to CM8* (covered later) is possible, but not to DB2 V8

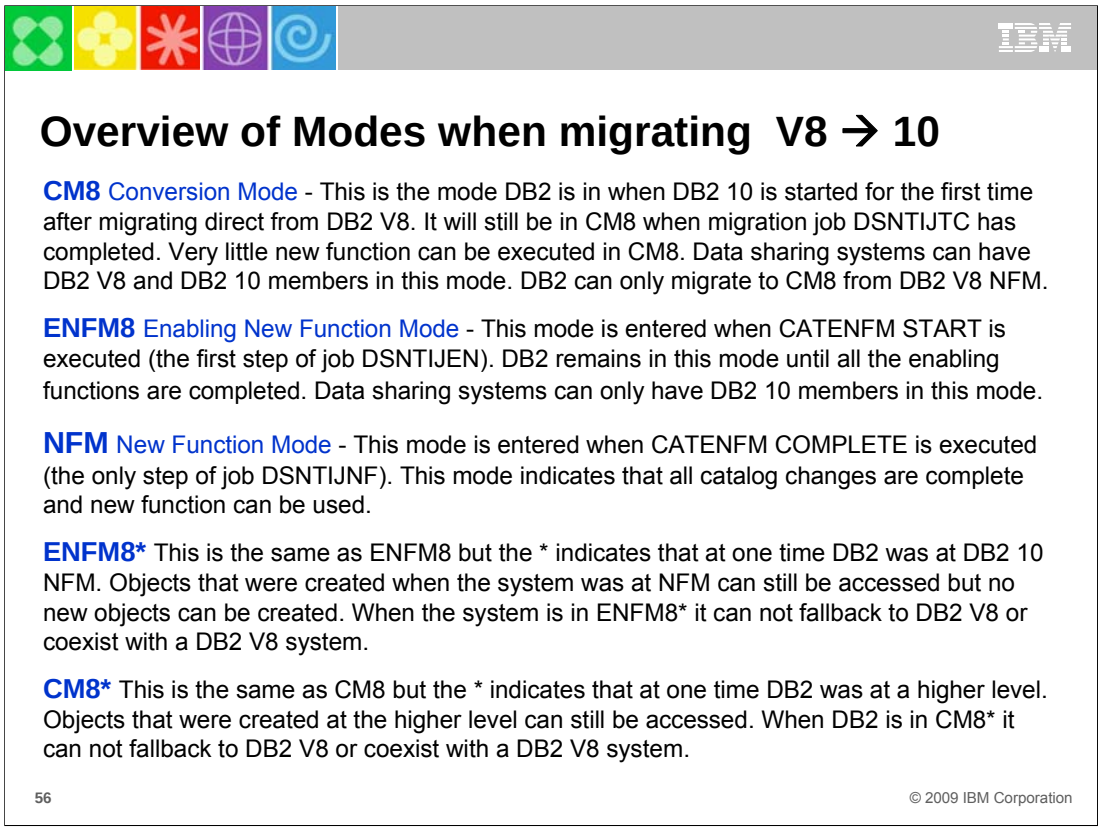

Enabling New Function Mode 8 (ENFM8) is entered when CATENFM START is executed, following a previous migration (CATMAINT) direct from DB2 V8 to DB2 10. This is the first step of job DSNTIJEN. DB2 remains in this mode until all the enabling functions are completed.

Data sharing systems can only have DB2 10 members in this mode. The enabling functions are:

•Place the DB2 subsystem in enabling-new-function mode 8 (ENFM8).

•The first time that you run job DSNTIJEN, DB2 saves the RBA or LRSN of the system log in the BSDS.

ENFM8* is the same as ENFM but the * indicates that the at one time the system was at NFM. Objects that were created when the system was at NFM can still be accessed but no new objects can be created. When the system is in ENFM8* it can not fallback to DB2 V8 or coexist with a DB2 10 system. When DB2 is in ENFM8* the mode can be changed with the following:

•Migration job DSNTIJCS takes the system to CM8*.

•Migration job DSNTIJNF takes the system to DB2 10 NFM.

CM8* is the same as CM but indicates that at one time DB2 was at a higher level. Objects created at the higher level can still be accessed.

When in CM8*, DB2 cannot fallback to DB2 V8, or coexist with a DB2 V8 system. CM8* is entered when DSNTIJC is run and DB2 is in ENFM8, ENFM8* or NFM. When DB2 is in CM8*, the mode can be changed as follows:

Migration job DSNTIJEN

•If DB2 went from NFM or ENFM8* to CM8*, DSNTIJEN takes it to ENFM8*.

•If DB2 went from ENFM8 to CM8*, it completes ENFM8 processing and leaves DB2 in ENFM8.

Migration job DSNTIJNF

•If DB2 went from NFM or ENFM8* to CM8*, DSNTIJNF takes it to NFM.

•If DB2 went from ENFM8 to CM8*, DSNTIJNF does one of two things:

- •If ENFM8 processing had completed it takes DB2 to NFM.
- •If not, DSNTIJNF indicates that DSNTIJEN must be run.

Migration job DSNTIJES

•If DB2 went from NFM or ENFM8* to CM8*, DSNTIJES takes it to ENFM8*.

•If DB2 went from ENFM8 to CM8*, DSNTIJES indicates that DSNTIJEN must be run.

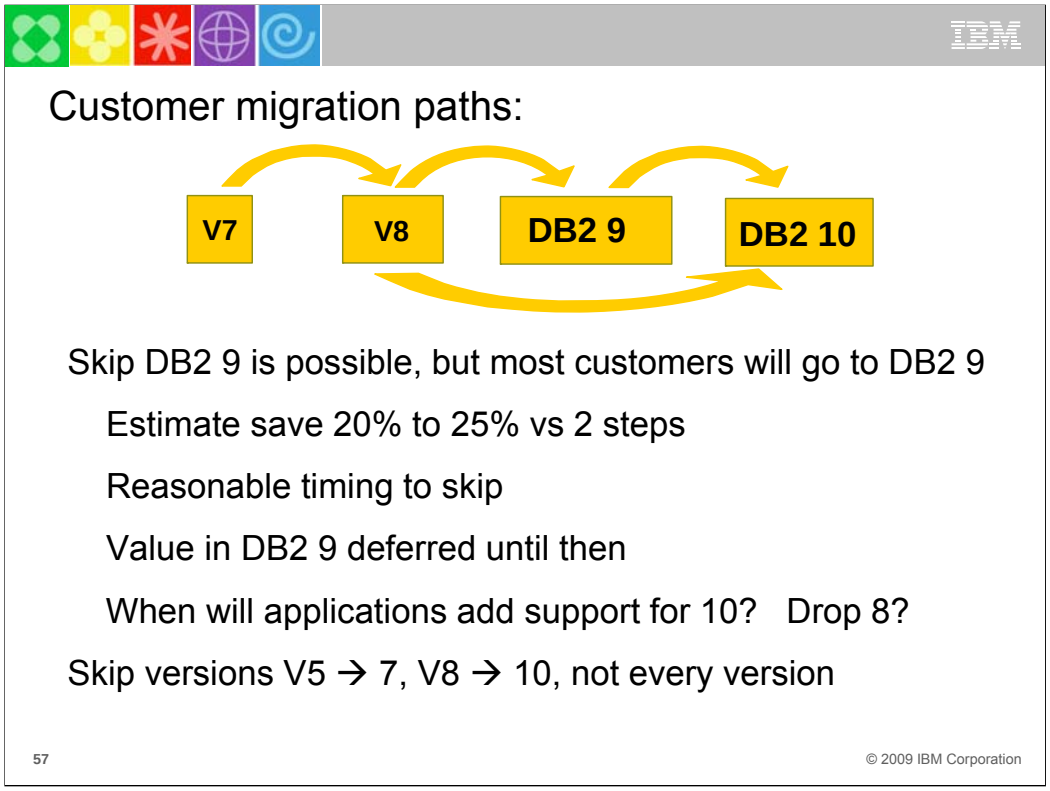

Do you want to migrate from V8 to DB2 10? The first question is the objective in terms of time you expect to save. The next questions are timing and service. Do you want support? DB2 10 will probably become generally available in 2010. When can you migrate to it? DB2 V8 end of marketing is Sept 2009. Migrating to DB2 9 is faster, easier and smoother now.

We have experience with skip release migration. We did provide skip level migration from V5 to V7, but not for any other recent releases. The ability to skip helps some customers who have not installed recent versions, but this work tends to help a minority of the DB2 customer base. The key question for skipping DB2 9 is, "How long can you wait for that function?" If you are on DB2 V8 today, how long after GA do you wait before moving to a new version? Support for the new version in applications and tools can delay the process. Waiting for others to deliver success stories adds more time. If you wait three years after GA, then you could start moving to DB2 10 in 2014, and probably get there about a year later. I expect DB2 V8 to be out of support before that, so you would probably have additional expense for extended service.

The savings for skipping a version migration are less than 50%, since the education and needed application and administration changes are about the same. Customers who do skip migration reported that the project takes longer, taking about 50% more work than a normal migration path. Changing from DB2 V8 or before to DB2 10 would require a cultural shift that some describe as culture shock. If customers spend the bulk of their migration project time in testing, then savings could be up to 40%, but most customer plans should show 20% to 25% reduction over two migrations. The cost for skipping is primarily later delivery of DB2 9 improvements by two to four years. Do you need the improvements in DB2 9 before that? CPU savings, especially in utilities. Disk savings via compression for indexes. Improved insert and update rates. Improved SQL and pureXML for developer productivity. Better availability.

If customers are on V7 or earlier, that version is out of service. Recommend migrating to V8 to be on a supported version and to allow a supported skip version migration.

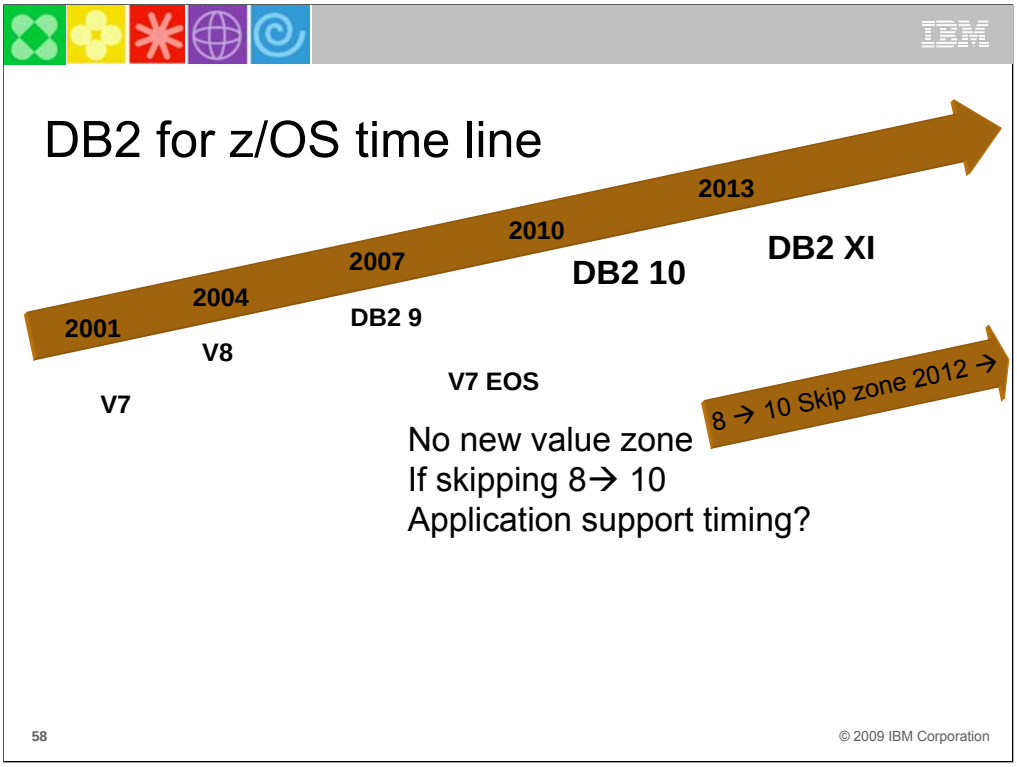

The key question for skipping DB2 9 is, “How long can you wait for that function?” If you are on DB2 V8 today, how long after GA do you wait before moving to a new version? Is it one year? Two? Or three? In those cases, you’ll start moving to DB2 10 in 2011, 2012 or 2013, and probably get there about a year later.

The savings for skipping a version migration are less than 50%, since the education and needed changes are about the same. If customers spend the bulk of their time in testing, then savings could be up to 40%, but most customer plans should show 20% to 25% reduction over two migrations. The cost for skipping is primarily later delivery of DB2 9 improvements by three or four years.

Do you need the improvements in DB2 9 before that? CPU savings, especially in utilities. Disk savings via compression for indexes. Improved insert and update rates. Improved SQL and pureXML for developer productivity. Better availability.

For customers who are still running V7, the option to skip V8 → 10 is very attractive and makes the current path clear. Customers who have just migrated to V8 may like this alternative for the short term.

DB2 9 allows more online schema changes and utilities that allow concurrent access. The BACKUP SYSTEM and RESTORE SYSTEM capabilities are improved.

Improved granularity for security and better network security come in DB2 9.

XML has become pervasive, and delivery in DB2 improves productivity while it avoids more data integration projects for the future.

DB2 9 is the best warehouse delivery in DB2 for z/OS ever, with improved SQL, optimization and surrounding products.

Index compression can reduce the disk space for indexes by half.

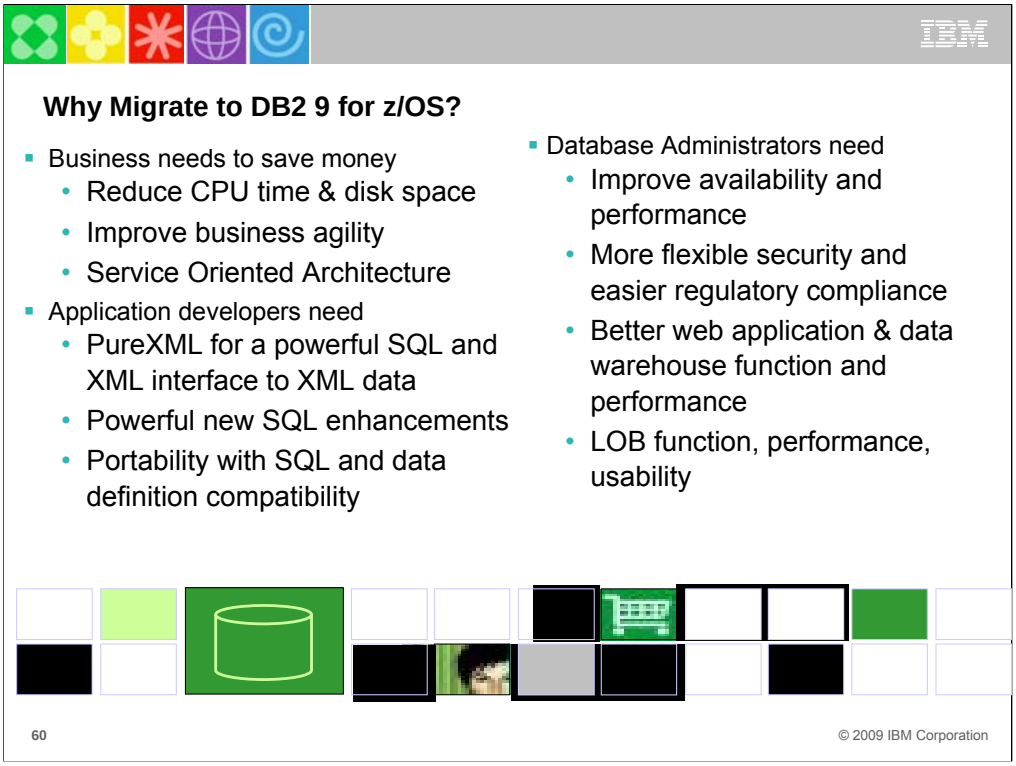

DB2 9 has a lot for everyone. Here are just a few of the strategic highlights.

The business needs include CPU cycle reductions that deliver in most utilities, disk space reduction from index compression, 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.

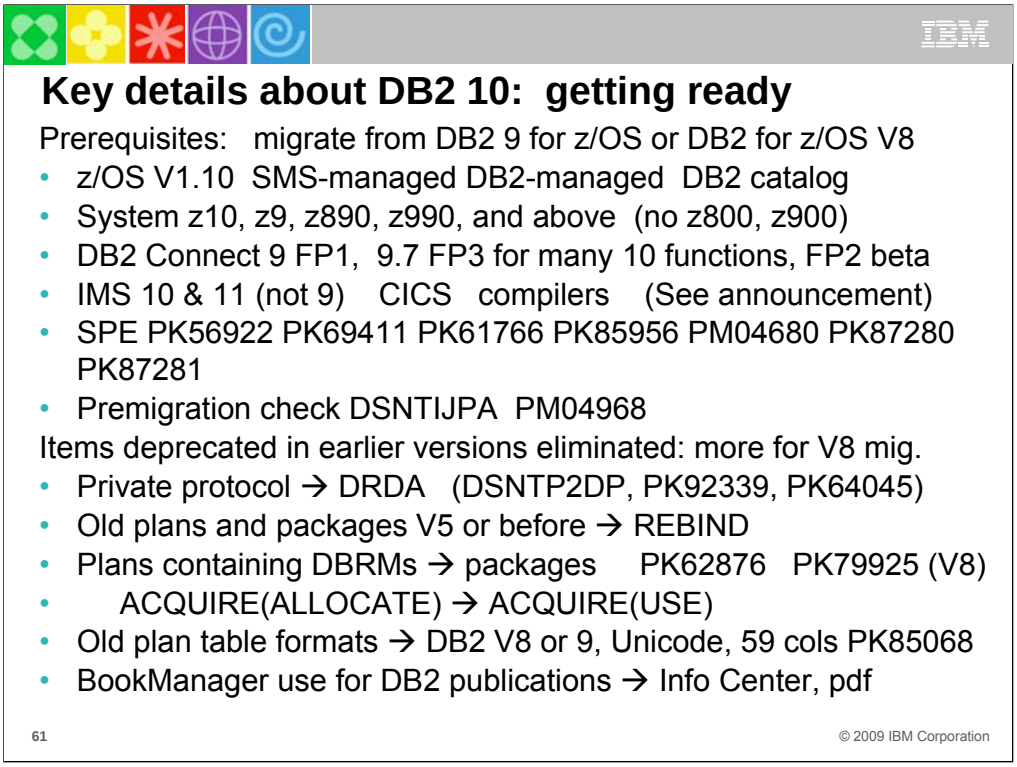

Here are a couple of thoughts about what might be required in hardware and software to run DB2 10. Much will depend upon the timing of the deliveries and market acceptance. Moving forward as quickly as possible means that some of the past must be left behind. See the list of deprecated functions from prior versions.

The above functions are still included in DB2 9, but are generally deprecated and may be removed from future versions. Note the direction indicated to the right of the arrows, as these are the functions provided to replace the existing function. If you are using any of these functions, you are advised to move to the new function.

See the Installation Guide section, “Functions that are deprecated” and the announcement material for more information on these changes.

http://www.ibm.com/common/ssi/rep_ca/8/897/ENUS206-098/ENUS206-098.PDF

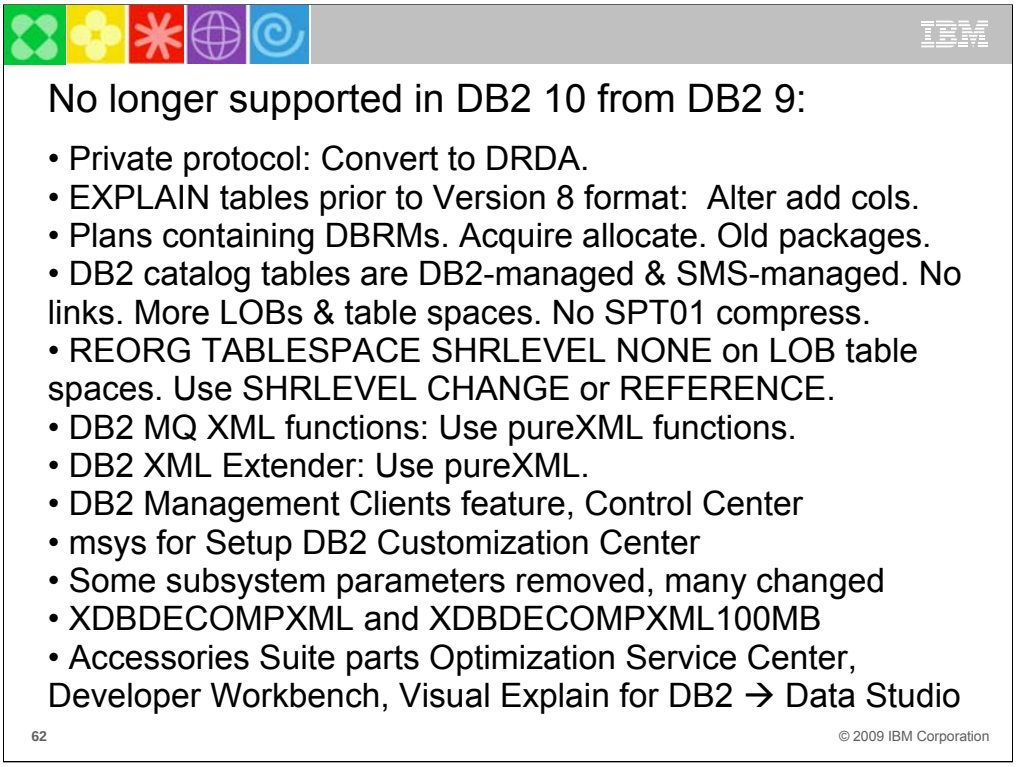

No longer supported, moving from DB2 9
Private protocol: Convert to DRDA.
EXPLAIN tables prior to Version 8 format: Alter tables to add the needed columns.
Plans containing DBRMs. Plans will contain packages, which have the DBRMs. Use the COLLID option of REBIND to make the needed changes. Acquire allocate is not a supported BIND option any longer. Use acquire use. Plans and packages should be rebound on current releases to benefit from optimizer enhancements. Packages from DB2 V5 or earlier will be automatically rebound when accessed by DB2 10.
DB2 catalog tables are DB2-managed and SMS-managed. Catalog and directory tables do not have links, but have more LOBs and more table spaces. Compression for table space SPT01 is not supported.
REORG TABLESPACE SHRLEVEL NONE on LOB table spaces. Use SHRLEVEL CHANGE or REFERENCE.
DB2 MQ XML functions: Use pureXML functions.
DB2 XML Extender: Use pureXML.
DB2 Management Clients feature is no longer available. Control Center does not support connections to DB2 10.
msys for Setup DB2 Customization Center no longer supported. Use installation panels instead.
Several subsystem parameters
• Annotated XML schema decomposition using XDBDECOMPXML and XDBDECOMPXML100MB
• Accessories Suite will no longer include Optimization Service Center component and the Data Studio Workbench feature. Visual Explain for DB2 for z/OS is not available for DB2 10. The recommended query optimization and service tools for DB2 for z/OS are Optim Query Tuner and Optim Query Workload Tuner. The recommended no-charge query optimization and service tool for DB2 for z/OS is Data Studio (http://www.ibm.com/software/data/studio/). These tools are based and built on the foundation of Optimization Service Center and Optimization Expert.

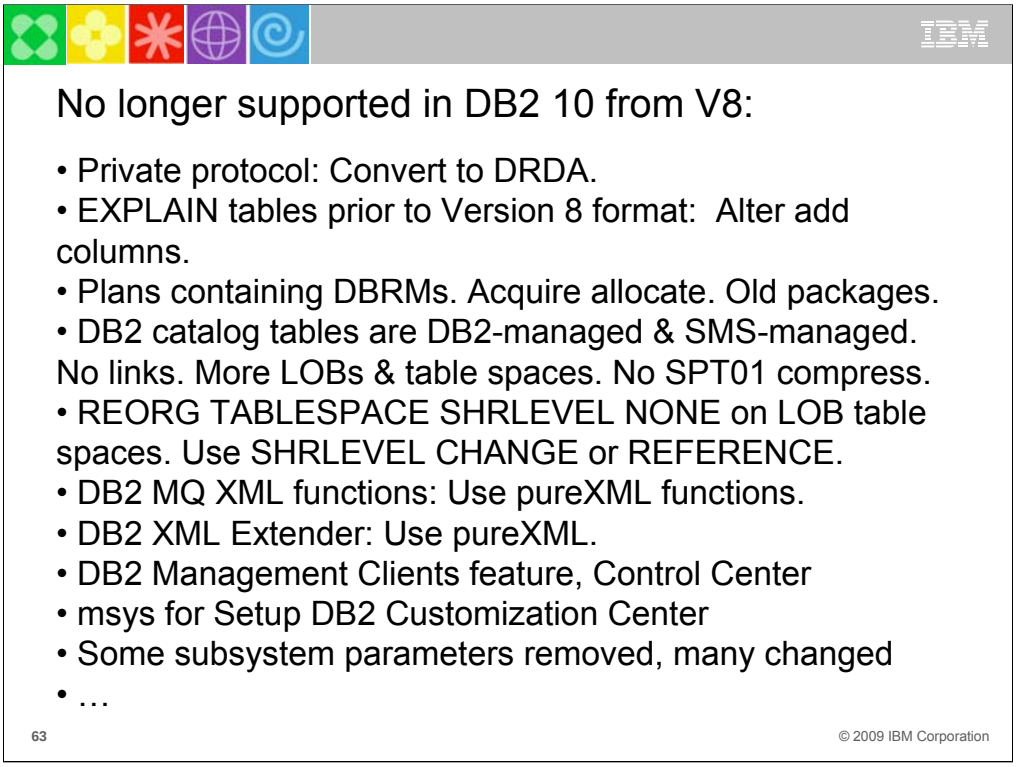

No longer supported. This list is about twice as long as the DB2 9 list.

Private protocol: Convert to DRDA.

EXPLAIN tables prior to Version 8 format: Alter tables to add the needed columns.

Plans containing DBRMs. Plans will contain packages, which have the DBRMs. Use the COLLID option of REBIND to make the needed changes. Acquire allocate is not a supported BIND option any longer. Use acquire use. Plans and packages should be rebound on current releases to benefit from optimizer enhancements. Packages from DB2 V5 or earlier will be automatically rebound when accessed by DB2 10.

DB2 catalog tables are DB2-managed and SMS-managed. Catalog and directory tables do not have links, but have more LOBs and more table spaces. Compression for table space SPT01 is not supported.

REORG TABLESPACE SHRLEVEL NONE on LOB table spaces. Use SHRLEVEL CHANGE or REFERENCE.

DB2 MQ XML functions: Use pureXML functions.

DB2 XML Extender: Use pureXML.

DB2 Management Clients feature is no longer available. Control Center does not support connections to DB2 10.

msys for Setup DB2 Customization Center no longer supported. Use installation panels instead.

Several subsystem parameters are removed, and many change default values.

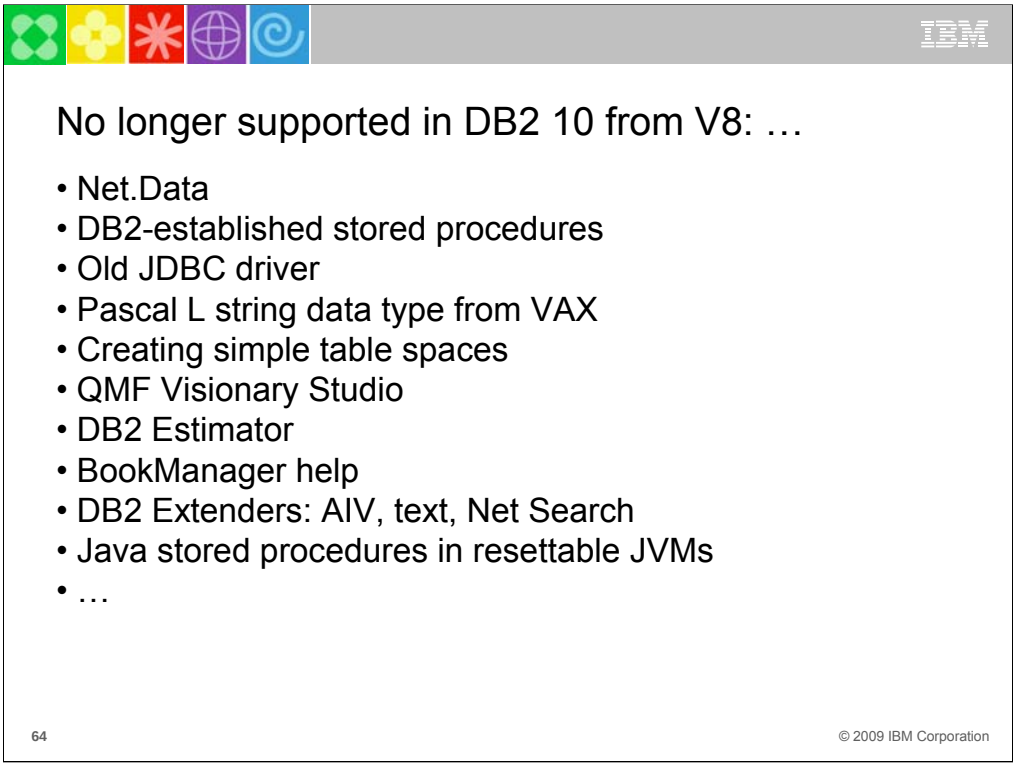

No longer supported …:
•Net.Data is removed. WebSphere is the strategic IBM solution for delivering DB2 data to Web applications.
•DB2-established stored procedure address spaces are no longer supported. Workload Manager (WLM) managed stored procedure address spaces is the strategic solution for stored procedure support, and migration to WLM managed stored procedure spaces is required for use of stored procedures in DB2 10.
•JDBC/SQLJ Driver for OS/390 and z/OS is no longer supported. All Java application programs and Java routines that are currently written to work with the JDBC/SQLJ Driver for OS/390 and z/OS need to be modified to work with the IBM DB2 Driver for JDBC and SQLJ (formerly known as the DB2 Universal JDBC Driver). The steps for migrating JDBC and SQLJ applications from the legacy JDBC/SQLJ Driver for OS/390 and z/OS to the IBM DB2 Driver for JDBC and SQLJ can be found in the Application Programming Guide and Reference for Java (SES1-3023). In addition, all WLM-managed stored procedures address spaces that are set up to execute Java routines must be modified to reference the IBM DB2 Driver for JDBC and SQLJ.
•Connections from VAX machines and the PASCAL L string data type are no longer supported.
•Creation of simple table spaces is no longer supported. DB2 10 for z/OS no longer implicitly creates simple table spaces nor allows customers to create simple table spaces. However, DB2 10 for z/OS continues to support simple table spaces created in previous versions.
•DB2 QMF Visionary Studio program is removed from DB2 QMF Enterprise Edition.
•DB2 Estimator is not available for DB2 10.
•BookManager-based online help has been removed. The prior help support has been replaced by the Information Management Software for z/OS Solutions Information Center (Information Center). The web-based Information Center is updated periodically during the life of each DB2 version, thus ensuring reader access to the most up-to-date information.
•AIV Extender, Text Extender, and Net Search Extender are removed.
•Java stored procedures no longer run in resettable JVMs.

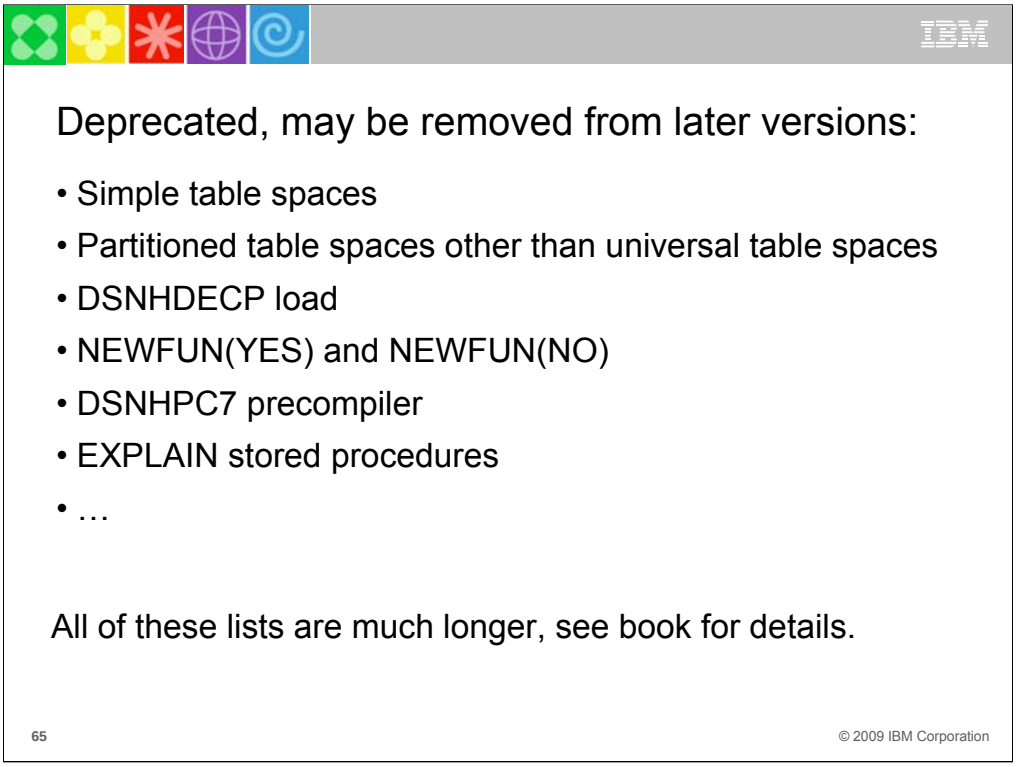

Deprecated and may be dropped from future versions: This list is for all customers on any version, but note items removed as well.

Simple and partitioned table spaces other than universal table spaces are deprecated. Use alter in new function mode to convert single-table table space to universal.

Some current use of DSNHDECP is deprecated. If you have code that loads DSNHDECP and maps it with macros, you should plan to change that code by using the new techniques. Customers who want to have one library for multiple DSNHDECP modules need to make this change.

SQL processing options NEWFUN(YES) and NEWFUN(NO) options are deprecated. Use NEWFUN(V10) rather than NEWFUN(YES). Use NEWFUN(V9) or NEWFUN(V8) rather than NEWFUN(YES).

The DSNHPC7 precompiler is deprecated, although it is still present in DB2 10. Use the current precompiler or coprocessor.

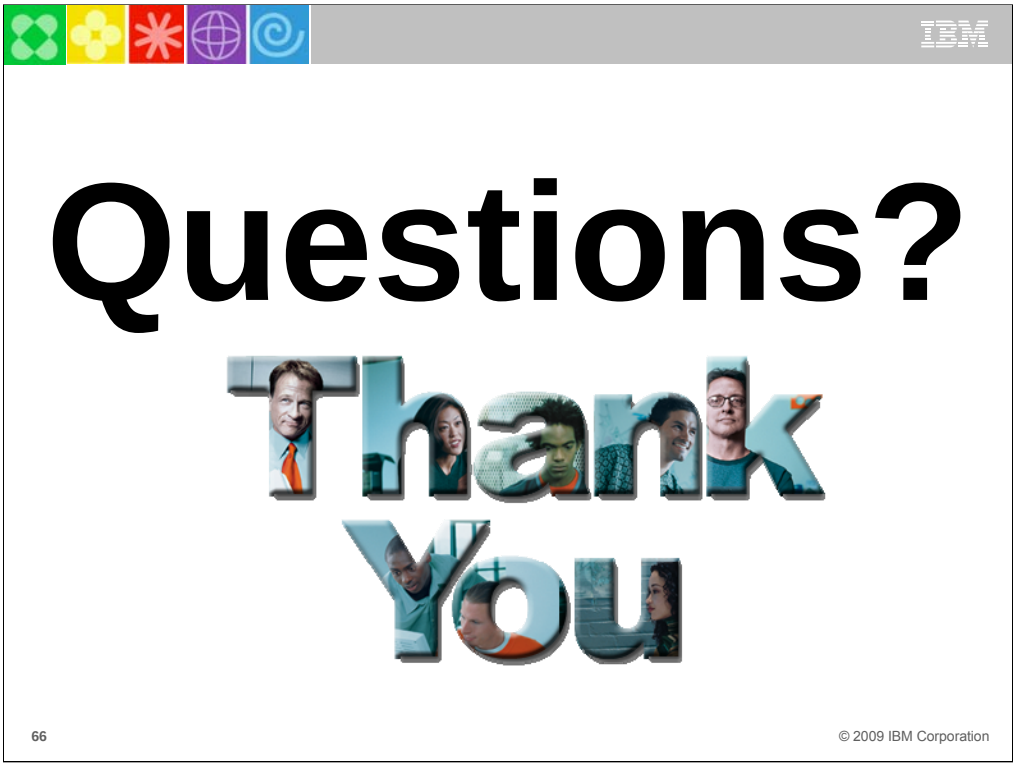

I hope we have lots of questions, but thank you in any case.

DB2 for z/OS Technical Strategy

- Application enablement
  - Apps can easily connect to DB2 from anywhere
  - Advanced SQL, XML capability, application portability
- Extend the lead in availability, scalability and performance.
  - Parallel Sysplex: the best scale-out solution in the industry
  - Tight integration between DB2 and the System z hardware and z/OS operating system
  - Advanced solutions for compliance with data security and privacy regulations
  - Workload consolidation: System z is the ultimate consolidation platform
  - Eliminate all causes of outages
- Reduce cost of ownership
  - DB technology that can handle large workloads with fewer people
  - Advanced autonomics to make the system more self-managing and self-tuning
  - Storage and CPU optimization, including specialty engines
- Improved data warehousing capabilities

© 2009 IBM Corporation

DB2 10 Major Focus Areas

- Performance, reduced CPU consumption
- Scalability, more concurrent active threads
- Continuous availability
- Contention reduction
- DBA / system administration productivity
- Advanced application functionality
- Security advancements

© 2009 IBM Corporation

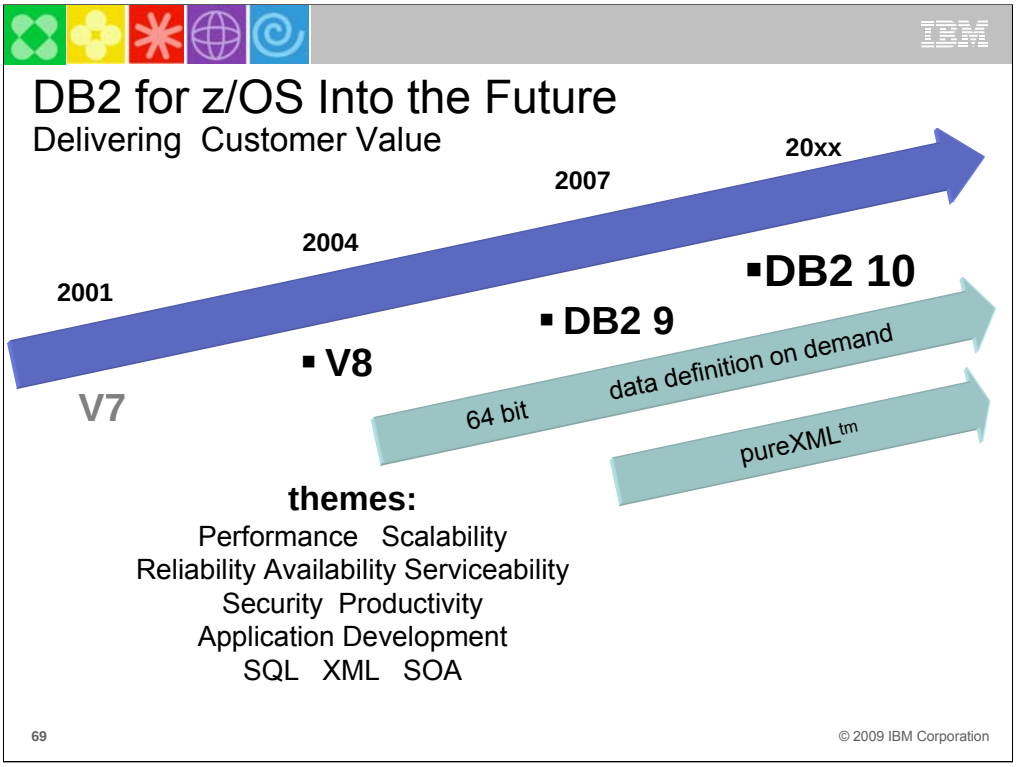

DB2 for z/OS V7 became generally available (GA) March 2001, and V8 delivered three years later. DB2 9 became generally available in March 2007, three more years. We expect the next version will be roughly 3 years from DB2 9 GA to DB2 10 or DB2 10 or whatever the name becomes.

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 10.

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.

DB2 9 and 10 in IBM Redbooks Publications

1. DB2 9 Technical Overview SG24-7330
2. DB2 9 Performance Topics SG24-7473 updated Dec. 2009
3. DB2 9 Stored Procedures SG24-7604
4. Index Compression with DB2 9 for z/OS redp4345
5. SQL Reference for Cross-Platform Development
6. Enterprise Database Warehouse, SG24-7637
7. 50 TB Data Warehouse on System z, SG24-7674
8. New Tools for Query Optimization SG24-7421
9. LOBs with DB2 for z/OS SG24-7270
10. Deploying SOA Solutions SG24-7663
11. Enhancing SAP - DB2 9 SG24-7239
12. SAP Application on Linux z SG24-6847
13. Best practices SAP BI - DB2 9 SG24-6489-01
14. Data Sharing in a Nutshell, SG24-7322
15. Securing DB2 & MLS z/OS SG24-6480-01
16. Data Sharing: Distributed Load Balancing & Fault Tolerant Configuration redp4449
17. Considerations on Small & Large Packages redp4424
18. Backup and Recovery Considerations redp4452
19. Powering SOA with IBM Data Servers SG24-7259
20. Packages Revisited, SG24-7688
21. Data Studio V2.1 Web Services redp4510
22. Ready to Access Solid-State Drives redp4537
23. Distributed Functions SG24-6952
24. Buffer Pool Monitoring & Tuning redp4604
25. Securing & Auditing Data SG24-7720
26. Serialization and Concurrency SG24-4725-01 new
27. Utilities SG24-6289-01 draft

© 2009 IBM Corporation

DB2 library more information http://www.ibm.com/software/data/db2/zos/library.html

Many IBM Redbooks publications, Redpapers and one cross-platform book on DB2 9 are published, in addition to the standard library, with more in the works. Check for updates.

http://www.redbooks.ibm.com/cgi-bin/searchsite.cgi?query=db2+AND+9+AND+for+AND+z/os

1. DB2 9 Technical Overview, SG24-7330 http://www.redbooks.ibm.com/abstracts/SG247330.html
2. DB2 9 Performance Topics, SG24-7473, http://www.redbooks.ibm.com/abstracts/SG247473.html
3. DB2 9 Stored Procedures, SG24-7604, http://www.redbooks.ibm.com/abstracts/SG247604.html
4. Index Compression DB2 9, REDP4345, http://www.redbooks.ibm.com/abstracts/redp4345.html
5. Deploying SOA Solutions SG24-7663, http://www.redbooks.ibm.com/abstracts/SG247259.html
6. Cross-Platform Development, http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html
7. Enterprise Data Warehousing, SG24-7637, http://www.redbooks.ibm.com/abstracts/sg247637.html
8. LOBs: Stronger & Faster SG24-7270, http://www.redbooks.ibm.com/abstracts/SG247270.html
9. Securing DB2 & MLS z/OS, SG24-6480-01, http://www.redbooks.ibm.com/abstracts/sg246480.html
10. Enhancing SAP, SG24-7239, http://www.redbooks.ibm.com/abstracts/SG247239.html
11. Best practices SAP BI, SG24-6489-01, http://www.redbooks.ibm.com/abstracts/sg246489.html
12. New Tools for Query Optimization, SG24-7421, http://www.redbooks.ibm.com/abstracts/sg247421.html
13. Data Sharing in a Nutshell, SG24-7322, http://www.redbooks.ibm.com/abstracts/sg247421.html
14. DB2 9 for z/OS Data Sharing: Distributed Load Balancing and Fault Tolerant Configuration http://www.redbooks.ibm.com/abstracts/redp4449.html
15. Considerations on Small and Large Packages redp4424 http://www.redbooks.ibm.com/abstracts/redp4424.html
16. Backup and Recovery Considerations redp4452 http://www.redbooks.ibm.com/abstracts/redp4452.html
17. Powering SOA IBM Data Servers, SG24-7259 http://www.redbooks.ibm.com/abstracts/SG247259.html
18. Packages Revisited, SG24-7688 http://www.redbooks.ibm.com/abstracts/SG247688.html
19. 50 TB Data Warehouse Benchmark on IBM System z http://www.redbooks.ibm.com/abstracts/sg247674.html
20. SAP on DB2 9 for z/OS: Implementing Application Servers on Linux for System z http://www.redbooks.ibm.com/abstracts/sg246847.html
21. IBM Data Studio V2.1: Getting Started with Web Services on DB2 for z/OS http://www.redbooks.ibm.com/abstracts/redp4510.html
22. Ready to Access DB2 for z/OS Data on Solid-State Drives http://www.redbooks.ibm.com/abstracts/redp4537.html
23. Parallel Sysplex Operational Scenarios http://www.redbooks.ibm.com/abstracts/sg242079.html
24. Distributed Functions http://www.redbooks.ibm.com/abstracts/sg246952.html
25. Buffer Pool Monitoring & Tuning http://www.redbooks.ibm.com/abstracts/redp4604.html
26. Securing and Auditing Data http://www.redbooks.ibm.com/abstracts/sg247720.html
27. Serialization & concurrency, SG24-4725-01 http://www.redbooks.ibm.com/abstracts/sg244725.html
28. Utilities SG24-6289-01 http://www.redbooks.ibm.com/Redbooks.nsf/RedpieceAbstracts/sg246289.html