We discuss the latest news about DB2 for z/OS, including the changes in disks, IBM System z9 and z10 Integrated Information Processor or zIIP and new z10 processors. Then the discussion will move DB2 9 for z/OS, with XML and SQL that is more consistent across the DB2 family. Utility enhancements help with new function, more LOB and XML support, better performance and improved availability, removing the BUILD2 step from online reorg. DB2 9 enhances DB2's ability to handle new web and enterprise applications. DB2 9 improves the ability to handle new applications with XML, large objects, and many SQL and security improvements. DB2 9 builds upon and extends DB2 traditional strengths and the ground-breaking V8 in many areas: online schema evolution, Unicode, XML, DB2 family SQL, utilities, security and 64-bit virtual storage.

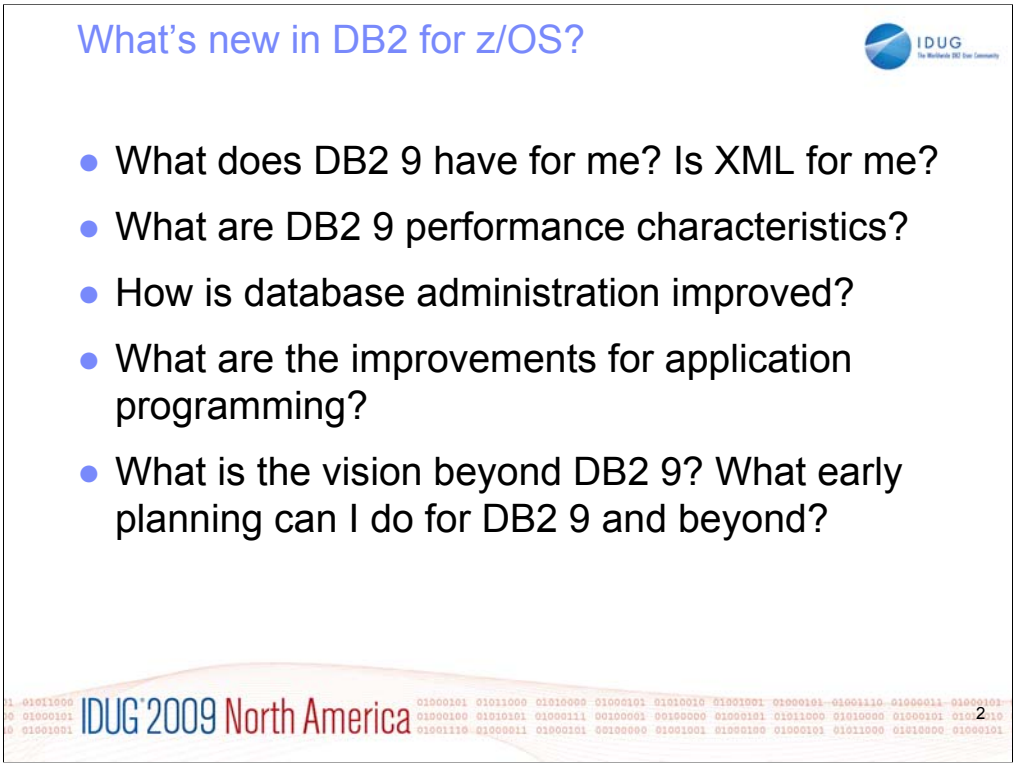

We discuss the latest news about DB2 for z/OS, including the changes in disks, IBM System z9 and z10 Integrated Information Processor or zIIP and new z10 processors. Then the discussion will move DB2 9 for z/OS, with XML and SQL that is more consistent across the DB2 family. Utility enhancements help with new function, more LOB and XML support, better performance and improved availability, removing the BUILD2 step from online reorg. DB2 9 enhances DB2's ability to handle new web and enterprise applications. DB2 9 improves the ability to handle new applications with XML, large objects, and many SQL and security improvements. DB2 9 builds upon and extends DB2 traditional strengths and the ground-breaking V8 in many areas: online schema evolution, Unicode, XML, DB2 family SQL, utilities, security and 64-bit virtual storage.

Innovating for System z Workloads

IBM Information Infrastructure is Optimized for Mainframe Environments

Latest changes

- System z10 EC and BC
- High performance FICON
- Solid state disk
- Disk encryption
- InfoSphere Warehouse on System z
- CICS V4

All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice. Availability, prices, ordering information, and terms and conditions will be provided when the product is announced for general availability.

The System z environment keeps getting bigger, better, faster, and cheaper. Today we'll talk about many changes, starting with the platform. System z10 processors offer a new range of faster processors with lower costs. Getting information into the processors has improved in a number of ways, with faster FICON channels, a solid state disk option for a new addition to the storage pyramid, and encryption for disk drives.

A new beta program has begun for InfoSphere Warehouse on System z, providing many of the same benefits delivered on other platforms. CICS Version 4 extends the scope and brings in the CICS Explorer, using an Eclipse framework.

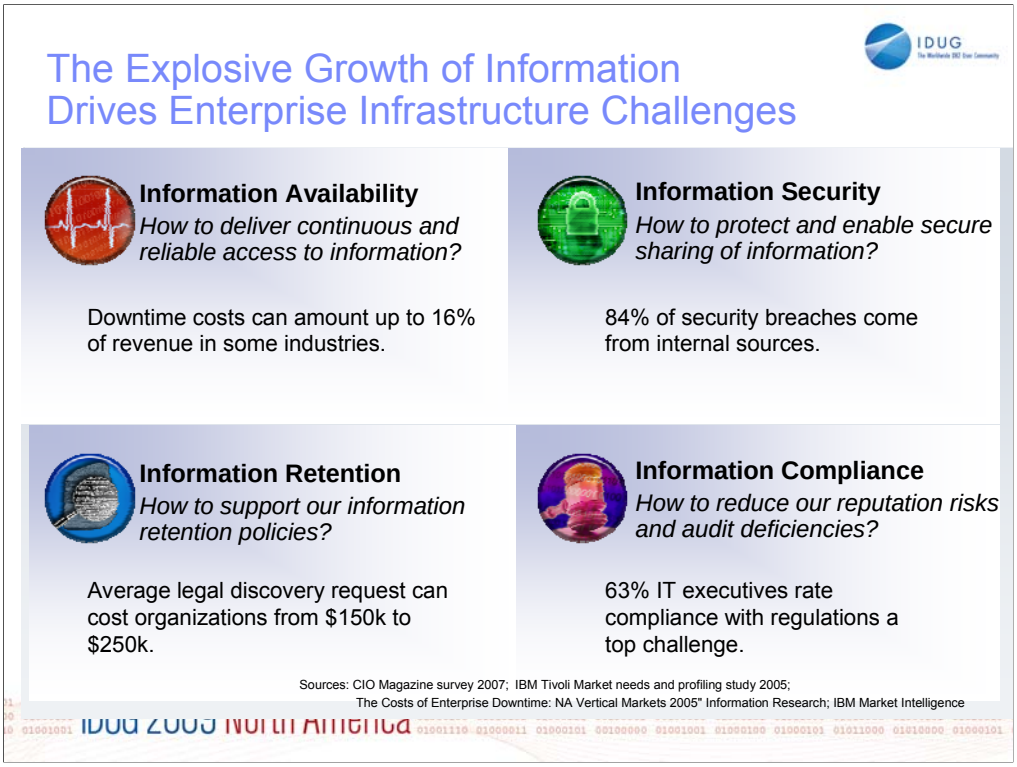

What we're proposing is that enterprises that want to succeed in such a challenging business climate need to focus on four key areas to ensure their information infrastructures can support the goals of the business. The 4 key areas are:

- Information Availability
- Information Security
- Information Retention
- Information Compliance

•What do we mean by Information Infrastructure? This is the overall infrastructure of integrated and optimized storage HW and SW, application servers, applications and middleware, networks, and endpoint devices that combine to deliver information across the extended enterprise.

The IBM Information Infrastructure aims to help businesses get the right information to the right people when they need it… in a secure manner.

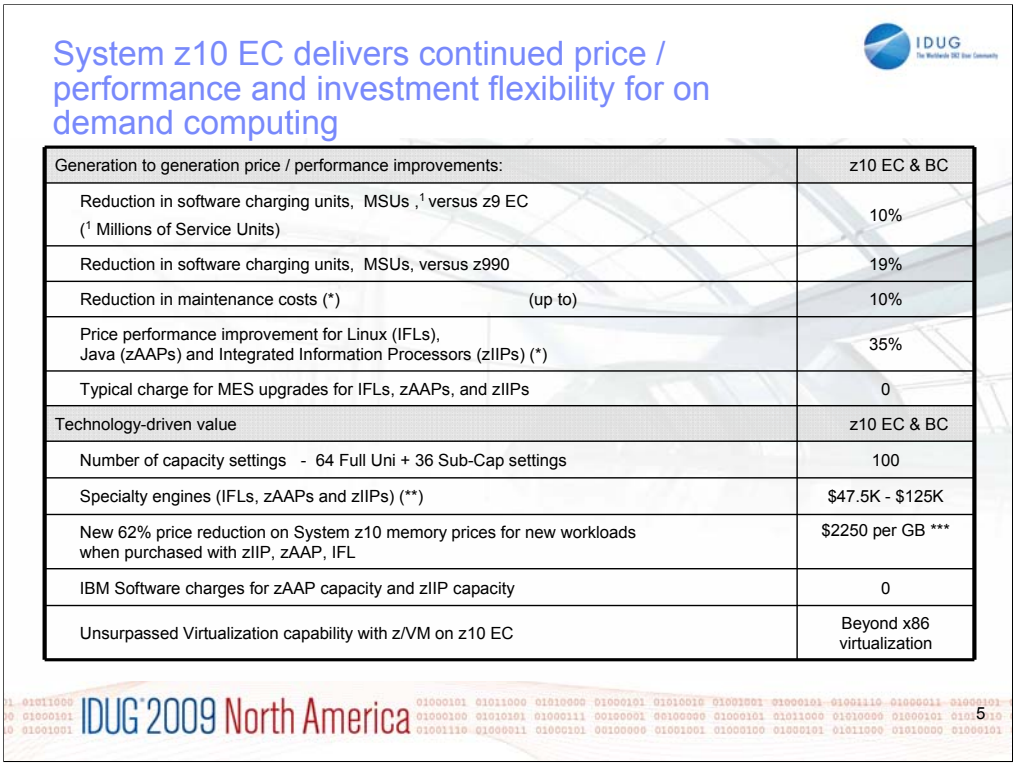

Many of the improvements help directly in price and performance. The dollars for MIPS, which I still think of in thousands of dollars per MIPS, is now about $66 on the z10 BC – not thousands, just dollars.

Memory has been reduced as well, so that you can purchase 16 gigabytes of memory per engine with a zIIP at about $2250 per gigabyte, a reduction of about 60%.

Plus there are System z New Application License Charge (zNALC) pricing metrics for New Workloads

On/Off Capacity on Demand (On/Off CoD) enhancements to better manage volatile business requirements

(*) – comparisons shown are z10 EC vs. z9 and z10 EC vs. z990

(**) Prices may vary by country $47.5K price is for z10 BC zIIP

(***) Limited to 16 GB per engine

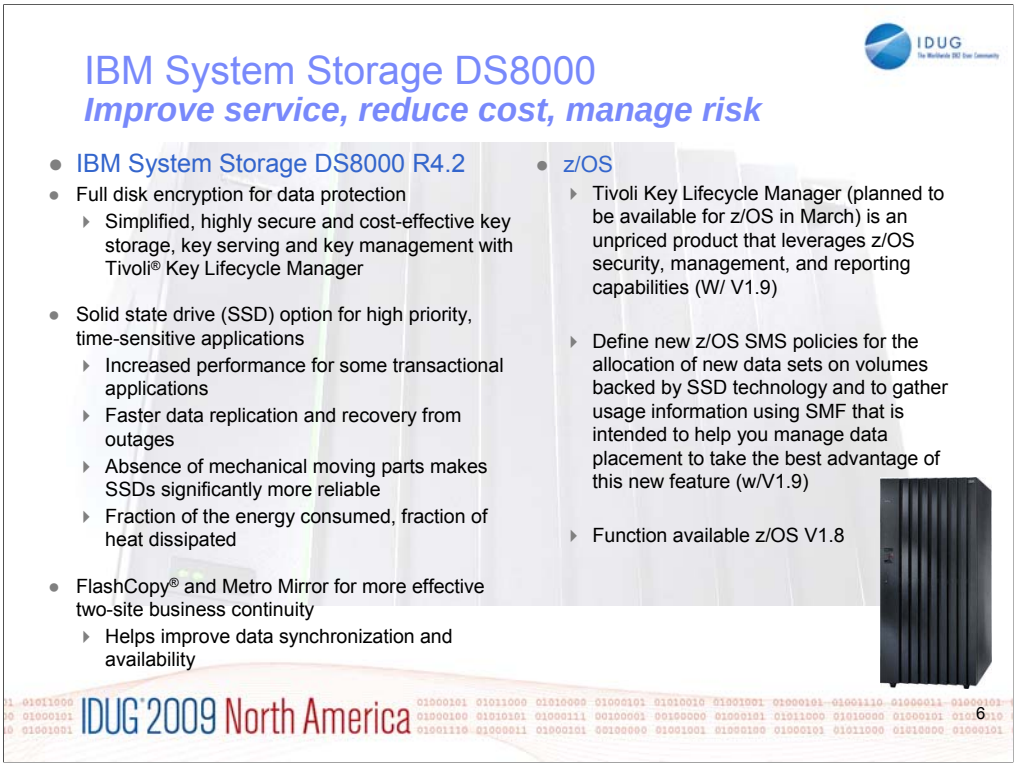

The IBM System Storage DS8000 R4.2 was announced Feb 9 2009 and will GA March 2009. The DS8000 R4.2 has many enhancements, the following are exclusive to z/OS environments:

•Solid state drives (SSD, also called flash memory) support. This new feature on DS8000 controllers introduces a powerful new technology to IBM enterprise storage by using solid state flash memory arrays in place of hard disk drives (HDD). Because it eliminates disk latency, SSD-based storage can provide a substantial improvement in storage response times compared to HDD. This technology provides the greatest performance advantage for high-use data that is written infrequently and processed randomly. Support is planned to allow you to define new SMS policies for the allocation of new data sets on volumes backed by SSD technology and to gather usage information using SMF intended to help you manage data placement to take the best advantage of these new features.

•Metro Mirror (PPRC) support for Flashcopy. This function will be designed to allow you to perform FlashCopy operations where Metro Mirror volumes are the destination while preserving the duplex state of the volumes, and to copy the data to secondary volumes during the FlashCopy operations. This is expected to help you improve storage availability and achieve more consistent recovery point objectives (RPO) by removing this reason for interrupting Metro Mirror duplexing.

•Drive-level encryption and advanced key management support. With this new functionality, the DS8000 secures sensitive data when drives are physically removed from the data center. Drive-level encryption has no performance impact and the solution is transparent to the application and to the server, which minimizes the costs associated with encrypting data at rest with a solution that is simple to deploy and simple to maintain.

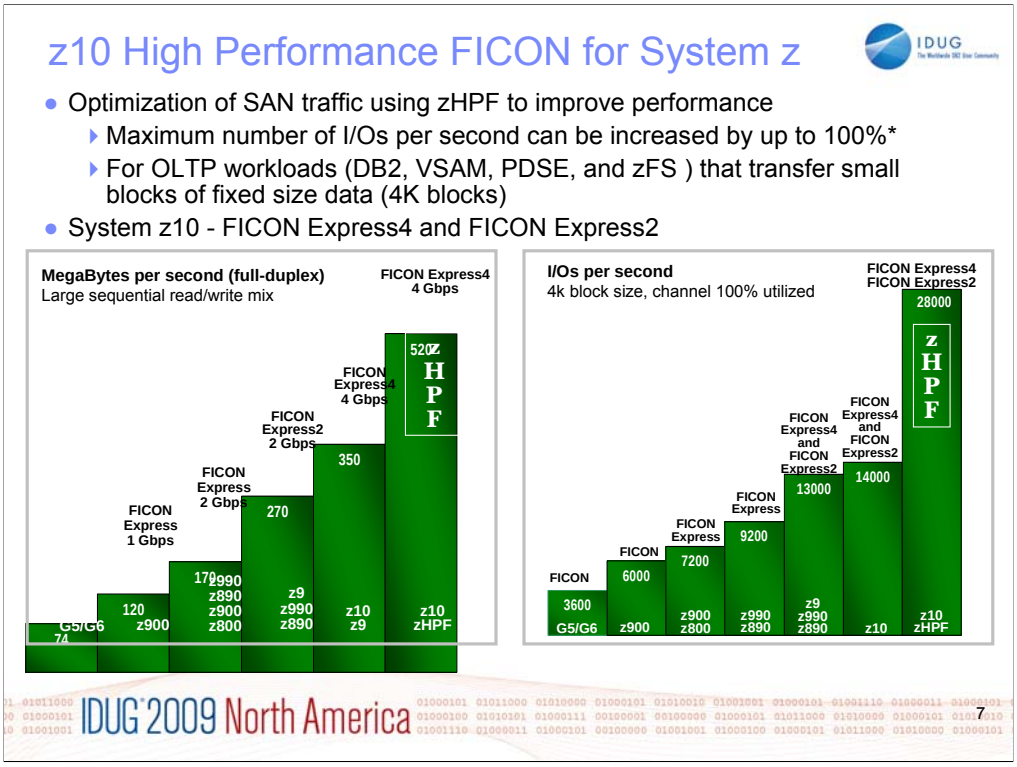

High Performance FICON and Solid State Disk more than doubles the random throughput per channel for the same amount of channel time.

Channel consolidation - fewer channels can manage more storage capacity.

Actual throughput depends on the percentage of I/Os eligible for zHPF. DB2 prefetch I/Os are not yet eligible.

Improves SSD response time and throughput by 20%

zHPF requires DS8000 Release 4.1, z10 processor and z/OS 1.10 or SPE to z/OS 1.8

SSD technology speeds up data access and removes bottlenecks imposed by spinning disks

– Other important performance features include MIDAW, HyperPAV, AMP, High Performance FICON, 4 gbps FICON links

Optimization of storage area network (SAN) traffic using zHPF to improve performance

- Maximum number of I/Os per second can be increased by up to 100%*
- For OLTP workloads (DB2, VSAM, PDSE, and zFS) that transfer small blocks of fixed size data (4K blocks)

Exclusive to System z10 - FICON Express4 and FICON Express2

Requires Control unit exploitation – IBM DS8000™ Release 4.1

- z/OS V1.7 with the IBM Lifecycle Extension for z/OS V1.7 (5637-A01), V1.8, V1.9, or V1.10 with PTFs

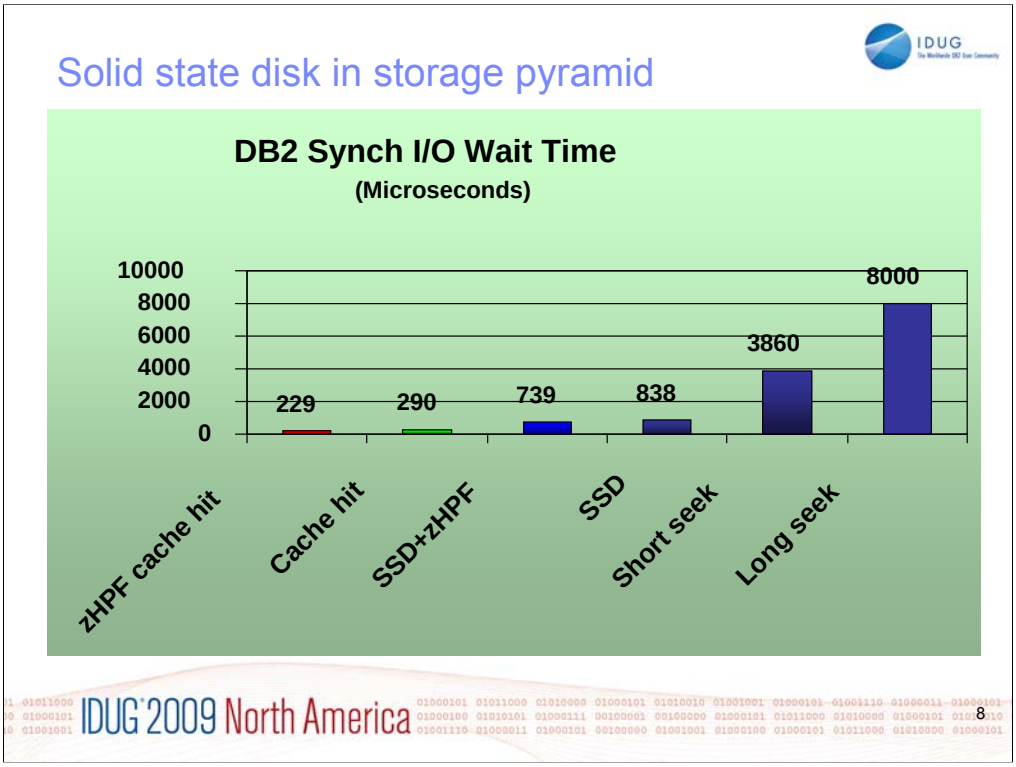

Understanding where solid state disk fits into the storage hierarchy is very important. This device is very fast for random IO, but it's more expensive than magnetic disk.

Solid state disk provide much faster random IO times than hard disk. The disk cache is the fastest, at 229 to 290 microseconds. The solid state disk times are 739 to 838 microseconds. Getting data from the spinning disk requires roughly 4 to 8 milliseconds.

If you have very effective bufferpools and disk cache today, then solid state disk might not be as useful. If your IO is more random, so that your average random IO is 4 milliseconds or more, then the solid state disk may be able to hold enough of the highest IO density information to reduce the IO times significantly.

While solid state disk is higher than magnetic disk today, the cost is rapidly lowering, and new technology promises even more reductions. See Jeff Berger's presentations for more detail.

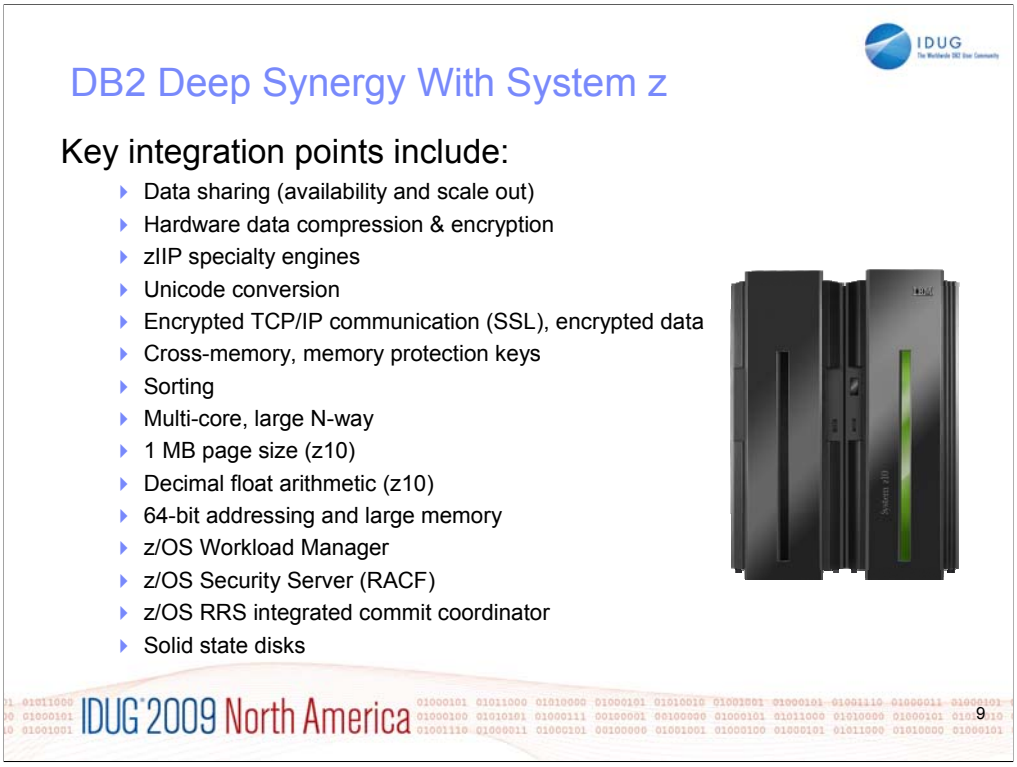

Data sharing is a prime example of deep synergy with System z. DB2 worked with the System z design team for nearly 10 years to produce a robust platform for horizontal scaling. The evolution has continued for 15 more years now after delivery.

Hardware data compression and encryption provides improved costs, easier management and robust resilience for the platform. Cross-memory and protection keys work with APF authorization and RACF for the underlying system integrity.

Specialty engines can reduce costs very substantially, reducing both hardware and software costs.

The z/OS workload manager (WLM) has changed in almost every release to improve work flow with DB2. DB2 has a dispatcher, the z/OS WLM.

Sorting, decimal arithmetic, decimal float, encryption, and Unicode conversions are examples of unique instructions in z/Architecture that DB2 uses.

System z10 Benefits for DB2

Taking System z to the next level

- Faster CPUs, more CPUs, more memory means better DB2 performance, scalability
  - 50% more n-way performance "on average"
  - About 60% more raw uniprocessor performance
    - Mileage will vary based on workload
  - 70% more server capacity (56->64 CPUs)
  - Up to 64 CPUs, z/OS 1.9 needed for 64-way in a single lpar
  - Up to 1.5 TB, z/OS 1.8 needed for >256G in a single lpar
- New OSA Express3, 10GbE for faster DRDA access
- Infiniband CF links
- HiperDispatch
- Hardware Decimal Floating Point facility
- 1MB page size (DB2 X plans to exploit)
- 50+ machine instructions added to improve compiled code efficiency (DB2 X plans to exploit)

The design of the IBM System z10™ processor chip is the most extensive redesign in over 10 years, resulting in an increase in frequency from 1.7 GHz (z9 EC) to 4.4 GHz on the z10 EC. The z10 BC processors run at 3.5 GHz. DB2 measurements found most workloads having a range of 1.4 to 2.1 times faster with z10 compared to z9.

Larger memory: DB2 users can potentially see higher throughput with more memory used for DB2 buffer pools, EDM pools or SORT pools.

Improved IO: improvements in the catalog and allocation can make the large number of data sets much faster and easier to manage. Disk IO times and constraints can be reduced.

System z is designed for secure data serving, yet also was enhanced to provide improvement enhances for CPU intensive workloads. The result is a platform that continues to improve upon all the mainframe strengths customers expect, yet opens a wider aperture of new applications that can all take advantage of System z10s extreme virtualization capabilities, and lowest TCO versus distributed platforms.

See section 4.3.1 z10 performance in the latest updates of DB2 9 for z/OS Performance Topics, SG24-7473 for additional detail.

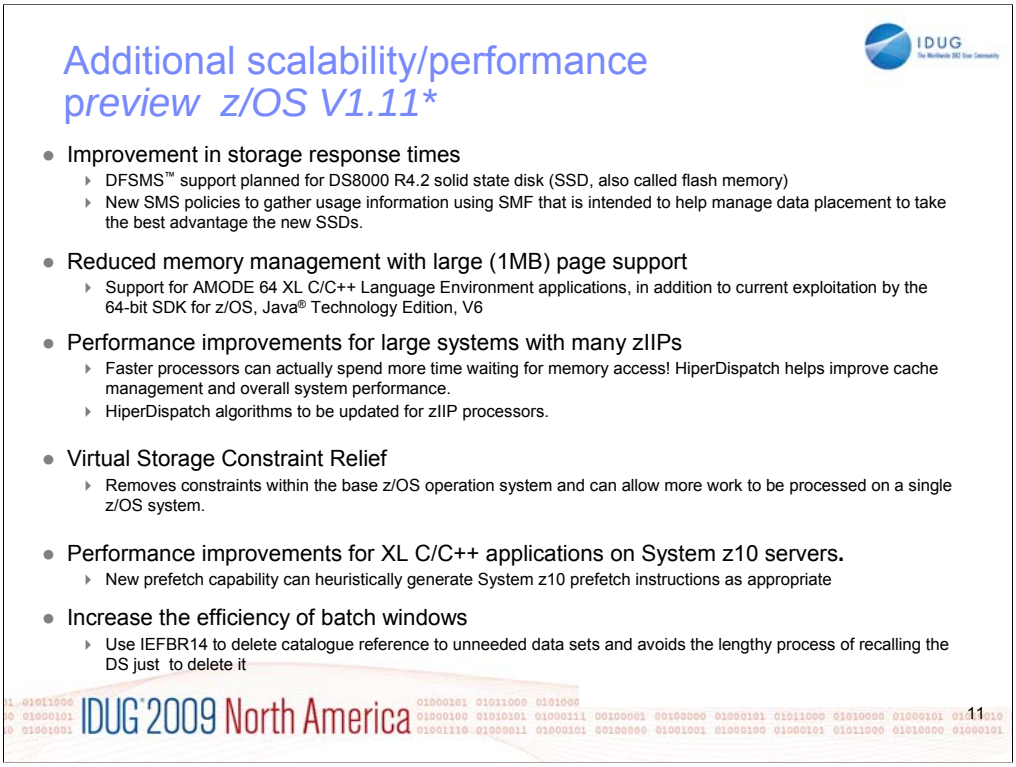

The industry has begun to hit fundamental physical limits for chip design. Large-magnitude CPU speed increases with each generation of chip are a thing of the past and capacity increases will increasingly come from higher n-way and more multi-threading. z/OS and its subsystems provide for scalability not only based on chip speeds, but on a single image, clustering, storage, and data handling basis as well. This holistic and balanced approach to scalability means your System z environment is capable of handling the growth of your user base, applications, business processes, and data processing needs.
Details on the scalability and performance enhancements intended for z/OS V1.11:
In z/OS V1.11, **DFSMS** support is planned for solid state drives (SSD, also called flash memory) on DS8000. This new DS8000 feature introduces a powerful new technology to IBM enterprise storage by using solid state flash memory arrays in place of hard disk drives (HDD). Because it eliminates disk latency, SSD-based storage can provide a substantial improvement in storage response times compared to HDD for high-use data that is written infrequently and processed randomly. Also, support is planned to allow you to define new SMS policies for the allocation of new data sets on volumes backed by SSD technology and to gather usage information using SMF that is intended to help you manage data placement to take the best advantage of these new features. This function is planned to be made available on z/OS V1.9 and V1.10 with PTF for APAR OA27703. For more information about the new DS8000 features, please refer to Hardware Announcement dated February 10, 2009.
Appropriate use of large (1 MB) pages can help reduce memory management overhead and increase translation lookaside buffer (TLB) hit ratios for exploiting programs.
z/OS V1.11 has designs for continued improvement to cache and memory management; this will continue to be true for the foreseeable future. The time it takes to retrieve data from memory, while progressively shorter in the absolute sense on newer server models, has become progressively longer when measured in processor cycle time increments. In other words, though memory access is much faster than it used to be, processors spend more cycles waiting for it than they ever did. **HiperDispatch** helps address the system's management of cache. In z/OS V1.10, HiperDispatch was directed toward improving cache management in multiprocessing LPARs. For z/OS V1.11 HiperDispatch design is planned to be changed to improve the performance for large scale z/OS systems that include zIIP processors. These changes will be intended to improve system performance for LPARs with a large number of zIIPs.

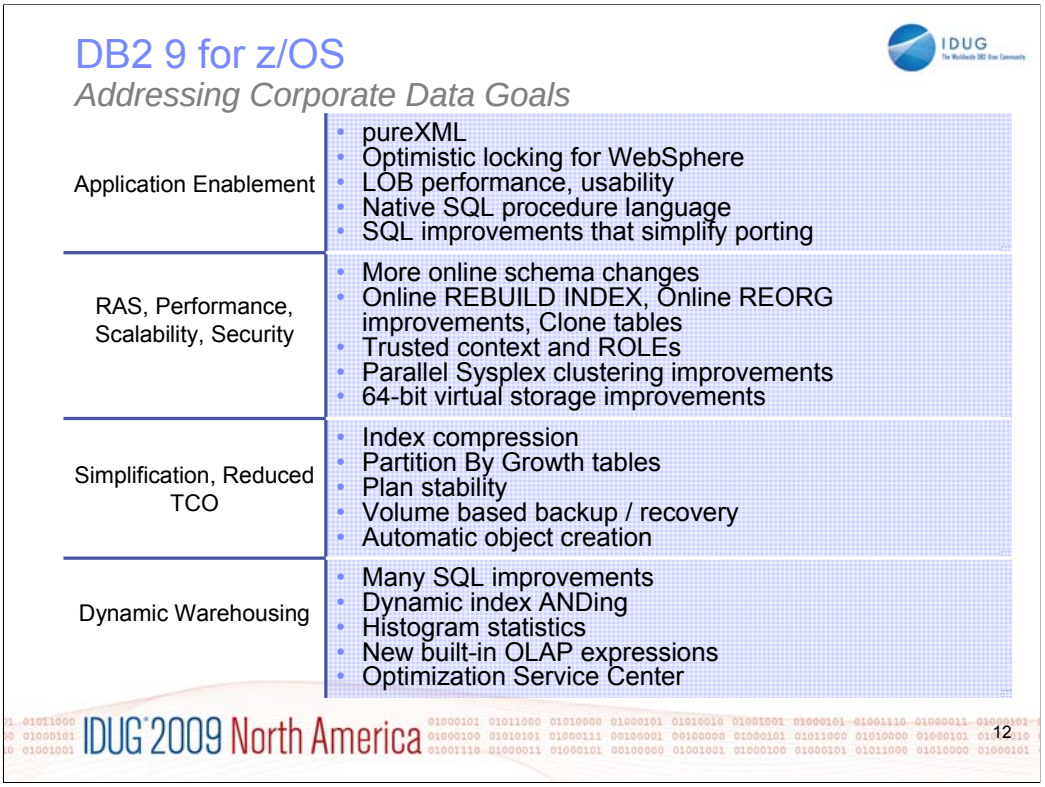

One of the key initiatives of V8 was online schema evolution, and that theme is expanding and changing to be data definition on demand. These are key improvements for resilience. One of the important changes is to be able to replace one table quickly with another. Another is to be able to rename a column 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. Changing DB2 early code does not require an IPL.

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.

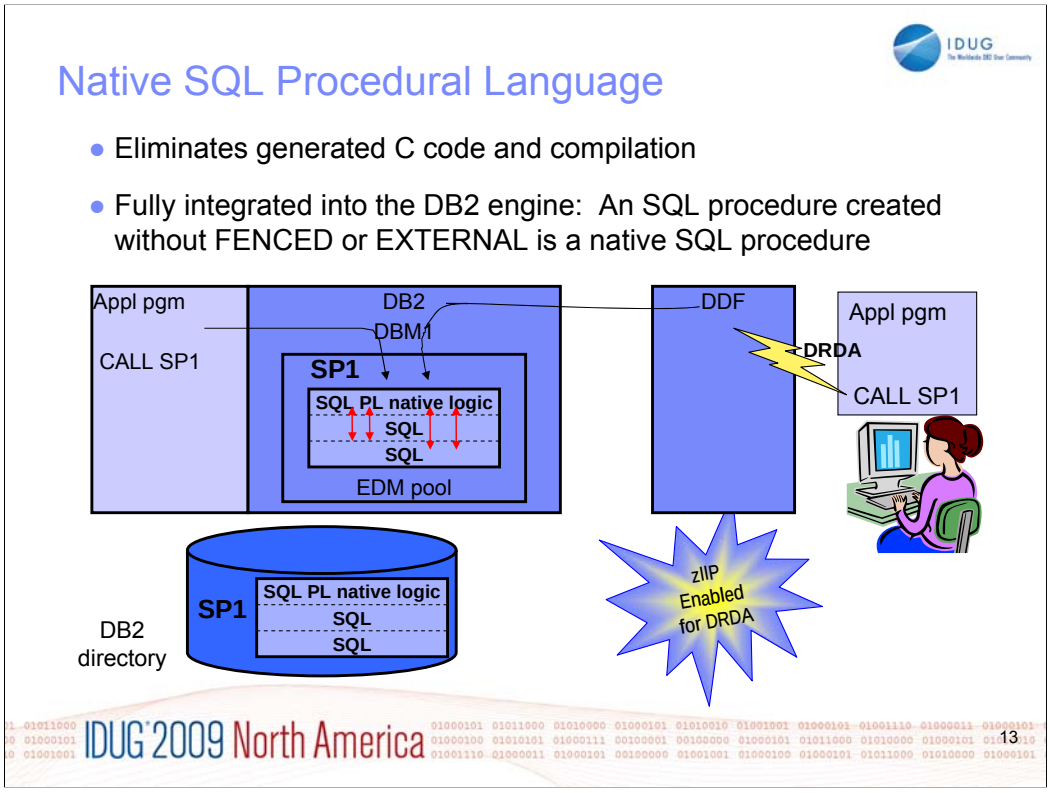

With DB2 9 new function mode, when you create a native SQL stored procedure, its procedural statements are now converted to a native representation that is stored in the DB2 catalog and directory, as it is done with other SQL statements. The parameter list and procedure options are stored in the database catalog tables as in the prior releases. When you call a native SQL procedure, DB2 loads the native representation from the catalog and the DB2 engine executes the procedure.

Extensive support for versioning: VERSION keyword on CREATE PROCEDURE
CURRENT ROUTINE VERSION special register ALTER ADD VERSION
ALTER REPLACE VERSION ALTER ACTIVATE VERSION

BIND PACKAGE with new DEPLOY keyword. Allow to deploy from Test to prod without doing a CREATE PROC stmt

Native SQL stored procedures

Stored procedures written in SQL procedure language enhance portability and ease of use when using DB2 for z/OS as your enterprise information source. This language is an ANSI standard language. It is similar to the proprietary stored procedure languages of several competitive databases, which assists in migrating and porting to DB2 for z/OS.

SQL stored procedures are supported by the DB2 Development Center tooling, providing an environment to code, test, and debug modules from your connected workstation. This language is currently converted to C when the CREATE PROCEDURE statement is executed. The C program is then automatically prepared, compiled, linked, and bound. The developer does not need to work with the C code.

SQL stored procedures code will be natively integrated into the DB2 engine, eliminating the conversion to C. Additionally, extensions to the bind command will allow for the promotion of the program and access paths between environments without needing to recreate the stored procedure.

•When native stored procedure requests are invoked from DRDA TCP/IP connections, the processing within the native stored procedure is eligible for zIIP specialty engine processing.

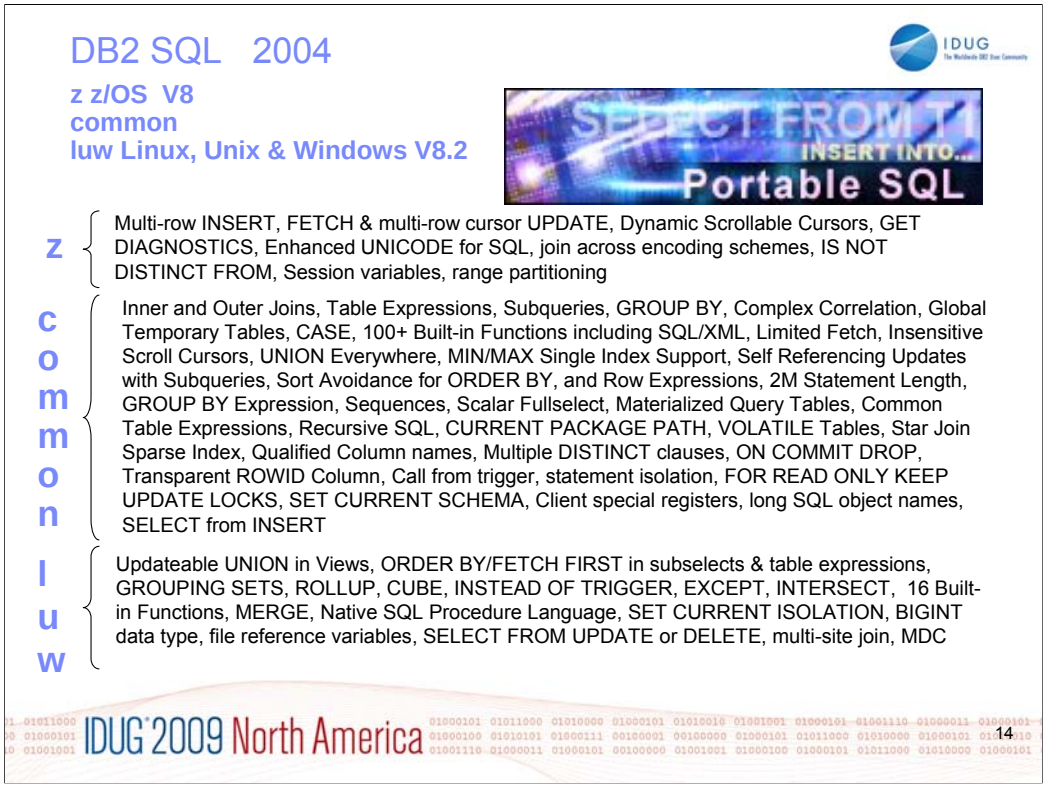

This chart shows the relationship of SQL in the DB2 family comparing DB2 for Linux, Unix & Windows with DB2 for z/OS for key language constructs. This chart compares the z/OS Version 8 from March 2004 with the LUW version from October 2004.

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

DB2 SQL 2007

z z/OS 9
common
luw Linux, Unix & Windows 9

z — Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE for SQL, join across encoding schemes, IS NOT DISTINCT FROM, Session variables, **TRUNCATE, DECIMAL FLOAT, VARBINARY, optimistic locking, FETCH CONTINUE, ROLE, MERGE, SELECT from MERGE**

common — Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions including SQL/XML, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index Support, Self Referencing Updates with Subqueries, Sort Avoidance for ORDER BY, and Row Expressions, 2M Statement Length, GROUP BY Expression, Sequences, Scalar Fullselect, Materialized Query Tables, Common Table Expressions, Recursive SQL, CURRENT PACKAGE PATH, VOLATILE Tables, Star Join Sparse Index, Qualified Column names, Multiple DISTINCT clauses, ON COMMIT DROP, Transparent ROWID Column, Call from trigger, statement isolation, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT, **UPDATE or DELETE, INSTEAD OF TRIGGER, Native SQL Procedure Language, BIGINT, file reference variables, XML, FETCH FIRST & ORDER BY in subselect and fullselect, caseless comparisons, INTERSECT, EXCEPT, not logged tables**, **range partitioning, compression**

luw — Updateable UNION in Views, GROUPING SETS, ROLLUP, CUBE, 16 Built-in Functions, SET CURRENT ISOLATION, multi-site join, MERGE, MDC, **XQuery**

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

DB2 SQL new cross platform reference

z z/OS **9**
common
luw Linux, Unix & Windows **9.5**

z Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE for SQL, join across encoding schemes, IS NOT DISTINCT FROM, TRUNCATE, VARBINARY, FETCH CONTINUE, MERGE, SELECT from MERGE, index compression

common Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions including SQL/XML, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index, Self Referencing Updates with Subqueries, Sort Avoidance for ORDER BY, and Row Expressions, 2M Statement Length, GROUP BY Expression, Sequences, Scalar Fullselect, Materialized Query Tables, Common Table Expressions, Recursive SQL, CURRENT PACKAGE PATH, VOLATILE Tables, Star Join Sparse Index, Qualified Column names, Multiple DISTINCT clauses, ON COMMIT DROP, Transparent ROWID Column, Call from trigger, statement isolation, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT, UPDATE or DELETE, INSTEAD OF TRIGGER, Native SQL Procedure Language, BIGINT, file reference variables, XML, FETCH FIRST & ORDER BY in subselect & fullselect, caseless comparisons, INTERSECT, EXCEPT, not logged tables, range partitions, data compression, **session variables, DECIMAL FLOAT, optimistic locking, ROLE**

luw Updateable UNION in Views, GROUPING SETS, ROLLUP, CUBE, **more** Built-in Functions, SET CURRENT ISOLATION, multi-site join, MERGE, MDC, XQuery, **XML enhancements, array data type, global variables, vendor syntax**

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

What You Can Do with pureXML

- Create tables with XML columns or alter table add XML columns
- Insert XML data, optionally validated against schemas
- Create indexes on XML data
- Efficiently search XML data
- Extract XML data
- Decompose XML data into relational data
- Construct XML documents from relational and XML data
- All the utilities and tools support for XML

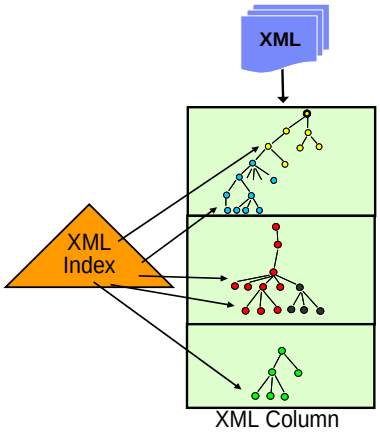

The amount of business information in XML form is already as great or greater than other forms and growing faster - failure to leverage efficiently as structured data means high cost and/or missed opportunity. DB2 9 provides the best of both worlds, pureXML[tm] for native storage and integrating XML with object-relational. Performance, integrity, protection, and scale from the proven DB2 infrastructure with the flexibility of XML/XPath and relational/SQL. This overcomes the complexity & limitations of prior models (shred, CLOB, or XML only). In 2006 IBM introduced a new generation data server with the availability of DB2 9. The explosive growth of XML based data standards in all industries means competitive advantage for those businesses that use it most effectively and efficiently. Client, policy and claims processing in Insurance; supply chain management in Retail; financial transactions and asset management in Banking; patient care in Healthcare; citizen service in Government; implementing Service Oriented Architectures (SOA) in Computing Software and Services - and many other processes across all industries - increasingly rely on information captured and exchanged in XML form. Our clients are increasingly managing XML format text documents in a content management system for proper governance and efficient use in the business process workflow. But few are realizing the full value of all the business data they possess that are in XML format.

The impact of pureXML is seen by a large banking client with a requirement to update over 500,000 XML data records per day. Attempts to use a competitors relational data server failed. Using DB2 9 with pureXML, the application was able to update more than half a million data records in less than an hour.

A large insurance client has seen the impact of pureXML to development time and cost with a 65% reduction in lines of code and more than 75% reduction in time required to develop services accessing XML data.

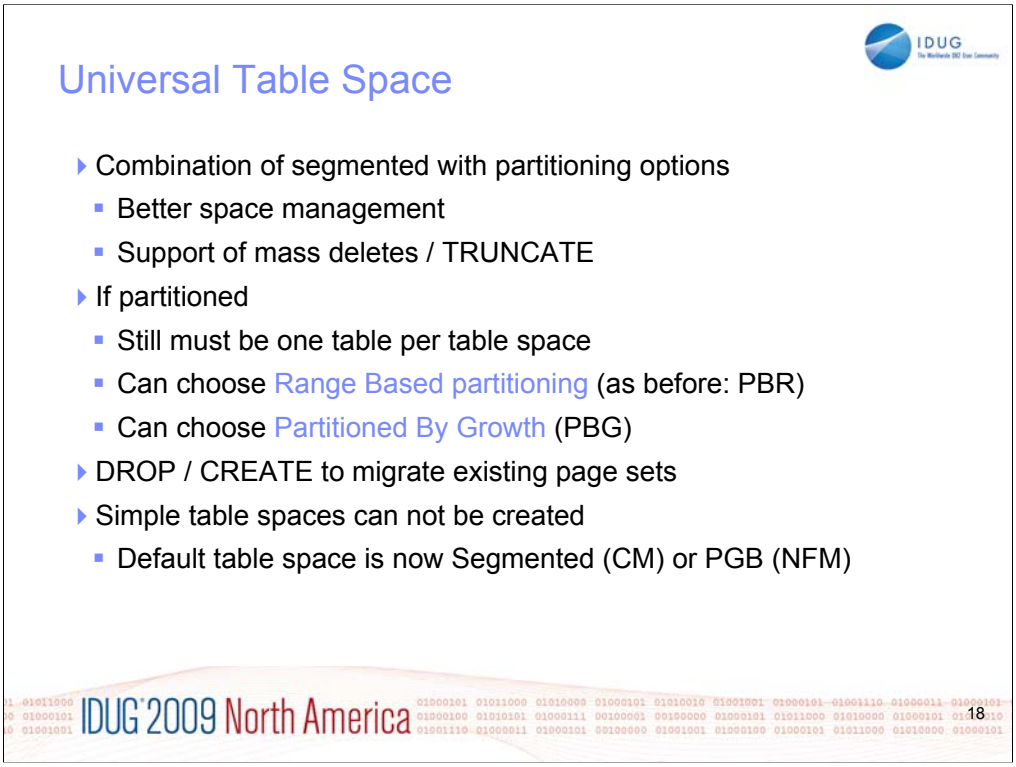

We need a partitioned table space with some of the segmented table space features. The advantages of segmented space maps for partitioned table spaces:
Universal Table space (Partition By Range): A partitioned segmented table space. Partitioning column required. One table per table space.
Universal Table space (Partition By Growth) A partitioned segmented table space. No partitioning column required. One table per table space.
Partition By Growth (PBG): Single-table table space, where each partition contains a segmented page set (allows segmented to increase from 64GB to 16TB or 128 TB with 32K pages). Eliminates need to define partitioning key and assign key ranges. Partitions are added on demand. A new partition is created when a given partition reaches DSSIZE. DSSIZE defaults to 64G. Up to MAXPARTITIONS Retains benefits of Utilities and SQL parallelism optimizations for partitioned tables SEGSIZE defaults to 4 & LOCKSIZE defaults to ROW Considerations: Single-table table space Always defines as LARGE (DSSIZE) Need PBR for query partition elimination No LOAD PART, ALTER ADD PART, or ROTATE PART All indexes are NPSIs
Automatic repositioning of Variable columns to end of row. Length attributes replaced with Indicators positioned after fixed length columns. Any table space created or reorged in DB2 9 NFM.
To Convert: REORG or LOAD REPLACE a table space or partition, or ADD PARTITION No EDITPROCs or VALIDPROCs PIT RECOVER will set the table space to the row format of the PIT
Catalog / Directory remains in Basic Row Format (BRF)

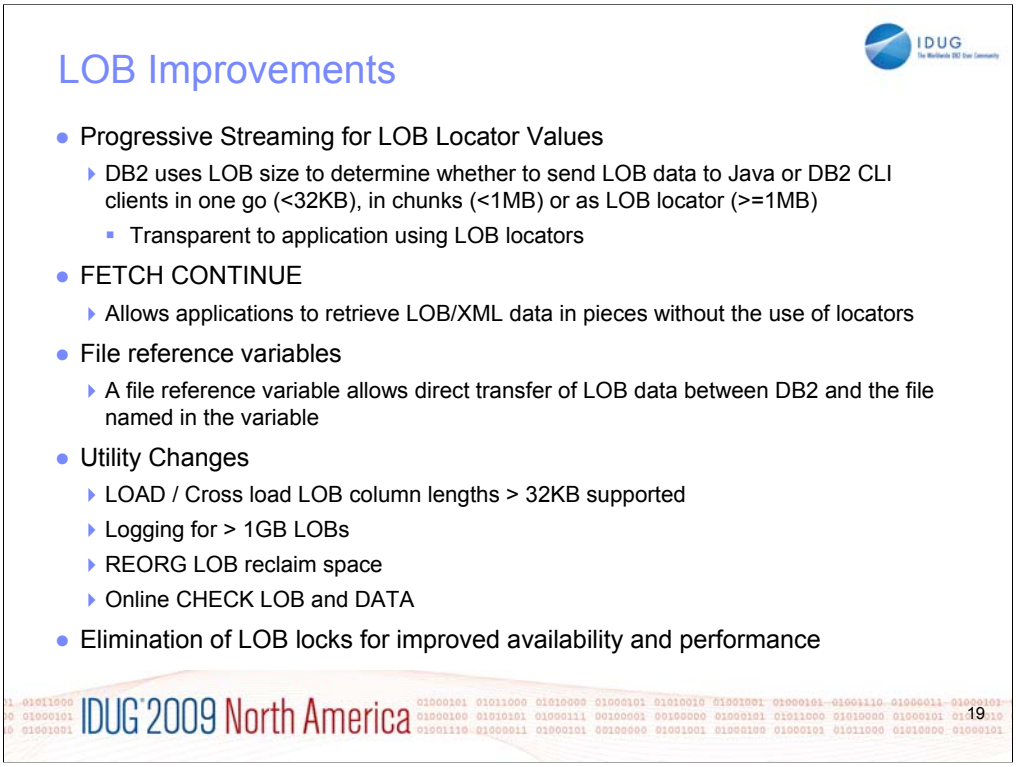

For Java and DB2 CLI programs that use locators with LOBs, Improves performance and less network traffic for LOBs that are less than 1MB
Default behavior if using DB2 9 for z/OS Requires DB2 Connect 9.1 FP 1 No changes required to programs using locator values
DB2 Client and Type-4 driver manage progressive streaming of data to program
DB2 for z/OS determines whether to flow LOB values or Locators to client based on size thresholds for JDBC, SQLJ, and CLI
For small LOBs, (Default <= 32KB) the performance should approximate that of retrieving a VARCHAR column of comparable size
Medium size LOBs (Defaults > 32KB and <= 1MB)
For large LOBs (Default over 1MB) locators are still used
Specific FETCH that contains LOB or XML columns Used with programs that materialize LOBs
Application uses a buffer that might not be large enough to hold the entire LOB or XML value. If any of the fetched LOB or XML columns do not fit, DB2 returns information about truncated columns and the actual length.
Retrieve LOB or XML data in multiple pieces without use of locators Must specify WITH CONTINUE on initial FETCH Subsequent fetches use FETCH CURRENT CONTINUE
Application must manage buffers & reassemble data Not required to fetch entire object before moving to next SQLCA indicates whether data is truncated
LOAD / Cross load LOB column lengths > 32KB supported
Logging for > 1GB LOBs
REORG LOB reclaim space SHRLEVEL(REFERENCE) Allows LOG NO SHRLVEL(NONE) is still an option.
Online CHECK LOB and DATA
Elimination of LOB locks Now using LRSN & page latching for consistency checks Prior to DB2 9, LOB locks were held until commit Even for UR
Space search for LOB allocation No LOB locks acquired for space search Read LSN granularity improved to page level in LOB table space
Improved availability & performance Particularly for UR readers
Requirements:: NFM “Locking protocol 3” GBP changes APAR Automatic in non-data sharing

Indexing Enhancements

- Larger index pages allow for more efficient use of storage
  - Fewer page splits for long keys
  - More key values per page
- Index compression provides page-level compression
  - Data is compressed to 4K pages on disk
  - 32K/16K/8K pages results in up to 8x/4x/2x disk savings
  - No compression dictionaries
    - Compression on the fly
    - No LOAD or REORG required
- Rebuild Index SHRLEVEL CHANGE
- Define RANDOM index keys to avoid hot spots with multiple processes inserting sequential keys

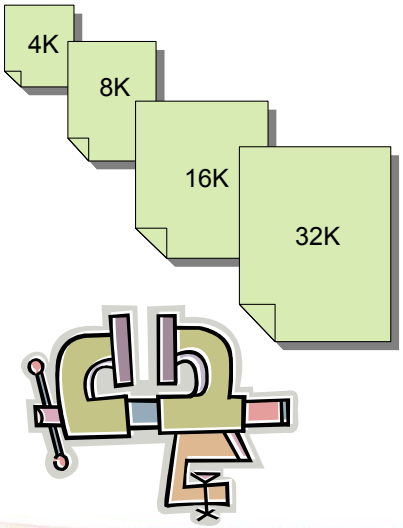

Indexing improvements contribute to the overall improvements in query performance. Specific improvements include index compression, index on expression, index key randomization, and larger index page sizes.

Larger index pages allow for more efficient use of storage, with fewer page splits for long keys and more key values per page.

Multiple processes inserting sequential keys can create hot spots on indexes. Randomized index keys avoid hot spots. Application insert throughput improved via avoidance of locking conflicts, but retrieval of sequential rows is likely to be slower.

Bigger index page: 4K, 8K, 16K, or 32K page → Up to 8 times less index split

Good for heavy inserts to reduce index splits. Especially recommended if high latch class 6 contention in data sharing. Two forced log writes per split in data sharing

Or high latch class 254 contention in non data sharing shown in IFCID 57

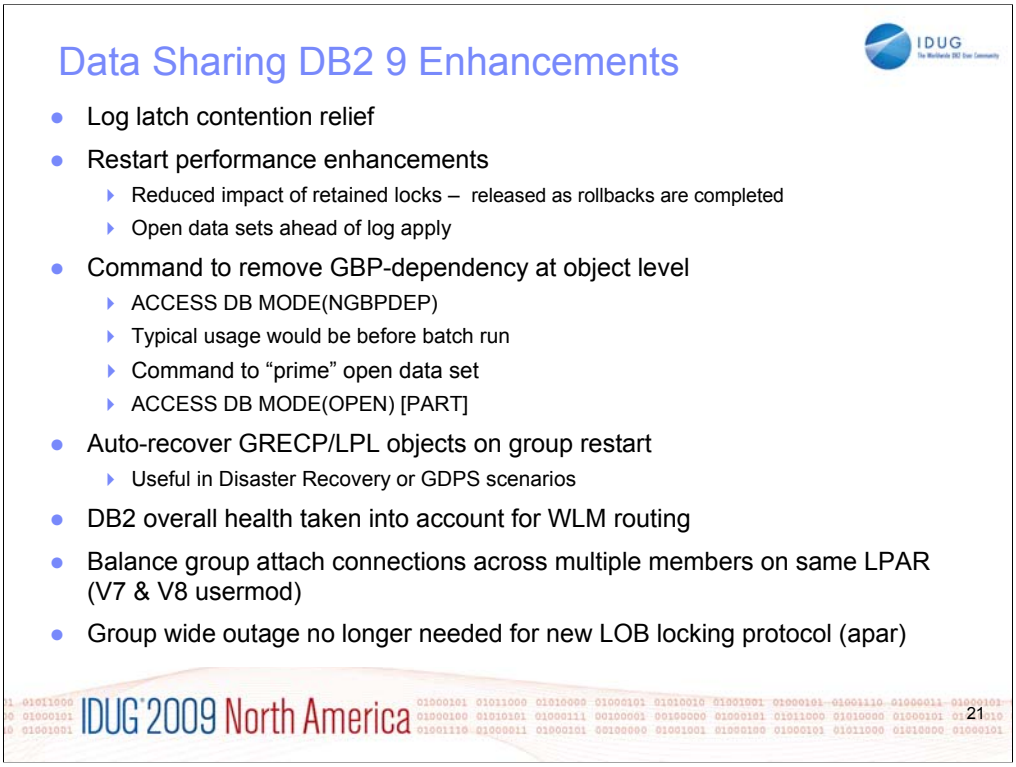

Many improvements help data sharing efficiency and usability in DB2 9. Logging was improved substantially for data sharing customers. Several enhancements help with faster restart, releasing retained locks faster and allowing data sets to be opened sooner.

-ACCESS DB…., watch for a PTF to allow ranges and wild cards in the –ACCESS command

Mode (OPEN) Forces the physical opening of the page set or partition on just the local member. This moves the overhead of the physical open from an SQL thread to the command thread. This improves the transaction rate for the first SQL thread to reference a given page set or partition.

Mode (NGBPDEP) Converts the specified page set or partition, non-disruptively to a non-group buffer pool dependent. You should use this before running large batch processes against a particular page set or partition to improve performance in a data sharing environment. Only issue this command to the member on which you plan to run the batch programs.

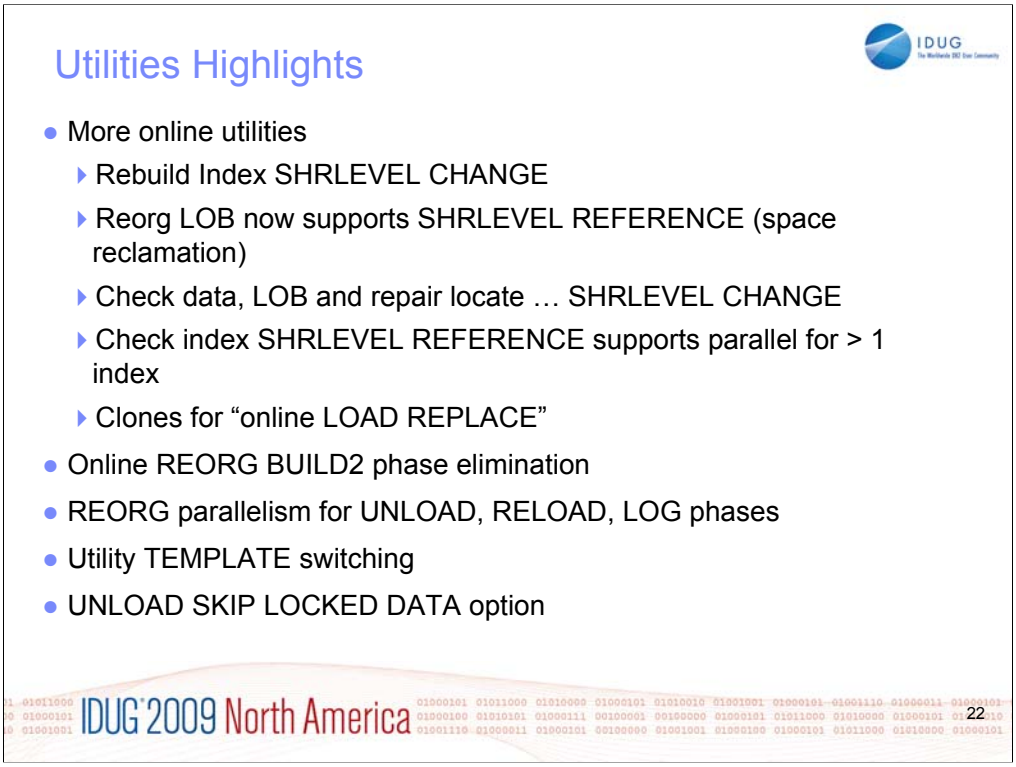

If you look at all the utility offerings, we have been adding more SHRLEVEL CHANGE and REFERENCE utilities to improve availability in every release. Now with DB2 9, the following utilities have SHRLEVEL CHANGE (RUNSTATS, COPY, REORG TABLESPACE, REORG INDEX, LOAD RESUME, REBUILD, UNLOAD, CHECK INDEX, CHECK DATA, CHECK LOB, and REPAIR). Cloned tables effectually function as LOAD REPLACE SHRLEVEL CHANGE. We have also improved availability for REORG TABLESPACE of a part/part range by removing the BUILD2 phase.

From a performance perspective, REORG of partitioned table spaces now has partition parallelism much like already existed in REBUILD INDEX, CHECK INDEX, COPY, RECOVER, and LOAD.

Utility template switching allows flexibility in which template applies to data sets based on attributes of the object processed.

UNLOAD SKIP LOCKED DATA gives another option besides just UR processing or CS processing. UR will not get locks, and therefore can unload “dirty” data. CS acquires locks and can therefore be slowed with lock contention. SKIP LOCKED is a CS option that will skip those rows/pages that are locked and give consistent data (but not all rows) without being slowed much by contention.

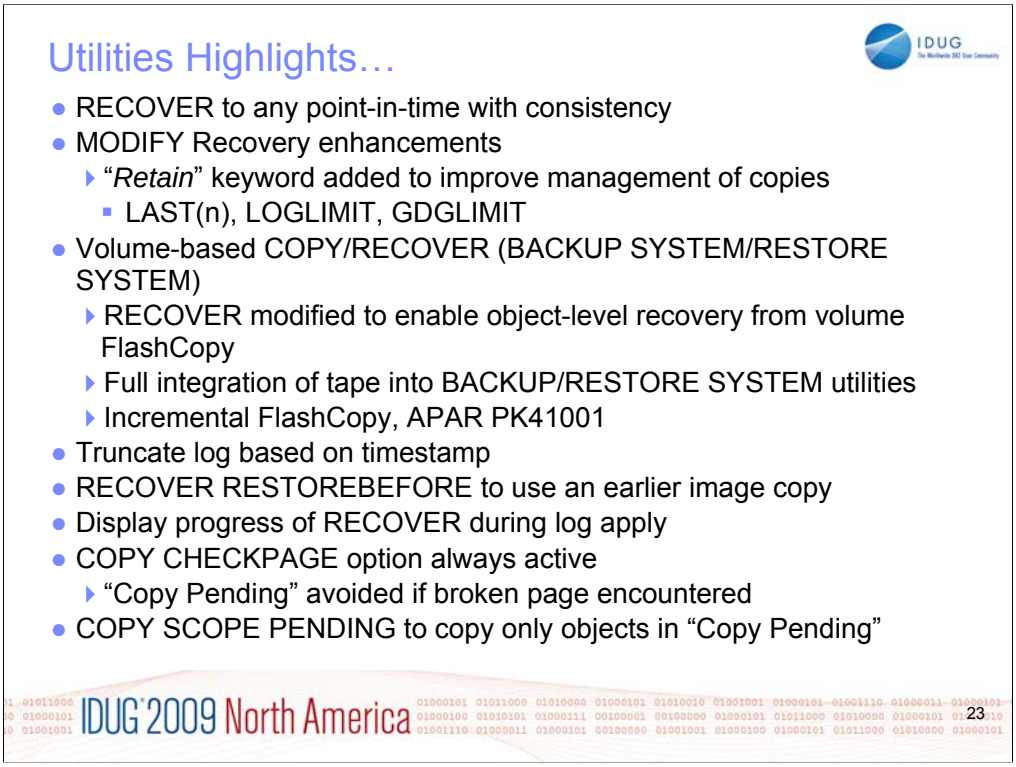

The COPY and RECOVER utilities each had several enhancements in DB2 9. One of the most important is RECOVER to point in time with consistency. We see that customers can't quiesce their workloads to provide a clean point for recovery, so RECOVER now recovers to a point on the log, then rolls back the in-flight work. The quiesce is not needed.

Prior to DB2 9 the user specified the deletion criteria either as before a specific date or by greater than a given age in days. DB2 9 has an alternative, by which instead of deletion criteria, retention criteria can be specified.

There are a number of improvements and extensions to the volume base utilities, or the BACKUP/RESTORE SYSTEM utilities.

If you need to do a conditional restart, DSNJU003 now allow you to specify a timestamp instead of an RBA/LRSN. The option is also available for the SYSPITRT for easier prep to truncate the log before running RESTORE SYSTEM.

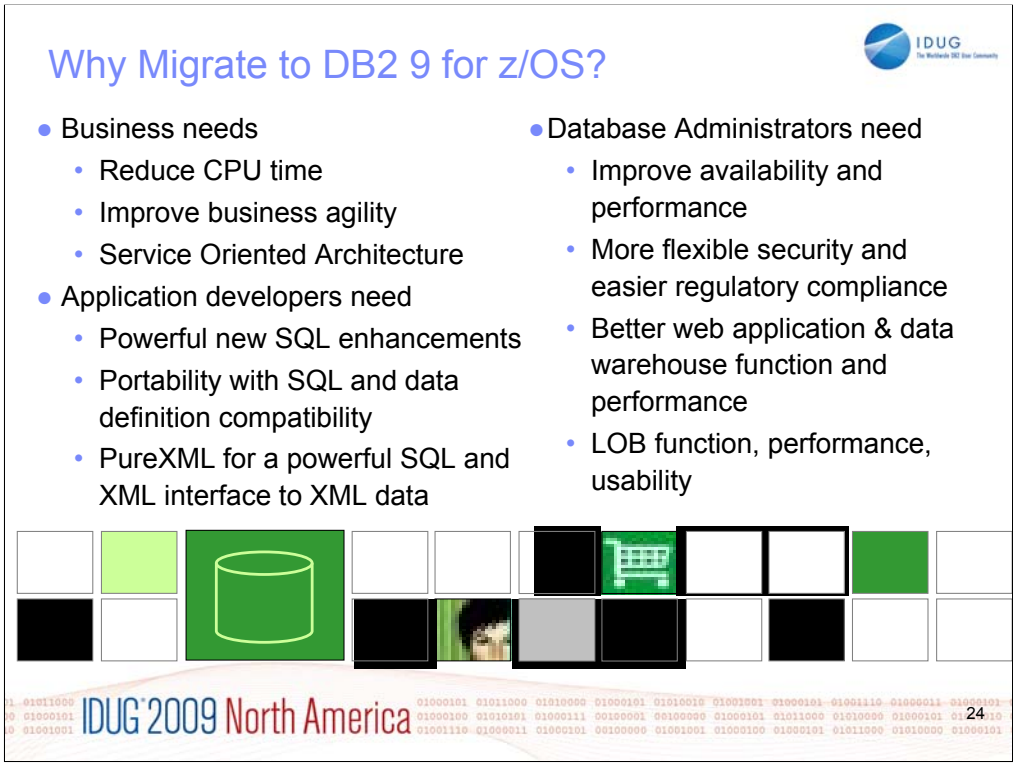

DB2 9 has a lot for everyone, unlocking the potential of V8. Here are just a few of the highlights. The business needs include CPU cycle reductions that deliver in most utilities, improved query optimization, improved business agility via faster implementation cycles, and new pureXML™ that builds a strong foundation for SOA and XML initiatives. Kevin Campbell, an Application Architect at Univar USA said it better than I can, “This is not a bolt-on or band-aid approach, DB2 9 for z/OS is XML without compromise.”

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

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

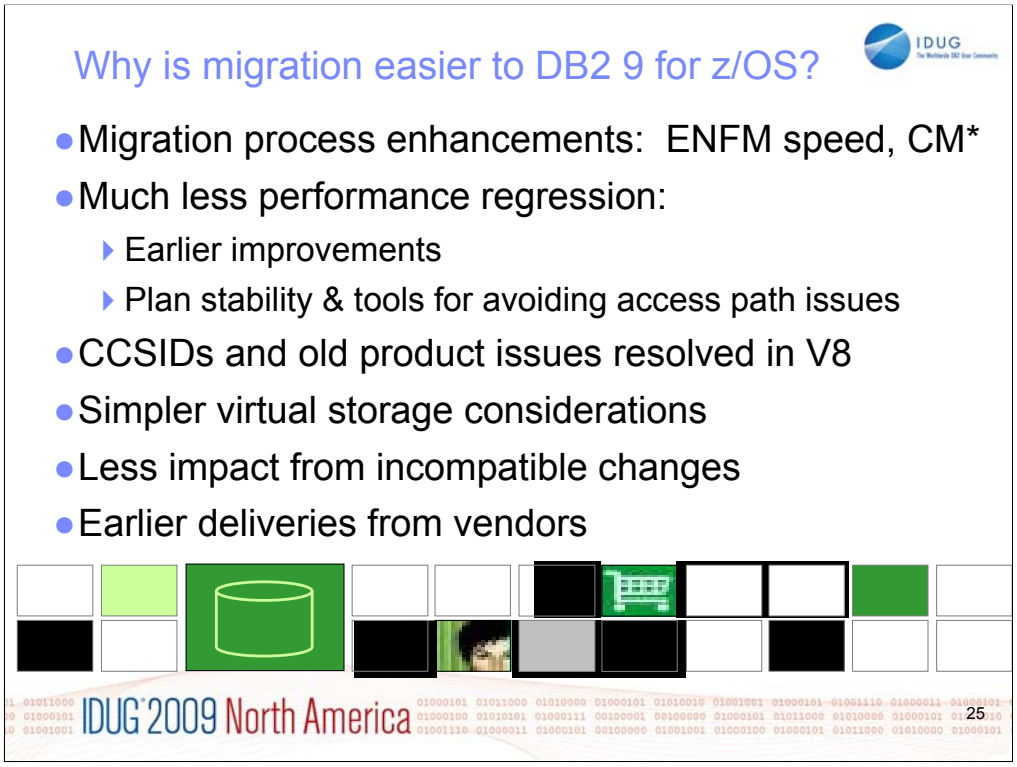

DB2 9 migration has been easier for many customers and tends to be faster than DB2 V8. Many customers note that the process and the quality are solid.

Migration process enhancements

ENFM is shorter (10 minutes versus an hour), as only two table spaces need to be reorganized in this phase (versus 18 in V8 including the largest ones).

Drop back to CM* if needed (no single steps without a drop back).

Much less performance regression

More consumable performance improvements

Utility CPU improvements in CM

Bind stability to reduce concern about access path regression.

CCSIDs and old product issues resolved in V8 migration, not a problem.

Simpler virtual storage considerations – incremental improvements.

Less impact from incompatible changes

Many vendors are ready today.

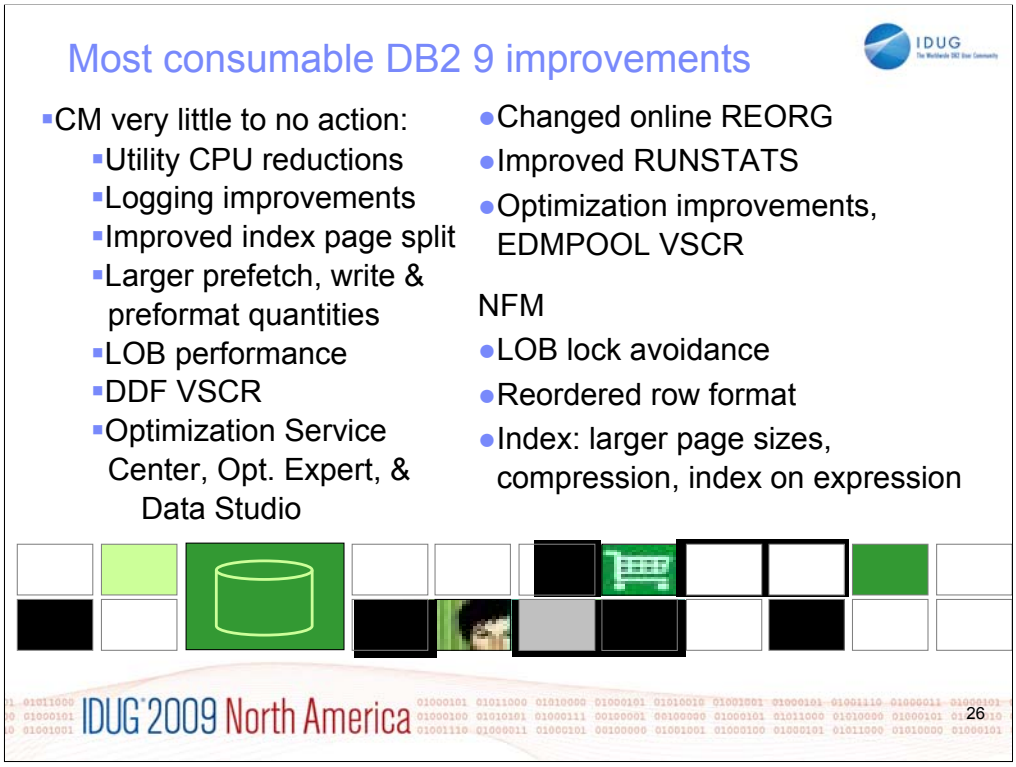

Here are some highlights for items that deliver the most quickly and easily:

Very little to no action is required for the utility CPU reductions, logging improvements, improved index page split, larger prefetch, write & preformat quantities, some LOB performance, DDF virtual storage constraint relief. The first group delivers in CM.

The next items require some work. Changed online REORG and other utility improvements require process changes and use of SHRLEVEL(CHANGE).

Improved RUNSTATS statistics needs some analysis to determine where the value is greater than the cost of gathering the new statistics.

Optimization improvements are automatic for dynamic SQL, but require work to REBIND for static SQL. In both cases, we need baselines to check for regression. REOPT(AUTO) for dynamic SQL needs analysis to be sure the improvement is working. EDMPOOL virtual storage constraint relief also requires a REBIND.

Optimization Service Center takes some learning, but should be fast for those who have used Visual Explain in the past. See the new redbook, SG24-7421, DB2 9 for z/OS: New Tools for Query Optimization.

LOB lock avoidance requires a quiesce of all subsystems in NFM until APAR PK62027. Reordered row format requires a REORG in NFM and varying length columns, and use with small columns can reduce compression. See APARs.

Index improvements for larger page sizes, compression, index on expression require database design work to determine where they are applicable. ALTERs, REORGs and creation of new indexes are needed.

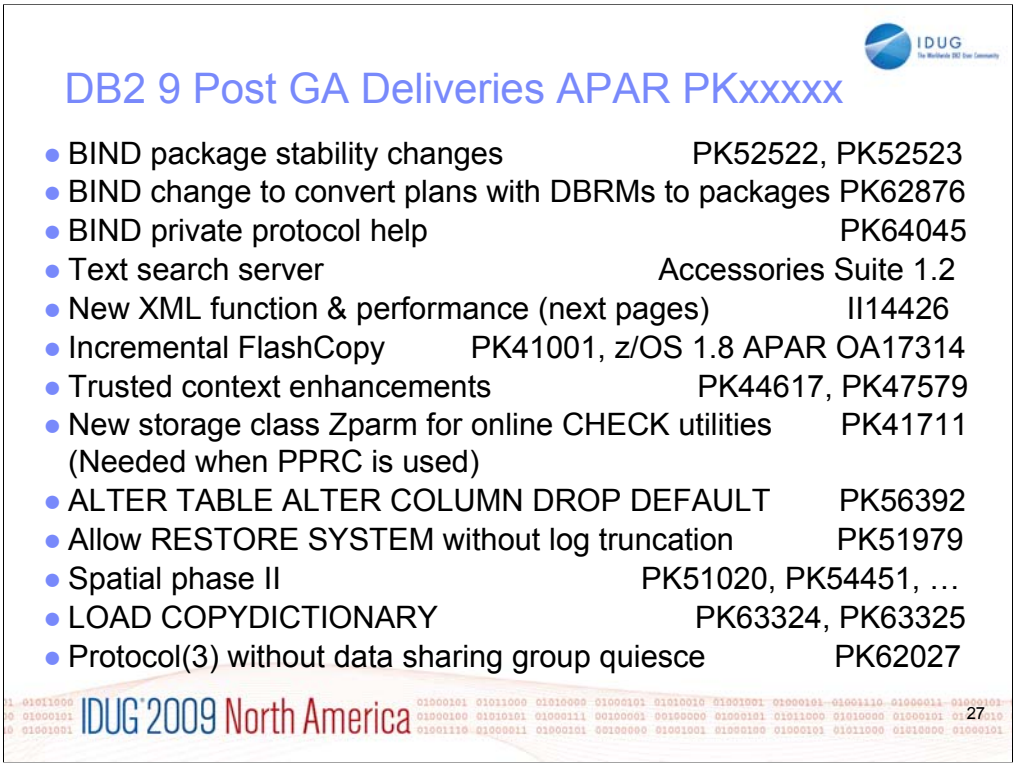

BIND stability changes deliver the capability to keep two or three versions of packages and the ability to switch from the current access paths to the older ones. This function can help you with regression from a REBIND.
Plans containing DBRMs, rather than using packages, are deprecated. That is to say we expect plans containing DBRMs to disappear some time in the future. To help with migration from DBRMs to packages, a change to BIND is shipped which can perform the conversion to packages.
A text search server delivered in Accessories Suite 1.2 at the end of 2007.
XML performance improvements and new XMLTABLE and XMLCAST functions are very important for early XML applications and consistency. Incremental FlashCopy makes use of the new z/OS capability. Trusted context enhancements are needed by early users. A new storage class parameter is needed when PPRC is used and is provided for online CHECK utilities.
ALTER TABLE ALTER COLUMN DROP DEFAULT extends the capability to ALTER a default. DB2 allows RESTORE SYSTEM recovery without requiring log truncation.
Many new spatial functions deliver compatibility with the DB2 for LUW function.
Changes come in DB2 to work with new fixpacks of OSC, OE and Data Studio.
The capability to copy a dictionary from one partition to another using the LOAD utility as added. APAR PK62027 avoids the need to quiesce the data sharing group to get the new LOB lock avoidance.

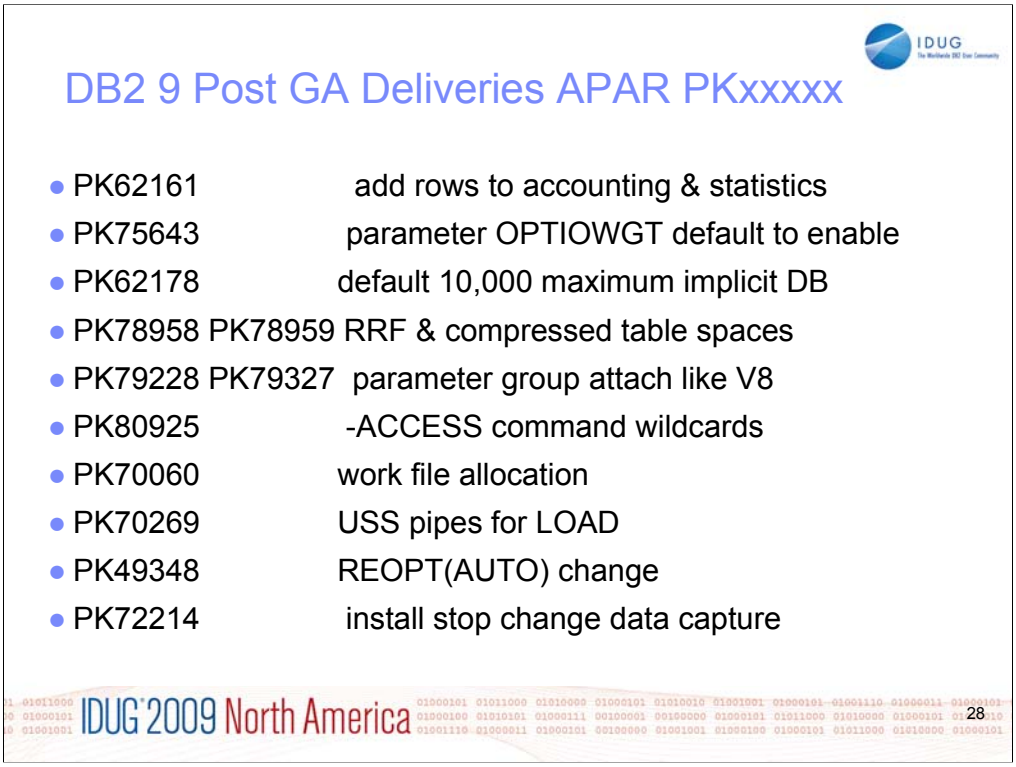

DB2 9 has been generally available for about two years now, and some new function has delivered in the service stream. Delivery in APARs is used when the need is urgent and the function can be delivered with a very small risk of disruption to existing customers. The delivery criteria become stronger as more customers move to the version.

Rather than creating 60,000 databases implicitly, changes in PK62178 allow default to 10,000 and allow you to set the number.

If you want to pipe information to the LOAD utility, look at PK70269.

PK72214 changes install to stop changed data capture for catalog tables as needed, asking you to restart them after switching to DB2 9 and ENFM.

Changes are in process for the combination of compressed table spaces and RRF. Some customers with many short varchars, e.g. varchar(1) found that compression was not effective with RRF, so the default is changed to not convert to RRF.

-ACCESS only worked for specific names, but it's being changed to allow wild cards and ranges.

OPTIOWGT was generally recommended, and becomes the default with PK75643.

If you have multiple members of a data sharing group in a single LPAR and prefer the V8 group attach logic, parameters are added to control the random group logic.

DB2 was designed to handle transactions and queries. In Version 1, the CPU times for transactions were roughly double those of IMS for similar transactions, so DB2 was positioned for the information center. Each release of DB2 has included substantial improvements for some types of queries: optimization improvements, SQL language, indexes, and avilit to manage the work. So DB2 has increased its ability to handle large, complex queries on large volumes of data over the years.

Customers have been using DB2 for queries over the time, and many do have warehouses on DB2 for z/OS. In most cases, the workload is shared with other work, rather than dedicated or standalone.

Current trends in business intelligence match the strengths of System z and DB2 for z/OS. Cost improvements with specialty engines and improvements in the latest versions make DB2 for z/OS even better for warehousing.

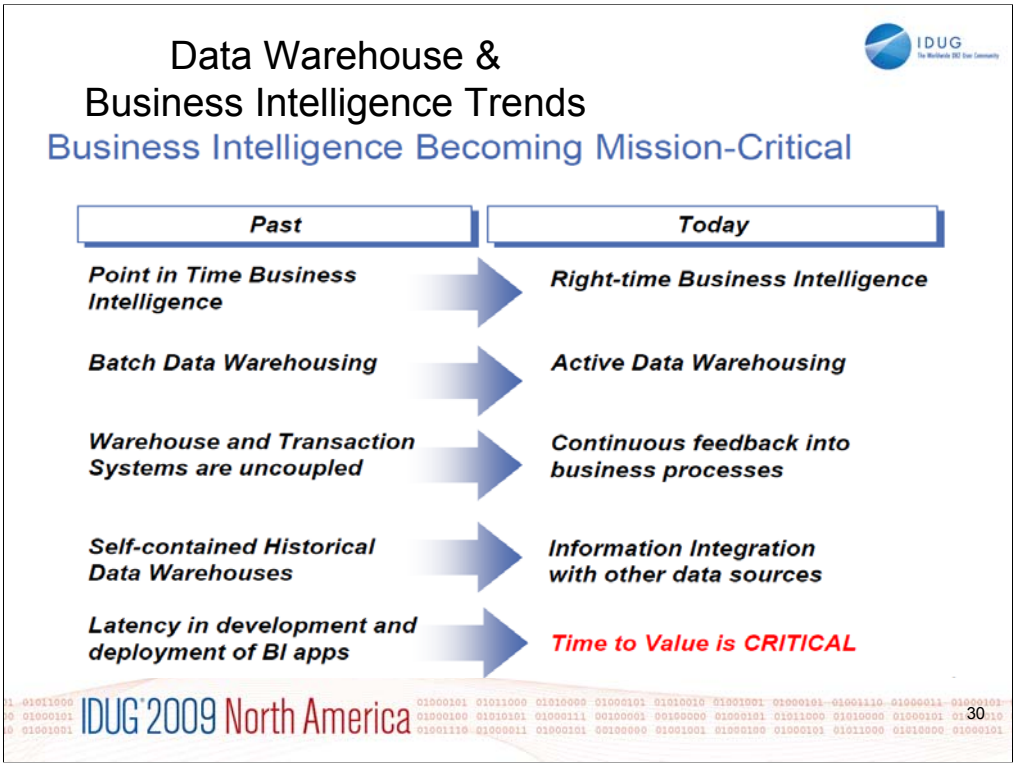

Many of the strongest trends in data warehousing and business intelligence address the key strengths of System z, with high quality of service, availability, resilience, scalability, security and the ability to share data and workloads.

Rather than being used as separate systems, business intelligence delivers more value by being integrated into the business processing, so that the business intelligence runs with more current information and the rest of the business can use the capabilities and data.

Integration and use of a wider range of data sources fits with the System z. Rather than needing to extract and duplicate the information, moving to another platform, processing power and disk can be shared. If data needs to be copied, the process is much simpler.

IBM announced a beta program for InfoSphere Warehouse on System z which will provide a highly scalable, lower cost way to design, populate and optimize a DB2 for z/OS data warehouse to support BI applications such as Cognos 8 BI. This new offering will simplify operational complexity with a single database for both operational and warehouse data - - reducing costs related to data duplication while providing more efficient access to DB2 data. Significantly improve query performance for users who want to drill down into specific data stored in DB2 for z/OS. (For example, a sales manager looking to build sales plans can look not only at revenue by year, but also data by quarter, month, and customer to make fast business decisions.) Give customers the ability to support near real-time decisions based on core business data managed in DB2 for z/OS, helping customers gain additional competitive advantage and value from their operational data.
InfoSphere Warehouse on System z will further strengthen the IBM Data Warehousing and Business Intelligence software solution for System z that today includes: DB2 for z/OS
•InfoSphere Information Server for System z
•InfoSphere MDM Server for System z
•Cognos 8 BI for System z
http://www.ibm.com/software/data/info/new-systemz-software/

DB2 Developer Workbench → Data Studio

before

IBM DB2 Developer Workbench V9.1

- SQL Query Editor
- SQLJ Editor
- SQL Builder
- XQuery Builder
- SQL Routine Debugger
- Java Routine Debugger
- XML Editor
- XML Schema Editor
- Data Management
- Visual Explain
- Project Management

Data Studio is a full replacement of DB2 Developer Workbench plus much more

- ***DB2 for Linux, Unix, Windows v8.x, v9.1.x, v9.5***
- ***DB2 for z/OS v7, v8, 9***
- ***DB2 for i v5r2, v5r3, v5r4, v6***
- ***Informix Dynamic Server (IDS) v9.x, v10.x, v11***

now

IBM Data Studio

- Integrated Query Editor – SQL + XQuery
- SQLJ Editor
- SQL Builder
- XQuery Builder
- SQL Routine Debugger
- Java Routine Debugger
- XML Editor
- XML Schema Editor
- Data Management
- Visual Explain
- Project Management

- ER Diagramming
- Data Distribution Viewer
- Object Management
- Browse & Update Statistics
- Security Access Control
- Connection Management integration with Kerberos and LDAP
- Data Web Services
- IDS Server Support
- Health Monitoring DB2 for LUW 9.5 and DB2 z/OS v9

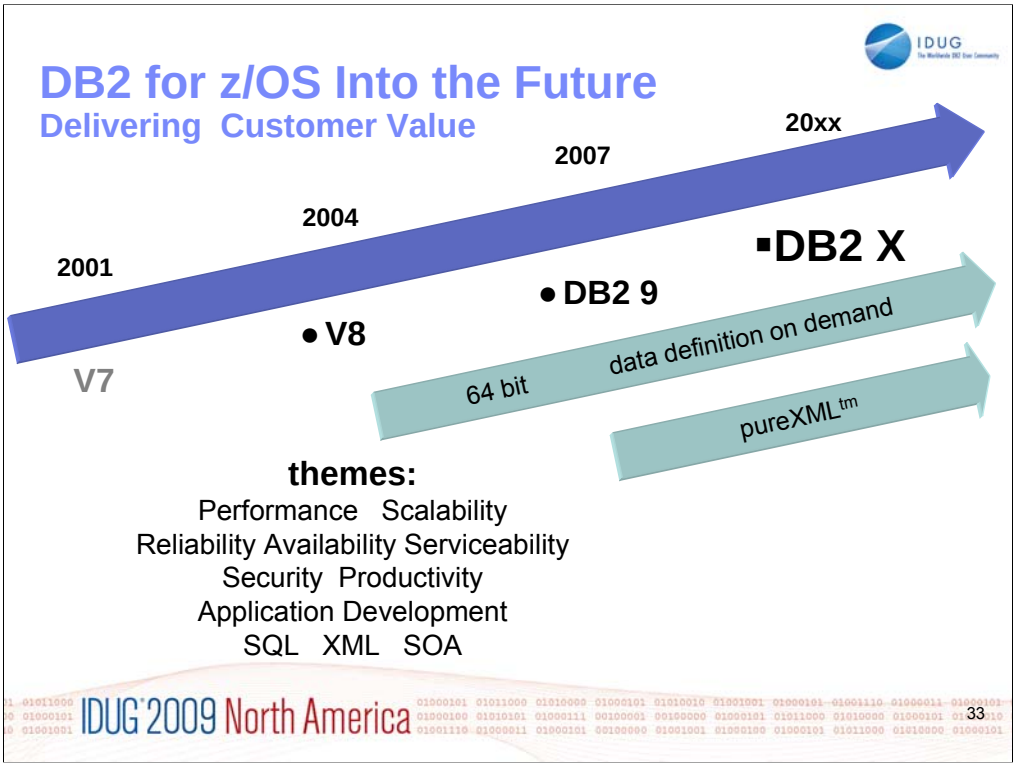

DB2 for z/OS V7 became generally available (GA) March 2001, and V8 delivered three years later. DB2 9 became generally available in March 2007, three more years. We expect the next version will be roughly 3 years from DB2 9 GA to DB2 10 or DB2 X or whatever the name becomes.

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

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

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

DB2 X: Preliminary Performance Plan → Significant CPU Reductions: Best on z10

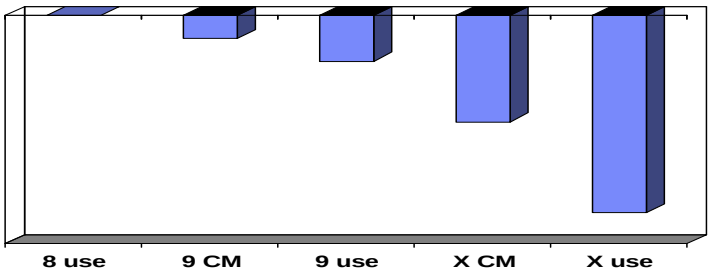

Transactions
Batch
REBIND

DB design changes
Hash access
application changes
SQL adjustments

Your situation will vary. Less CPU is better.
Processors z890, z990, z9 and later z/OS 1.10 and later

Reducing CPU from DB2 9 to DB2 X 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 X 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 X 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 X.

We expect DB2 X 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.

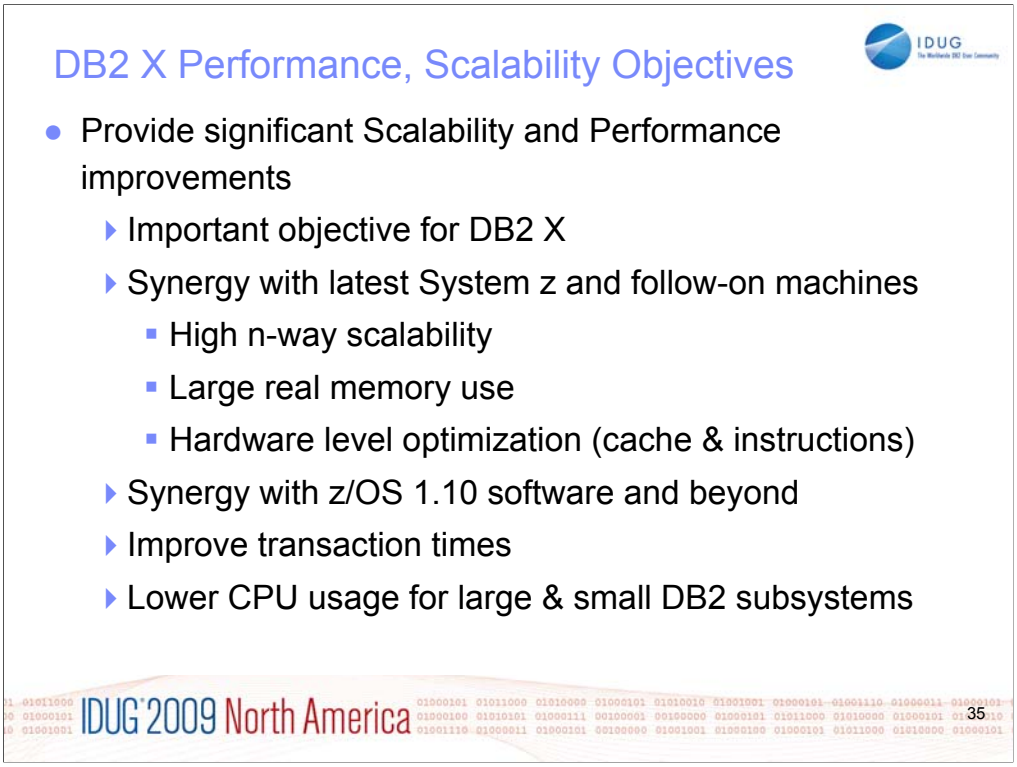

Providing significant scalability and performance improvements is an important DB2 X objective. Synergy with the latest System z10 processors and follow-on machines provides part of the improvements. Being able to deliver high scalability for increasing numbers of processors is important for growth and 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.

See the next pages for detailed changes.

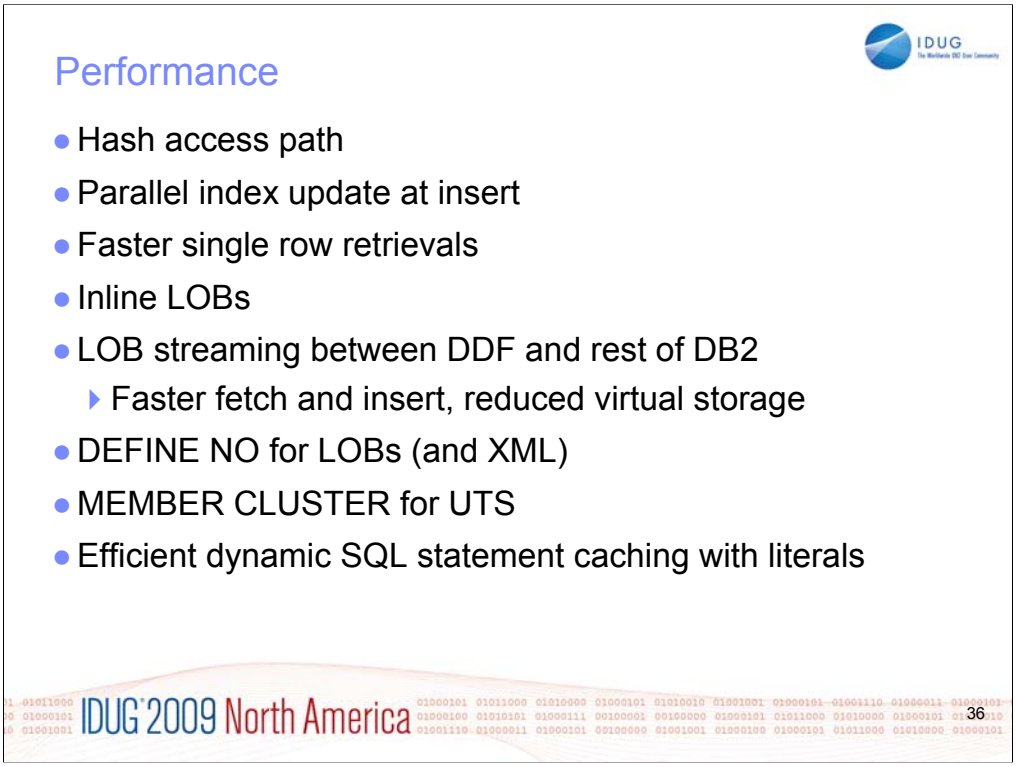

Hashing is an faster alternative to index access for a key = value lookup, not needing to touch the index pages.

Indexes can be updated in parallel to speed the process when many indexes must be maintained.

A single row select is another common process that is speeded up in this release.

Inline LOBs resolve the common performance problem of having large objects where most LOBs are small, by keeping part of the LOB on the page with the rest of the data. Streaming for LOBs improves distributed processing. Define NO lets DB2 avoid data set definitions.

Member clustering is now allowed for data sharing situations with frequent updating for universal table spaces.

Even if the dynamic SQL has literal values, rather than host variables or parameter markers, DB2 can use the dynamic statement cache effectively.

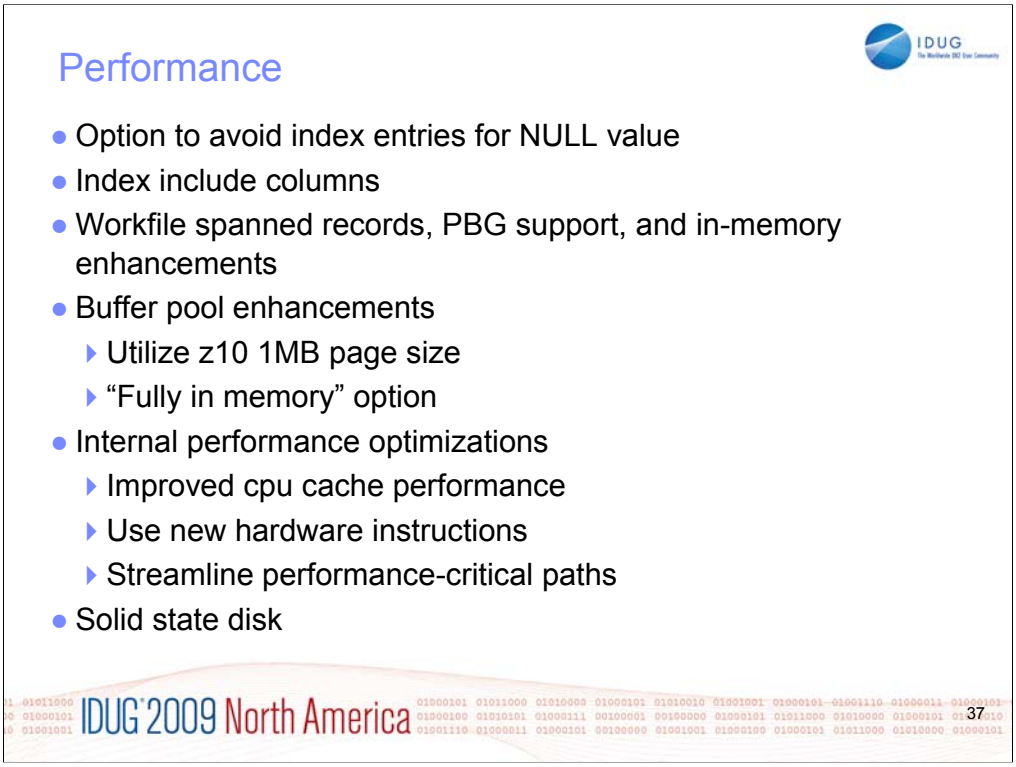

Being able to have a column that is unique, but includes extra columns can improve performance substantially in some situations, and makes the DB2 family more consistent.

If you want to have a table be completely in memory, then the new option can make the process easier. System z10 and z/OS 1.10 allow a 1 megabyte page size, instead of a 4 kilobyte page size, improving efficiency for large amounts of memory.

Many of the performance do not require changes in applications or administration. The CPU cache performance and new hardware instructions can reduce CPU time without customer action, other than moving to this version.

Solid state disk is working today with DB2, and integration can improve the performance and value. Solid state disk is expected to improve substantially in size and value in the next few years.

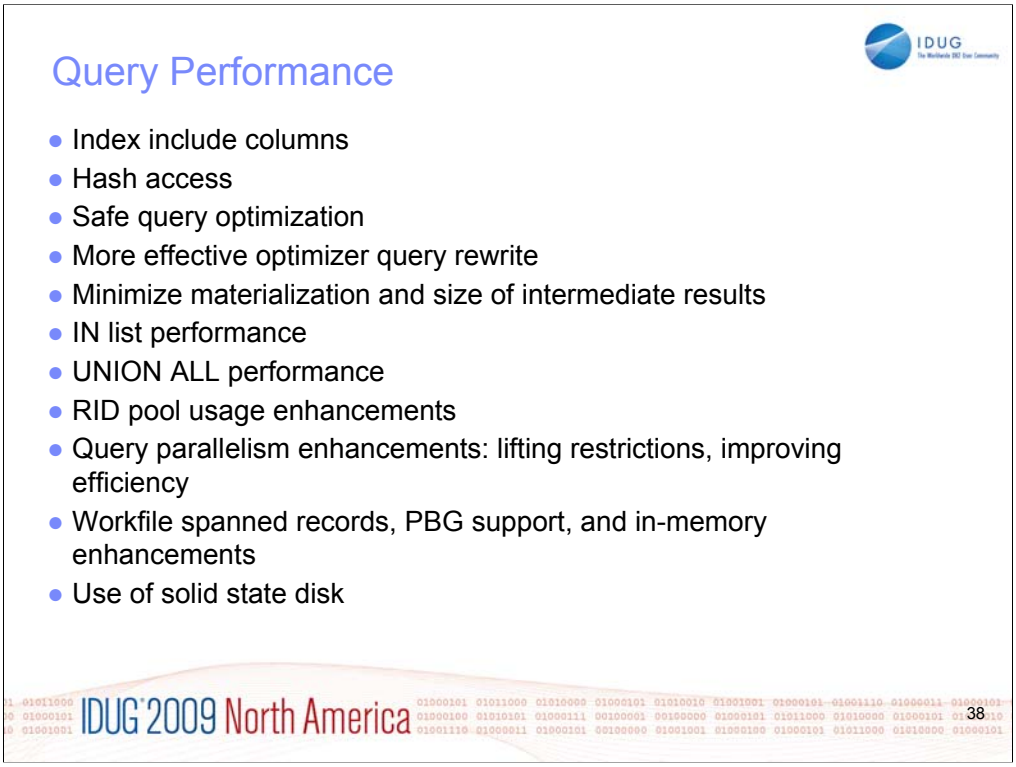

A wide range of optimization improvements will help with DB2 X query performance. New techniques for access data more quickly are the first step. The ability to have a unique index with additional columns delivers the ability to use indexes more effectively. Where the full key is equal to a value, a hash can reduce CPU and memory access substantially.

Optimization techniques include better ability to avoid bad access paths where an average cost is a little better, but some access will be much worse. Improved optimizations can reduce the query processing by a substantial amount.

Increasing parallel processing by lifting restrictions and improving efficiency will return results faster and enable more use of zIIP.

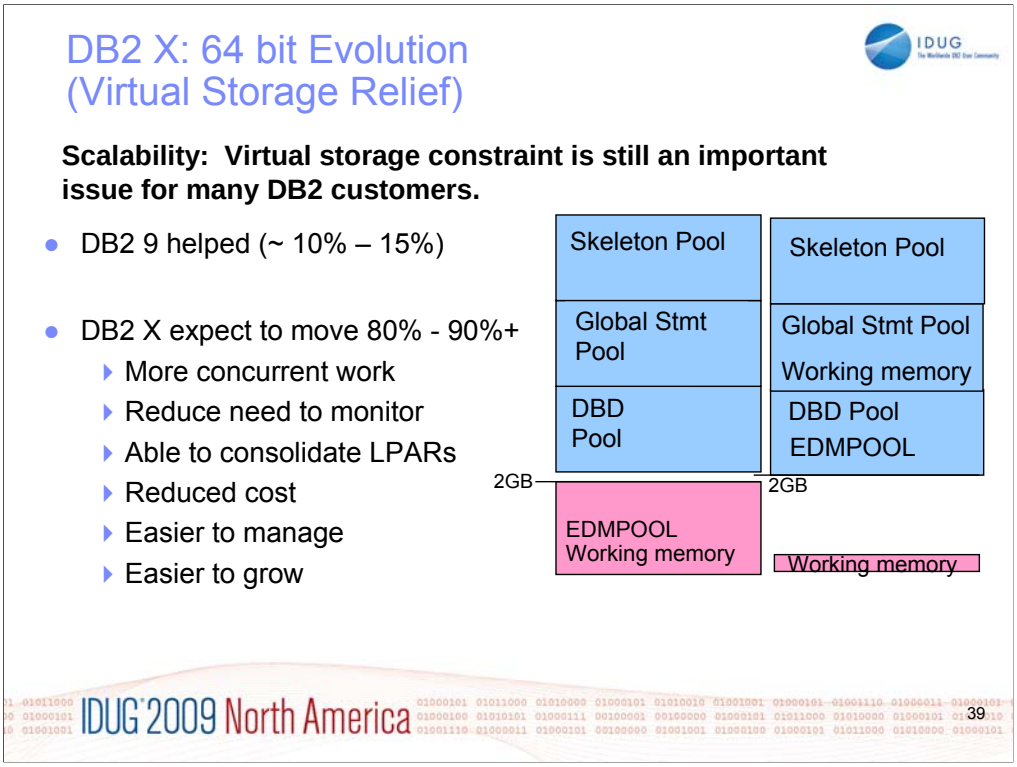

Virtual storage is most common constraint for large customers. Virtual storage can limit the number of concurrent threads for a single member or subsystem.

The DB2 9 virtual storage objective was 10-15% relief. The DB2 X 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 X will address a number of the next items, such as utility locking, catalog concurrency.

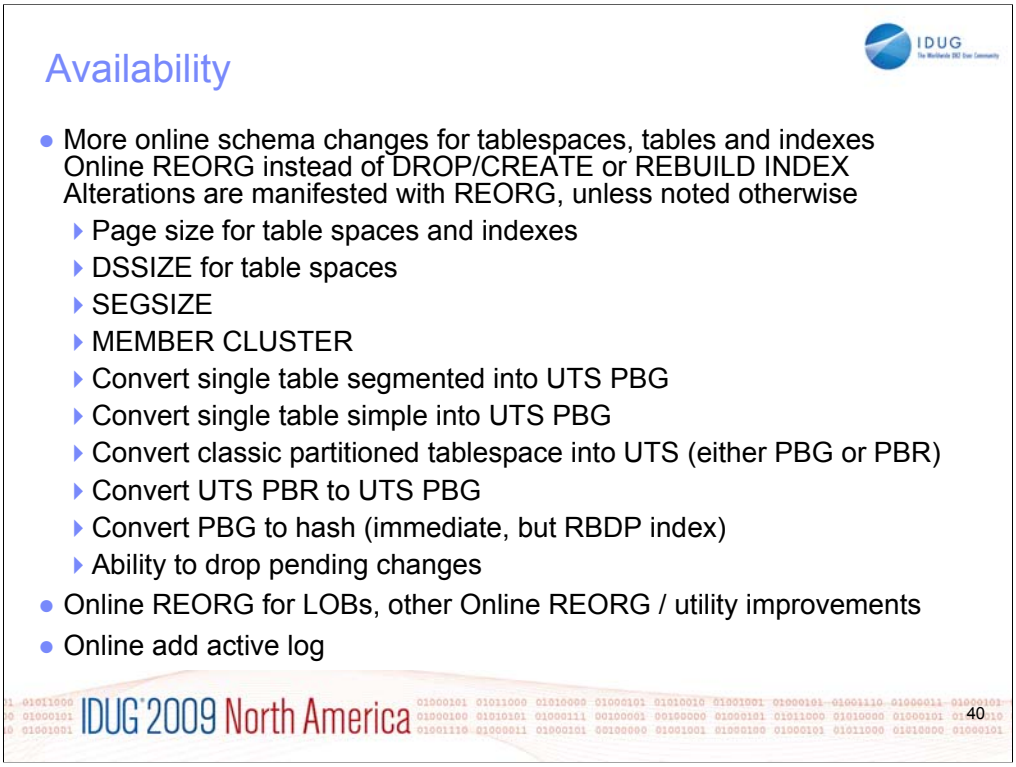

Continuous availability requirements continue to escalate. Large batch and maintenance windows are in the past. Those windows are being closed on the fingers of DBAs. DBAs increasingly need the ability to make all changes and to do all maintenance activities online or around the clock.

DB2 X 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.

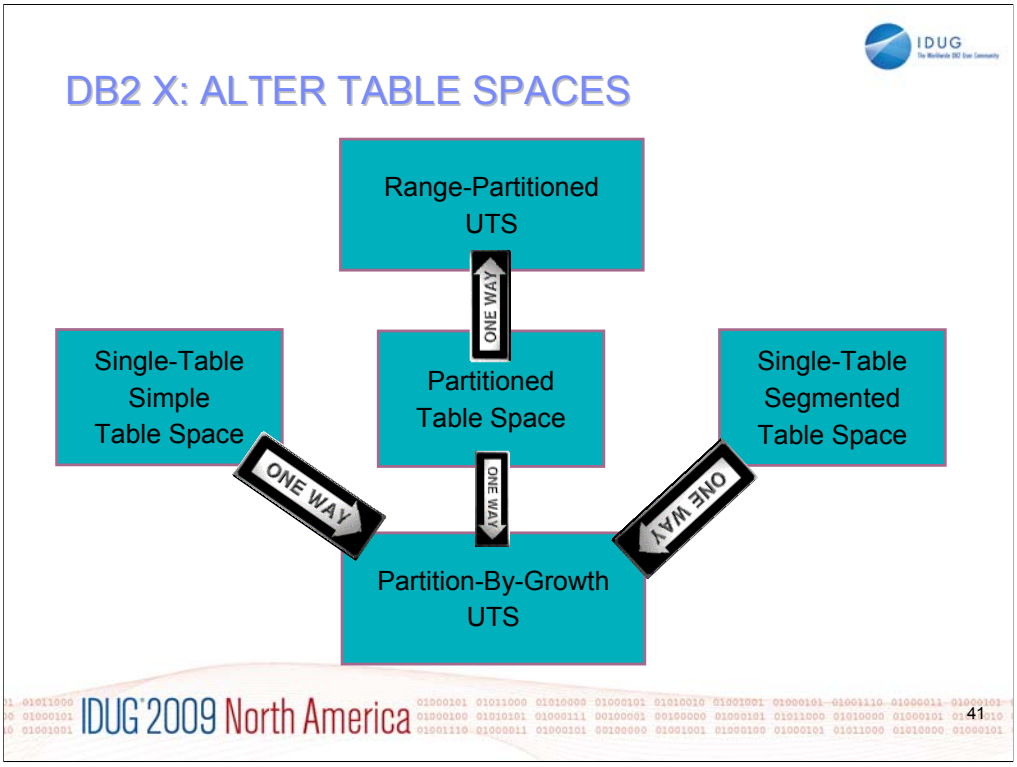

These are the changes in table space type in diagram form.

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.

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

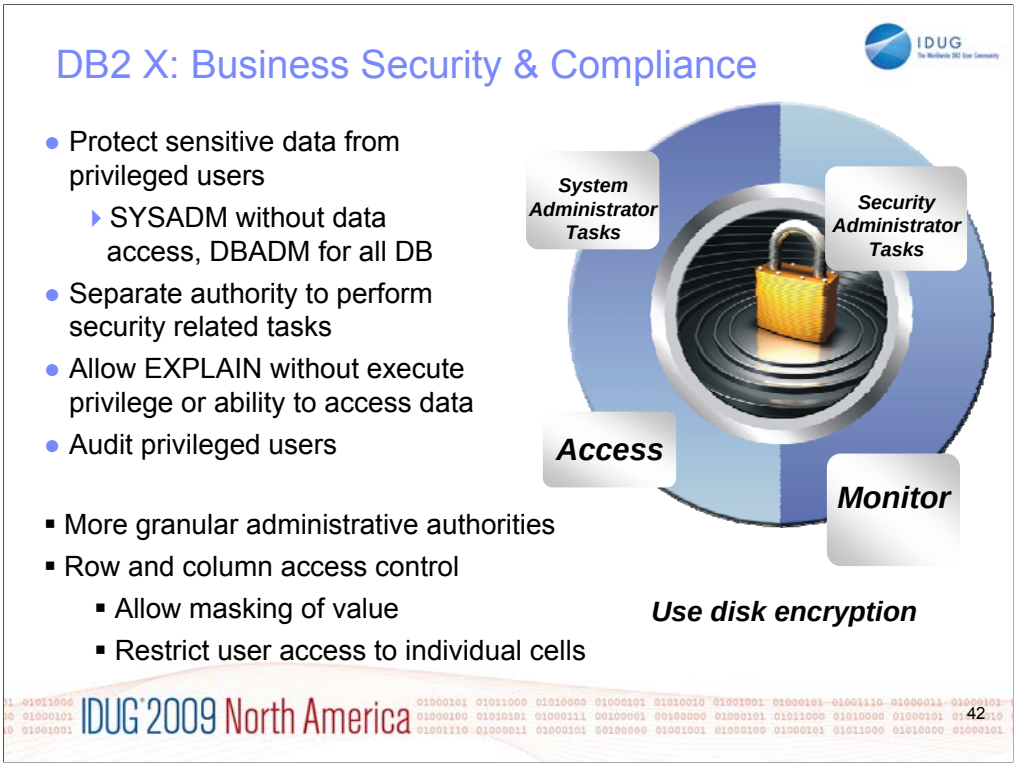

Customers are being pressed for a wide range of improved security and compliance. Data retention is a growing need. Protecting sensitive data from the privileged users and administrators is required. Separation of authority for security, access, and some common tasks, like EXPLAIN will help. Auditing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

In DB2 X, we expect to have a form or temporal data or the ability for a table to contain both current and historical data, and to query the information as of a specific point in time.

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.

DB2 X: Productivity – Doing More with Less!

- Auto statistics collection
- Compress 'on the fly'
  - Avoid need to run utility
- Reduce contention, more online processing
- Automatic config of IBM supplied UDFs and SPs
- Enhancements for Data Studio

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

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.

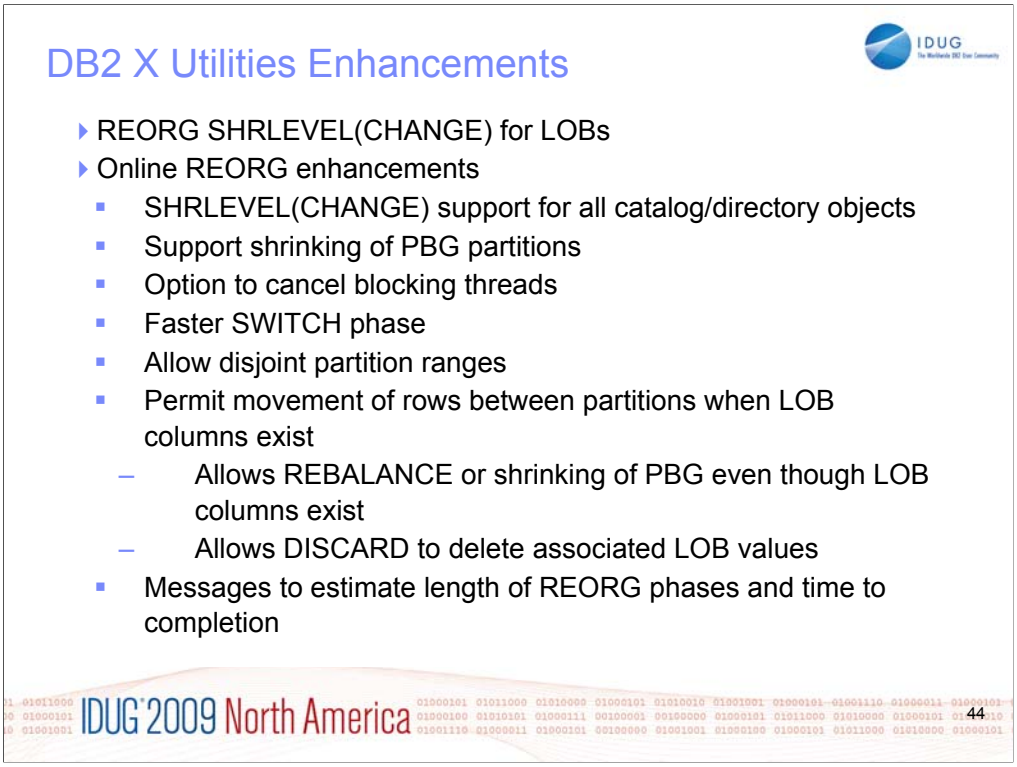

Continuous availability requirements continue to escalate. Large batch and maintenance windows are in the past. Those windows are being closed on the fingers of DBAs. DBAs increasingly need the ability to make all changes and to do all maintenance activities online or around the clock.

DB2 X 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.

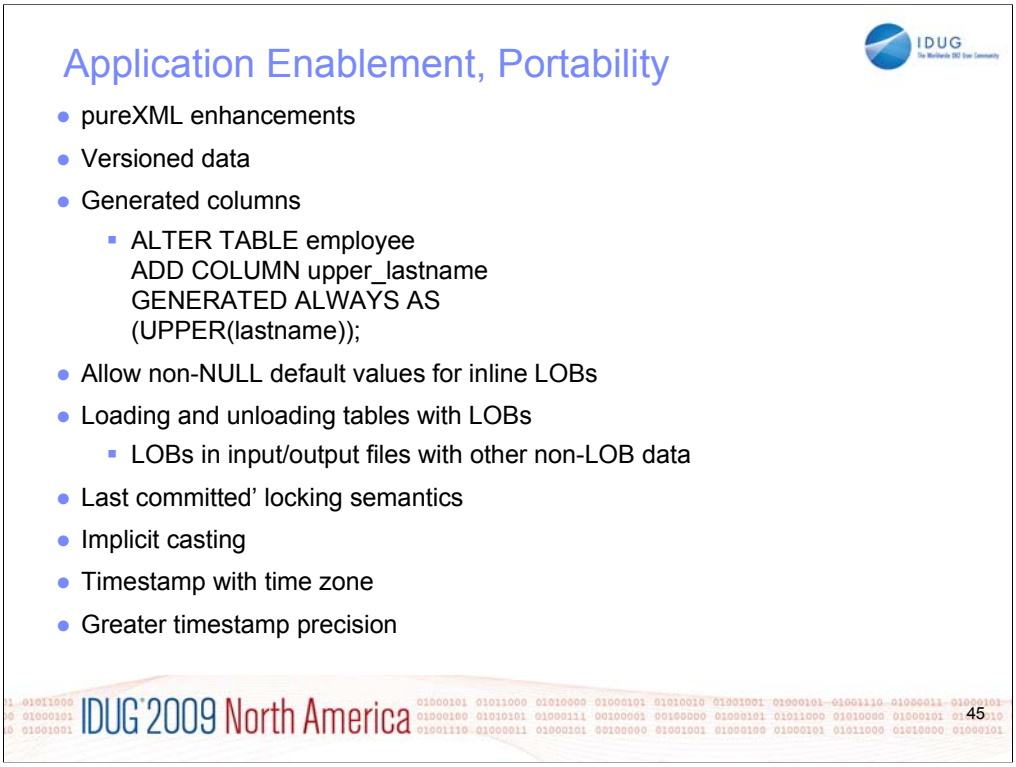

The pureXML improvements are noted two slides later.

Generated columns allow a column to be controlled by the specification, rather than needing all update applications to do the work.

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).

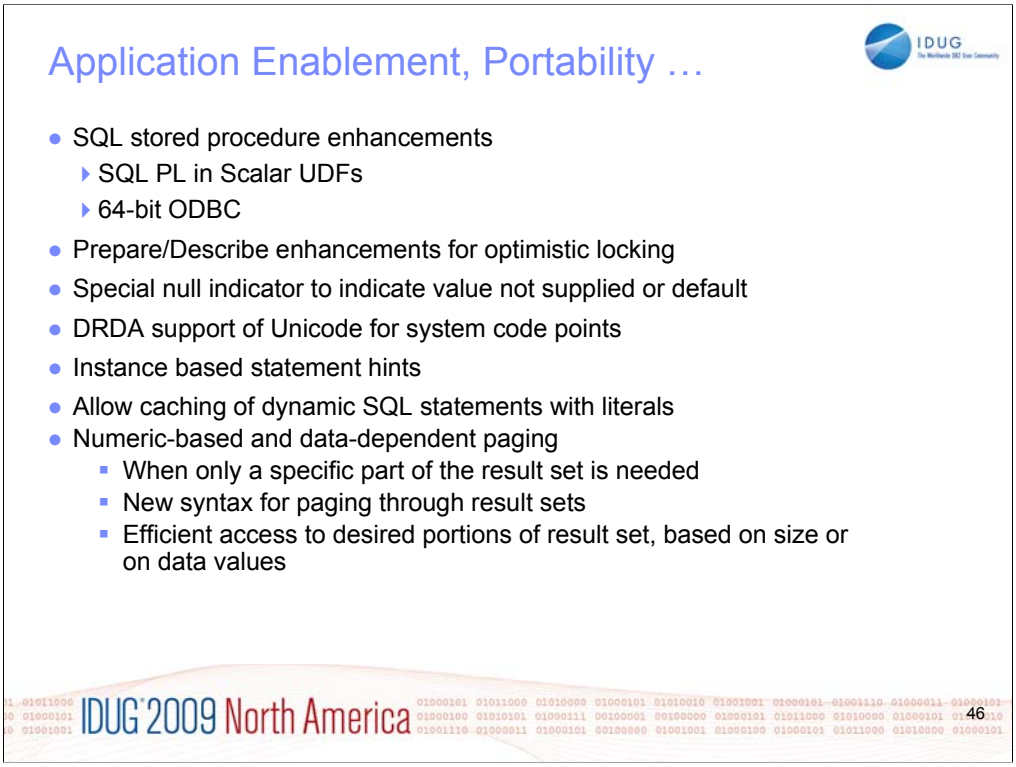

SQL procedure language is now allowed in scalar user-defined functions. Applications can use data above the bar in the new ODBC structures.

Optimistic locking can be managed better in dynamic SQL. 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 either to skip unused part of result set or to retrieve the interesting part of result set only. Numeric-based pagination allows applications to access part of a result set based on an absolute position. Data-dependent pagination lets applications access part of a result set based on a logical key value.

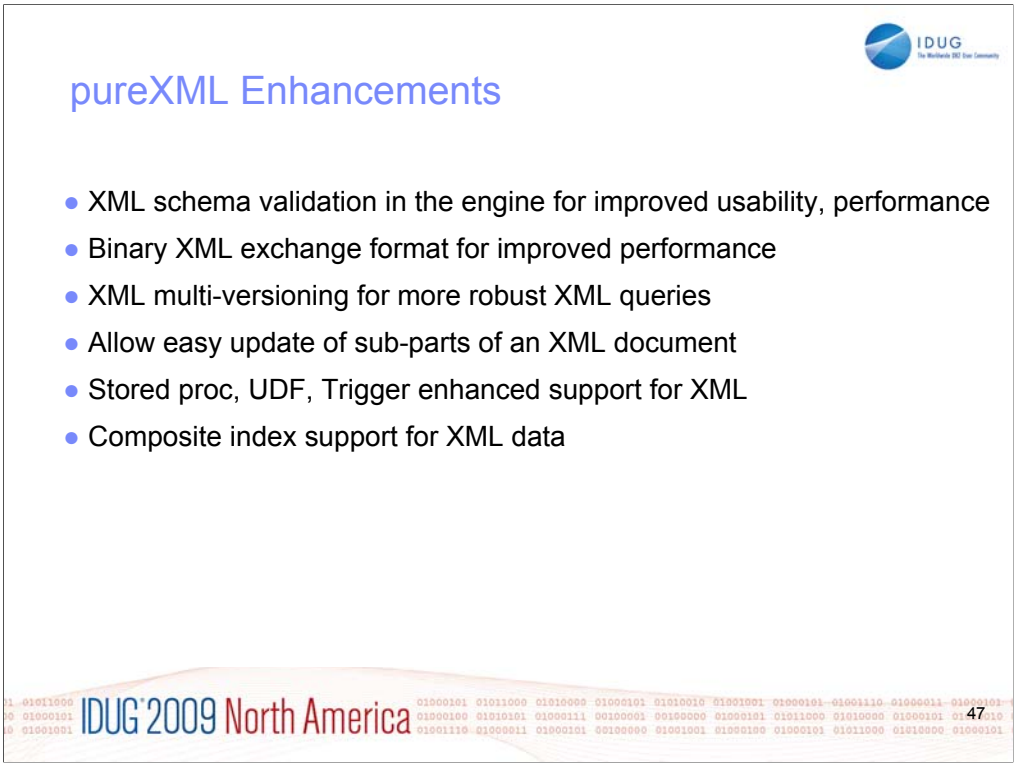

A range of XML improvements delivers a strong release 2 of the pureXML function. Customers use of DB2 9 pureXML shaped this delivery of improved performance and usability.

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.

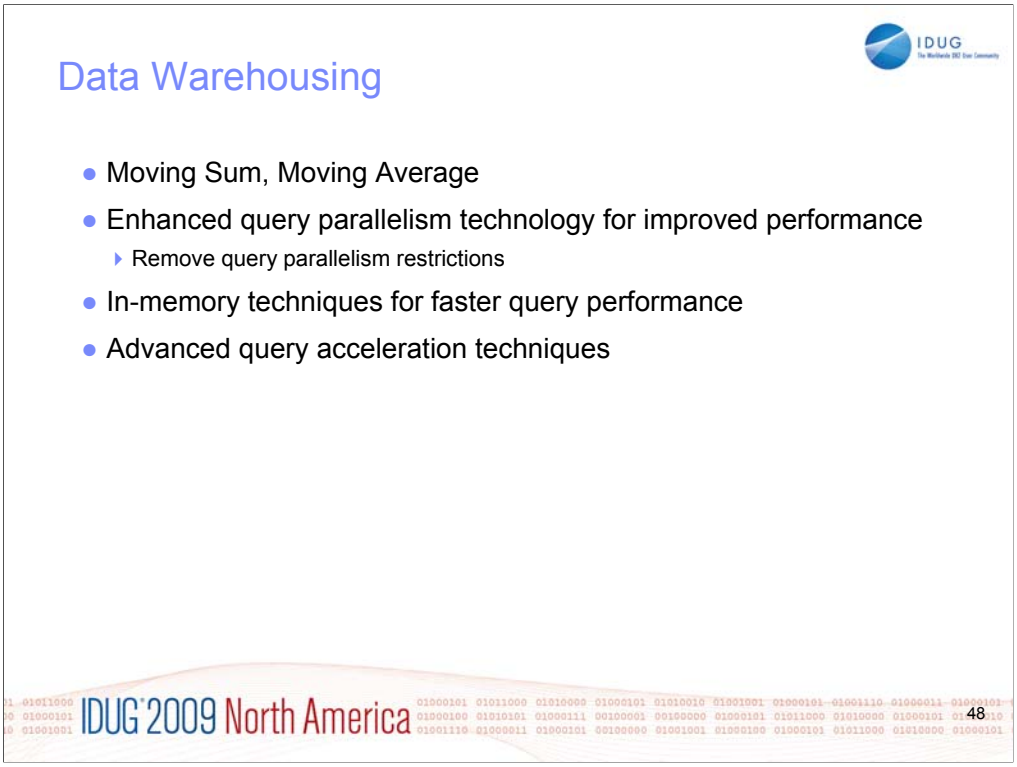

DB2 warehousing capabilities continue to be enhanced at a rapid pace. Improvements in the SQL and XML were noted earlier. Improvements in optimization and additional parallel processing allow faster query processing.

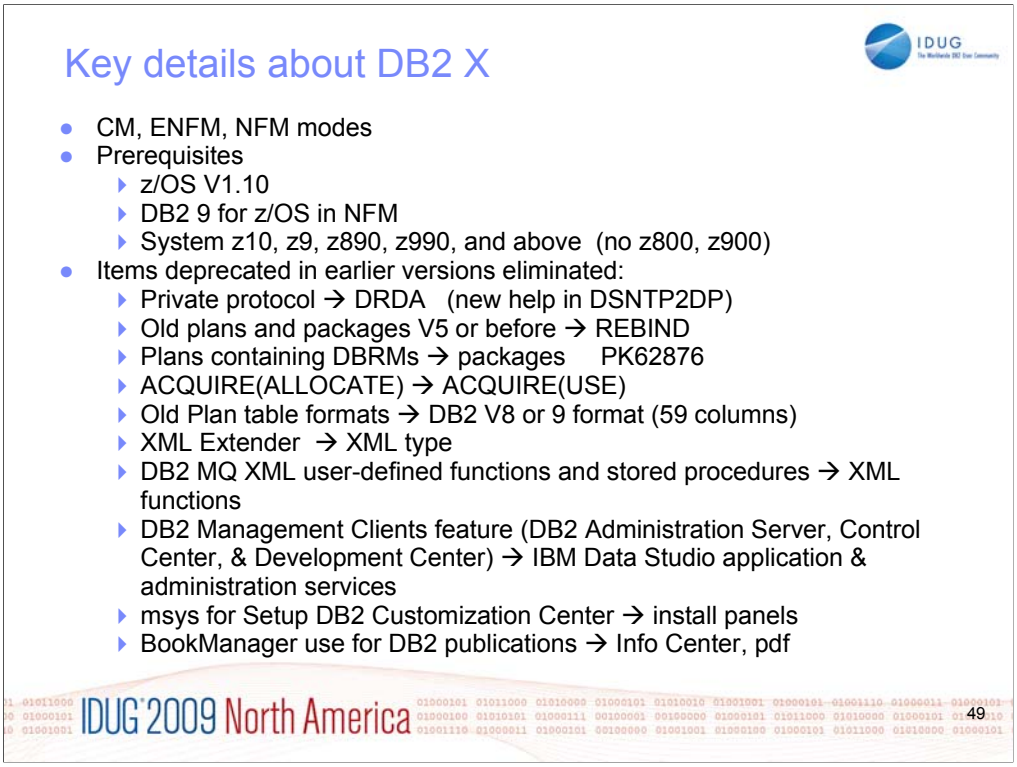

Here are a couple of thoughts about what might be required in hardware and software to run DB2 X. Much will depend upon the timing of the deliveries and market acceptance. Moving forward as quickly as possible means that some of the past must be left behind. See the list of deprecated functions from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

The above features are still included in DB2 9 and may be dropped 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

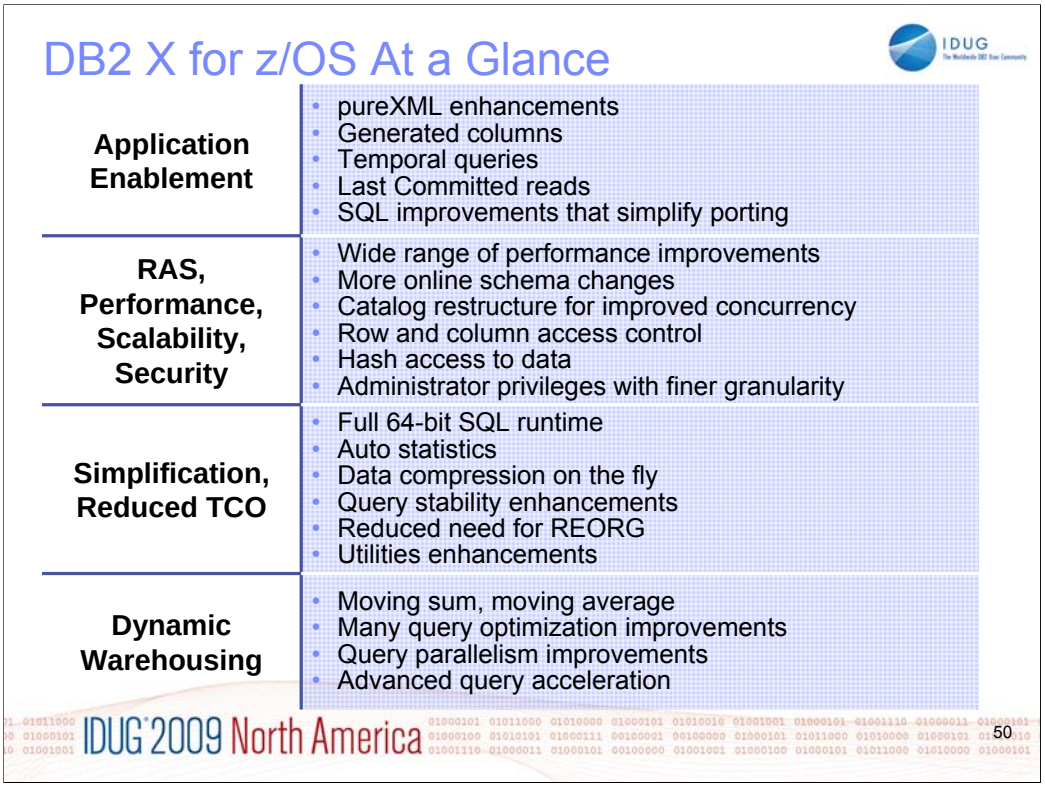

DB2 X builds upon many themes: core platform strengths of performance, scalability, reliability, stability, availability, resilience, and security. All of this work helps with total cost. PureXML and Schema evolution or data definition on demand enhancements are included. Most of the 64 bit work should be completed in DB2 X.

XML, SQL, web services and other programming interfaces extend for usability. 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. Warehousing continues to evolve, with key trends matching DB2 for z/OS strengths.

Get updated books **latest February 28, 2009**
http://publib.boulder.ibm.com/infocenter/imzic/

- ✓Administration Guide
- ✓Data Sharing: Planning and Administration
- ✓Performance Guide
- ✓Utility Guide and Reference
- ✓Application Programming Guide & Reference
- ✓Application Programming Guide & Reference for JAVA™
- ✓ODBC, Spatial, …
- ✓Redbooks
- ✓Installation Guide
- ✓RACF Access Control Module Guide
- ✓Messages
- ✓Codes

You can get most of the books from the Information Center or from the DB2 Library web page. The books continue to be updated, so get the latest ones. Some of the IBM Redbooks publications will be helpful. You may need books from the z/OS Library as well.
http://publib.boulder.ibm.com/infocenter/imzic/

http://www.ibm.com/support/docview.wss?rs=64&uid=swg27011656

http://www.ibm.com/systems/z/os/zos/bkserv/r9pdf/

Be sure to use the latest information to save time and problems. Some of the IBM Redbooks publications have added lately or updated (next page).

DB2 9 in IBM Redbooks Publications

1. DB2 9 Technical Overview SG24-7330
2. DB2 9 Performance Topics SG24-7473
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. DB2 9 Optimization Service Center 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

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.

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 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
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. Optimization Service Center, 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. DB2 for z/OS: Considerations on Small and Large Packages redp4424 http://www.redbooks.ibm.com/abstracts/redp4424.html
16. DB2 9 for z/OS: 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. DB2 9 for z/OS 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/redpieces/abstracts/sg247674.html
20. SAP on DB2 9 for z/OS: Implementing Application Servers on Linux for System z http://www.redbooks.ibm.com/redpieces/abstracts/sg246847.html
21. IBM Data Studio V2.1: Getting Started with Web Services on DB2 for z/OS http://www.redbooks.ibm.com/redpieces/abstracts/redp4510.html
22. Parallel Sysplex Operational Scenarios http://www.redbooks.ibm.com/redpieces/abstracts/sg242079.html
23. Watch for titles on DB2 distributed; serialization & concurrency; utilities

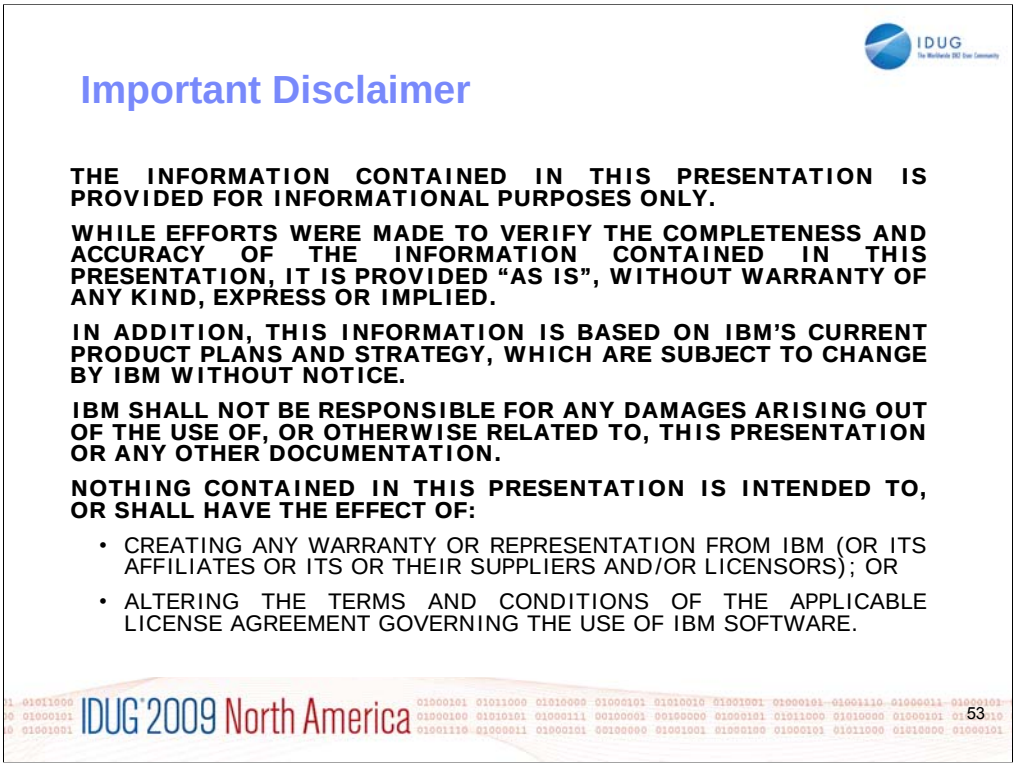

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice. Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in the operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only. 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.

IBM MAY HAVE PATENTS OR PENDING PATENT APPLICATIONS COVERING SUBJECT MATTER IN THIS DOCUMENT. THE FURNISHING OF THIS DOCUMENT DOES NOT IMPLY GIVING LICENSE TO THESE PATENTS.

TRADEMARKS: THE FOLLOWING TERMS ARE TRADEMARKS OR ® REGISTERED TRADEMARKS OF THE IBM CORPORATION IN THE UNITED STATES AND/OR OTHER COUNTRIES: AIX, AS/400, DATABASE 2, DB2, e-business logo, Enterprise Storage Server, ESCON, FICON, OS/400, Netfinity, RISC, RISC SYSTEM/6000, System i, System p, System x, IBM, Lotus, NOTES, WebSphere, z/Architecture, z/OS, System z,

The FOLLOWING TERMS ARE TRADEMARKS OR REGISTERED TRADEMARKS OF THE MICROSOFT CORPORATION IN THE UNITED STATES AND/OR OTHER COUNTRIES: MICROSOFT, WINDOWS, ODBC

Session: A04

What's New in DB2 9 for z/OS?

Questions?

Roger Miller

IBM Silicon Valley Lab

millerrl@us.ibm.com

Thanks for coming.

Title – DB2 for z/OS technical evangelist, strategist, architect, designer, developer, writer, service, DB2 factotum (from the Latin for does everything or jack of all DB2 trades, master of several).

Current Projects - Roger is working to roll out DB2 9 for z/OS, to design the next improvements in DB2.

Technical accomplishments/education - Roger Miller is a DB2 for z/OS technical evangelist, architect and designer who worked on many facets of DB2, ranging from overall design issues to SQL, languages, install, security, audit, standards, performance, concurrency, and availability. He has worked for 30 years on DB2 development, product design and strategy. He often helps customers to use the product, answers many questions and presents frequently to user groups.

Fun facts - Roger likes hiking, bicycling, reading, Shakespeare, Yosemite and bears. He has learned to like working out with a personal trainer.