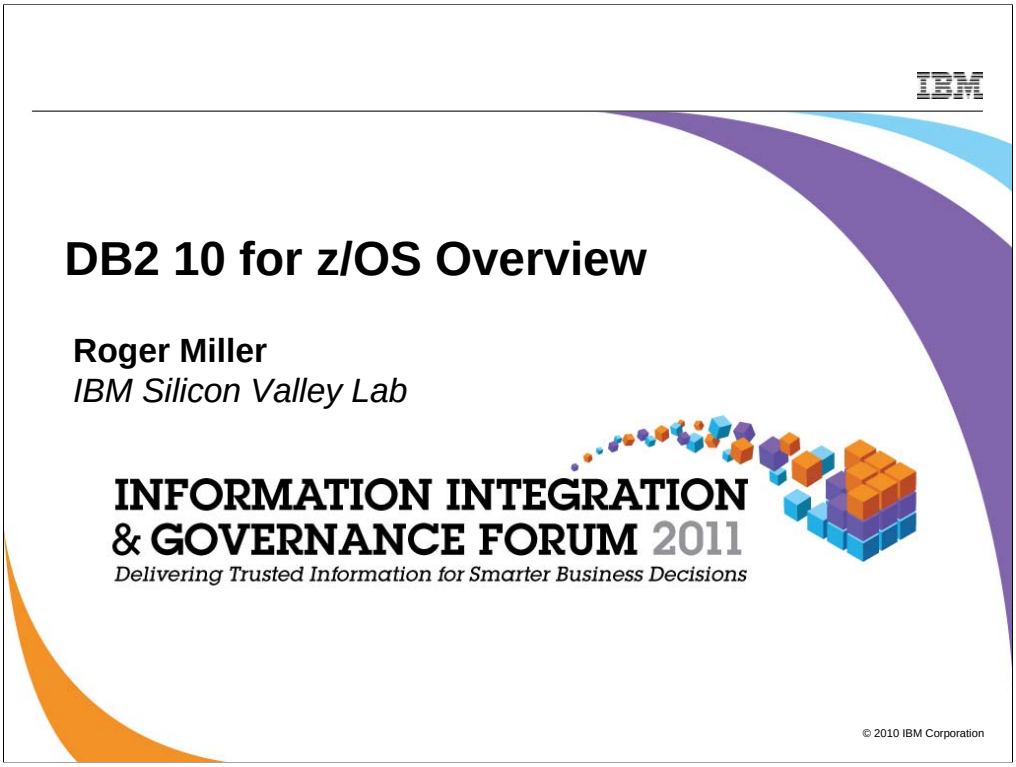

DB2 10 for z/OS – More for less

DB2 10 for z/OS provides the best reduction in CPU for transactions and batch in 22 years. 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 and pureXML improvements 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.

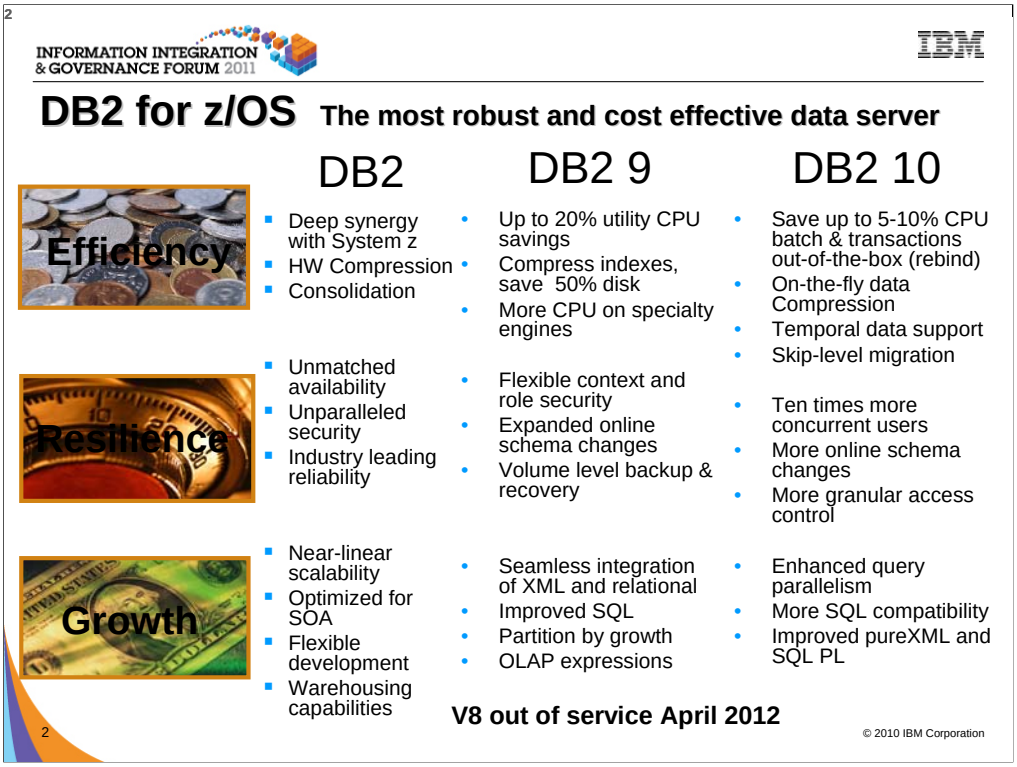

Customers have come to know DB2 as the most robust and cost effective data server. With every version of DB2, we are focused on the needs of our customers to operate efficiently, to be up and running 24x7, and to grow with their business. With V9, customers get CPU and disk savings as well as a boost 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, after rebinding static SQL applications. 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.

INFORMATION INTEGRATION & GOVERNANCE FORUM 2011

IBM

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 change: data definition, utilities, & subsystem
5. Improved security with more granularity
6. Temporal or versioned data
7. SQL enhancements improve portability
8. pureXML performance and usability
9. Hash, index include columns, skip migration, …
   Pick your favorite!
10. Productivity improved for database & systems administrators, and application programmers

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.

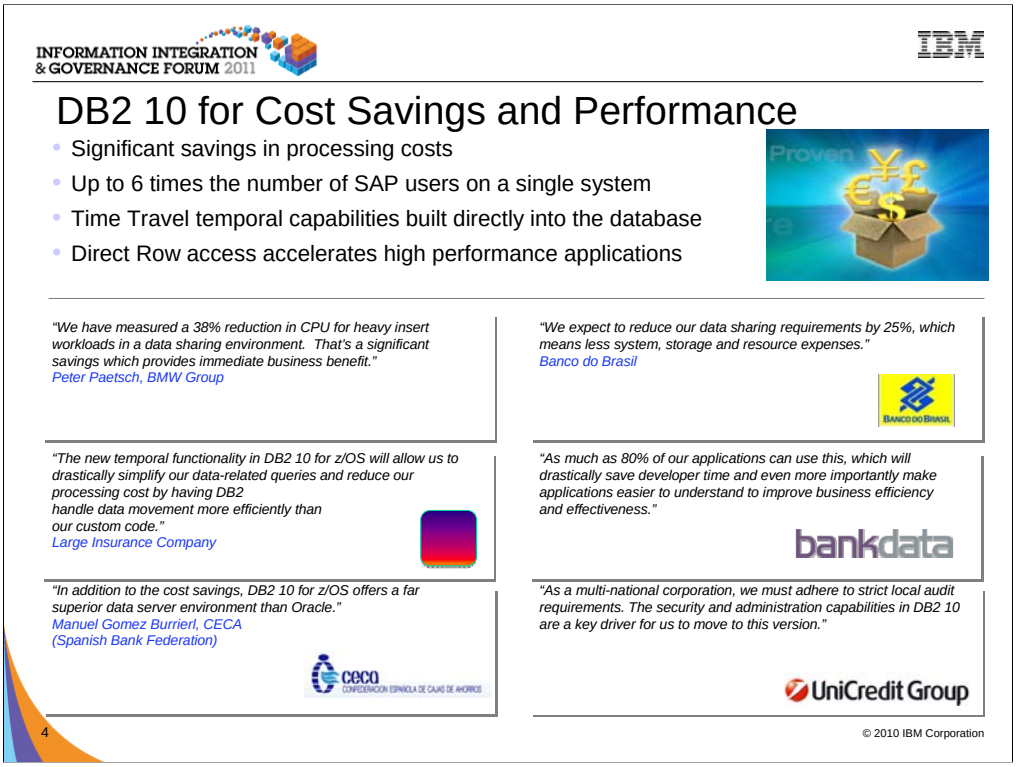

It's easy to see what customers are saying about DB2 10. Their words are on the web pages, in list servers and blogs.

"Over the past several months, BMW has tested the new version of DB2 10 for z/OS, focusing on specific features and comparing these directly to the same features in DB2 9 for z/OS. One of the IBM design goals expected a general improvement in massive parallel SQL-insert performance, where we achieved close to 40% CPU improvement and significant elapse time reduction in direct comparison to DB2 9 for z/OS. For all of our critical tested selects statements, the version 10 optimizer chose the optimal access path, sometimes even improving previous access path choices in version 9. Overall, we are very pleased with the added functionality and architectural enhancements, and are looking forward to this exciting release." Philipp Nowak, DB2 Product Manager, BMW Group

"With the scalability improvements in DB2 10, we expect to be able to quickly reduce our production data sharing group from 20 members to 15", said Paulo Sahadi, Senior Production Manager at Banco do Brasil. With DB2 10 able to handle 5-10 times as many threads as the previous version, the upgrade will immediately give the bank some much-needed room for future workload growth while simultaneously reducing their data sharing overhead. "We will also save some CPU and storage from removing the five DB2 systems, and we will have to spend a lot less time monitoring our virtual storage." Paulo Sahadi, Senior Production Manager, Information Management Division, Banco do Brasil

See much more on the web.

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

http://www.ibm.com/software/data/db2/zos/testimonials.html

Some Beta Customer Performance Feedback

Workload	Results
Customer1: Distributed Concurrent Insert	50% DB2 elapsed time reduction; 15% chargeable CPU reduction after enabling high perf DBAT
Customer2: CICS online transactions	Approx. 7% CPU reduction in DB2 10 CM after REBIND, Another 4% reduction with 1MB page usage
Customer3: CICS online transactions	Approx 5% CPU reduction
Customer4: Data sharing heavy concurrent insert	38% CPU reduction
Customer5: Queries	Average CPU reduction 28% from V8 to DB2 10 NFM
Customer6: Batch	Overall 28% CPU reduction after rebind packages
Customer7: DDF OLTP	40% CPU reduction for JDBC stored procedures workload, 15% CPU reduction for securities trading

These figures show some customer performance measurements from the beta program. These measurements reflected various customer work better than a benchmark, but were usually less repeatable. Customers could not generally have dedicated resources, so they measured multiple times and checked for consistency of the runs. Most of the customer information showed the ability to get improvements similar to those in the benchmark measurements, with a wider range of work and results.

Beta Customer Feedback on Selected New Functions

Workload	Results
Multi row insert (data sharing)	33% CPU reduction from DB2 9, 4x improvement from V8 due to LRSN spin reduction
Parallel Index Update	30-40% Elapsed time improvement with class 2 CPU time reduction
Inline LOB	SELECT LOB shows 80% CPU reduction
Include Index	17% CPU reduction in insert after using INCLUDE INDEX
Hash Access	20-30% CPU reduction in random access 16% CPU reduction comparing Hash Access and Index-data access. 5% CPU reduction comparing Hash against Index only access

These measurements are for some specific function improvements measured by customers, rather than for a broader workload.

Top Items Driving DB2 10 Decisions

- CPU / Performance improvements
- Virtual storage enhancements
  - Reduce number of members, save money
- Stability & regression
- Security enhancements
  - Built-in security, trace & audit features, new roles, end-to-end auditing
  - Cleaner/safer environment; Better audit/compliance
- Temporal
- Skip-level migration DB2 V8 → DB2 10

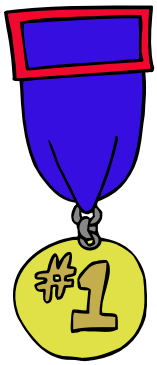

Vast majority of beta customers plan production in 2011

Customers have been enthusiastic about DB2 10 for z/OS performance and scalability. Some customers are able to simplify their structure and many are seeing better productivity.

The security and temporal function are seeing stronger than expected early acceptance. A safer infrastructure with better audit function can help customers avoid the need for a new security structure.

Customers are signing up for the V8 to DB2 10 migration to save time and to get these improvements faster. The first production has started. Most of the 23 beta customers plan to move to production in 2011.

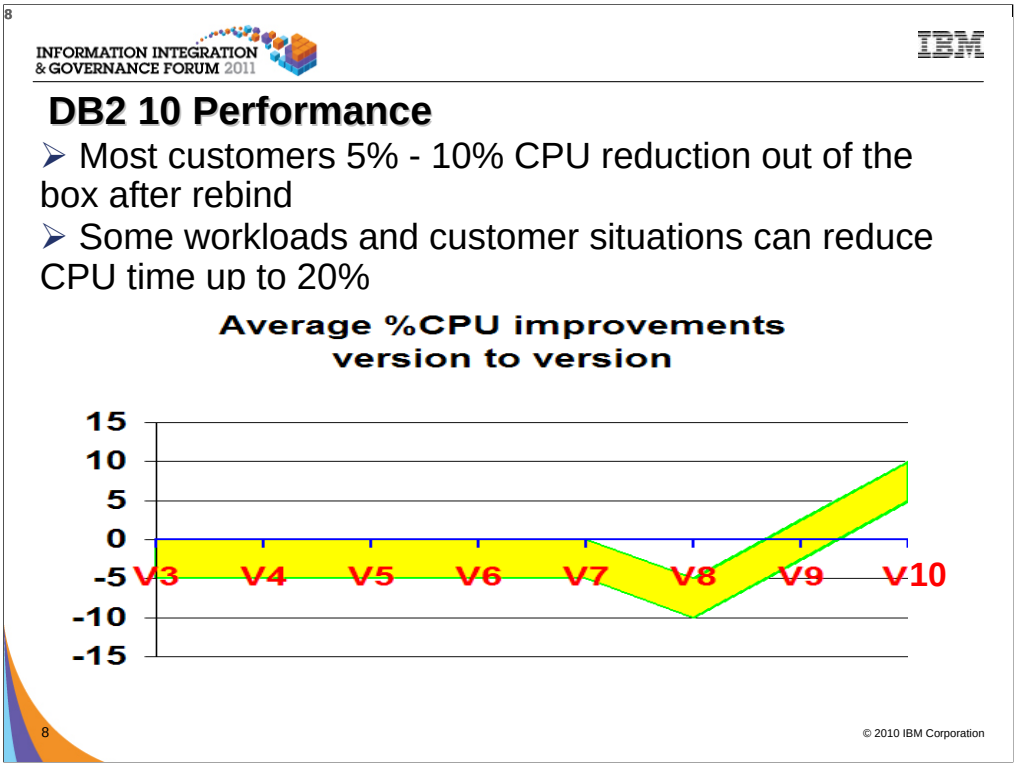

The objective for general transaction and batch performance has been to minimize the regression. Version 2 in 1988 provided a substantial Improvement in transaction and batch work, but the past 21 years have seen the focus on removal of bottlenecks, scalability, query performance, and minimizing performance 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 and rebinding the applications. 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.

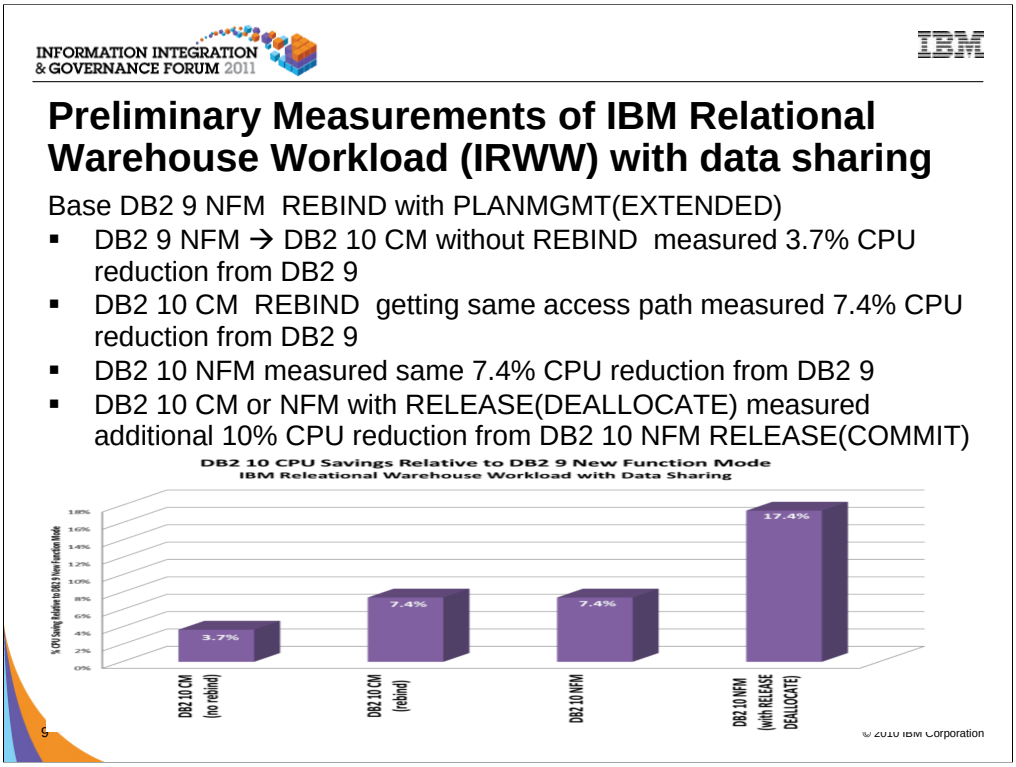

This is the scenario for a benchmark transaction that is run on DB2 9, then on DB2 10. This scenario uses some new function in DB2 9 to BIND or REBIND a package with access control management to allow three copies. These are fairly light CICS transactions that have been used for many DB2 transaction benchmarks.

- In step 1, this application is moved to DB2 10 CM without a REBIND, and the result is a 3.7% reduction in CPU time.
- In step 2, still in DB2 10 CM, a REBIND is performed but with exactly the same access path. With the REBIND, the CPU savings over DB2 9 was 9.4%, double that without the REBIND.
- In step 3, moving to NFM, the CPU time is the same.
- In step 4, these transactions are changed to use RELEASE(DEALLOCATE), saving an additional 10% of the CPU time compared to the prior RELEASE(COMMIT).

So this scenario demonstrates the runtime improvements and CPU value of REBIND and RELEASE(DEALLOCATE)

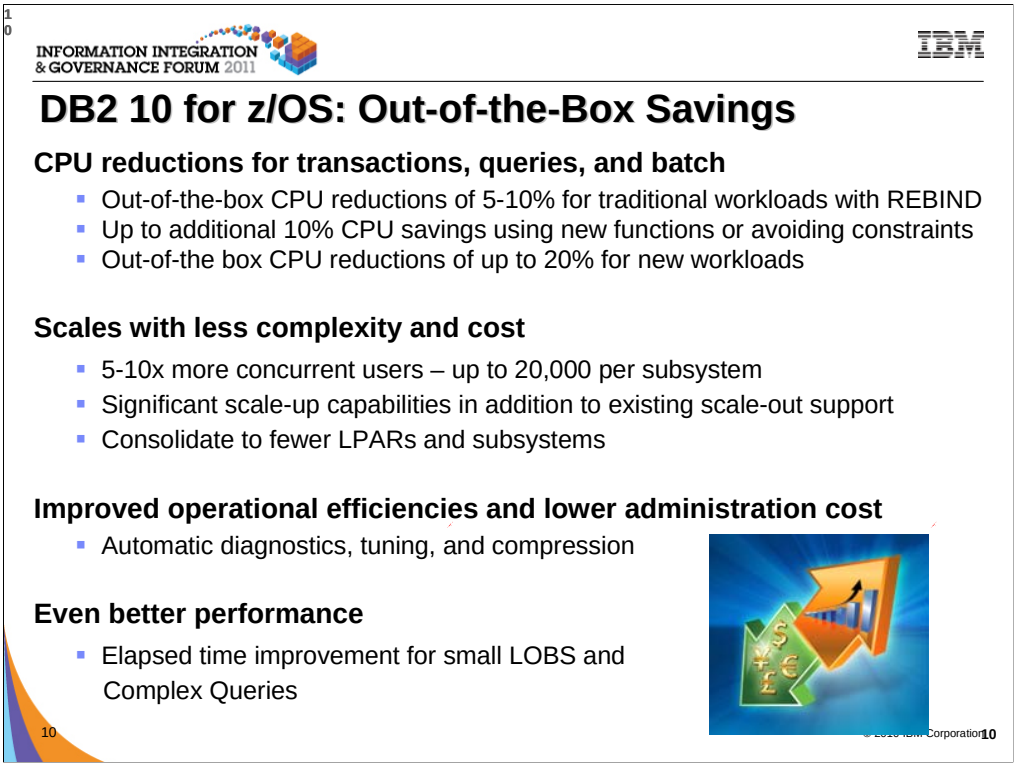

Improved operational efficiency for out-of-the-box savings Version 10 delivers great value by reducing CPU usage. Compared to previous releases of DB2 for z/OS, most customers can achieve out-of-the-box CPU savings of five to ten percent for traditional workloads and up to 20 percent for some workloads. DB2 reduces CPU usage by optimizing processor times and memory access, leveraging the latest processor improvements, larger amounts of memory, solid-state drives, and z/OS enhancements. Improved scalability and constraint relief can add to the savings. Productivity improvements for database and systems administrators can drive even more savings.

In Version 10, performance improvements focus on reducing CPU processing time without causing significant administration or application changes. Most performance improvements are implemented by simply migrating to Version 10 and rebinding. You gain significant performance improvements from distributed data facility (DDF) optimization, buffer pool enhancements, parallelism enhancements, and more.

Most customers should see 5% - 10% CPU reduction out of the box after rebinding. Some workloads and customer situations can reduce CPU time more. While versions 3, 4, 5, 6, and 7 generally increased CPU times by a small amount, less than 5%, version 8 increased CPU time by 5% to 10% for most customers. DB2 9 often reduced CPU a little or increased very little (less than 2%). New function, improved scalability, and faster hardware compensated for the increases in CPU time. Using the new function could change the increases into reductions, particularly with DB2 V8 and multi-row fetch.

Early DB2 10 performance benchmarking and customer experience has shown that most customers can expect to get 5% to 10% CPU reduction after rebinding. Some customers will get more and some less. Some situations can reduce CPU time more than that. Customers who have scalability issues, such as virtual storage constraints or latching can see higher improvements. Opportunities for tuning can take advantage of memory improvements. High volume, short-running distributed transactions can take advantage of CPU reductions, using release deallocate. Concurrent sequential insert can be reduced from 5% - 40%. Queries can be improved as much as 20% without access path change, and more for better access paths. A workload with native SQL procedures has shown up to 20% CPU reduction. For DB2 utilities, customers moving from DB2 9 should expect a small (0% to 7%) reduction in CPU times varying by utility, while customers moving from DB2 V8 will see larger CPU reductions in the range of 20%.

Productivity improvements: Improvements in SQL and XML improve productivity for those who develop new applications and for those who are porting from other platforms. Automating, reducing, or eliminating tasks, and avoiding manual invocation improves productivity and can help avoid problems. Resiliency improvements for virtual storage and availability increase productivity. DB2 10 improvements make the install, migration, and service processes faster and more reliable, including the ability to skip from V8 to DB2 10.

Innovations in Version 10 drive new value in resiliency through scalability improvements and fewer outages, whether those outages are planned or unplanned. Scalability delivers the ability to handle five to ten times more concurrent users in a single DB2 subsystem than in previous releases of DB2 for z/OS (as many as 20,000 concurrent threads). Improved availability is supported by schema evolution, or data definition on demand, and manageability enhancements for query performance.

INFORMATION INTEGRATION & GOVERNANCE FORUM 2011

IBM

DB2 Deep Synergy With System z

Key integration points include:

- Data sharing (availability and scale out)
- zIIP and other specialty engines
- Unicode conversion
- Encrypted communication & data
- Hardware data compression & encryption
- Cross-memory, memory protection keys
- Sorting
- Multi-core, large N-way
- 64-bit addressing and large memory
- z/OS Workload Manager
- z/OS Security Server (RACF)
- z/OS RRS integrated commit coordinator
- System z10 1 MB page size, decimal float
- Solid state disks
- zEnterprise z196, zBX, z10, …

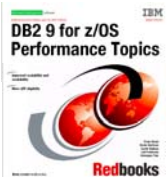

11

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.

DB2 has unique ways to use the z10 and zEnterprise to deliver additional value.

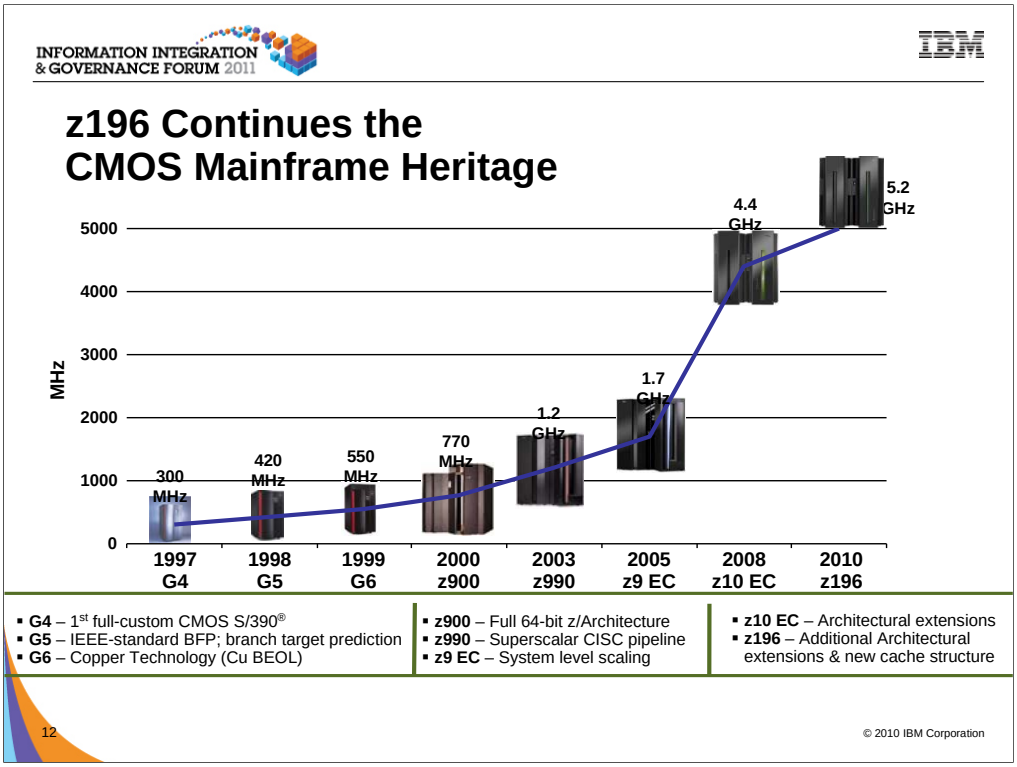

The zEnterprise system offers substantial improvements in the base z196 and a new hybrid structure.

The design of the IBM System z10™ processor chip wass 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. The average performance increase for the z10 EC over the z9 EC is about 58%, but we see substantial variation in that ration as workloads change, from 40% to 80% for most workloads, but some improve by a factor of 2.1 times faster, while some can run at very close to the same speed. The number of cycles per instruction increases to roughly 5 cycles per instruction.

It 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

INFORMATION INTEGRATION & GOVERNANCE FORUM 2011

IBM

DB2 and zEnterprise 196 performance

Measurement data currently available shows

- DB2 OLTP workloads observing 1.3x to 1.6x DB2 CPU reduction compared to z10 processors
- Higher DB2 CPU reduction can be achieved as number of processors per LPAR increases
- With DB2 10 and zEnterprise, CPU reduction can be up to 1.8x compared to DB2 9 and z10 with many processors per LPAR

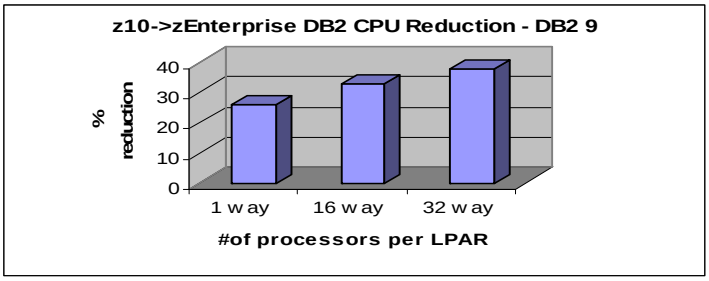

LSPR measurements of DB2 9 with the zEnterprise z196 show substantial reductions when compared to z10 processors. This transaction workload showed a range of 1.3 to 1.6 times CPU time reduction, with the best CPU reductions when more processors per LPAR are used. Including the larger number of faster processors (80 vs 64) and DB2 10 can mean fewer footprints.

https://www-304.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindexpdf/$file/SC28118714_20100714.pdf

https://www-304.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprwork?OpenDocument&pathID=

INFORMATION INTEGRATION & GOVERNANCE FORUM 2011

IBM

System zEnterprise Benefits for DB2

Taking System z synergy to the next level

- Faster CPUs, more CPUs, more memory → better DB2 performance, scalability
- Compression hardware expected to increase DB2 data compression performance
- Cache optimization, 192M L4 Cache expected to benefit DB2 work
- Hybrid architecture query performance acceleration with IBM Smart Analytics Optimizer
- Excellent synergy with DB2 10 → significant CPU reduction and scalability increase
  - CPU reductions
  - Remove key single system scaling inhibitors: virtual storage, latching, catalog, utilities, …
  - Translation Lookaside Buffer changes expected to improve performance for 1MB page sizes
  - Buffer pool management

14

The zEnterprise z196 and DB2 10 take synergy to the next level. See Chris Crone's presentation for more detail. Faster CPUs, more CPUs, and more memory means better DB2 performance and scalability.

Compression hardware improvements are expected to increase DB2 data compression performance.

192M L4 Cache is expected to benefit DB2 workloads, as DB2 uses the memory. DB2 can take an advantage of cache optimization on zEnterprise.

Hybrid architecture delivers new opportunities for DB2 query performance acceleration with IBM Smart Analytics Optimizer.

Excellent synergy with DB2 10 removes many single system scaling inhibitors – virtual storage, latching, catalog concurrency, and utility concurrency.

The Translation Lookaside Buffer Changes are expected to improve DB2 10 performance for 1MB page sizes. Buffer pool improvements for large buffers will provide additional help.

Combined with DB2 10 improvements CPU reduction, buffer pool management, relief for virtual storage constraint and latch contention, DB2 applications can enjoy significant cost reduction and scalability improvement on zEnterprise.

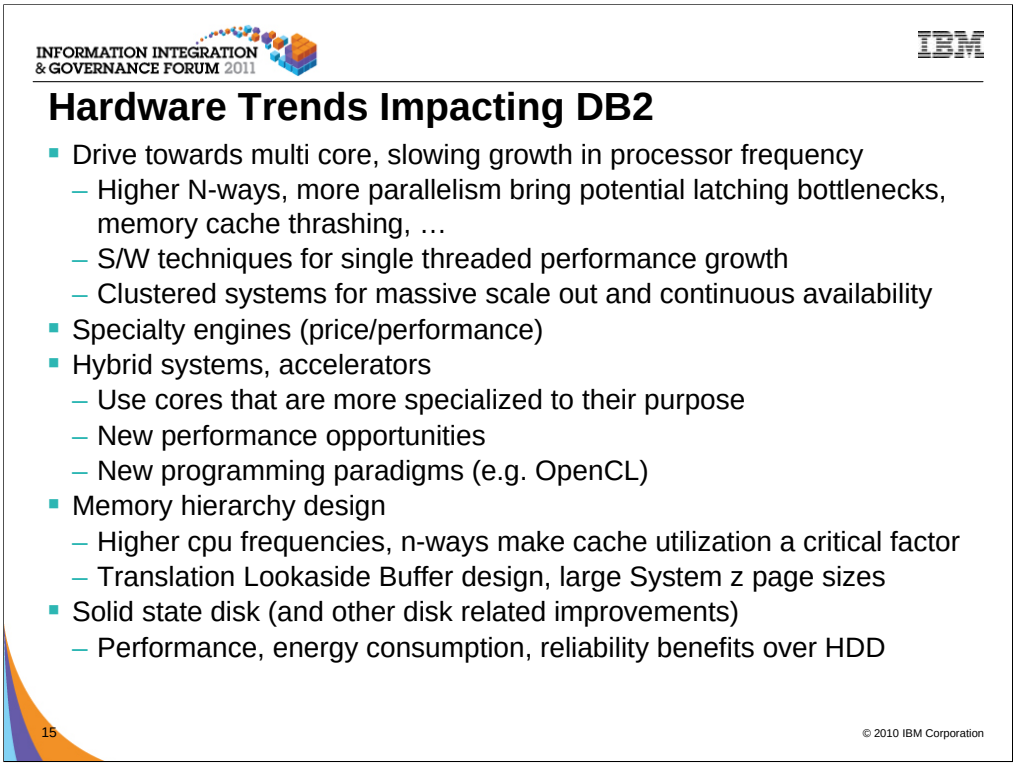

Multi-core processors, alone, will be insufficient to enable application-level systems price & performance improvements at historical rates. Semiconductor scaling, which has provided the foundation for system-level improvements in cost and performance, is slowing down. While densities are expected to continue to improve, power efficiency is expected to lag. Transistor performance at constant power density has been close to constant since the 90nm node. The resulting lack of processor frequency growth has forced the industry to seek performance improvements through the introduction of multiple cores on a processor, and the parallel software to exploit them. In contrast to nodes prior to 32nm and 22nm, power density at constant frequency (and constant switching factor) will increase node-to-node past 22nm. This means that if a core was migrated from 22nm to 15nm and replicated to fill a chip of similar size, the operating frequency of these cores would have to be reduced in order to use the chip in a system of similar power supply and cooling capacity.

INFORMATION INTEGRATION & GOVERNANCE FORUM 2011

IBM

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

The IBM Smart Analytics Optimizer was announced in July 2010. This technique works with the z10 and zEnterprise to provide to a hybrid structure delivering lower cost query on System z.

http://www.ibm.com/software/data/infosphere/smart-analytics-optimizer-z/

IBM purchased Netezza, one of the leading warehouse appliance vendors. IBM adds to the lines of hardware, data management and integration, warehousing and analytics that includes Cognos, SPSS, and now Netezza.

http://www.eweekeurope.co.uk/news/ibm-looks-ahead-to-future-with-netezza-13361

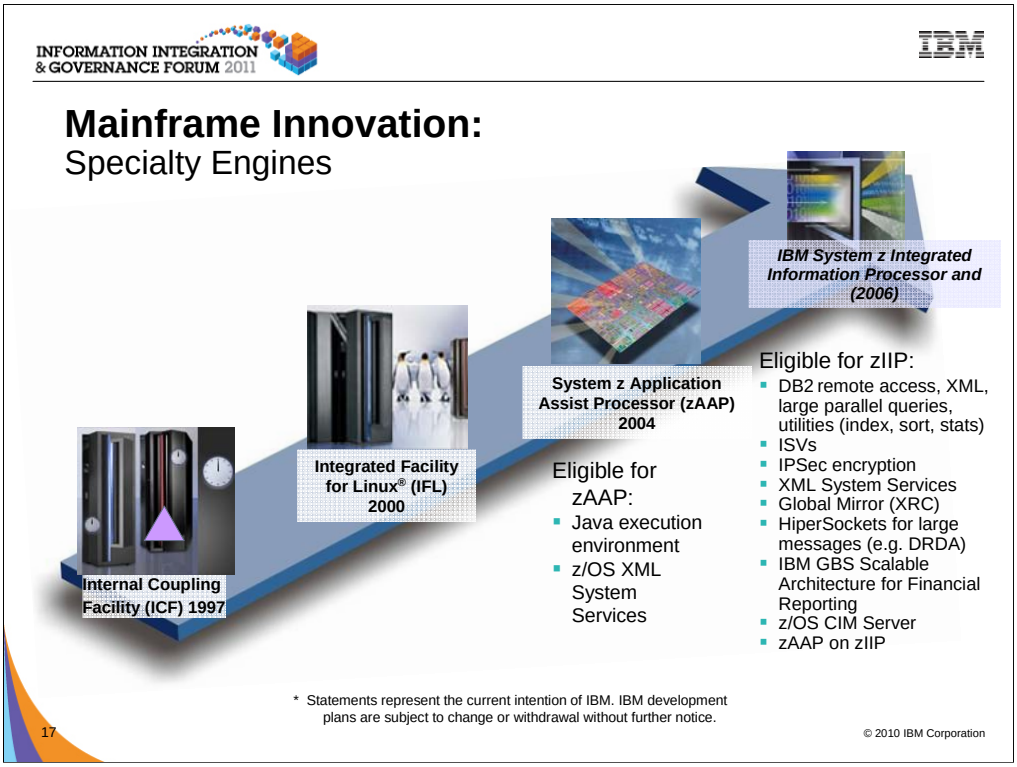

The IBM family of specialty engines has evolved. The zIIP uses have multipled since its 2006 introduction, with changes in 2006, 2007, 2008, 2009, and 2010 so far.

Announced August 2009:

•z/OS CIM Server

•DB2 sort utility

•zAAP on zIIP

The changes in 2010 include improvements in service to add a little more remote processing and improvements in DB2 10.

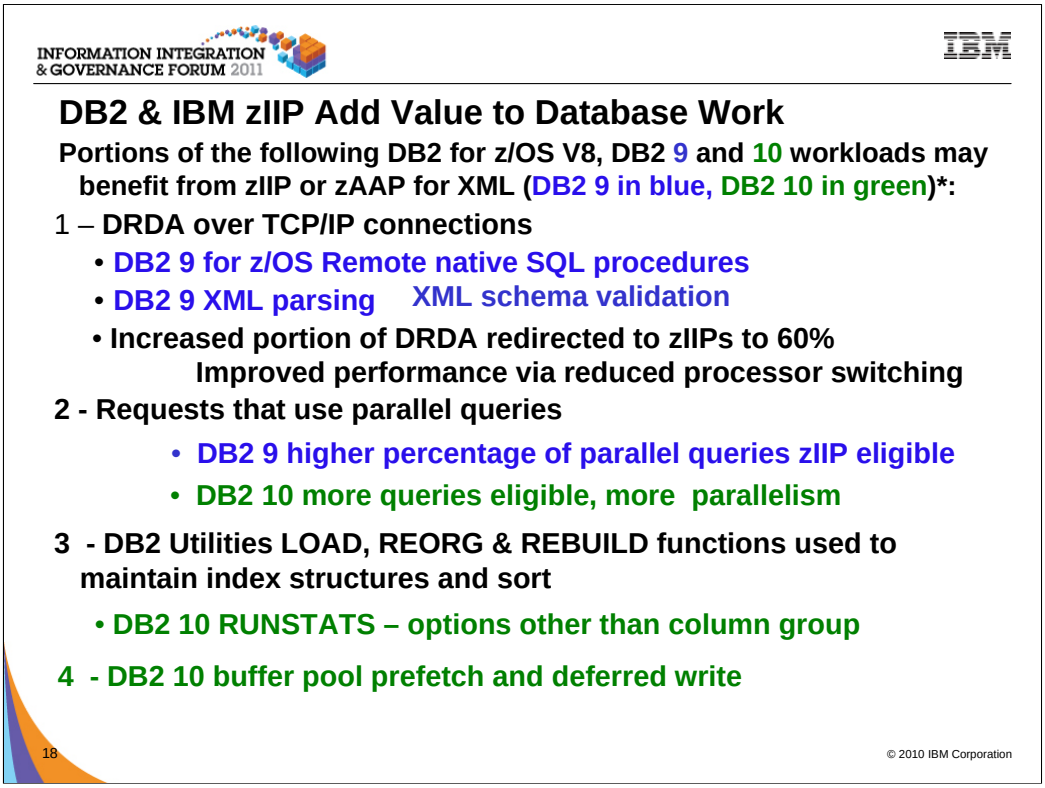

*** zIIP allows a program working with z/OS to have all or a portion of its enclave Service Request Block (SRB) work directed to zIIP. Above types of DB2 work are those running in enclave SRBs, of which portions can be sent to zIIP.**

The zIIP is designed so that a program can work with z/OS to have all or a portion of its enclave Service Request Block (SRB) work directed to the zIIP. The above types of DB2 V8 work are those executing in enclave SRBs, of which portions can be sent to the zIIP. Not all of this work will be run on zIIP. z/OS will direct the work between the general processor and the zIIP. The zIIP is designed so a software program can work with z/OS to dispatch workloads to the zIIP with no anticipated changes to the application – only changes in z/OS and DB2.

IBM DB2 for z/OS version 8 was the first IBM software able to take advantage of the zIIP. Initially, the following workloads can benefit:

• SQL processing of DRDA network-connected applications over TCP/IP: These DRDA applications include ERP (e.g. SAP), CRM (Siebel), or business intelligence and are expected to provide the primary benefit to customers. Stored procedures and UDFs run under TCBs, so they are not generally eligible, except for the call, commit and result set processing. DB2 9 remote native SQL Procedure Language is eligible for zIIP processing. BI application query processing utilizing DB2 parallel query capabilities; and functions of specified DB2 utilities that perform index maintenance.

•For more, see http://www.ibm.com/systems/z/ziip/

2010 New method to control the portion of SQL requests that are authorized to be diverted to zIIP engines with improved performance via reduced processor switching. This change also increases portion of DRDA that is authorized to run on zIIPs to 60%. APAR PM12256 for V8 & DB2 9. Included in DB2 10 base.

DB2 10 improvements include increased parallel processing, the RUNSTATS utility and buffer pool prefetch.

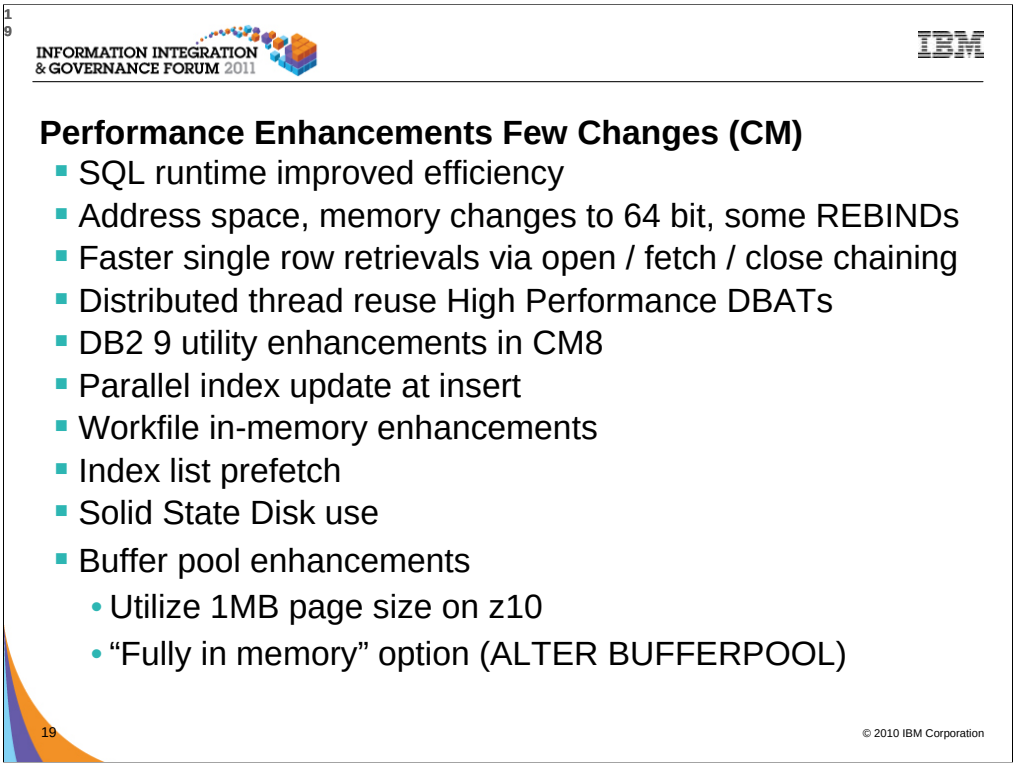

These are the improvements which we expect almost every customer to see as soon as DB2 10 is running, even in conversion mode.
Many paths within DB2 processing leverage better 64-bit memory capabilities. This results in better SQL performance for many existing SQL access plans.
All of the memory improvements provide immediate relief for all of many memory constrained systems. With some REBINDS, memory usage is reduced, allowing you to use memory more effectively for example in buffer pools for performance in your environment.
Better handling of singleton Cursor Selects through chaining, combines the open fetch and close work ONLY once across the network improving network bandwidth efficiency and overall performance.
Also for DDF transactions, there is increased DDF thread reuse. This enhancement starts to handle and reuse DDF threads though the same methods that we have had for a long time through CICS thread interfaces for robust consistent type transactions.
For people that are directly coming to DB2 10 through migrations from Version 8 they immediately get long list of the Version 9 enhancements especially all the Utility performance improvements. Some are experiencing 20% elapse time savings in Version 9.
Updates to index columns are done in parallel in DB2 10 improving insert performance. This out of the box enhancement along with DB2 10 List Pre-fetch capabilities improve all existing applications that use list pre-fetch activities across all existing accesses paths and especially when using indexes that may be a little disorganized and in need of a reorg.
DB2 can now use solid state disk devices. These devices are great for the workfiles, GTTs and other high performance table spaces within your environment.
Also, the enhanced way DB2 10 uses in-memory Workfiles and the improvements related to RID pool overflows helps all application systems avoid the deadly table space scan at the peak processing times.
DB2 can utilize the new bigger 1MB page size on z10 and provide additional buffer pool options to put a table fully in memory with an easy simple table space ALTER.
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.

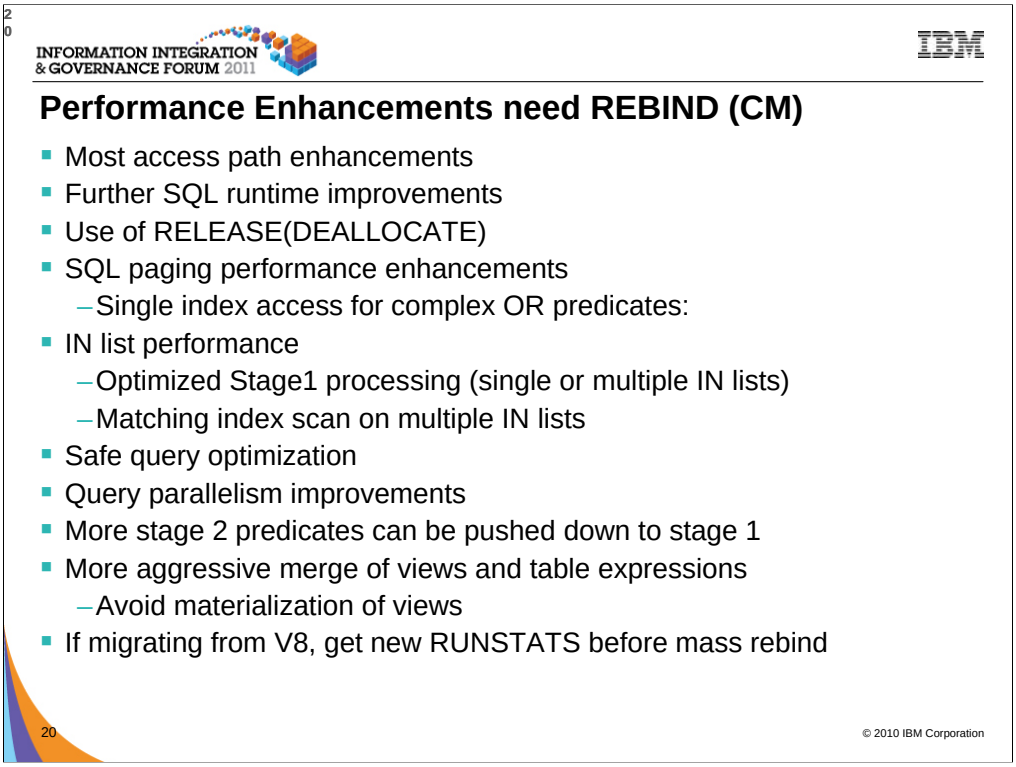

Rebind is required for a long list of improvements in optimization and parallelism. The key improvements from REBIND in CM include SQL inlist improvements, SQL paging enhancements, query parallelism improvements, and more aggressive view and table expression merge.

Improvements with WHERE OR clauses that have columns that all reference the same index can be optimized for a single Stage 1 evaluation instead of being evaluated multiple times and then retrieved multiple times. This consolidates and improves performance dramatically for processing.

IN-List predicates are now evaluated through Stage 1 processing and provide a matching index access of multiple IN-list Where clause criteria. This is common and great for applications that have cursor pagination within their application.

More SQL query parallelism offloads precious CPU main engine cycles and pushes more processing into your specialty zIIP engines.

REBINDS also let DB2 push down more predicate evaluations from Stage 2 to Stage 1 during data retrieval. This helps reduce the rows evaluated in each step of the SQL access path improving performance significantly for complex access paths.

This also comes into play with View and expression materialization. Since the amount of data through the materialized steps can be sometimes big, cutting down the amount of data in Stage 1 processing interim result sets can really help performance of these complex SQL statement and their processing situations.

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.

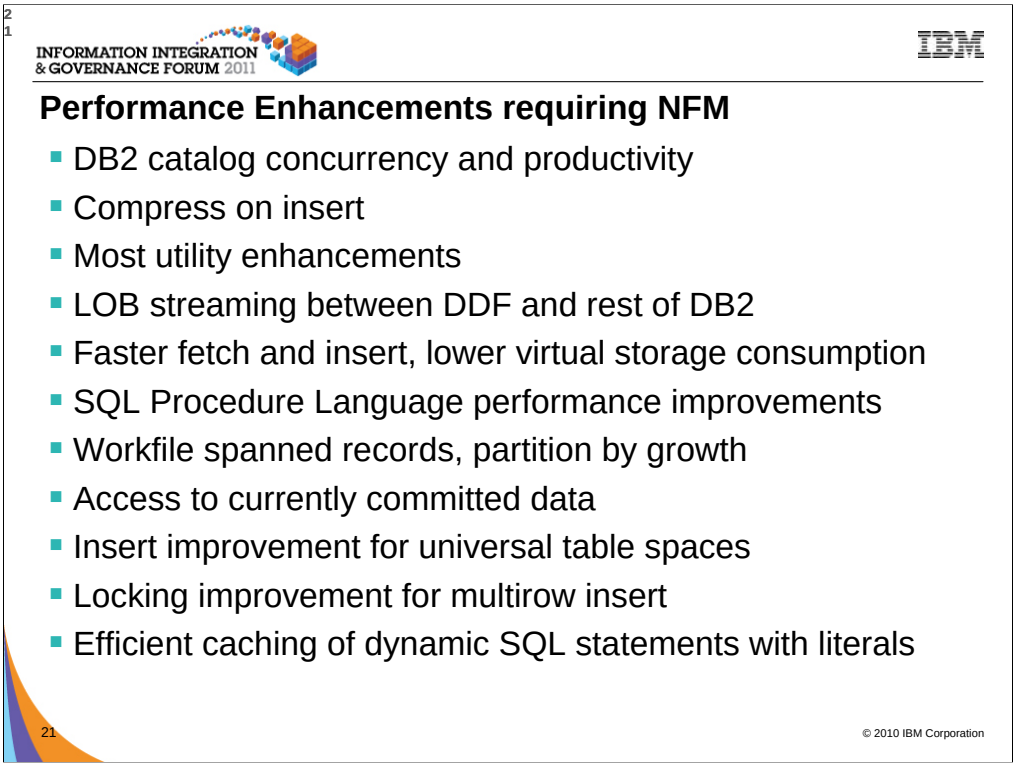

Some of the performance improvements require new function mode and some work by database administrators to tune the database design and often to REBIND. Efficient caching for literals needs a rebind in NFM.
One of the enhancements that will help everyone with large dynamic SQL applications are the improved processes that handle the Dynamic Statement Cache. They now consolidate SQL statements that are the same but have different literals. This reduces the SQL statement space used in the Dynamic Statement Cache and reuses the security and object verification and access path already developed for the SQL statement. This dramatically improves Dynamic Statement Cache cache hits, reduces the duplicate SQL statement previously held, makes room for more SQL improving the overall workload performance.
Also all the concurrency and performance enhancements through the DB2 10 Utilities improvements are available for more on-line ALTER and maintenance activities.
Better streaming and minimized LOB default sizes helps improve when LOB materialization is happening with the system and this is especially important for DDF type applications because of the network impact of large objects.
Small WORKFILEs are now available for simple predicate evaluation for improving performance DB2 provides native support for the SQL procedural language eliminating the cumbersome requirement to generate a C program from the SQL procedure that would then execute as an external stored procedure. DB2 10 SQL procedures are better optimized to execute more efficiently more common constructs are optimized within the DB2 code making SQL procedures very efficient for performance within the SQL procedure language.
WORKFILE can have expanded records up to 65K so larger Joins and answers set can be generated from DB2.
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.

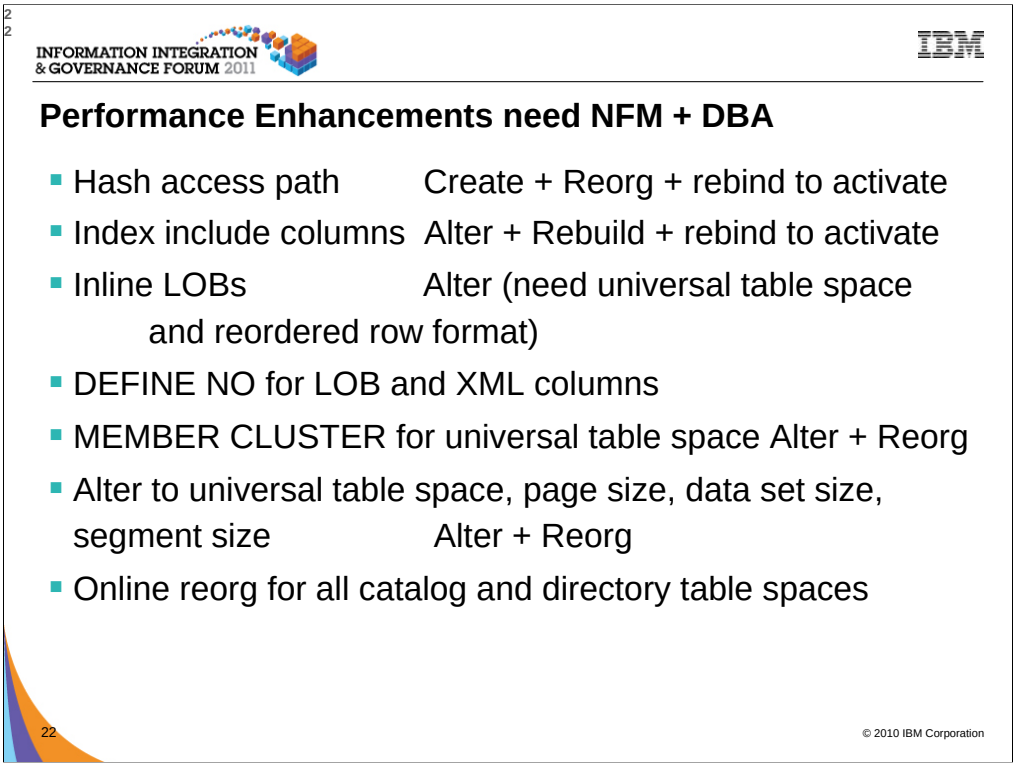

Some of the performance improvements require new function mode and some work by database administrators to tune the database design and often to REBIND. The DBAs can improve on their direct keys access and implement the new Hash access table space type reducing a key access to potentially a single I/O. This is great for the customer or product id lookouts that are done millions or billions of times a day. The Index include non-unique columns within a UNIQUE index definition provides great way to eliminate indexes and consolidate other indexes for better optimization and overall performance. This is great also for better access paths because all the columns are now in the single index for better access paths and fewer indexes to reference.

Inline LOBs use the reorder row format and handle the LOB better for overall streaming and application performance. This along with the DEFINE NO feature allows the row to be used and the data set for the LOB not to be defined. The application can still reference the row and get the proper results. The system doesn't define the LOB data set until the LOB is saved which can save a lot of space for tables that only sometimes had LOB entries.

The MEMBER CLUSTER setting can help eliminate contention and the clustering requirements within a data sharing environment table. This boosts performance and relieves contention across your data-sharing environment.

And now almost any attribute within the definition of the table space can be Altered within DB2 and then applied through an on-line reorg. This is great for application availability to keep your business rolling with DB2.

Hashing Summary

- Provides fast, direct location of most rows
  - Reduces I/O and CPU in most cases
  - Can replace an existing Primary or Unique Key Index
    - Faster Insertion/Deletion
- Size of Fixed Size Hash Area is important
  - Too small and performance degrades, too large and space is wasted
- DB2 helps you manage the size
  - REORG AUTOESTSPACE YES
  - RTS tracks the number of overflowed entries
- If clustering is important for query performance, then be aware that Hash will eliminate these benefits
- LOAD performance is slower with hash

Hash access is particularly valuable for applications that randomly access rows in a table. When an application uses equality predicates on a unique key to locate a row, the savings from using hash access can be substantial, particularly in cases where the index access would have to traverse an index that is many levels deep. Some examples of these kinds of accesses are random look-ups by bank or insurance policy account numbers. Because the base technology of hash access results in randomization of the rows within a fixed size hash data area, hash access is not intended for accesses that currently scan ranges of index keys, or that access tables that vary drastically in size. When used appropriately, hash access provides very fast access to random individual rows in a table.

IBM performance analysis has shown that a table that is organized by hash spaces can provide significant cost savings when applied in appropriate situations. For example:

13% class 2 CPU reduction was observed for a SELECT statement that retrieved all columns of a single row from a table that is organized by hash by using a fully qualified key, as compared to access through an index (with 3 index levels) on the same columns that defined the hash key.

37% class 2 CPU reduction was measured for 50,000 executions of a SELECT statement that retrieved all columns of single rows from a table that is organized by hash of the statement, as compared to access through an index (with 3 index levels) on the same columns that defined the hash key.

9% class 2 CPU reduction was measured for a SELECT statement that retrieved a subset of columns from a table that is organized by hash, when compared to index only access (with 3 index levels) on all columns.

In each case reduced get page and synchronous I/O operations were also observed.

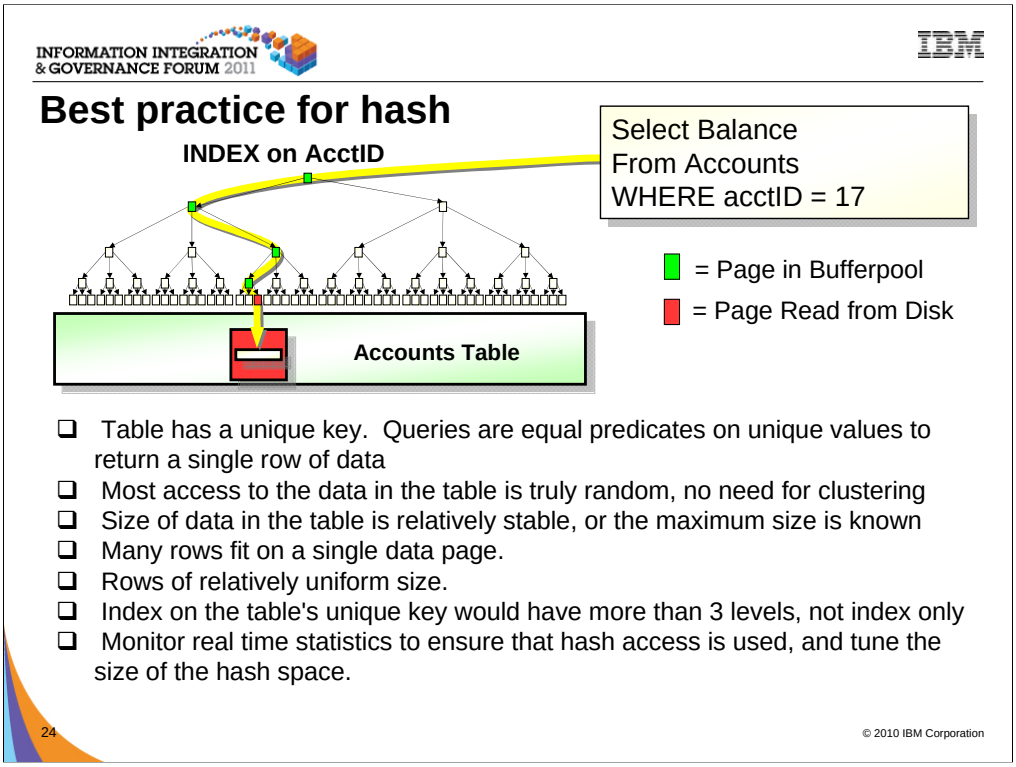

Some hash organization best practices

Evaluate the applications and workload thoroughly before adopting hash organization. Hash organized tables deliver the most reductions and response time improvements in certain specific situations:

•The table has a unique key

•Queries that access the table specify equality predicates on unique values to return a single row of data.

•Most access to the data in the table is truly random. Applications that use range scans, or that depend on clustered data, do not perform optimally with hash organized tables. You can use IFCID199 to verify that access is truly random.

•The size of the data in the table is relatively stable, or the maximum size of the data is known. The amount of space that must be dedicated to a hash organized table is fixed.

•Many rows fit on a single data page. When too few rows fit within a single data page, additional space might be required to achieve the benefits of hash organization.

•The tables contains rows of relatively uniform size.

•The benefits of hash access are greatest when an index on the table's unique key would have more than 3 levels.

•After adopting hash organization, monitor real time statistics to ensure that hash access is used, and tune the size of the hash space.

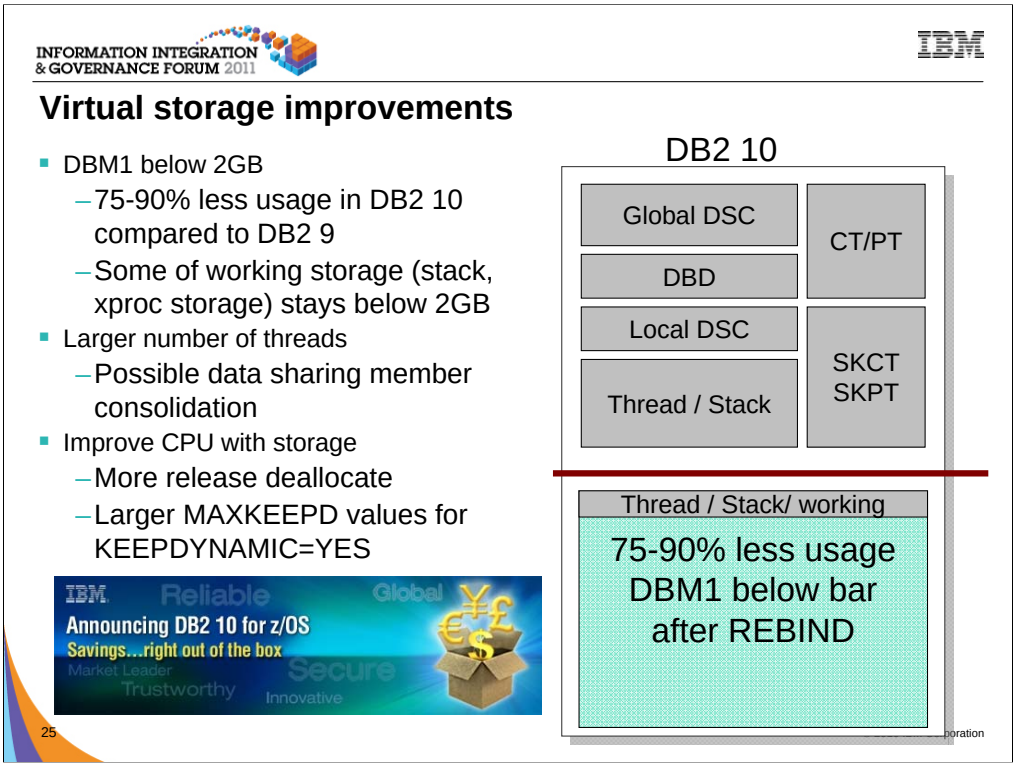

Laboratory measurements and early customer experience have shown substantial savings in the primary constrained address space, DBM1. Most measurements have shown 75% to 90% savings for the virtual storage in that address space below the 2 GB bar. Some EDMPOOL and some working storage remains below the bar.

This storage relief allows many more threads or concurrent users in a DB2 subsystem, allowing new possibilities for optimization.

Some customers will be able to consolidate data sharing members, saving on memory, CPU and administration time.

Other customers will be able to use the storage to improve service or to reduce CPU time more. Some common examples are expected to be use of RELEASE(DEALLOCATE) and larger amounts of dynamic statement cache.

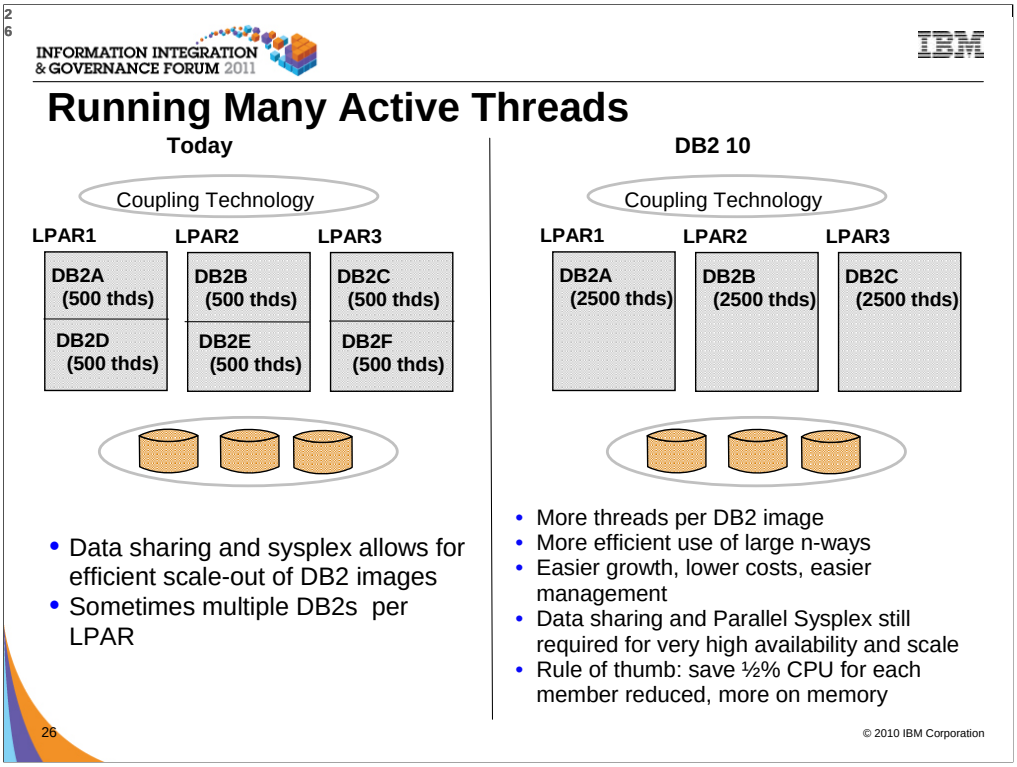

Customers are constrained by virtual memory to various degrees. This slide shows a relatively extreme situation experienced by some customers today. With a maximum of 500 threads (very dependent upon workload) in a DB2 subsystem, this customer is using two DB2 subsystems in the same data sharing group 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 11 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.

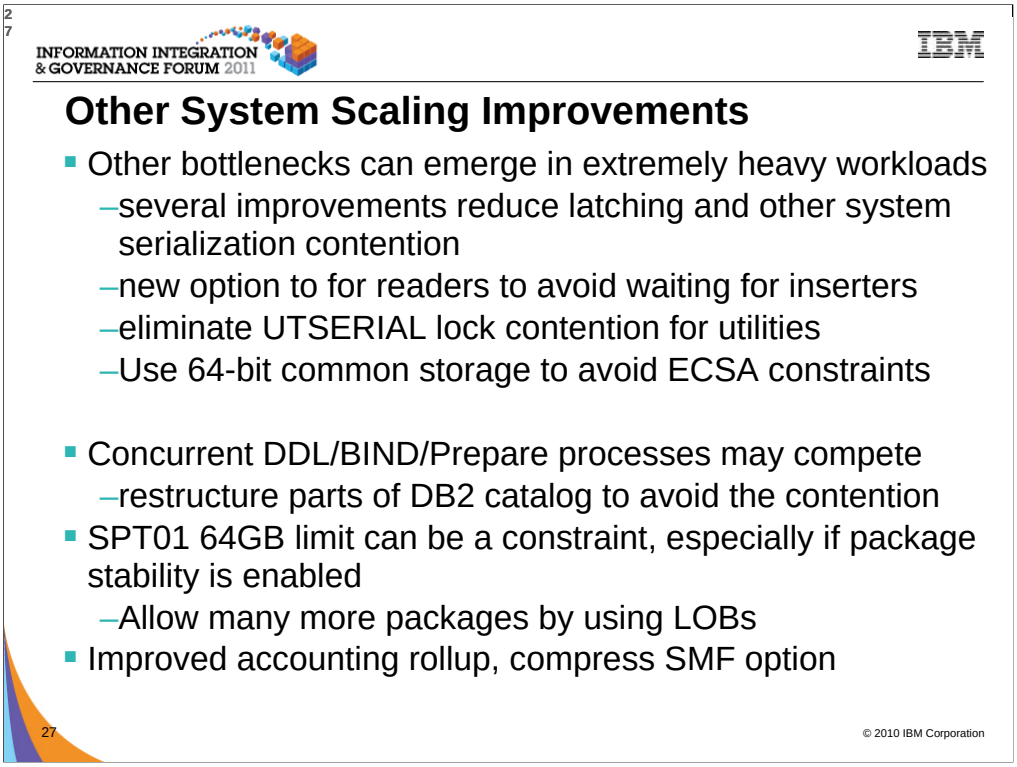

Increasing the number of concurrent threads will expose the next tier of constraints. DB2 10 will address a number of the next items, such as utility locking and catalog concurrency.

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.

Major changes in DB2 10 catalog & directory

- Improve availability and productivity
- Increase maximum size substantially
- Reduce contention: BIND, DDL, utilities
- Catalog changes: Remove links
  - Many more table spaces, partition by growth
  - Row level locking, reordered row format
  - CLOB and BLOB columns for long strings
    - Inline for performance
  - Online reorganization and check
  - More automatic: DB2-managed SMS-controlled

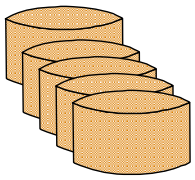

The DB2 catalog and directory are restructured in DB2 10 ENFM to improve productivity and availability. You'll see these improvements in NFM. The current size limits are increased substantially and contention among process like BIND, dynamic SQL, data definition and utilities is reduced. With more table spaces and more structures, more work is required for some process, such as BIND.

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. Inline LOBs are used for the performance improvements. 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.

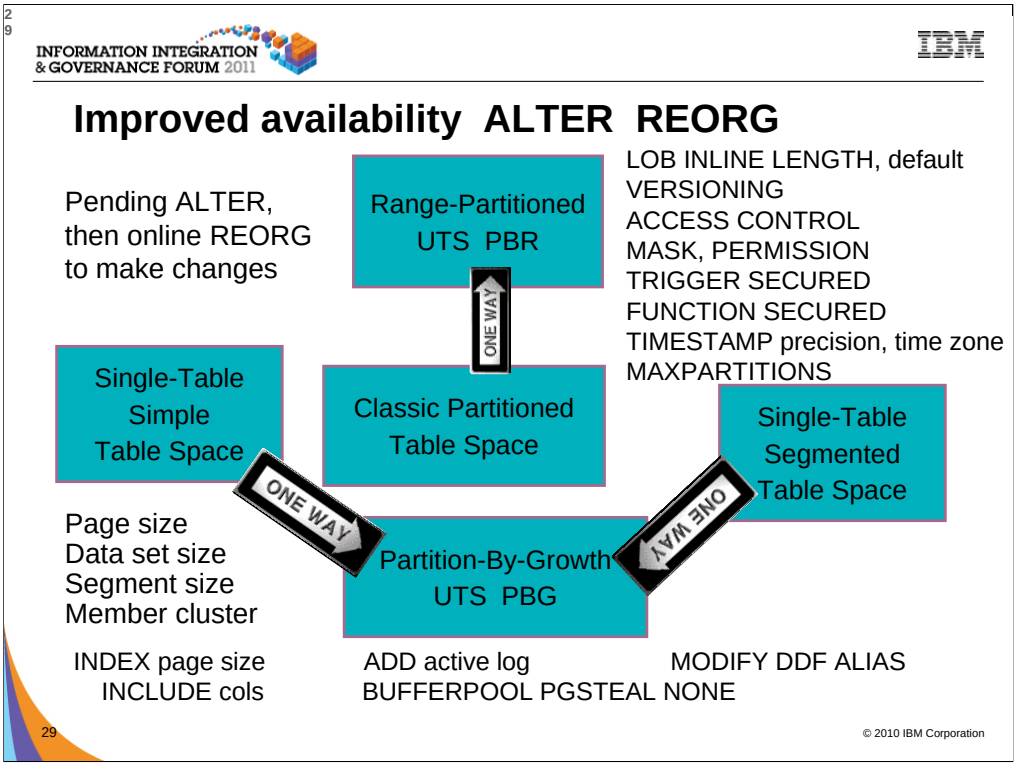

DB2 10 brings many new options for ALTER BUFFERPOOL, TABLE, INDEX, and TABLE SPACE. These are the changes in table space type in diagram form, adding the ability to change from single table segmented, simple or partitioned table spaces to universal table spaces. This release also adds the ability to 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.

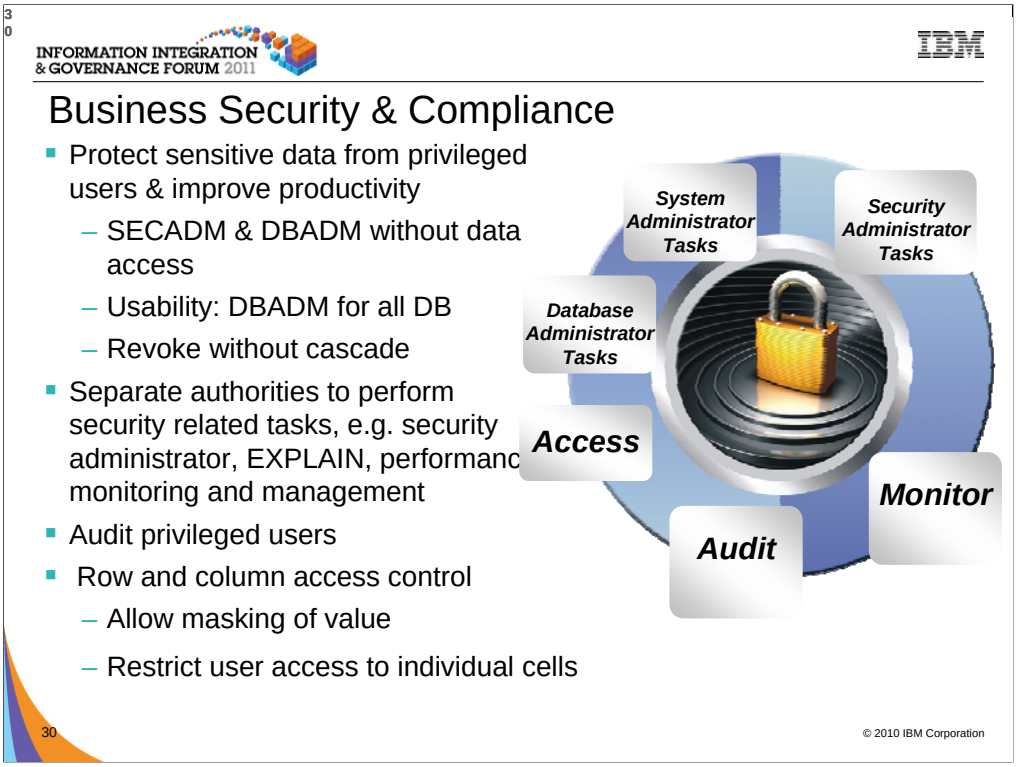

Customers are being pressed for a wide range of improved security and compliance. Data retention is a growing need. Protecting sensitive data from the privileged users and administrators is required. Separation of authority for security, access, and some common tasks, like EXPLAIN will help. Auditing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

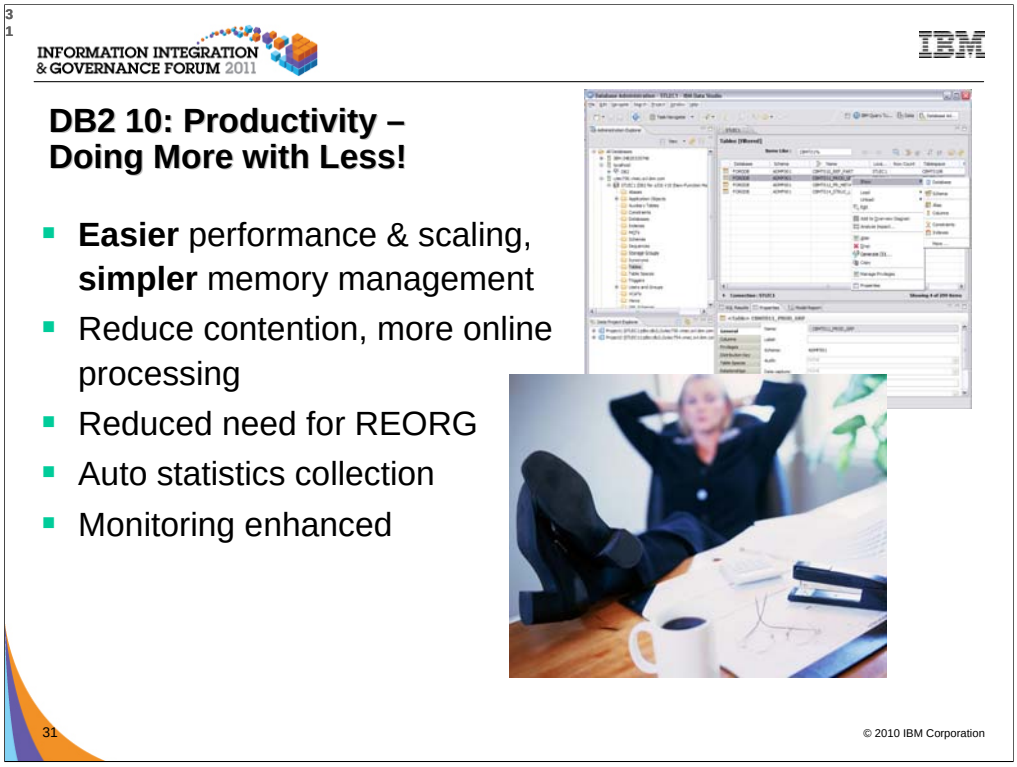

Some of the improvements come with Data Studio for application programming and administration – stronger cross-platform graphical interfaces, better integration with Java, improvements in the ability to develop and debug.

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.

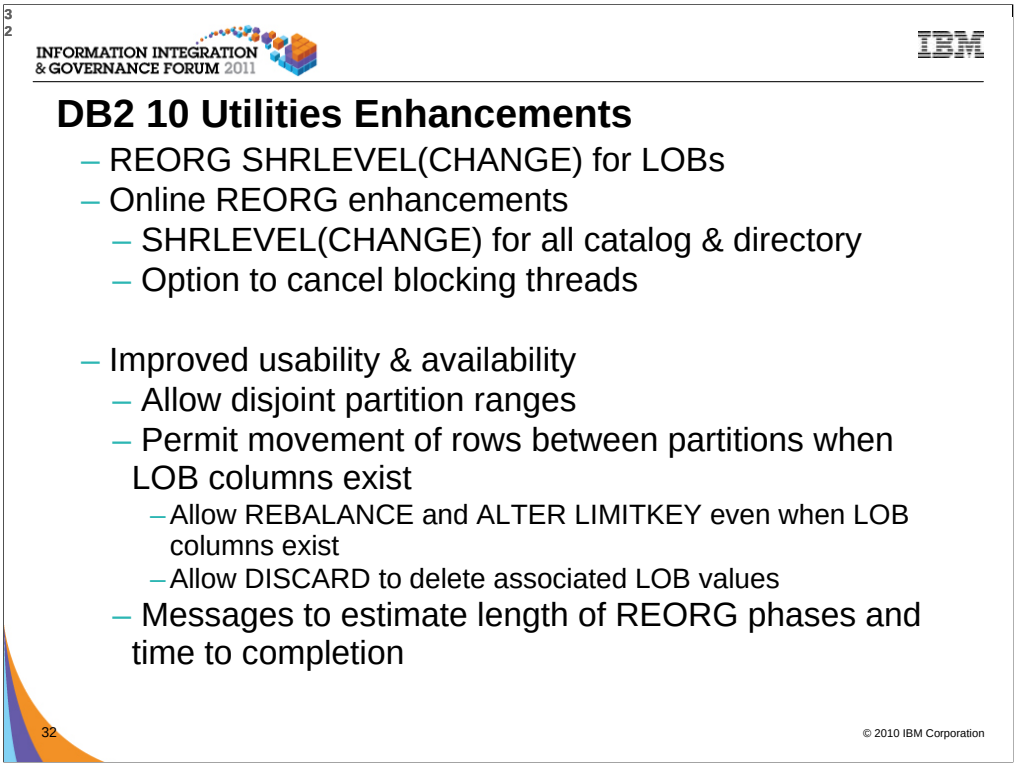

Continuous availability requirements continue to escalate. Large batch and maintenance windows are in the past. Those windows are being closed on the fingers of DBAs. DBAs increasingly need the ability to make all changes and to do all maintenance activities online or around the clock.

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.

Query Processing Enhancements

- **Performance Improvements**
  - Improved caching of dynamic SQL with literals
  - Safe Query Optimization
  - Aggressive View Merge
  - IN List Processing
  - SQL Pagination
  - Parallelism Enhancements
- **Access Path Stability**
  - Relief from package REBIND regression

Access path improvements deliver improved response time and reduced resources and simpler management for packages and queries. The performance improvements include better optimization for some common situations, improved caching for dynamic SQL statements which contain literals instead of parameter markers, and increased parallel processing. Removing some parallel restrictions provides faster response times and allows more use of zIIP.

Access path stability improvements help eliminate regression from a REBIND. The DB2 9 package management changes helped many customers reduce their fear of REBIND, and these improvements take the next step, extending and expanding the capability. New capabilities make the processing more efficient and easier to manage.

Query Enhancements

- CPU time reductions for queries, batch, & transactions
  - Complex predicate processing improvements
- SQL enhancements: Moving Sum, Moving Average, temporal, timestamp, implicit cast, SQL PL, …
- pureXML improvements
- Access improvements: Index include columns, Hash
- Optimization techniques
  - Remove parallelism restrictions; more even parallel distribution
  - Scalability: memory and latching relief allow more parallel
  - Increased zIIP use – parallel, prefetch, RUNSTATS
  - In-memory techniques for faster query performance
- Analysis: instrumentation, Data Studio & Optim Query Tuner

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.

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.

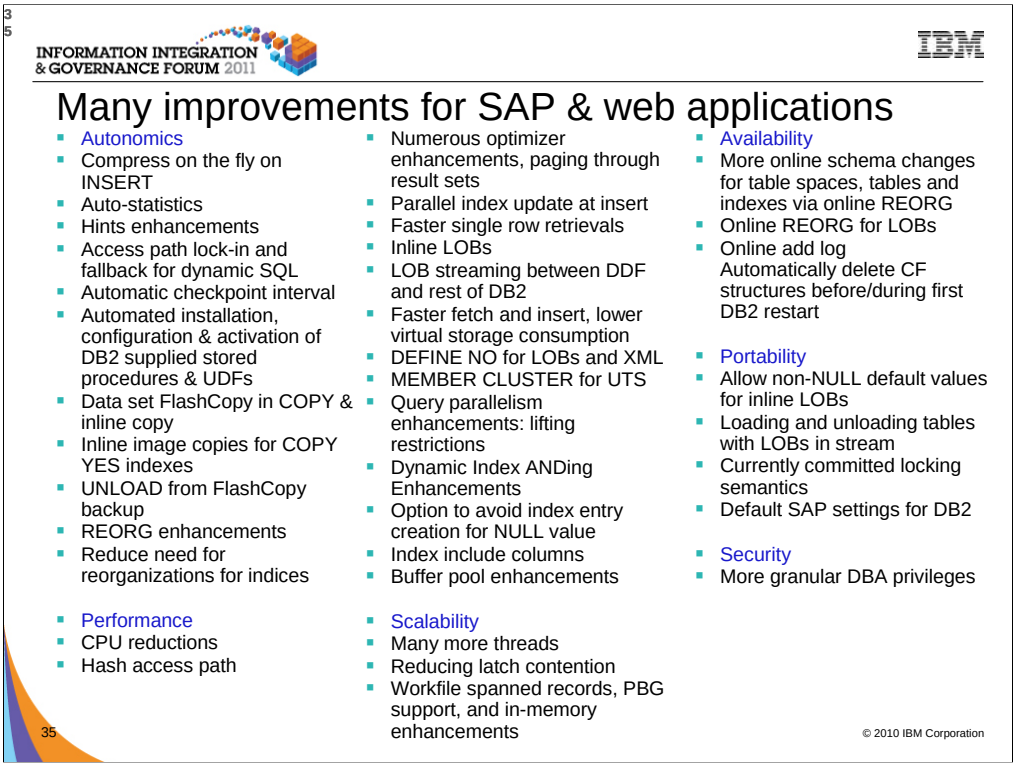

The enhancements cover many aspects of the database technology including new applications support, SQL enhancements, performance and scalability, continuous availability, data warehousing improvements as well as reducing the total cost of ownership.

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.

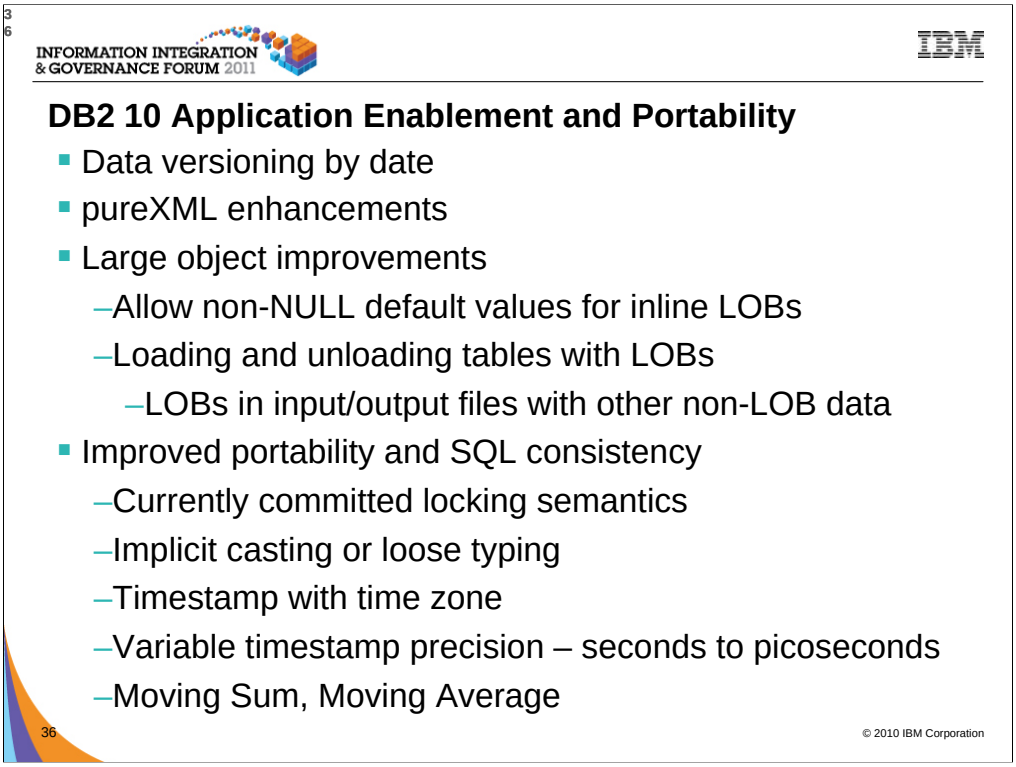

DB2 10 provides the ability to version your data by date. Using the timestamp picoseconds enhancements all the data within a table can have unique timestamps. This enhancement along timestamp time zone capabilities provides the application designer with options for tables that hold global data activities. This is an important feature for global financial and other global industry companies.

Large object are also improved by being able to put smaller LOBs that will fit on the same data page within the data row. This allows the row and the related LOB data to be on the same row and eliminate the extra I/Os to get the LOB data from its other table. This feature combined with the ability to use NULL or default values for the LOB gives additional reasons to have small LOBs in-line with their associated row data.

LOB handling is also improved as the LOB data can be included in the standard input/output files with other non-LOB data. This eliminates the hassles of the large amount of extra LOB files previously needed to support the loading and unloading of LOB data within your system.

DB2 10 also improves its compatibility and SQL consistency within the DB2 family and with other DBMS vendors. This allows any other DBMS vendor systems to be more easily ported to DB2 10 on System z eliminating availability, scalability and performance problems. This feature can quickly resolve your performance, maintenance and scalability for some of your UNIX systems, some SAP or other packaged software systems.

This compatibility also extends to the ability to implicitly cast unlike data types for easily moving or integrating data across application data types, program languages and platforms.

This portability is also reflected in the new package level parameters to control whether the application looks at only currently committed data or not. This improves application concurrency and provides flexibility within your application design for when the system should ignore rows that are in the process of being inserted and only use currently committed rows

The Timestamp with Time Zone and the pico seconds of the timestamps features are a great way to set up a fact table within a data warehouse or business intelligence database. These components along with the new SQL capabilities for calculating a moving sum or moving average are additional DB2 10 capabilities that make it easier for operational business intelligence applications.

INFORMATION INTEGRATION & GOVERNANCE FORUM 2011

IBM

Versioned data or Temporal Data

- Table-level specification to control data management based upon time
- Two notions of time:
  - System time: notes the occurrence of a data base change
    - "row xyz was deleted at 10:05 pm"
    - Query at current or any prior period of time
    - Useful for auditing, compliance
  - Business time: notes the occurrence of a business event
    - "customer xyz's service contract was modified on March 23"
    - Query at current or any prior/future period of time
    - Useful for tracking of business events over time, application logic greatly simplified

- New syntax in FROM clause to specify a time criteria for selecting historical data

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.

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

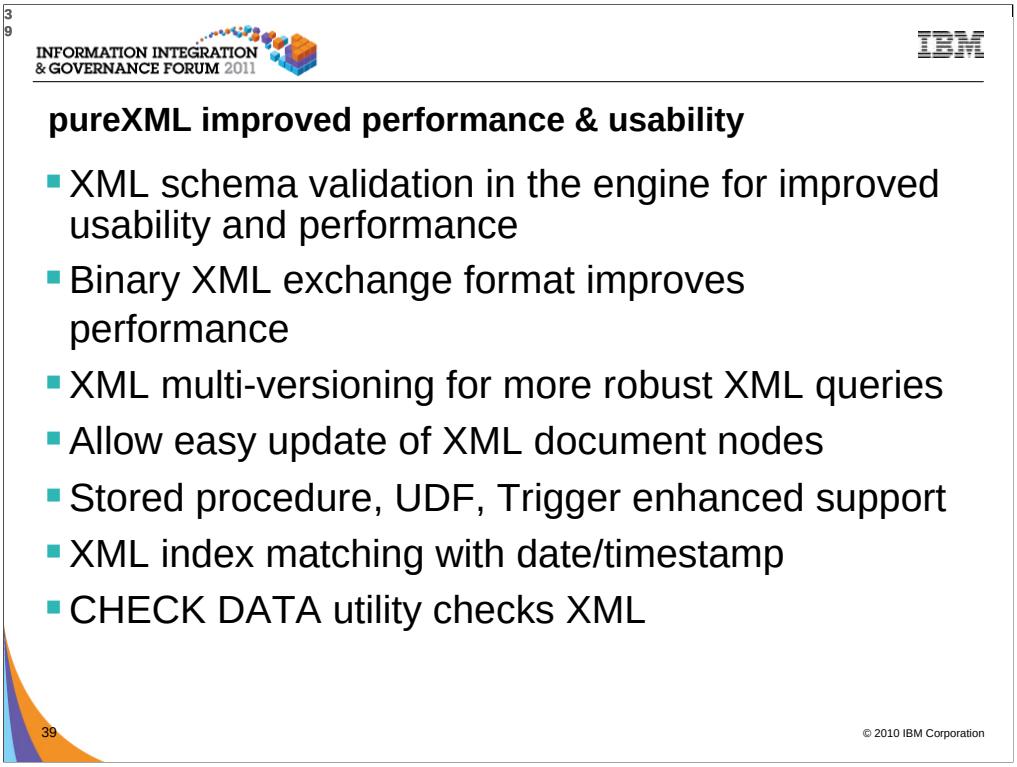

A range of XML improvements delivers a strong release 2 of the pureXML function. Customers use of DB2 9 pureXML shaped this delivery of improved performance and usability.

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.

During the early programs for DB2 10 for z/OS, more than 100 companies were involved, getting their applications and tools ready for customers. Most vendors are ready today, but please talk to your vendors to find out what releases and fixes you need.

INFORMATION INTEGRATION & GOVERNANCE FORUM 2011

IBM

Jump into DB2 10! The water's fine.

DB2 V8 DB2 9 DB2 10

Key Questions are WHEN? and HOW?

41

Is your current fish bowl getting constrained? What is limiting you? Is it CPU? Virtual storage? Latching? DB2 catalog and directory? Concurrent utilities? Are you currently running DB2 9? V8? V7? Should you migrate to DB2 10?

The answer is a definite Yes. The question is not so much whether to migrate as when and how to migrate. If you are running DB2 9 today, then DB2 10 is in your future, giving you more room to grow, with fewer limits, lower costs, and more for less. If you are running DB2 V8 today, then you have a choice of jumping to DB2 9 or directly to DB2 10. So the key question is, "When should I migrate to DB2 10?"

See the migration paper and presentation for a lot more information.

ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/migration/upgrading-to-db2z10-miller.pdf
ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/migration/db2-10-migration-planning-miller.pdf

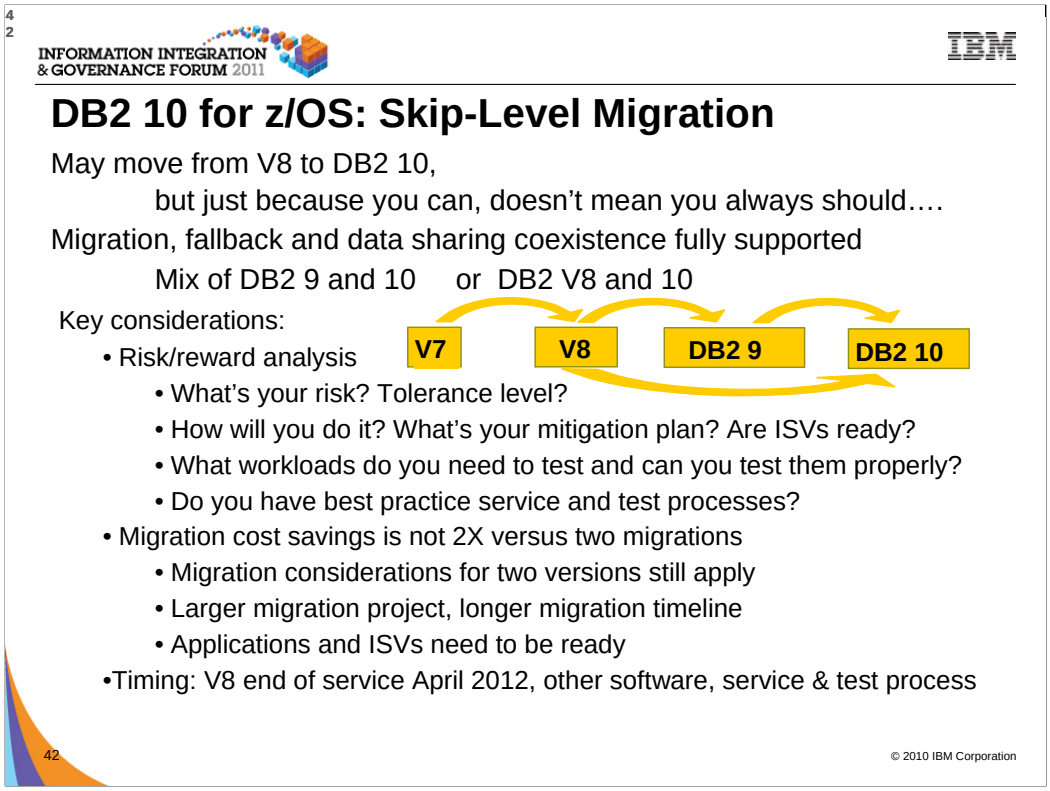

DB2 10 supports migration from DB2 9 NFM or from V8 NFM. Customers not yet running V8 or DB2 9 should plan to migrate to DB2 for z/OS V8 NFM as preparation for a migration to DB2 10. We estimate that about one customer in five migrated using a skip version technique for V5 to V7, and we'll see a similar fraction this time. The key value for skip customers on V7 or new on V8. DB2 V8 end of service is announced as April 2012.

When should I migrate to DB2 10? At this point, DB2 10 is in beta. Some of the key information for making this decision is not yet known. They include the date for DB2 10 general availability, V8 extended service, and pricing. This information will come in announcements. While DB2 10 is expected to be better than prior versions, it will have maturity and service delivery like other software, with more defects at first fewer as the software matures. Determining when the software is ready for a specific customer and when the customer is ready for the software depends upon the specific customer resources for testing, prior experience, and the value for the improvements versus the need for stability. Many customers depend upon tools or other software, and having that software work with DB2 is a prerequisite. When this information is known, we can answer the question.

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 than for a single version. 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 extended service on V8.

http://www.ibm.com/support/docview.wss?uid=swg21006951

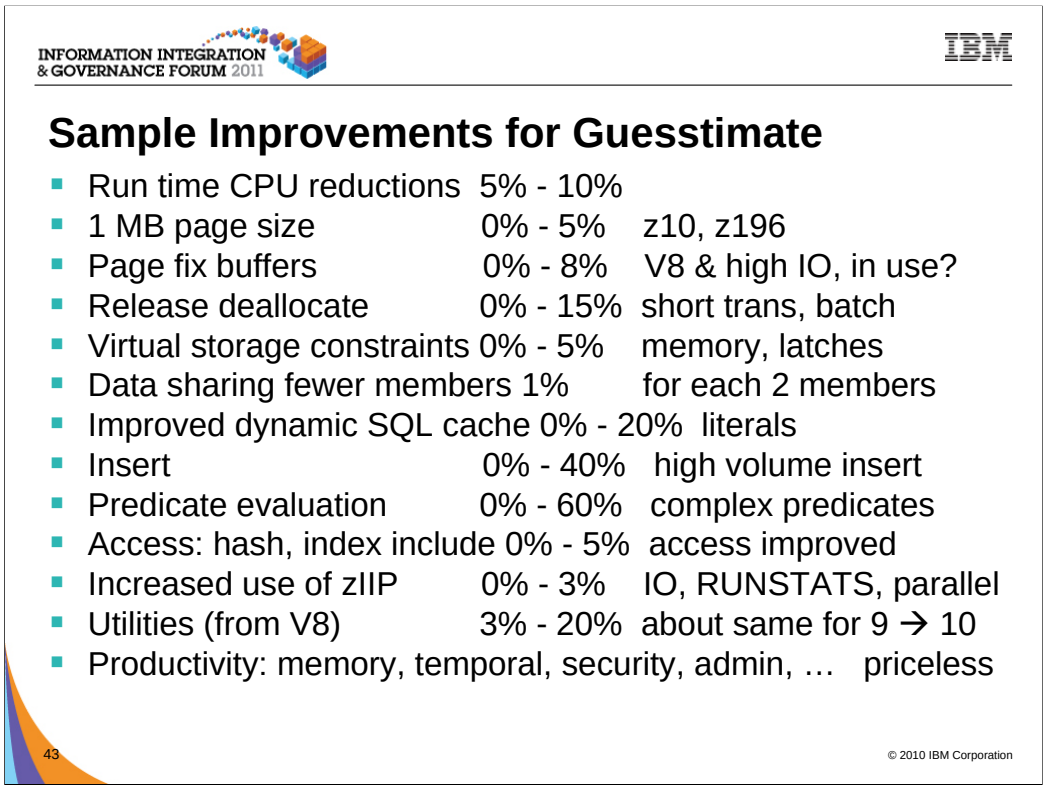

The common range for CPU reductions is very wide. Understanding the magnitude of the gains for each individual customer and the breadth of applicability are important. Which gains provide the peak workload, determining the charges in most pricing options?

For most customers looking at a general workload, the expectation ranges from 5% to 10%. Transactions with only a few SQL statements don't get the above, but can benefit from the increased ability to use release(deallocate). The change to use 1 MB hardware page sizes can be up to 5%, if you have a z10 or z196 and configure the LFAREA. 1 MB page sizes also depends upon page fixed buffers. Many customers have not taken advantage of the V8 function, while it can save up to 8% of the CPU time if the amount of IO is high.

Virtual storage constraint relief is generally up to 5%, but extreme cases can save much more. Estimate saving ½% of CPU for each active member removed from data sharing.

Queries with many predicates can improve up to 60%. Many customers reported insert improvements up to 40%, and larger for V8.

Increased use of zIIP comes from prefetch read, deferred write, most options of the RUNSTATS utility, and increased parallelism, and can give up to 3%.

The DB2 10 utilities are roughly the same as those in DB2 9 for CPU time, but much better than DB2 V8 utilities.

The biggest benefits from DB2 10 are the productivity improvements in memory management, temporal SQL, security and administration.

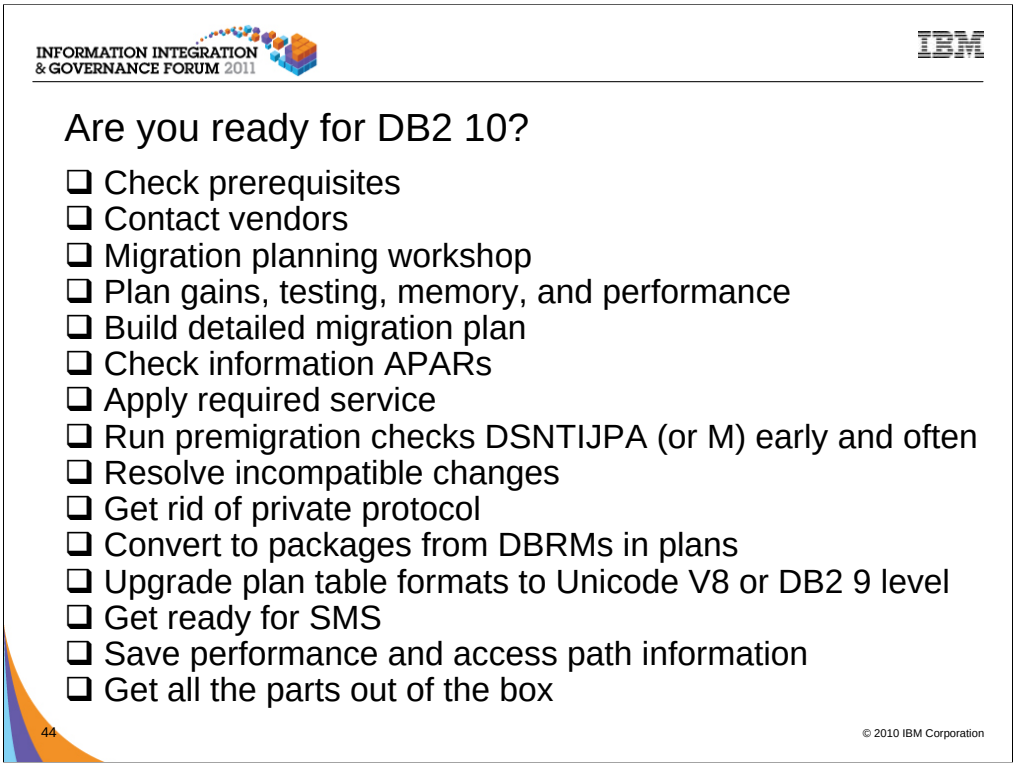

This is a checklist on getting ready for DB2 10. You will want to get this work done before starting to work with the new code. When this checklist is completed, you can move to the checklists in the DB2 Installation Guide.

See the checklist in the DB2 10 Upgrade paper and presentation.

ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/migration/upgrading-to-db2z10-miller.pdf
ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/migration/db2-10-migration-planning-miller.pdf

The Migration Planning Workshop (MPW) offering enables customers to understand the breadth of features delivered in DB2 for z/OS versions. Many resources are brought together in a single offering. The migration process is explained. Customers leave the session with materials they can use to start their installation / migration immediately, or in the future. Their questions are given attention, typically in a small group setting. While most MPWs are small, multi-company events, they can be delivered for individual companies or even user groups. Get 2 page trifold:

ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/migration/db2-10-migration-planning-workshop-trifold.pdf
ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/migration/db2-9-migration-planning-workshop-trifold.pdf

Who should attend? An MPW contains content that is appropriate to: Application Developers, Database Administrators, System Administrators, Architects, IT decision makers, Project Managers. Customers should contact their IBM representative for more information about the Migration Planning Workshop.

The morning session provides information for all parties. The afternoon is more migration focused, and therefore a better fit for Database Administrators, System Administrators, and Project Managers. A typical agenda provides a DB2 for z/OS overview in the morning and DB2 migration (preparations, planning, process) in the afternoon.

What you can expect? • An understanding of the features delivered with DB2 and how they can benefit your enterprise. • Clarity of the migration process. • References for many subjects, including: Migration, Fallback, Prerequisites & preparations. You will leave with Presentation materials, Checklists, Project plan framework, Related documentation, Networking and Contacts.

Important features of information center

- Find helpful usage instructions in the “Information center home” section.
- Easily send feedback by clicking the **Feedback** link at the bottom of any topic.
- Conveniently download the PDF version of the information from the link at the bottom of any topic. *(Look for the PDF icon!)*
- Efficiently search for the information you need by using the search features.

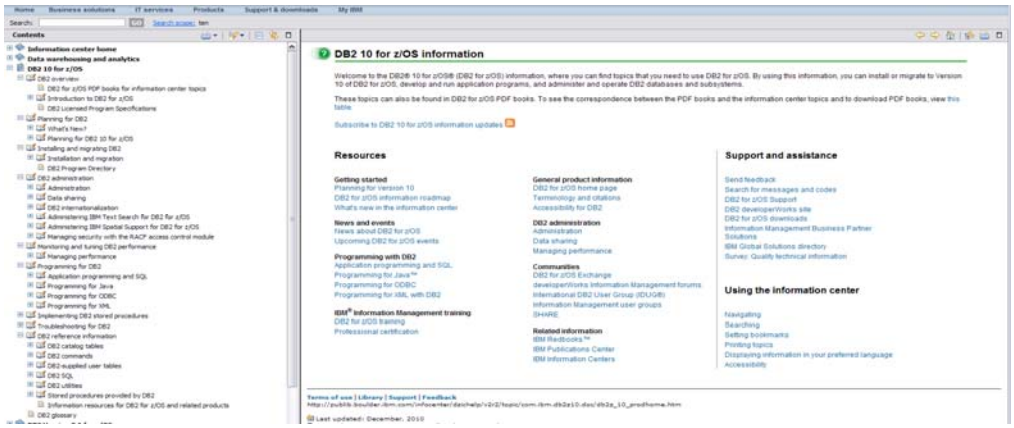

I am a big fan of the Information Center and use it many times a day. I know the books very well, but I find more and find it faster with the search. It’s easy to answer questions with an RTFW by sending the URL of the page with an answer.

http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp

The Information Center is also easy to improve if it’s not right, or you can’t find what you need.

We welcome your feedback about the DB2 10 for z/OS information. Send documentation feedback by clicking the **Feedback** link at the bottom of any information center topic. Feedback is received directly by the information development team.

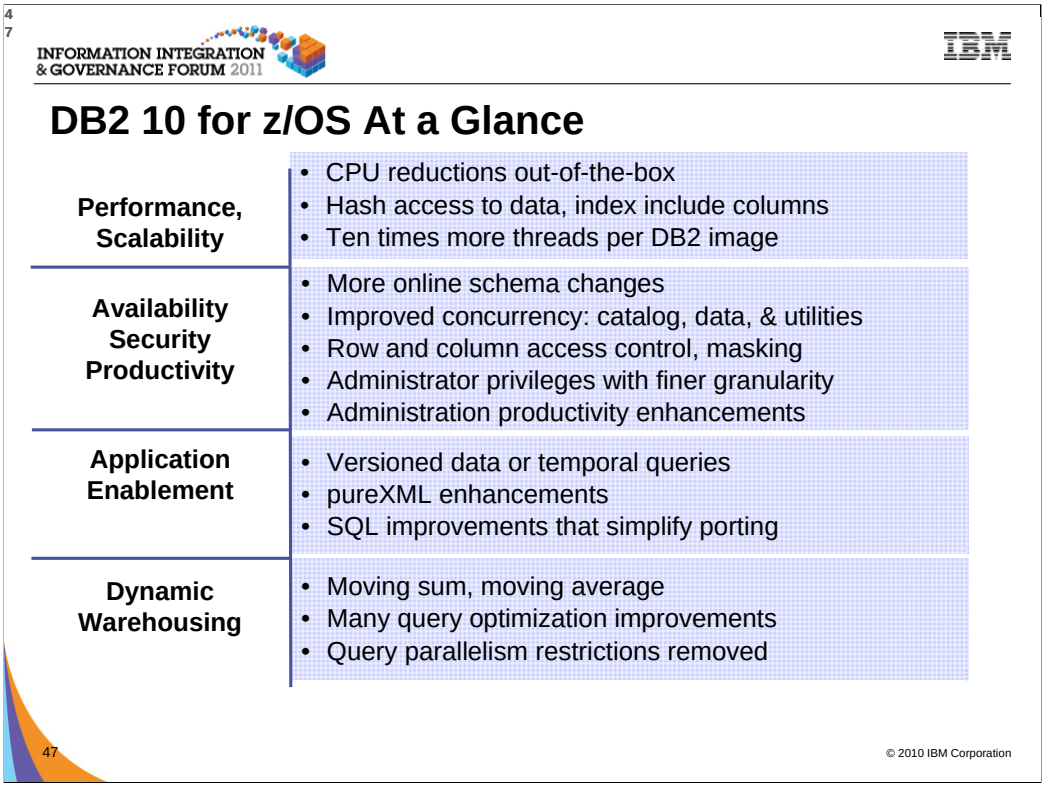

DB2 10 for z/OS provides the best reduction in CPU for transactions and batch for 22 years, since V2R1. We expect most customers to reduce CPU times between 5% and 10%. Applications which can take advantage of additional benefits, such as hash access, can have larger CPU and memory reductions. Scalability 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, 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.

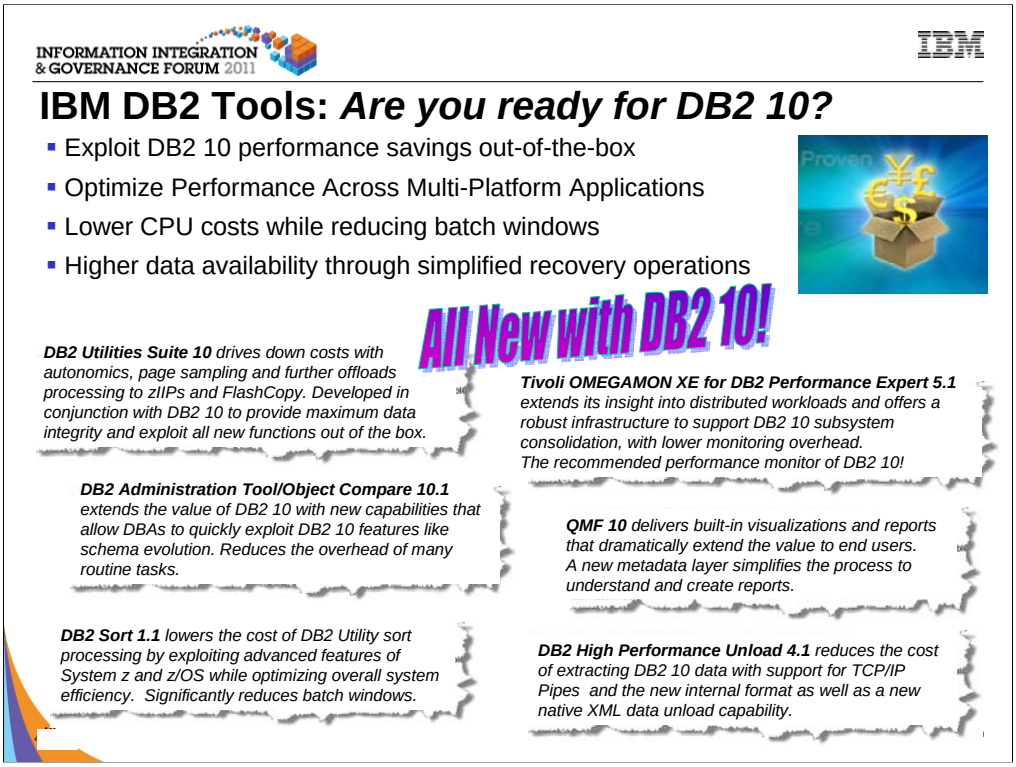

Utilities – Focus on eliminating outages, improving performance, reducing resource consumption, reduce complexity and improve automation. Day 1 utility support for DB2 10 function. Some of the key enhancements are:
•Flashcopy support at data-set level for COPY, RECOVER, REORG, LOAD, & REBUILD INDEX, For ex: can create an image copy data set to be a transaction-consistent image copy data set with no application outage
•Significant enhancement to REORG to reduce outages, including Support new Online Schema in DB2 10 and hash table, both before and after conversion to hash format, Improved performance for part-level REORG w/ non-partitioning indexes and REORG INDEX to reduce ET, SHRLEVEL CHANGE support for LOB table space for both LOY YES/NO w/ no mapping table required
•Major enhancements to RUNSTATS incl. zIIP support, autonomic features, page & auto-sampling rates
Admin Tool: DBA-managed performance improvements - Easily migrate existing tables to hash access, Manage new Security models and autonomic statistics collection, Rollout Application BiTemporal Data – "as of"
HPU - new internal format to UNLOAD (perf. Feature) to as well as TCPIP support via USS pipe and native XML UNLOAD support
Tivoli OMEGAMON XE for DB2 Performance Expert, v5.1 W5655-W37 introduces an end-to-end response time measurement capability surfacing DB2 for z/OS SQL metrics, making it IBM's most comprehensive DB2 application performance assessment tool. Use it to evaluate the efficiency of, and optimize performance of your DB2 for z/OS DBs across your heterogeneous application environment. It offers all the capabilities of Tiv. OMEGAMON XE for DB2 PM on z/OS and Buffer Pool Analyzer for z/OS, in addition to its own unique capabilities.
Utilities – Focus on eliminating outages, improving performance, reducing resource consumption, reduce complexity and improve automation. Day 1 utility support for DB2 10 function. Some of the key enhancements are:
•Flashcopy support at data-set level for COPY, RECOVER, REORG, LOAD, & REBUILD INDEX, For ex: can create an image copy data set to be a transaction-consistent image copy data set with no application outage
•Significant enhancement to REORG to reduce outages, including Support new Online Schema in DB2 10 and hash table, both before and after conversion to hash format, Improved performance for part-level REORG w/ non-partitioning indexes and REORG INDEX to reduce ET, SHRLEVEL CHANGE support for LOB table space for both LOY YES/NO w/ no mapping table required
•Major enhancements to RUNSTATS incl. zIIP support, autonomic features, page & auto-sampling rates
Admin Tool: DBA-managed performance improvements - Easily migrate existing tables to hash access, Manage new Security models and autonomic statistics collection, Rollout Application BiTemporal Data – "as of"
HPU - new internal format to UNLOAD (perf. Feature) to as well as TCPIP support via USS pipe and native XML UNLOAD support
Tivoli OMEGAMON XE for DB2 Performance Expert, v5.1 W5655-W37 introduces an end-to-end response time measurement capability surfacing DB2 for z/OS SQL metrics, making it IBM's most comprehensive DB2 application performance assessment tool. Use it to evaluate the efficiency of, and optimize performance of your DB2 for z/OS DBs across your heterogeneous application environment. It offers all the capabilities of Tiv. OMEGAMON XE for DB2 PM on z/OS and Buffer Pool Analyzer for z/OS, in addition to its own unique capabilities.

INFORMATION INTEGRATION & GOVERNANCE FORUM 2011

IBM

Questions?

Watch for upcoming white papers and conferences

IOD IDUG Share

I hope we have lots of questions, but thank you in any case. We have provided many additional notes in the slides and additional slides at the end of this presentation.

We have a new white paper on DB2 10 from Dave Beulke now and expect to have another from Julian Stuhler of Triton Consulting in the future.

https://www14.software.ibm.com/webapp/iwm/web/signup.do?source=sw-infomgt&S_PKG=db2_zos_reduce_costs

DB2 10 is a hot topic at upcoming conferences, so please plan to attend IOD, IDUG or Share.

http://www.ibm.com/software/data/db2/zos/events.html

http://www.idug.org

http://www.share.org

http://www.ibm.com/software/uk/data/conf/

INFORMATION INTEGRATION & GOVERNANCE FORUM 2011

IBM

DB2 9 and 10 IBM Redbooks Publications

1. DB2 10 Technical Overview SG24-7892 new
2. Extremely pureXML DB2 10 & 9 SG24-7915 new
3. DB2 10 Performance Topics coming soon
4. DB2 9 Technical Overview SG24-7330
5. DB2 9 Performance Topics SG24-7473
6. DB2 9 Stored Procedures SG24-7604
7. Serialization and Concurrency SG24-4725-01
8. Distributed Functions SG24-6952
9. Utilities SG24-6289-01
10. DB2 and Storage Management, SG24-7823
11. Index Compression with DB2 9 for z/OS redp4345
12. SQL Reference for Cross-Platform Development
13. Enterprise Database Warehouse, SG24-7637
14. 50 TB Data Warehouse on System z, SG24-7674
15. LOBs with DB2 for z/OS SG24-7270
16. Deploying SOA Solutions SG24-7663
17. Enhancing SAP - DB2 9 SG24-7239
18. Best practices SAP BI - DB2 9 SG24-6489-01
19. Data Sharing in a Nutshell, SG24-7322
20. Securing DB2 & MLS z/OS SG24-6480-01
21. Data Sharing: Dist Load Balancing & config. redp4449
22. Packages Revisited, SG24-7688
23. Ready to Access Solid-State Drives redp4537
24. Buffer Pool Monitoring & Tuning redp4604
25. Securing & Auditing Data SG24-7720

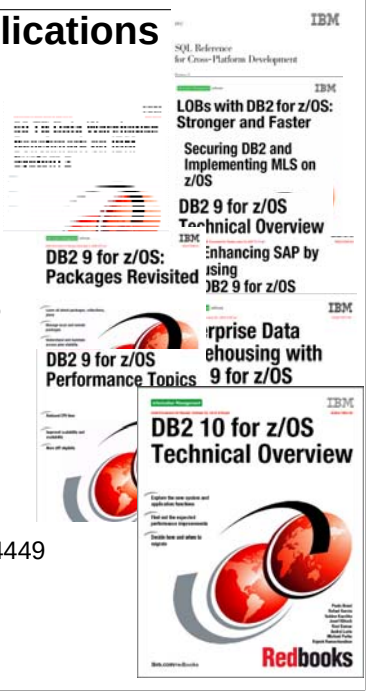

50

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+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 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. 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 Architecture 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/abstracts/sg246289.htm
29. DB2 9 and Storage Management SG24-7823 http://www.redbooks.ibm.com/abstracts/sg247823.htm

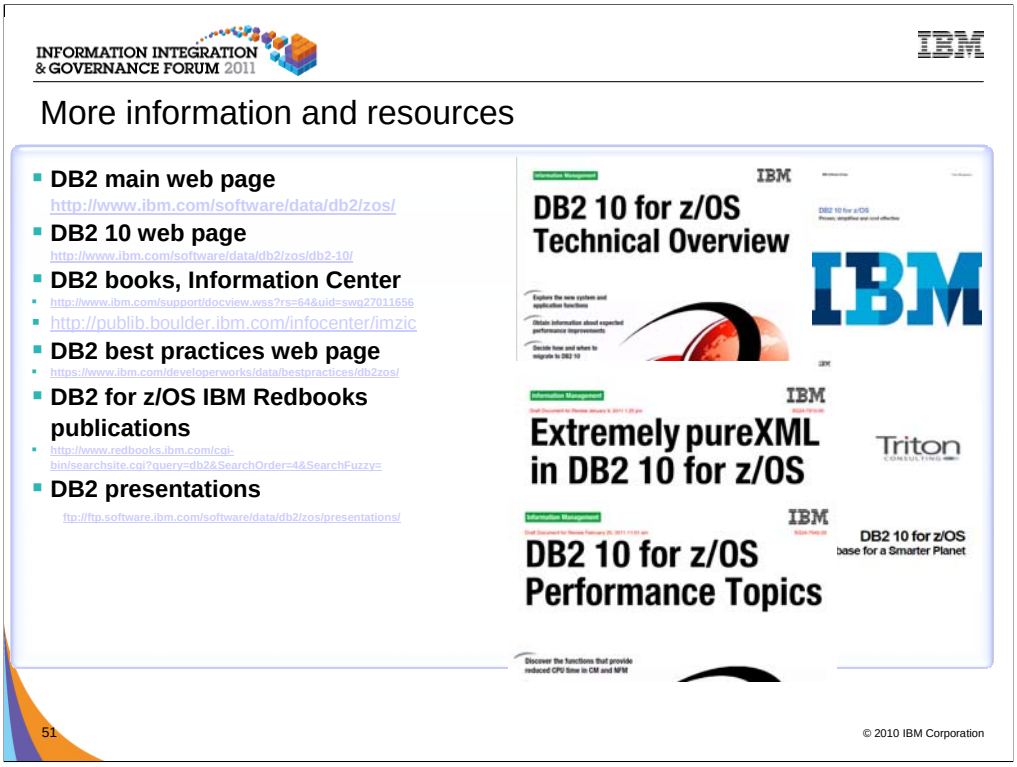

Here are some resources for a discussion of business value in DB2. White papers are useful for a quick summary, and IBM Redbooks publications provide more detailed technical discussion.

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

ftp://ftp.software.ibm.com/software/data/pubs/papers/DB2_for_zOS_V9_Business_Value_White_Paper.pdf
http://www.ibm.com/software/os/systemz/newsletter/mainstreamed11_uk.html

DB2 10 Technical Overview SG24-7892,
http://www.redbooks.ibm.com/abstracts/sg247892.html

Extremely pureXML in DB2 10 SG24-7915,
http://www.redbooks.ibm.com/abstracts/sg247915.html

The Business Value of DB2 for z/OS, SG24-6763,
http://www.redbooks.ibm.com/abstracts/sg246763.html

DB2 9 for z/OS Technical Overview, SG24-7330, chapter 2 on System z synergy, http://www.redbooks.ibm.com/abstracts/sg247330.html

DB2 9 for z/OS Performance Topics, SG24-7473, performance on z10,
http://www.redbooks.ibm.com/abstracts/sg247473.html

See more on the next pages …

Here are the latest presentations and papers on the web.

DB2 10 overview

ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/overview/db2-10-overview-miller.pdf

ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/overview/db2-10-questions-answered.pdf

DB2 10 upgrade paper and presentation

ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/migration/upgrading-to-db2z10-miller.pdf

ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/migration/db2-10-migration-planning-miller.pdf

DB2 10 performance

ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/performance/db2-10-performance-share-2011-miller.pdf

SAP Whitepapers DB2 10 for z/OS is certified for SAP NetWeaver 7.30 and SAP R/3 4.6

- **SAP article on DB2 10 *(published by SAP)*** http://www.sdn.sap.com/irj/sdn/db2
- **SAP Best Practice Guide for Migrating to DB2 10 for z/OS *(published by SAP)*** https://websmp207.sap-ag.de/~sapidb/011000358700001414122010E
- **(Updated) Business Continuity Guide for Running SAP on System z – based on DB2 10 for z/OS, DB2 Connect 9.7 FP3a, SAP NetWeaver 7.10 and Tivoli Automation for z/OS V3.3** http://publibfp.dhe.ibm.com/epubs/pdf/iapacs03.pdf
- **DB2 10 for z/OS with SAP on IBM System z Performance Report** – new techdocs white paper http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101845
- **DB2 10 for z/OS – Optimized for SAP** – http://cattail.boulder.ibm.com/cattail/?source=s#view=andreas.r.mueller@de.ibm.com/files/3198290001883DDBA202FBE4093F23B6
- **SAP on DB2 10 for z/OS - Being More Productive, Reducing Costs and Improving Performance** – http://www.sdn.sap.com/irj/sdn/db2?rid=/library/uuid/005c6b33-aaf0-2d10-fcbb-b42e89ac5791

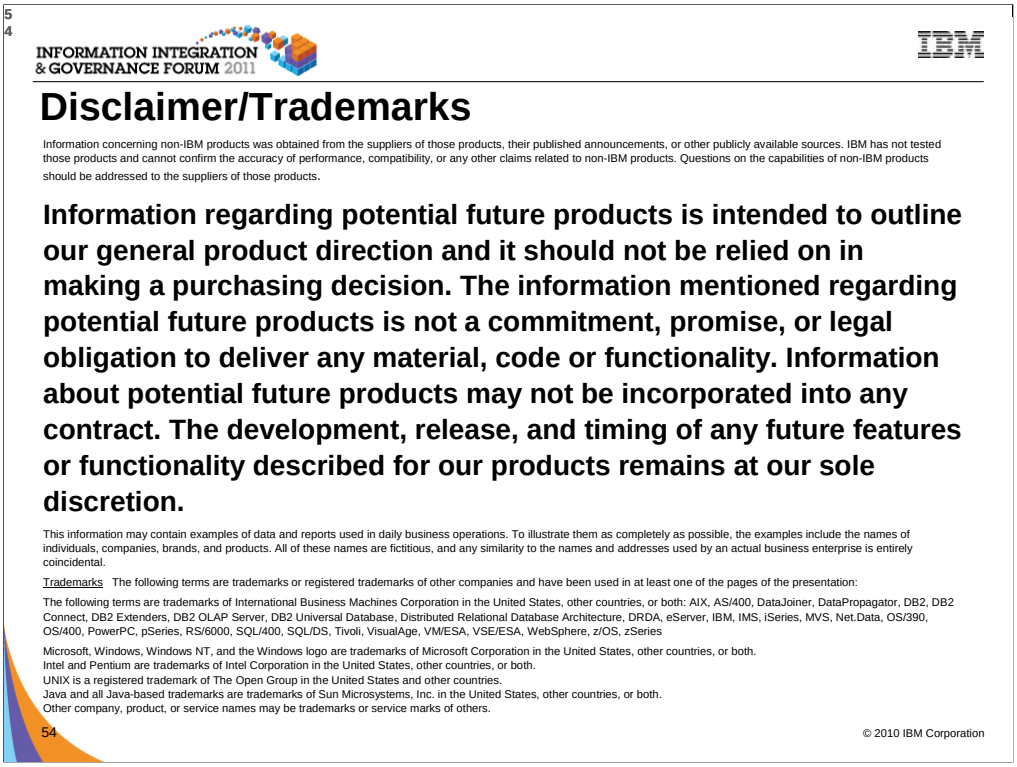

See the web for current information, the announcement, …

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://www.ibm.com/support/docview.wss?uid=swg21006951

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

This slide stack has too many slides for a one hour presentation. Here are some suggestions.

1 hour slides 1 – 7, 9, 13 - 15, 18 - 25, 29 – 31, 37, 41, 47, 48
30 minute overview slides 1 – 7, 13, 14, 25, 29 – 31, 37, 41, 47, 48
skip 13 and 14 if zEnterprise is not useful
1 slide slide 2, 3, 4, or 47

Outline add from these sections where the customer is interested.

Overview and customer experience	2 – 7
Performance	8 – 10
System z synergy	11 – 18
Modes and performance	19 – 22
Hash	23 – 24
Scalability	25 – 28
Availability, Security , Productivity	29 – 32
Application & SQL	33 – 40
Migration strategy	41 – 46
At a glance and tools	47 – 48
Questions and pointers for more	49 – 55