



What's New in DB2 10 for z/OS?

Monday, February 28, 2011: 1:30 PM-2:30 PM

Room 211A (Anaheim Convention Center)


Roger Miller and Ann Hernandez *IBM Silicon Valley Lab*






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.

DB2 for z/OS
The most robust and cost effective data server



	DB2	DB2 9	DB2 10
 <p>Efficiency</p> <ul style="list-style-type: none"> Deep synergy with System z HW Compression Consolidation 	<ul style="list-style-type: none"> Up to 20% utility CPU savings Compress indexes, save 50% disk More CPU on specialty engines 	<ul style="list-style-type: none"> Save up to 5-10% CPU batch & transactions out-of-the-box (rebind) On-the-fly data Compression Temporal data support Skip-level migration 	
 <p>Resilience</p> <ul style="list-style-type: none"> Unmatched availability Unparalleled security Industry leading reliability 	<ul style="list-style-type: none"> Flexible context and role security Expanded online schema changes Volume level backup & recovery 	<ul style="list-style-type: none"> Ten times more concurrent users More online schema changes More granular access control 	
 <p>Growth</p> <ul style="list-style-type: none"> Near-linear scalability Optimized for SOA Flexible development Warehousing capabilities 	<ul style="list-style-type: none"> Seamless integration of XML and relational Improved SQL Partition by growth OLAP expressions 	<ul style="list-style-type: none"> Enhanced query parallelism More SQL compatibility Improved pureXML and SQL PL 	

V8 out of service April 2012



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.

5




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.

DB2 10 for z/OS - Feedback from Beta Customers



We are really thrilled about "Temporal Data" feature – this feature has the potential to significantly reduce overheads. We have estimated that 80% of our existing temporal applications could have used "the DB2 10 temporal features" instead of application code - this feature will drastically save developer time, testing time – and even more importantly make applications easier to understand so improve business efficiency and effectiveness

Frank Petersen -System Programmer BANKDATA

bankdata

"We use very large buffer pools – some of them up to 3.2GB in size. We rely on efficient access to buffered data and any saving in the cost of accessing that data will be very beneficial."

Philipp Nowak, BMW DB2 Product Manager

We have measured a 38% reduction in CPU and a 7% reduction in suspend time for some heavy insert workloads in a data sharing environment. That's a significant saving which provides immediate business benefit."

Philipp Nowak, BMW DB2 Product Manager

SHARE
In Anaheim
2011

What Customers are Saying



"DB2 10 enhances our ability to support our rapidly growing workloads while delivering some very valuable new function with immediate business benefits."

Paulo Sahadi, Senior Production Manager, Information Management Division, Banco do Brasil

As a multi-national corporation, we must adhere to strict local audit requirements. The security and administration capabilities in DB2 10 are a key driver for us to move to this version.



"In addition to the cost savings, DB2 10 for z/OS offers a far superior data server environment than Oracle database"

**Manuel Gomez Burrierl
CECA (Spanish Bank Federation)**



Some Beta Customer Performance Feedback

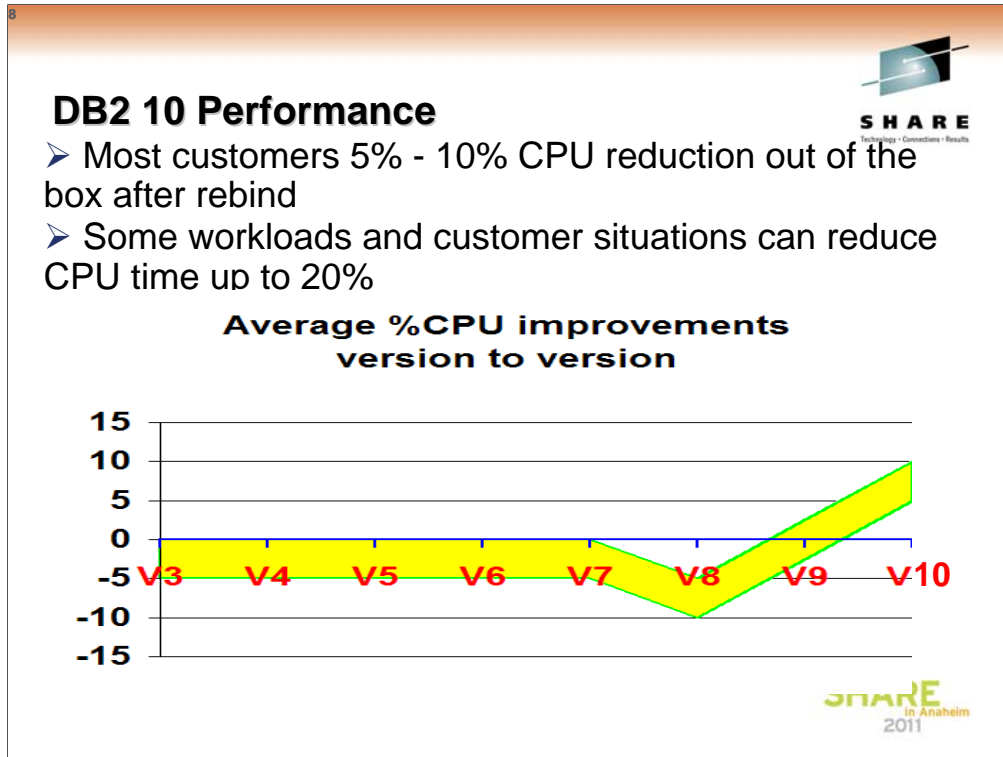


Workload	Results
Customer1: Distributed Concurrent Insert	50% DB2 elapsed time reduction; 15% chargeable CPU reduction after enabling high perf DBAT
Customer2: CICS online transactions	Approx. 7% CPU reduction in DB2 10 CM after REBIND, Another 4% reduction with 1MB page usage
Customer3: CICS online 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

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



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 its 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.

Preliminary Measurements of IBM Relational Warehouse Workload (IRWW) with data sharing

Base DB2 9 NFM REBIND with PLANMGMT(EXTENDED)

- DB2 9 NFM → DB2 10 CM without REBIND measured 3.7% CPU reduction from DB2 9
- DB2 10 CM REBIND getting same access path measured 7.4% CPU reduction from DB2 9
- DB2 10 NFM measured same 7.4% CPU reduction from DB2 9
- DB2 10 CM or NFM with RELEASE(DEALLOCATE) measured additional 10% CPU reduction from DB2 10 NFM RELEASE(COMMIT)

DB2 10 CPU Savings Relative to DB2 9 New Function Mode
IBM Relational Warehouse Workload with Data Sharing



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)

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



CPU reductions for transactions, queries, and batch

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

Scales with less complexity and cost

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

Improved operational efficiencies and lower administration cost

- Automatic diagnostics, tuning, and compression

Even better performance

- Elapsed time improvement for small LOBS and Complex Queries



10

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.

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, ...



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.

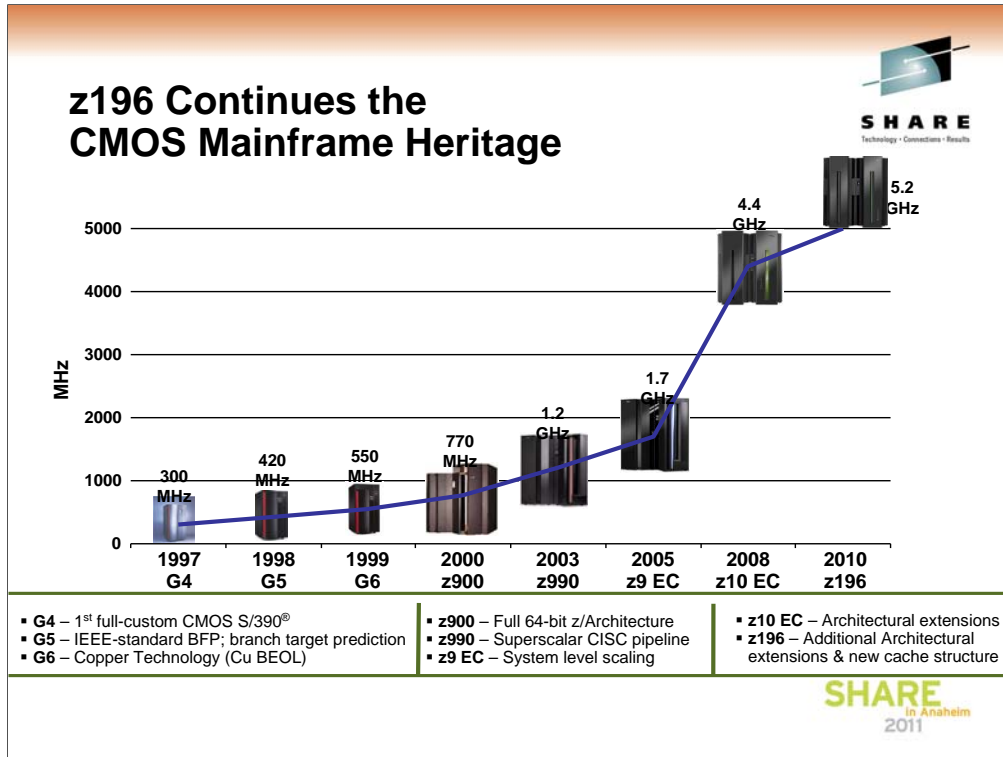
Hardware data compression and encryption provides improved costs, easier management and robust resilience for the platform. Cross-memory and protection keys work with APF authorization and RACF for the underlying system integrity.

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 was 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 ratio 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

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



SHARE
In Anaheim
2011

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.

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



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.

Hardware Trends Impacting DB2

- Drive towards multi core, slowing growth in processor frequency
 - Higher N-ways, more parallelism bring potential latching bottlenecks, memory cache thrashing, ...
 - S/W techniques for single threaded performance growth
 - Clustered systems for massive scale out and continuous availability
- Specialty engines (price/performance)
- Hybrid systems, accelerators
 - Use cores that are more specialized to their purpose
 - New performance opportunities
 - New programming paradigms (e.g. OpenCL)
- Memory hierarchy design
 - Higher cpu frequencies, n-ways make cache utilization a critical factor
 - Translation Lookaside Buffer design, large System z page sizes
- Solid state disk (and other disk related improvements)
 - Performance, energy consumption, reliability benefits over HDD



SHARE
in Anaheim
2011

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.

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



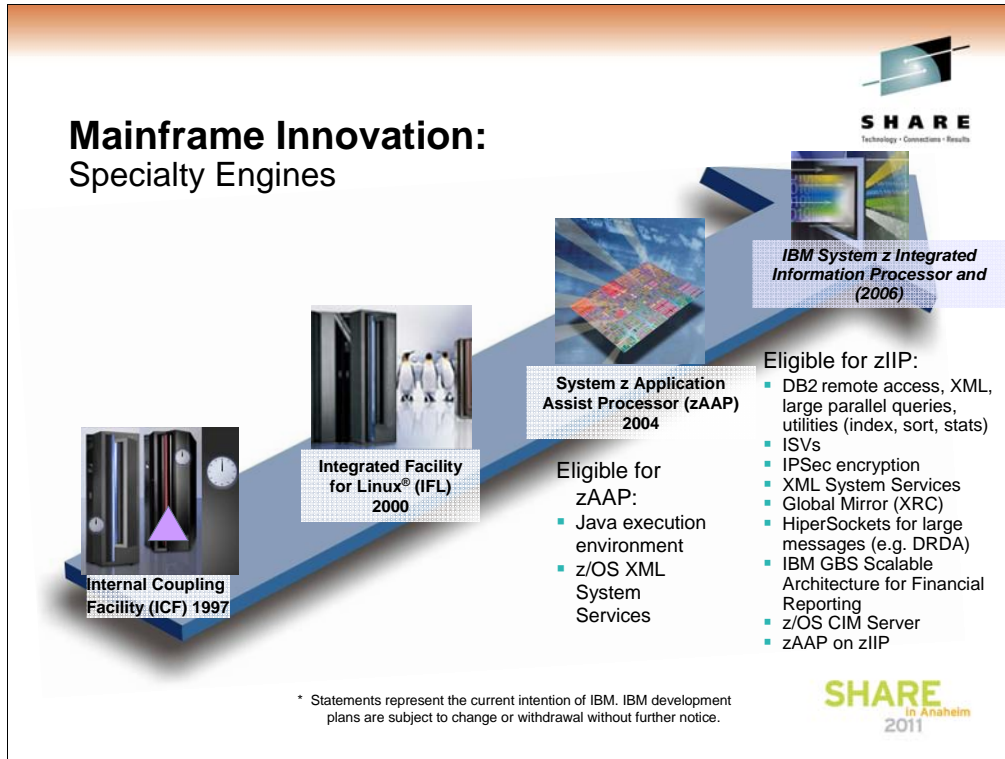
SHARE
In Anaheim
2011

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.eewekeurope.co.uk/news/ibm-looks-ahead-to-future-with-netezza-13361>




The IBM family of specialty engines has evolved. The zIIP uses have multiplied 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.



DB2 & IBM zIIP Add Value to Database Work

Portions of the following DB2 for z/OS V8, DB2 9 and 10 workloads may benefit from zIIP or zAAP for XML (DB2 9 in blue, DB2 10 in green)*:

- 1 – DRDA over TCP/IP connections
 - DB2 9 for z/OS Remote native SQL procedures
 - DB2 9 XML parsing XML schema validation
 - Increased portion of DRDA redirected to zIIPs to 60%
Improved performance via reduced processor switching
- 2 - Requests that use parallel queries
 - DB2 9 higher percentage of parallel queries zIIP eligible
 - DB2 10 more queries eligible, more parallelism
- 3 - DB2 Utilities LOAD, REORG & REBUILD functions used to maintain index structures and sort
 - DB2 10 RUNSTATS – options other than column group
- 4 - DB2 10 buffer pool prefetch and deferred write



* 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.

9



Performance Enhancements Few Changes (CM)

- SQL runtime improved efficiency
- Address space, memory changes to 64 bit, some REBINDs
- Faster single row retrievals via open / fetch / close chaining
- Distributed thread reuse High Performance DBATs
- DB2 9 utility enhancements in CM8
- Parallel index update at insert
- Workfile in-memory enhancements
- Index list prefetch
- Solid State Disk use
- Buffer pool enhancements
 - Utilize 1MB page size on z10
 - “Fully in memory” option (ALTER BUFFERPOOL)



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 through 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.

20



Performance Enhancements need REBIND (CM)

- Most access path enhancements
- Further SQL runtime improvements
- Use of RELEASE(DEALLOCATE)
- SQL paging performance enhancements
 - Single index access for complex OR predicates:
- IN list performance
 - Optimized Stage1 processing (single or multiple IN lists)
 - Matching index scan on multiple IN lists
- Safe query optimization
- Query parallelism improvements
- More stage 2 predicates can be pushed down to stage 1
- More aggressive merge of views and table expressions
 - Avoid materialization of views
- If migrating from V8, get new RUNSTATS before mass rebind



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.

Performance Enhancements requiring NFM

- DB2 catalog concurrency and productivity
- Compress on insert
- Most utility enhancements
- LOB streaming between DDF and rest of DB2
- Faster fetch and insert, lower virtual storage consumption
- SQL Procedure Language performance improvements
- Workfile spanned records, partition by growth
- Access to currently committed data
- Insert improvement for universal table spaces
- Locking improvement for multirow insert
- Efficient caching of dynamic SQL statements with literals



SHARE
In Anaheim
2011

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 enhancement reduces the CPU time for workloads that execute queries that require the use of small work files.


1
2



SHARE
Technology • Connections • Results

Performance Enhancements need NFM + DBA

- Hash access path Create + Reorg + rebind to activate
- Index include columns Alter + Rebuild + rebind to activate
- Inline LOBs Alter (need universal table space and reordered row format)
- DEFINE NO for LOB and XML columns
- MEMBER CLUSTER for universal table space Alter + Reorg
- Alter to universal table space, page size, data set size, segment size Alter + Reorg
- Online reorg for all catalog and directory table spaces




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.



SHARE
Technology • Connections • Results

Virtual storage improvements

- DBM1 below 2GB
 - 75-90% less usage in DB2 10 compared to DB2 9
 - Some of working storage (stack, xproc storage) stays below 2GB
- Larger number of threads
 - Possible data sharing member consolidation
- Improve CPU with storage
 - More release deallocate
 - Larger MAXKEEPD values for KEEP DYNAMIC=YES



DB2 10

Global DSC	CT/PT
DBD	
Local DSC	SKCT SKPT
Thread / Stack	
Thread / Stack / working	
75-90% less usage DBM1 below bar after REBIND	

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.


2
4

Running Many Active Threads

Today

Coupling Technology

LPAR1	LPAR2	LPAR3
DB2A (500 thds)	DB2B (500 thds)	DB2C (500 thds)
DB2D (500 thds)	DB2E (500 thds)	DB2F (500 thds)




- Data sharing and sysplex allows for efficient scale-out of DB2 images
- Sometimes multiple DB2s per LPAR


DB2 10


Coupling Technology

LPAR1	LPAR2	LPAR3
DB2A (2500 thds)	DB2B (2500 thds)	DB2C (2500 thds)



- More threads per DB2 image
- More efficient use of large n-ways
- Easier growth, lower costs, easier management
- Data sharing and Parallel Sysplex still required for very high availability and scale
- Rule of thumb: save ½% CPU for each member reduced, more on memory





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.

Other System Scaling Improvements



- Other bottlenecks can emerge in extremely heavy workloads
 - several improvements reduce latching and other system serialization contention
 - new option to for readers to avoid waiting for inserters
 - eliminate UTSERIAL lock contention for utilities
 - Use 64-bit common storage to avoid ECSA constraints
- Concurrent DDL/BIND/Prepare processes may compete
 - restructure parts of DB2 catalog to avoid the contention
- SPT01 64GB limit can be a constraint, especially if package stability is enabled
 - Allow many more packages by using LOBs
- Improved accounting rollup, compress SMF option




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.


2
6



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

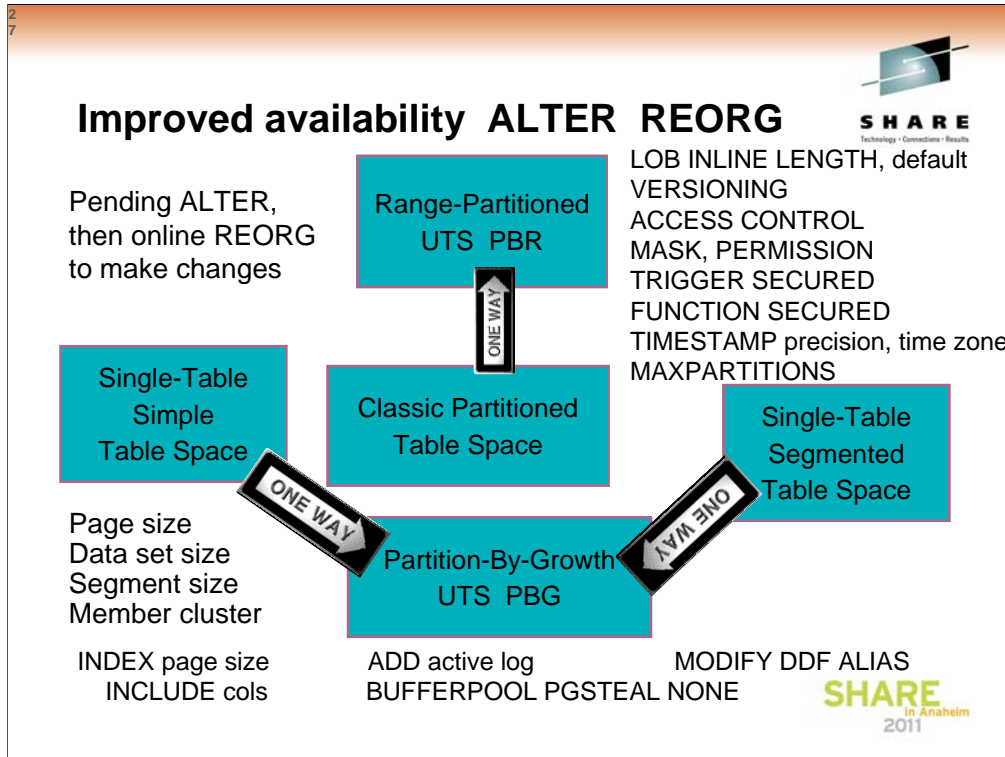




The DB2 catalog and directory are restructured in DB2 10 ENFM to improve productivity and availability. You'll see these improvements in NFM. The current size limits are increased substantially and contention among process like BIND, dynamic SQL, data definition and utilities is reduced. With more table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

2
8

Business Security & Compliance

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




SHARE
Technology • Connections • Results



SHARE
In Anaheim
2011

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.

DB2 10: Productivity –
Doing More with Less!



- **Easier** performance & scaling, **simpler** memory management
- Reduce contention, more online processing
- Reduced need for REORG
- Auto statistics collection
- Monitoring enhanced



SHARE
in Anaheim
2011


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.

50



DB2 10 Utilities Enhancements

- REORG SHRLEVEL(CHANGE) for LOBs
- Online REORG enhancements
 - SHRLEVEL(CHANGE) for all catalog & directory
 - Option to cancel blocking threads
- Improved usability & availability
 - Allow disjoint partition ranges
 - Permit movement of rows between partitions when LOB columns exist
 - Allow REBALANCE and ALTER LIMITKEY even when LOB columns exist
 - Allow DISCARD to delete associated LOB values
 - Messages to estimate length of REORG phases and time to completion



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.

5
2

Many improvements for SAP & web



- **Autonomics**
 - Compress on the fly on INSERT
 - Auto-statistics
 - Hints enhancements
 - Access path lock-in and fallback for dynamic SQL
 - Automatic checkpoint interval
 - Automated installation, configuration & activation of DB2 supplied stored procedures & UDFs
 - Data set FlashCopy in COPY & inline copy
 - Inline image copies for COPY YES indexes
 - UNLOAD from FlashCopy backup
 - REORG enhancements
 - Reduce need for reorganizations for indices
- **Performance**
 - CPU reductions
 - Hash access path
- Numerous optimizer enhancements, paging through result sets
- Parallel index update at insert
- Faster single row retrievals
- Inline LOBs
- LOB streaming between DDF and rest of DB2
- Faster fetch and insert, lower virtual storage consumption
- DEFINE NO for LOBs and XML
- MEMBER CLUSTER for UTS
- Query parallelism enhancements: lifting restrictions
- Dynamic Index ANDing Enhancements
- Option to avoid index entry creation for NULL value
- Index include columns
- Buffer pool enhancements
- **Scalability**
 - Many more threads
 - Reducing latch contention
 - Workfile spanned records, PBG support, and in-memory enhancements
- **Availability**
 - More online schema changes for table spaces, tables and indexes via online REORG
 - Online REORG for LOBs
 - Online add log
 - Automatically delete CF structures before/during first DB2 restart
- **Portability**
 - Allow non-NULL default values for inline LOBs
 - Loading and unloading tables with LOBs in stream
 - Currently committed locking semantics
 - Default SAP settings for DB2
- **Security**
 - More granular DBA privileges



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 Application Enablement and Portability



- Data versioning by date
- pureXML enhancements
- Large object improvements
 - Allow non-NULL default values for inline LOBs
 - Loading and unloading tables with LOBs
 - LOBs in input/output files with other non-LOB data
- Improved portability and SQL consistency
 - Currently committed locking semantics
 - Implicit casting or loose typing
 - Timestamp with time zone
 - Variable timestamp precision – seconds to picoseconds
 - Moving Sum, Moving Average

SHARE
in Anaheim
2011

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.

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



SHARE
in Anaheim
2011

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.

DB2 SQL 2010

z z/OS 10 common

luw Linux, Unix & Windows 9.8

z { Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE SQL, join across encoding schemes, IS NOT DISTINCT FROM, VARBINARY, FETCH CONTINUE, MERGE, SELECT from MERGE, **data versioning, access controls**

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, OmniFind, spatial, range partitions, data compression, session variables, DECIMAL FLOAT, optimistic locking, ROLE, TRUNCATE, index & XML compression, created temps, **inline LOB, administrative privileges, implicit cast, date/time changes, currently committed, moving sum & average, index include columns, PureScale**

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, even more vendor syntax, temp table compression

SHARE
Technology • Connections • Results

In Anaheim
2011

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>

5
6

pureXML improved performance & usability

- XML schema validation in the engine for improved usability and performance
- Binary XML exchange format improves performance
- XML multi-versioning for more robust XML queries
- Allow easy update of XML document nodes
- Stored procedure, UDF, Trigger enhanced support
- XML index matching with date/timestamp
- CHECK DATA utility checks XML

SHARE
Technology • Connections • Results

SHARE
In Anaheim
2011

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

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

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

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

The CHECK DATA utility is extended to check XML data.



Query Enhancements

- 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

SHARE
in Anaheim
2011

Query enhancements in DB2 build on the improvements in DB2 V8 and 9. The CPU reductions can make a difference 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.



- Continued growth of System z platform:
 - As of 2Q 2009 - Over 6,200 applications, from over 1,600 ISVs enabled on the System z platform
 - Over 2,200 applications on z/OS 1.8 and above
 - More than 3,800 for all z/OS releases
 - Over 3,000 applications on Linux on z
 - More than 500 new Linux applications added in 2008
 - Represents 18% growth from 2007
 - Over 250 already added in 2009
- ISV Partner loyalty
 - Over 1,500 ISVs building applications for System z
 - Over 1,800 applications on z/OS 1.8 and above (over 3,500 for all z/OS releases)
 - **2,800+ applications for Linux on System z**
 - 86% of our ISVs maintain OS currency

Jump into DB2 10! The water's fine.



DB2 V8

DB2 9

DB2 10

Key Questions are WHEN? and HOW?

SHARE
In Anaheim
2011

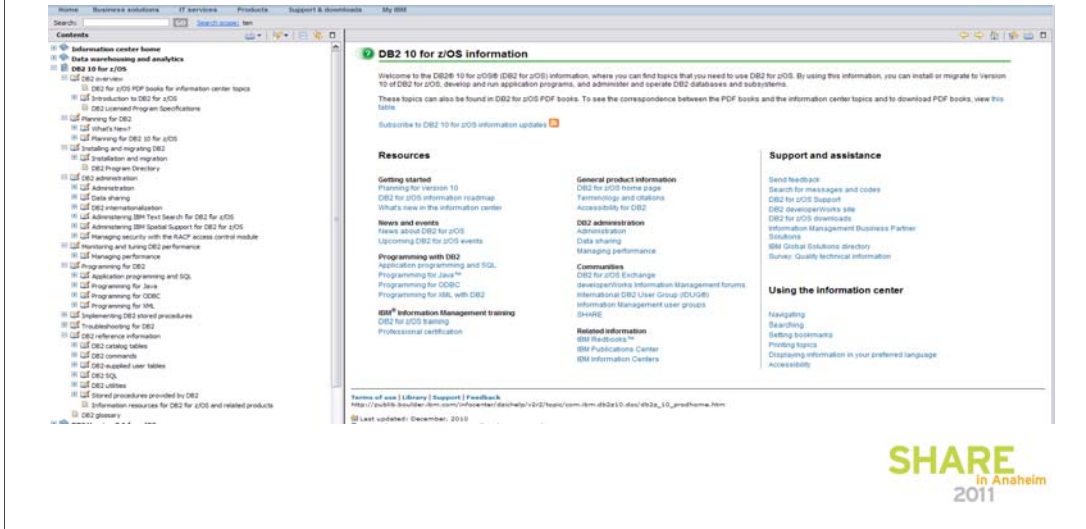
Is your current fish bowl getting constrained? What is limiting you? Is it CPU? Virtual storage? Latching? DB2 catalog and directory? Concurrent utilities? Are you currently running DB2 9? V8? V7? Should you migrate to DB2 10?

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.

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.

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.

Improvements to installation & migration information based on customer feedback

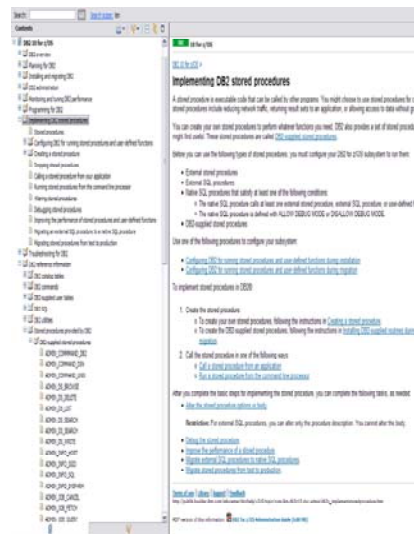
- Pre-migration and migration checklists were added to help plan for and keep track of the migration process. The checklists include links to the pre-migration and migration steps. Access the checklists from the following locations:
 - Introduction to migration from DB2 Version 8:
http://publib.boulder.ibm.com/infocenter/ecsimzic/v1r0/topic/com.ibm.db2z10.doc.inst/db2z_intro2migfromv8.htm
 - Introduction to migration from DB2 Version 9.1:
http://publib.boulder.ibm.com/infocenter/ecsimzic/v1r0/topic/com.ibm.db2z10.doc.inst/db2z_intro2migfromv9.htm
- Installation and migration steps clearly define tasks to complete. Related concept and reference information are included as links from the tasks.
- Several new jobs simplify the setup and installation of DB2-supplied stored procedures and user-defined functions. The process is documented in [Installation step 20: Set up DB2-supplied routines \(optional\)](#) and [Migration step 25: Set up DB2-supplied routines \(optional\)](#).

For the V10 development cycle, the information development team focused on addressing solutions to top customer pain points. Based on customer feedback, one of the areas we focused on improving is the installation and migration information. Try the checklists. Those who used them noted that they helped.

Stored procedure information improvements



- Stored procedure information easier to find
- Implementing stored procedures
- DB2 reference information added
- Added more information:
 - Setting up WLM
 - Granting privileges
 - Moving stored procedures from test to production.
- Clearer information



The information development team focused on addressing solutions to top customer pain points. Based on customer feedback, another area we focused on improving is the stored procedure information.

We made the stored procedure information easier to find.

“Implementing stored procedures” explains all of the tasks involved in setting up, creating, and calling stored procedures with links to the detailed steps for each task.

We added all of the stored procedure reference topics to the “DB2 reference information” section.

We added missing information, such as setting up WLM environments, granting privileges, and migrating stored procedures from test to production.

We also rewrote information to make it clearer what you need to do to accomplish a particular goal. The improvements were implemented in DB2 10 and DB2 9 information.

DB2 10 for z/OS: Skip-Level Migration



May move from V8 to DB2 10,

but just because you can, doesn't mean you always should....

Migration, fallback and data sharing coexistence fully supported

Mix of DB2 9 and 10 or DB2 V8 and 10

Key considerations:

-
- Risk/reward analysis
 - What's your risk? Tolerance level?
 - How will you do it? What's your mitigation plan? Are ISVs ready?
 - What workloads do you need to test and can you test them properly?
 - Are you missing out on DB2 9 value in the meantime?
 - Migration cost savings is not 2X versus two migrations
 - Migration considerations for two versions still apply
 - Larger migration project, longer migration timeline, more risk
 - Applications and ISVs may not be ready
 - Timing: GA, V8 end of service, other software, ability to test early software



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>

4


4

DB2 10 for z/OS At a Glance



SHARE
Technology - Connections - Results


Performance, Scalability	<ul style="list-style-type: none"> • CPU reductions out-of-the-box • Hash access to data, index include columns • Ten times more threads per DB2 image
Availability Security Productivity	<ul style="list-style-type: none"> • More online schema changes • Improved concurrency: catalog, data, & utilities • Row and column access control, masking • Administrator privileges with finer granularity • Administration productivity enhancements
Application Enablement	<ul style="list-style-type: none"> • Versioned data or temporal queries • pureXML enhancements • SQL improvements that simplify porting
Dynamic Warehousing	<ul style="list-style-type: none"> • Moving sum, moving average • Many query optimization improvements • Query parallelism restrictions removed



SHARE
In Anaheim
2011

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.

IBM DB2 Tools: Are you ready for DB2 10?



SHARE

- Exploit DB2 10 performance savings out-of-the-box
- Optimize Performance Across Multi-Platform Applications
- Lower CPU costs while reducing batch windows
- Higher data availability through simplified recovery operations



All New with DB2 10!

DB2 Utilities Suite 10 drives down costs with autonomies, page sampling and further offloads processing to zIIPs and FlashCopy. Developed in conjunction with DB2 10 to provide maximum data integrity and exploit all new functions out of the box.

Tivoli OMEGAMON XE for DB2 Performance Expert 5.1 extends its insight into distributed workloads and offers a robust infrastructure to support DB2 10 subsystem consolidation, with lower monitoring overhead. The recommended performance monitor of DB2 10!

DB2 Administration Tool/Object Compare 10.1 extends the value of DB2 10 with new capabilities that allow DBAs to quickly exploit DB2 10 features like schema evolution. Reduces the overhead of many routine tasks.

QMF 10 delivers built-in visualizations and reports that dramatically extend the value to end users. A new metadata layer simplifies the process to understand and create reports.

DB2 Sort 1.1 lowers the cost of DB2 Utility sort processing by exploiting advanced features of System z and z/OS while optimizing overall system efficiency. Significantly reduces batch windows.

DB2 High Performance Unload 4.1 reduces the cost of extracting DB2 10 data with support for TCP/IP Pipes and the new internal format as well as a new native XML data unload capability.

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 TCP/IP 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 TCP/IP 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.

Questions?

Thank YOU

Watch for upcoming white papers and conferences

IOD IDUG Share

SHARE
Technology • Connections • Results

SHARE
In Anaheim
2011

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>

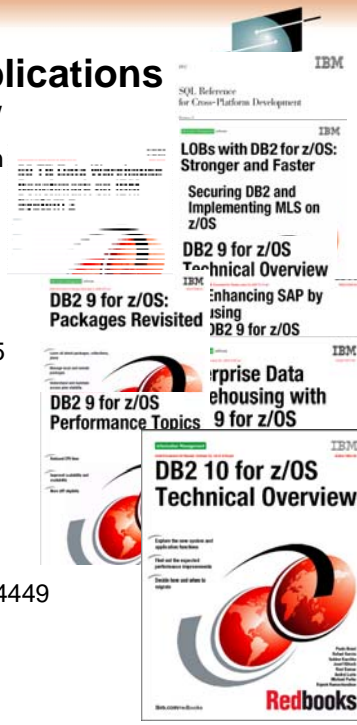
<http://www.idug.org>

<http://www.share.org>

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

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



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>
http://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.html>
29. DB2 9 and Storage Management SG24-7823 <http://www.redbooks.ibm.com/abstracts/sg247823.html>

More information and resources



- **DB2 main web page**
<http://www.ibm.com/software/data/db2/zos/>
- **DB2 10 web page**
<http://www.ibm.com/software/data/db2/zos/db2-10/>
- **DB2 books, Information Center**
 - <http://www.ibm.com/support/docview.wss?rs=64&uid=swg27011656>
 - <http://publib.boulder.ibm.com/infocenter/imzic>
- **DB2 best practices web page**
<https://www.ibm.com/developerworks/data/bestpractices/db2zos/>
- **DB2 for z/OS IBM Redbooks publications**
 - <http://www.redbooks.ibm.com/cgi-bin/searchsite.cgi?query=db2&SearchOrder=4&SearchFuzzy=>
- **DB2 presentations**
<ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/>



SHARE
in Anaheim
2011

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>

What DB2 Means to Business Value, Talk DB2 in business terms the suits can understand.
<http://www.ibm.com/databasemag.com/story/showArticle.jhtml?articleID=199201819>

System z Strengths and Values, SG24-7333,
<http://www.redbooks.ibm.com/abstracts/sg247333.html>

A01: DB2 10 for z/OS Overview




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

DB2 10 for z/OS

A perfect 10

SHARE
In Anaheim
2011

50



SHARE
in Anaheim
2011

Disclaimer/Trademarks

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements, or other publicly available sources. IBM has not tested these 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.

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

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

Trademarks The following terms are trademarks or registered trademarks of other companies and have been used in at least one of the pages of the presentation:
The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both: AIX, AS/400, DataJoiner, DataPropagator, DB2, DB2 Connect, DB2 Extenders, DB2 OLAP Server, DB2 Universal Database, Distributed Relational Database Architecture, DRDA, eServer, IBM, IMS, iSeries, MVS, Net.Data, OS/390, OS/400, PowerPC, pSeries, RS/6000, SQL/400, SQL/DS, Tivoli, VisualAge, VM/ESA, VSE/ESA, WebSphere, z/OS, zSeries
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Other company, product, or service names may be trademarks or service marks of others.

DB2 10 for z/OS is now generally available, but some specific items in this presentation are not. One example of function that is not in DB2 10 is the ability to migrate from Sybase.

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>

Questions and Answers about DB2 10 for z/OS

March 3, 2011

This article is a pointer to resources and information about DB2 10 for z/OS. Customers have asked many questions on webcasts and at conferences. Many of the pointers have more extensive information to respond to the questions.

Even though this product is generally available, changes are likely to occur for some of these answers. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

How can I get more detailed information about DB2 10? Are there any IBM Redbooks publications out now? How can I get the DB2 books?

More details are provided on the web:

DB2 main web page <http://www.ibm.com/software/data/db2/zos/>
DB2 10 web page <http://www.ibm.com/software/data/db2/zos/db2-10/>
DB2 books, Information Center

<http://www.ibm.com/support/docview.wss?uid=swg27019288>

DB2 best practices web page. These best practices present advice on the optimal way to use DB2 for z/OS to satisfy key business data processing needs. These presentations and articles are authored by leading experts in IBM's development and consulting teams.

<https://www.ibm.com/developerworks/data/bestpractices/db2zos/>

DB2 for z/OS IBM Redbooks publications <http://www.redbooks.ibm.com/cgi-bin/searchsite.cgi?query=DB2+AND+z/OS&SearchOrder=4&SearchFuzzy=>

DB2 10 booklet

<ftp://public.dhe.ibm.com/common/ssi/ecm/en/imm14075usen/IMM14075USEN.PDF>

The first DB2 10 IBM Redbooks publications are DB2 10 for z/OS Technical Overview and Extremely pureXML, discussing XML in DB2 9 and 10. Watch for Performance Topics, anticipated in the second quarter of 2011.

You can also learn more at upcoming conferences, including Share, IDUG, and IOD. Get information from the recent conferences from their web sites.

IDUG	www.idug.org
IOD	www.ibm.com/software/data/2010-conference/
Share	www.share.org
IBM Education	www.ibm.com/training/us/db2zospath
DB2 events	www.ibm.com/software/data/db2/zos/events.html

We'll be getting more information out at the conferences: Share, GSE, IOD, and IDUG. See the IDUG conference content if you have a membership. Register to join IDUG if you don't. Register for the virtual IOD conference.

Here are some sessions from the IDUG EMEA 2010 track for DB2 10:

- A01 - DB2 10 for z/OS Overview
- A02 - DB2 10 for z/OS Migration Planning
- A03 - DB2 for z/OS Trends and Future Direction
- A04 - DB2 9 & 10 for z/OS Enhancements for System Programmers
- A05 - DB2 10 for z/OS Security Enhancements
- A06 - DB2 10 Availability Enhancements
- A07 - DB2 10 - What's in it for SAP and Enterprise Applications
- A08 - What's new from the optimizer in DB2 10 for z/OS?
- A09 - DB2 for z/OS and Websphere Update
- A10 - Hash Access to DB2 Data - Faster, Better, Cheaper
- A11 - DB2 and System z Synergy
- A12 - DB2 10 Application topics – A sneak preview
- A13 - DB2 10 for z/OS Performance Preview
- A14 - Utilities Update - DB2 10 for z/OS
- A16 - The Next Level: Data Warehousing with DB2 10 and zEnterprise

How does DB2 10 use the zEnterprise?

Faster CPUs, more CPUs, and more memory means better DB2 performance and scalability. 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. This has excellent synergy with DB2 10, which removes many single system scaling inhibitors.

Compression hardware improvements in zEnterprise are expected to increase DB2 data compression performance. Larger cache (192 MB in level 4) is expected to benefit DB2 workloads. DB2 for z/OS can take an advantage of cache optimization on zEnterprise. Translation Lookaside Buffer changes are expected to improve DB2 10 performance for 1MB page sizes. The 1 MB hardware page sizes improve DB2 performance and are only on z10 and zEnterprise. Hybrid architecture delivers new opportunities for DB2 query performance acceleration with IBM Smart Analytics Optimizer.

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.

Should I upgrade to DB2 10? When and how should I migrate to DB2 10?

The answer to upgrading to DB2 10 is a definite Yes. The question is not so much whether to upgrade as when and how to upgrade. If you are running DB2

9 today, then DB2 10 is in your near future, giving you more room to grow, with higher limits, lower CPU, 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. Both choices are fully supported with migration and fallback.

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, then 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. Customers who migrate early will need more robust test and service processes.

See the Upgrade to DB2 10 paper at the end of this presentation for much more on this topic.

<ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/migration/db2-10-migration-planning-idug-emea-2010-nov-miller.pdf>

When will DB2 10 for z/OS be generally available or GA?

DB2 10 for z/OS is generally available world wide as of October 22, 2010.

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

When will DB2 V8 go out of support?

DB2 end of service is scheduled for April 2012, as announced in August 2010. Customers should be getting ready for their next migration if they are on DB2 V8. Any customers who are on V7 or lower should be migrating to V8 as soon as possible, leaving the versions which are out of service. Extended service can be provided for a fee, if you can't migrate and need support. Contact your IBM representative. Click on Product Support Lifecycle from this page to see the service cycle:

http://www-947.ibm.com/support/entry/portal/Planning/Software/Information_Management/DB2_for_z-OS

<http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=AN&subtype=CA&appname=g pateam&supplier=897&letternum=ENUS910-169&pdf=yes>

V8 has been withdrawn from marketing. If we're still on V7, what are our options?

The withdrawal from marketing means that new customers cannot order V8. As an existing customer, you can order and migrate to DB2 V8. Talk to your IBM representative to get DB2 software or migration support. Ask about the Migration Planning Workshop as well.

How should I start getting ready for DB2 10?

Are you ready for DB2 10? Are you running in DB2 9 NFM or DB2 V8 NFM? Have you looked through the information APARs? Have you contacted your vendors? Do you have the required service applied to your current system? Have you run the premigration job and worked to resolve the incompatibilities?

Have you gotten rid of private protocol? Have you converted to packages instead of DBRMs directly in plans? Those are first tasks, along with planning what you expect to gain and the process for testing the new version.

- Plan the migration process. Make sure that you have the prerequisites and current software that works with the new version. Check with your vendors. See the planning for migration presentations referenced below. Migration planning workshops are available at no charge for DB2 V8, DB2 9, and DB2 10 and can save you time, providing experience from other customers and IBM.
- Check for incompatible changes. As you go through the process, check for the incompatible changes and deprecated function in DB2 10 as well. That way, you can avoid some of these problems and make the changes in the normal course of work, so there is less effort when you upgrade to DB2 9 or DB2 10. Run the premigration jobs (DSNTIJP9 and DSNTIJPA) early and often.
- Read the Installation Guide, the Preventive Service Planning (PSP) bucket and Authorized Program Analysis Report (APAR) installation text. The Consolidated Service Test (CST) can provide a much more stable level of service for a broad stack of products. CST has worked for very small customers, very large customers, and those in between.
- Develop a test plan for your unique workload and environment. CST can help for more general processing work. Collect performance information and access path information. DB2 9 and 10 package stability can help you get back to the original access paths and are generally used in a release migration.
- Put plans in place to ensure that each step is tested and successful before continuing. Go through the next checklist in the DB2 upgrade paper and then use the more detailed checklists in the DB2 Information Center.

Do I need to REBIND? When?

REBIND is not required for migration to DB2 10, but REBIND is strongly recommended. Getting the best performance improvements and eliminating regression does depend upon rebind in most situations: getting current structures, better access paths, and reusing threads. Eliminating performance regression may depend upon REBIND. Storage constraint relief depends upon REBIND. Changing to use release deallocate requires a REBIND. The migration process has been smoother when customers REBIND, with fewer problems. All plans containing DBRMs must be rebound. All packages that were last bound on V5 or lower must be rebound. Static SQL with DEGREE(ANY) for parallel processing should be rebound, or it will be sequential. Other REBINDs can be staged over weeks of time, and REBIND is only needed once per package for the migration. Improvements in access paths can be very significant, such as stage 2 predicates that can become stage 1. REBIND in DB2 10 takes more CPU and

elapsed time than in prior versions, but more concurrent REBINDs are possible in NFM. So please be kind. REBIND.

What else is needed to get performance out-of-the-box?

Some customers have needed to use RELEASE(DEALLOCATE) to get better performance for short running transactions. With the IRWW workload, the improvement from use of this parameter was about 10%.

RELEASE(DEALLOCATE) could not be used by distributed applications, before DB2 10, and improvements have been implemented to recycle the thread every 200 transactions. RELEASE(DEALLOCATE) is best for high volume batch or transactions with few SQL statements in each COMMIT. The transactions also need to be well-behaved for locking. Local transactions need some mechanism to end the thread, so that utilities, data definition changes and other processes can be performed.

Can you explain more about the changes for RELEASE(DEALLOCATE) in DB2 10?

RELEASE(DEALLOCATE) has been part of DB2 for a long time, but DB2 10 makes the function more useful. Using RELEASE(DEALLOCATE) requires much more memory. The dramatic memory improvements in DB2 10 that you can achieve with rebind makes it possible to use RELEASE(DEALLOCATE) more. This change saves significant CPU time for high volume transactions with few short running SQL statements, without changing applications or DDL.

RELEASE(DEALLOCATE) works in CM, but does require a REBIND unless the packages already use RELEASE(DEALLOCATE). For DDF work, after rebinding with Release(Deallocate), the customer must issue the MODIFY DDF PKGREL(BINDOPT) command. By default JDBC is changed to use Release Deallocate. Additionally, even if packages are not rebound, a new DDF option called High Performance DBAT can be set that will cause resources to be periodically released after 200 Commits. So RELEASE(DEALLOCATE) is more applicable and safer to use in many more situations. RELEASE(DEALLOCATE) depends upon having very well debugged, well behaved applications that are careful with locking and commit frequently.

Some customers are getting a lot of performance improvements that can be made in CM on the first day. The improvements do require a REBIND in most situations, and that does mean testing, but DB2 version changes also take testing, so combining the work for a dramatic improvement will work for many customers. This change can be implemented for a few very high volume transactions and provide a great return.

The DB2 10 performance presentation shows a fairly common scenario. Using more CICS protected ENTRY (persistent) threads with RELEASE (DEALLOCATE) provides the opportunity for significant price/performance

improvement. High-volume, short running distributed transactions can take advantage of CPU reductions, using RELEASE(DEALLOCATE). 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 7.4%, 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). This change could also be made in CM with step 2.

So this scenario demonstrates the runtime improvements and CPU value of REBIND and RELEASE(DEALLOCATE) for a high volume, short transaction.

What does DB2 10 do to enable RELEASE(DEALLOCATE)? Why can't I rebind packages now in DB2 9 with this option and realize less CPU?

RELEASE(DEALLOCATE) takes memory, which is not available for many customers on DB2 9 or V8. In DB2 9, you get RELEASE(COMMIT) for DDF work, even if packages are bound RELEASE(DEALLOCATE). DB2 10 allows distributed RELEASE(DEALLOCATE). DB2 10 also changes to end the thread every 200 transactions. If your distributed transactions are already bound with RELEASE (DEALLOCATE), then they would not need to be rebound, just change the DDF setting.

What negative effects do I need to be aware of in using RELEASE (DEALLOCATE)?

If your application is not well behaved, then you can get into problems, deadlocks, timeouts, and inability to run utilities. If your process is missing commits or takes gross table space locks, then it's not a good candidate for DEALLOCATE.

Did some customers find performance regression?

Some customers found longer BIND times in DB2 10. The default for access path management has changed in DB2 10 from none to EXTENDED. Customers who move from DB2 9 and have used access path management have some improvements. Customers who did not have access path management will find increases in BIND CPU time from this change. If you want to reduce the time for BIND, then change the subsystem parameter back to none. Use access path management where you have noted problems with new access paths.

What is DBAT reuse? How does it better perform from DB2 Connect concentrator features?

DBAT is a DDF thread, which does not permit thread reuse in DB2 9, and will in DB2 10. Thread reuse with small, high volume transactions can improve CPU times significantly.

What happens to DSNZPARM COMPRESS_SPT01? I have seen recommendations to compress in DB2 9, but since compression in DB2 10 is unnecessary - is there something in the DB2 10 migration process to revert back to uncompressed?

The DB2 10 ENFM migration process changes the catalog and directory into the new format, including this change.

Is the combination of hash access and unique index access just overhead or can they complement each other?

The access techniques can be complementary. If you have fully qualified keys with equal predicates, then the hash can be used. Otherwise, for range predicates as an example, the indexes are needed. Hash means that you can't cluster, so hash will be used where fully qualified key access is the primary access and where clustering is not needed. The best candidates for hash access are random single row access into a large table with a fixed, known size, with many rows per page, small variation in row sizes, and number of index levels greater than 4.

Is memory used moved above the 2GB bar evenly spread through DB2 addresses spaces or are it concentrated in DBM1, for instance?

Most of the savings are in DBM1, as that is the primary constraint for customers. REBINDs are required for the EDMPOOL improvements and some working storage. The shared area above the bar can help with other address spaces. Some ECSA for instrumentation is also moved above the bar.

I thought DB2 would never have another skip a release migration. Why the change for V8 to V10?

Skip release is costly for developers and complex, so it's not for every version. DB2 delivered skip for V5 to V7 and now for V8 to V10. I think that most customers are best served by staying current with DB2 versions, but some customers get behind. Skip migration is great for customers running V7 today or for those who just finished migrating to V8. Customers who plan to migrate to DB2 9 in 2011 should generally do so. The presentation and paper on migration have much more information.

Does a statement concentrator need to be enabled for caching of dynamic SQL with literals?

Caching for dynamic SQL statements is part of DB2 10 NFM. You need to allow caching to ignore literals with a prepare attribute or an ODBC or JDBC attribute.

Will EXPLAIN show hash access if chosen over index access?

EXPLAIN shows hash access. See the presentations and documentation on hash access. The IDUG EMEA conference in Vienna included a full session on hash access, with best practices for establishing the hash access. Please get it from the IDUG web site under content for that conference.

Can you please expand on hash access and how it is useful?

Hash access is a new direct access technique within DB2 10. DB2 uses the direct keys defined and an algorithm to access the data. This direct access is sometimes done in a single I/O making it very fast and efficient for common customer id and product id access types. The direct single I/O access will improve any application that performs these keyed access types that are sometime done millions of times a day.

How is the hash access path established?

The IDUG EMEA conference in Vienna included a full session on hash access, with best practices for establishing the hash access. Please get it from the IDUG web site under content for that conference.

Is there any improvement compared for MEMBER CLUSTER compared to DB2 9?

MEMBER CLUSTER is not permitted for universal table spaces (UTS) in DB2 9 and is for DB2 10. With DB2 10, you can ALTER to a UTS, then REORG to change the table space to MEMBER CLUSTER, using a pending ALTER. Insert performance improvements build upon those in DB2 9. Insert rates improve a lot and CPU time reductions are very large in DB2 10.

What's the recommendation for IRLM storage? PC = YES / NO?

PC=YES has been the only option since V8.

The DB2 10 overview indicates that most utility enhancements in DB2 10 are available in NFM, rather than CM. Previous DB2 releases provided utility enhancements in CM. Can you explain this change in behavior please?

Delivering new function in new function mode means a better ability to fall back cleanly and better consistency for DB2 and utilities. Customers are spending less time in Conversion Mode and getting more value in New Function Mode.

During PENDING ALTER, is the object is accessible?

Yes, objects are accessible, as long as applications commit to allow the switch for the online reorg.

With DB2 9 and 10, we hear more and more about stabilizing access paths through hints and many more options. Is this another way of saying that the cost based optimizer has somehow reached a limit and that optimizer capabilities are no longer sufficient?

No. DB2 9 and 10 have many query access path improvements, but access paths for transactions and making the maintenance and release migration process simpler are also important.

Does use of the z/OS 1MB page size change the maximum DB2 data page size? Are we still stuck with 32K maximum, or are we going to be allowed to grow to a 1MB page size (and buffer pool) in DB2 10?

Page sizes are an overloaded term and often cause confusion. Hardware and z/OS have page sizes or frames of 4KB – and now add 1MB. DB2 page sizes remain the same: 4KB, 8KB, 16KB, and 32KB.

Has the global buffer pool contention seen with very large buffer pools been addressed in V10? We see contention with pseudo close causing all the buffer pools to be scanned.

This overhead was addressed in DB2 10, avoiding buffer pool scans.

Does DB2 10 support private protocol?

DB2 10 has only DRDA. Warnings have been delivered for 15 years, and conversion is necessary before moving to DB2 10. See the upgrade presentation and paper.

Timestamps in DB 10 provide 12 digits beyond the second or picoseconds. Do existing applications need to change to accommodate all 12 digits instead of microseconds?

Existing applications don't need to change if they don't need to change from using microseconds. The new timestamp capabilities within DB2 10 provide the database designer with the flexibility to define timestamp precision as application requirements dictate. DB2 10 allows 0 to 12 digits for timestamps with a compatible default of 6. So you could have picoseconds with 12 digits or nanoseconds with 9 digits, more precise than earlier DB2 versions. You could define timestamps with 0 digits, if seconds are precise enough. Existing applications do not need to change if the current timestamp precision is adequate for the application. A general purpose application which will process any DB2 data needs to change to handle this variation and timestamps with time zones.

How small is a small LOB?

A small LOB is able to fit on the same page size along with its other associated data column information. This avoids creating a separate table space for the LOB data. The desired is to fit these small LOBs within their 4k, 8k, 16k or 32k page sizes so that no additional I/O, memory, and CPU are needed to store or retrieve the LOB data. Typically these small LOBs can be anywhere from hundreds of bytes to 31k, so that they fit on the same page with the rest of the row.

What are the improvements for BLOB and CLOB?

Many improvements were made in DB2 10 for LOBs. Inline LOBs, improved LOB streaming to and from DDF are key performance improvements. Indexes can be

defined on the inline LOB. Better streaming and minimized LOB default sizes help improve performance when LOB materialization occurs. This is especially important for DDF applications because of the network impact of large objects. DEFINE NO for LOB and XML columns makes definition easier. REORG SHRLEVEL(CHANGE) for LOBs improves availability. Online REORG permits movement of rows between partitions when LOB columns exist, and allows DISCARD to delete associated LOB values. Non-NULL default values can be defined for inline LOBs. Now you can stream LOBs into one data set, rather than taking a data set for each LOB.

Is it feasible to convert Oracle or Sybase to DB2 10?

DB2 10 has even more DB2 family and other DBMS compatibility functions. We've seen many customers able to convert with DB2 9. The new timestamp capabilities and precision definitions, the stored procedure language enhancements and many other features that first came in DB2 LUW 9.7 are now offered in DB2 10 for z/OS. It is possible to convert Oracle to run on DB2 10 and z/OS. The Oracle complexities will determine the level of effort but DB2 10 has the majority of the functionality needed for a quick migration effort. SQL Skin works with DB2 for LUW to ease Sybase migration now, and might work with DB2 for z/OS in the future.

Will this change the DB2 Utilities?

Changes in utilities are an integral part of DB2 for z/OS. See utilities presentation on the IDUG web site.

Are there any changes on SQL HINTS topic included in DB2 10?

The infrastructure around the SQL BIND process is enhanced in DB2 10 not just the SQL Hints. DB2 10 provides the capability to provide hints globally to all users for an SQL statement and provide certain optimizer zparm settings such as parallelism star join and other BIND parameter settings for specific package or specific SQL statements. This provides extreme granularity and flexibility for turning on specific parameters for certain databases, applications, tables or even certain SQL statements within your environment. DB2 10 provides unprecedented levels of package optimization, stability and control through the basic, extended and versioning of packages within your DB2 system. This feature provides the flexibility to choose the package with the desired access path and lock it down. By locking down the access path, administrators can error code out a REBIND that doesn't have the same access path that we desired and prevent it from going into our production environment. This helps guarantee performance and stability within our production systems.

Can solid state drives be used with hardware replication?

Yes. Here are some pointers:

<http://www.ibm.com/systems/storage/solutions/ssd/>

ftp://public.dhe.ibm.com/storage/whitepaper/disk/ds8000_taneja.pdf

ftp://service.boulder.ibm.com/storage/software/virtualization/clipper_svc_5.pdf

<ftp://ftp.software.ibm.com/common/ssi/sa/wh/n/tsw03044usen/TSW03044USEN.PDF>
<https://www.ibm.com/developerworks/mydeveloperworks/blogs/InsideSystemStorage/tags/announcements?lang=en>

Can you explain the common SDSNEXIT in DB2 10?

Allowing tailored names for DSNHDECP will permit many subsystems to share the SDSNEXIT data set. Having more than one name means that programs cannot just load the DSNHDECP module to use the data. Note the changes in DB2 9 to prepare for this.

http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db29.doc.inst/db2z_functionsdeprecated.htm

Why one would need to select from SYSLGRNX?

SYSLGRNX has been needed for some recovery situations.

Will R10 restrict BINDING a plan to a single DBRM?

BIND for a DBRM is always to a package in DB2 10. Every DBRM will be bound to a package, and plans will need one or more packages for execution. See the migration presentation for more.

What happens for existing DBRMs bound into plans?

The best practice is to convert all plans to use packages on DB2 V8 or DB2 9. If you don't convert plans to use packages, they will be rebound when the plans run, using a default collection id. See the next question.

Will a formal conversion method be provided to convert stand alone DBRMs to packages for DB2 10?

The process was provided in DB2 9 and V8, about a year ago. See the migration presentation for the process. The Packages Revisited chapter 4 discusses the changes and process.

Packages Revisited, SG24-7688

<http://www.redbooks.ibm.com/abstracts/SG247688.html>

Does DB2 10 have any features like multirow fetch or insert that applications need to modify to improve performance?

Applications don't require changes to have performance features as dramatic as V8 multi row fetch and insert, but do have many improvements in SQL and XML in DB2 10. See the overview presentation for the application improvements, ranging from temporal or versioned data, new timestamp options, and improved ability to port from other platforms, and many XML enhancements. Many of the improvements depend upon implementing the latest clients, drivers, and DB2 Connect 9.7 fixpack 3a. See the migration presentation and paper for more.

We have currently disabled RRF in our DB2 subsystems since we have encountered numerous problems with it. Is RRF usage a must from the start in DB2 10?

No, the option still exists in DB2 10, but RRF is the default, and required for most new improvements. Customers can use this parameter to control conversion to RRF, but should transition to RRF.

What are the major software/hardware requirements for DB2 V10?

See the migration presentation and paper and the general availability announcement. Check the Program Directory on the web.

http://www.ibm.com/common/ssi/rep_ca/0/897/ENUS210-380/ENUS210-380.PDF

<http://www.ibm.com/support/docview.wss?uid=swg27019288>

<http://publibfp.boulder.ibm.com/epubs/pdf/i1088290.pdf>

We are running IMS V9. In DB2 10 must we upgrade our IMS to 10 or 11? We are currently starting to upgrade to IMS v10.

IMS V9 is out of service and is not supported with DB2 10. Migrate to IMS 10 or 11 soon. See the announcements for prerequisites. IMS 12 is already in beta.

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

What are the major conversion tasks for applications (primarily binds)?

Migration to DB2 10 has many steps similar to V8 and DB2 9. REBINDs are necessary to leverage many of the performance and memory enhancements. New Function Mode is necessary for other new enhancements. Review the migration presentation and Installation Guide to understand the feature requirements. REBINDs and adjusting for the function removed are the primary application programming tasks. See the migration presentation and paper.

Does BIND time increase in DB2 10?

BIND times do increase in moving to DB2 10. This increase is often offset by the improved concurrency for BIND. A substantial part of the increase has been the change to default to use of access path stability, which is an option.

Could you recap index include please? What is index include designed to do?

The new DB2 10 INCLUDE COLUMNS on a unique index feature provides the ability to add additional non-unique columns to a unique index definition. This new feature provides the ability to consolidate and eliminate other indexes on your tables leading to improved access and reduced I/Os by having fewer indexes to maintain. This also shortens DB2 access paths in some cases, from multiple index access to a single index access. This function is also part of DB2 for LUW.

Given potential data sharing member consolidation with storage relief, do you think open dataset might become more of a bottleneck/concern?

Some customers may find data set open a bottleneck, but most will not. DB2 9 handles roughly 60 thousand concurrent allocations for some customers today. DB2 9 has increased that number to about 100K. Changes in z/OS V1.12 improve allocation for very large numbers of data sets.

Will the DB2 catalog / directory have to be SMS managed? Will IBM supply a conversion guide or procedure?

Yes, the new data sets for the DB2 catalog change from using manual definition and extension to DB2 managed data sets under SMS control. See the migration presentation and paper for a start. The DB2 catalog and directory have some table spaces defined as DSSIZE 64G, which requires SMS managed storage with extended format and extended attributes.

How do the utilities work with LOB sizes greater than 32K?

The utilities have been working with large LOBs since DB2 V7 with APARs by having separate data sets or members for each LOB. New changes allow LOBs to span records instead of having separate data sets or members.

Will dynamic statement cache handle literals for select, insert, update and delete?

The DB2 10 dynamic statement cache (DSC) enhancements consolidate all types of SQL statements and their literals. So if the same statement is used within the application with different literals DB2 will consolidate it within the DSC leaving more room for other statements and better caching overall. This behavior requires a parameter to be set, as literals can provide important information for optimization.

There was mention of compress SMF option. How does this differ from ACCUMACC?

ACCUMACC, commonly called rollup, combines multiple accounting records into one. The rollup technique is being improved to work better with packages, allowing more customers to use this option. Compression uses an SMF compression option, often reducing data volumes by 70% to 80%, while increasing CPU time by less than 1%. Customers can use either one or both techniques to reduce SMF accounting volume.

What is different between MLS and Row (and column) access control?

The primary design for MLS is strict security labels, similar to the needs for military confidentiality. MLS builds upon the concepts in z/OS and RACF. Row and column access control is more flexible, with capabilities similar to those of a view, but oriented to security.

COBOL and C Stored procedures perform better than Native Stored Procedures in DB2 9. Will native stored Procedures perform better than C and COBOL Stored procedures in DB2 10?

Native SQL procedures came very close to COBOL and C in IBM laboratory measurements on DB2 9, within about 15%. DB2 9 remote native SQL procedures used much less standard CPU than COBOL or C by using zIIP engines. A DB2 10 native SQL procedure workload has shown up to 20% CPU reduction using SET statements, IF statements and SYSDUMMY1, making them very close to COBOL and C for local calls, and able to use zIIP for remote calls.

How much more zIIP offload is expected in DB2 10? Is there a way to utilize zIIP more?

See the performance presentation for an overview. The primary changes in DB2 10 for additional use of zIIP are removing some parallel restrictions, expanded XML validation, RUNSTATS options other than column groups and inline RUNSTATS, prefetch reads and deferred writes. Some customers will find that the new zAAP on zIIP facility will reduce complexity or make zIIP more useful.