



Database Project Opening and What is New in DB2 10 for z/OS and Beyond?

Roger Miller
IBM Silicon Valley Lab

Monday, August 2, 2010: 11:00 AM-12:15 PM
Session 7959 Room 103 (Hynes Convention Center)



DB2 10 for z/OS - Helping you improve operational efficiencies and gain competitive advantage


Broadcast Date: April 13, 2010 at 11:00 a.m. Eastern Daylight Time, 4:00 p.m. UTC
Developed for: IT Managers, Application Architects, Database Administrators, System Administrators, System Programmers, and System Architects.

There is something for everyone in DB2 10 for z/OS. DB2 10 for z/OS provides the best reduction in CPU for transactions and batch for 22 years, since V2R1. We expect most customers to reduce CPU times between 5% and 10% initially, with opportunity for more. Applications which can take advantage of additional benefits, such as hash access, can have larger CPU and memory reductions. Scalability is the second major benefit, with the ability to run five to ten times as many threads in a single subsystem by moving 80% to 90% of the virtual storage above the bar. Schema evolution or data definition on demand enhancements improves availability. SQL, pureXML, and web services extend usability and application portability for this platform. Productivity improvements for application developers and for database administrators are very important as data grows in scale and complexity. Warehousing continues to evolve, with key trends matching System z and DB2 for z/OS strengths of performance, scalability, reliability, stability, availability, resilience, and security. DB2 10 has a lot for everyone. Your organization needs to lower operating costs thru CPU cycle reductions while still building a strong foundation for SOA and XML initiatives. Database Administrators (DBAs) will find improved database performance, scalability, and availability. Memory management is dramatically reduced, so growth is much simpler. DBAs also get more flexible security to help with regulatory compliance. A wide range of enhancements improve ERP application and data warehouse functionality and performance. Large object (LOB) performance and flexibility are improved. Many DBA tasks are simplified.

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

Speakers: Roger Miller, DB2 for z/OS technical evangelist, IBM Software Group and David Beulke, Dave Beulke & Associates.

2



Disclaimer/Trademarks

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

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.

SHARE in Boston

At this point, the beta announcement for DB2 10 was made February 9, 2010. The general availability announcement has not been made, so this is not an available product. We have used the code name DB2 X for z/OS up to this point. The beta program started March 12, 2010.

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>



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"> • 20%-30% Utility CPU savings • Compress indexes, save 50% disk • More CPU on specialty engines 	<ul style="list-style-type: none"> • Save up to 20% CPU batch & transactions • 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 	

SHARE in Boston

**Beta started:
Mar 12, 2010**

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 DB2 9, 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.

Top 10 in DB2 10 for z/OS



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



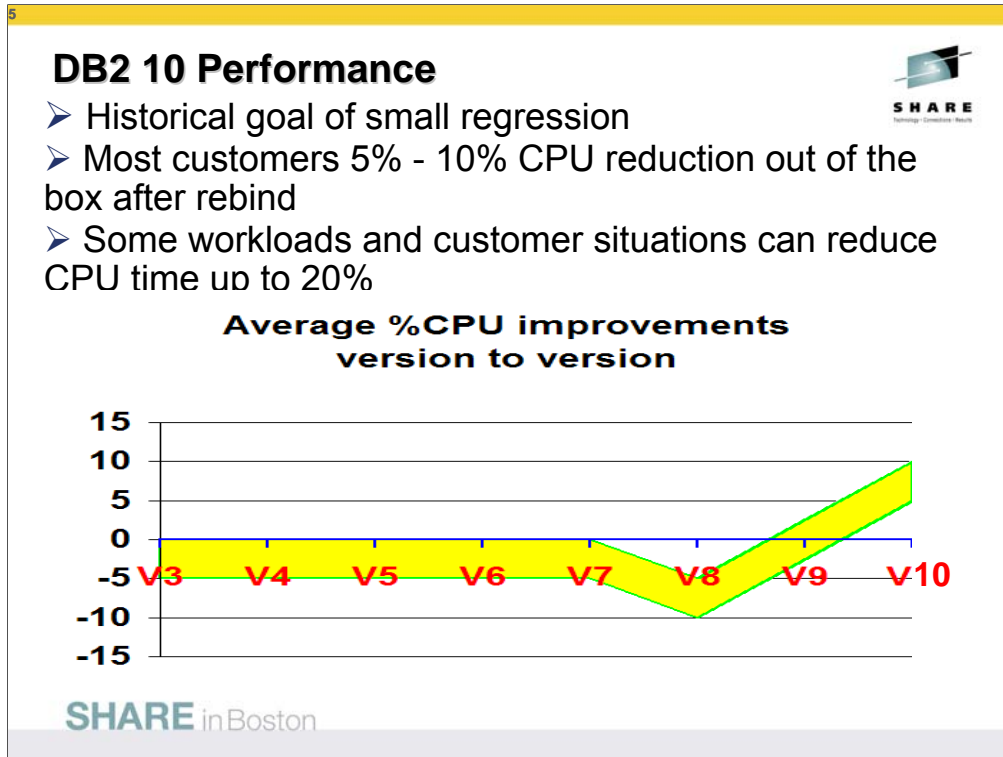
SHARE in Boston

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.



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



1. DB2 9 NFM REBIND with PLANMGMT(EXTENDED)
2. DB2 9 NFM → DB2 10 CM without REBIND
 - Measured 3.7% CPU reduction from DB2 9
3. DB2 10 CM REBIND getting same access path
 - Measured 7.4% CPU reduction from DB2 9
4. DB2 10 NFM
 - Measured same 7.4% CPU reduction from DB2 9
5. DB2 10 NFM with RELEASE (DEALLOCATE)
 - Measured additional 10% CPU reduction from DB2 10 NFM RELEASE(COMMIT)

SHARE in Boston

1. 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.
 2. In step 2, this application is moved to DB2 10 CM without a REBIND, and the result is a 3.7% reduction in CPU time.
 3. In step 3, 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.
 4. In step 4, moving to NFM, the CPU time is the same.
 5. In step 5, 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.

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
- 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



SHARE in Boston

7

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 196, zBX, z10, ...



SHARE inBoston

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.

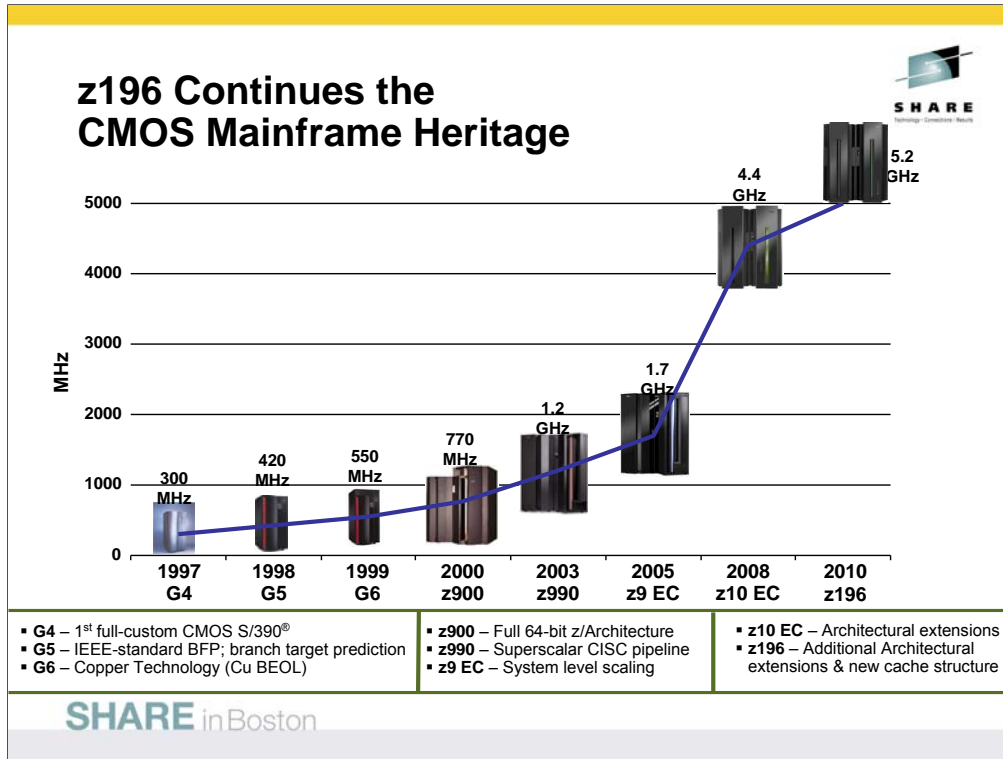
Hardware data compression and encryption provides improved costs, easier management and robust resilience for the platform. Cross-memory and protection keys work with APF authorization and RACF for the underlying system integrity.

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



System zEnterprise Benefits for DB2

Taking System z 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 cost 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



SHARE inBoston

Faster CPUs, more CPUs, more memory means better DB2 performance, scalability

Compression hardware improvements expected to increase DB2 data compression performance

192M L4 Cache expected to benefit DB2 workloads

DB2 for z/OS 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, which will remove many single system scaling inhibitors

Translation Lookaside Buffer Changes expected to improve DB2 10 performance for 1MB page sizes

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.

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 Boston

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.

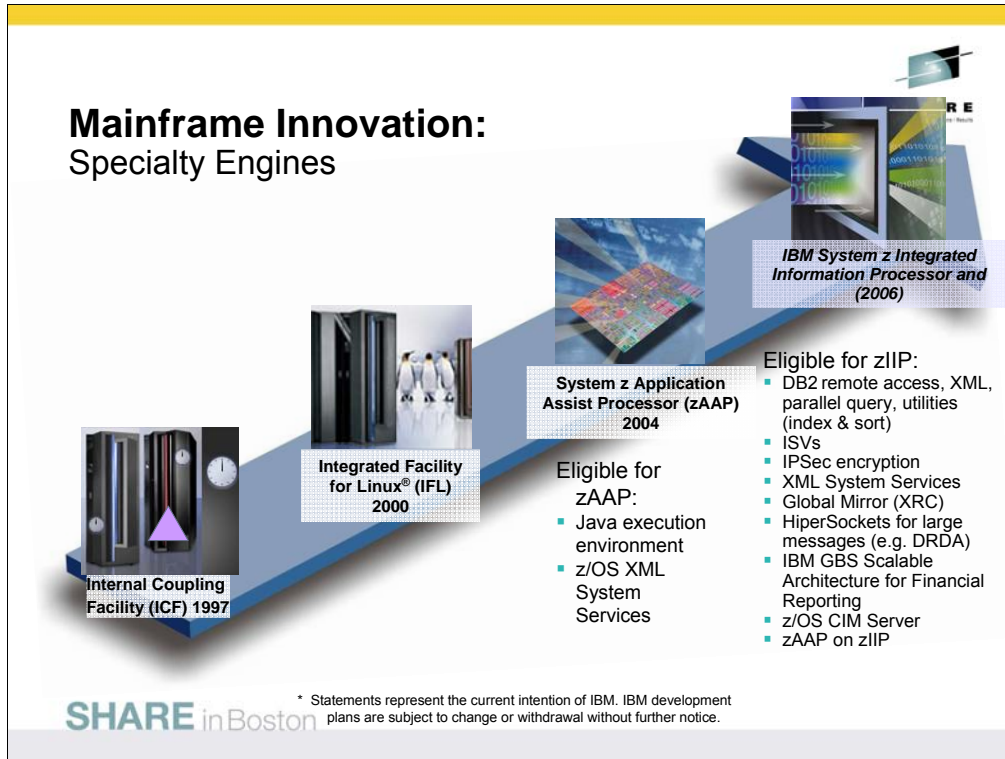
Hardware Trends Impacting DB2



- 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
- Phase change memory may be more scalable, lower cost than DRAM

SHARE in Boston

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.



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 (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 via DRDA to fully utilize zIIP
 - 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 still more queries eligible for parallelism
- 3 - DB2 Utilities LOAD, REORG & REBUILD
 - DB2 utility functions used to maintain index structure and sort
 - DB2 10 RUNSTATS
- 4 - DB2 10 Buffer Pool Prefetch

SHARE in Boston

* 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.



Performance Enhancements Requiring 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)

SHARE in Boston

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

Many paths within DB2 processing leverage better 64-bit memory capabilities. This results in better SQL performance for many existing SQL access plans.

All of the memory improvements provide immediate relief for all of many memory constrained systems. With some REBINDS, memory usage is reduced, allowing you to use memory more effectively for example in buffer pools for performance in your environment.

Better handling of singleton Cursor Selects through chaining, combines the open fetch and close work ONLY once across the network improving network bandwidth efficiency and overall performance.

Also for DDF transactions, there is increased DDF thread reuse. This enhancement starts to handle and reuse DDF threads though the same methods that we have had for a long time through CICS thread interfaces for robust consistent type transactions.

For people that are directly coming to DB2 10 through migrations from Version 8 they immediately get long list of the Version 9 enhancements especially all the Utility performance improvements. Some are experiencing 20% elapse time savings in Version 9.

Updates to index columns are done in parallel in DB2 10 improving insert performance. This out of the box enhancement along with DB2 10 List Pre-fetch capabilities improve all existing applications that use list pre-fetch activities across all existing accesses paths and especially when using indexes that may be a little disorganized and in need of a reorg.

DB2 can now use solid state disk devices. These devices are great for the workfiles, GTTs and other high performance table spaces within your environment.

Also, the enhanced way DB2 10 uses in-memory Workfiles and the improvements related to RID pool overflows helps all application systems avoid the deadly table space scan at the peak processing times.

DB2 can utilize the new bigger 1MB page size on z10 and provide additional buffer pool options to put a table fully in memory with an easy simple table space ALTER.

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

5

Performance Enhancements requiring REBIND (CM)

- Most access path enhancements
- 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
- REBIND enables further SQL runtime improvements
- If migrate from V8, get new RUNSTATS before mass rebind

SHARE inBoston

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

- Efficient caching of dynamic SQL statements with literals
- 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
- Insert improvement for universal table spaces

SHARE in Boston

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

One of the enhancements that will help everyone with large dynamic SQL applications are the improved processes that handle the Dynamic Statement Cache. They now consolidate SQL statements that are the same but have different literals. This reduces the SQL statement space used in the Dynamic Statement Cache and reuses the security and object verification and access path already developed for the SQL statement. This dramatically improves Dynamic Statement Cache cache hits, reduces the duplicate SQL statement previously held, makes room for more SQL improving the overall workload performance.

Also all the concurrency and performance enhancements through the DB2 10 Utilities improvements are available for more on-line ALTER and maintenance activities.

Better streaming and minimized LOB default sizes helps improve when LOB materialization is happening with the system and this is especially important for DDF type applications because of the network impact of large objects.


Small WORKFILES are now available for simple predicate evaluation for improving performance DB2 provides native support for the SQL procedural language eliminating the cumbersome requirement to generate a C program from the SQL procedure that would then execute as an external stored procedure. DB2 10 SQL procedures are better optimized to execute more efficiently more common constructs are optimized within the DB2 code making SQL procedures very efficient for performance within the SQL procedure language.

WORKFILE can have expanded records up to 65K so larger Joins and answers set can be generated from DB2.

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

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

9



Performance Enhancements need NFM + DBA work

- 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

SHARE inBoston


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.

2
0



DB2 10: 64 bit Virtual Storage Relief

DB2 9 helped (~ 10% – 15%)

DB2 10: 5 to 10 times more threads, up to 20,000

- Move 80% - 90% above bar
- More concurrent work
- Reduce need to monitor
- Able to consolidate LPARs
- Reduced cost
- Easier to manage
- Easier to grow

Skeleton Pool
Global Stmt Pool
DBD Pool EDMPOOL

Skeleton Pool
Global Stmt Pool Working memory
DBD Pool EDMPOOL

2GB

EDMPOOL Working memory

EDMPOOL Working memory

2GB bar

REBIND is needed for much of savings.

SHARE in Boston

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

The DB2 9 virtual storage relief was 10-15% relief. The DB2 10 relief is 80% to 90% of the DBM1 address space. The result is the ability to run much more concurrent work, with many cases of 5 to 10 times more threads. Some ECSA for instrumentation moves above the bar.


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


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

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

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
--

SHARE in Boston

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

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)

DB2 10

Coupling Technology

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

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

- 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 1/2% CPU for each member reduced, more on memory

SHARE in Boston

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

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



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

SHARE in Boston

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

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

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

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

Major changes in DB2 10 catalog & directory



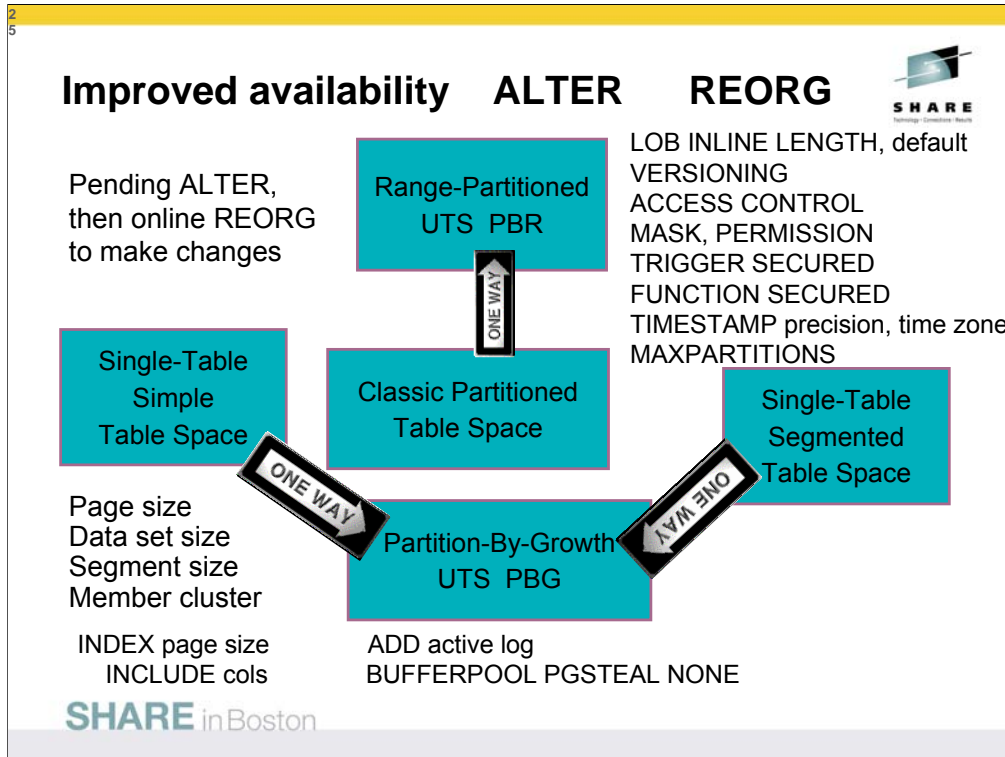
- Improve availability and productivity
- Increase maximum size substantially
- Reduce contention: BIND, DDL, utilities
- Allow SELECT from SYSLGRNX
- Catalog changes: Remove links
 - Many more table spaces, partition by growth
 - Row level locking, reordered row format
 - CLOB and BLOB columns for long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

SHARE in Boston

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

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

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





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
6

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 in Boston


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.

2
7

DB2 10: Productivity – Doing More with Less!



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



Manual invocation of

- RUNSTATS
- COPY/BACKUP SYSTEM
- QUIESCE
- MODIFY RECOVERY
- REORG

SHARE in Boston

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.



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

SHARE in Boston

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
0

Many improvements for SAP & web applications

- **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

SHARE in Boston

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 Boston

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 Boston

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 in Boston

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>

3
4



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 in Boston

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
- 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
 - increased zIIP use
 - In-memory techniques for faster query performance
- Analysis: instrumentation, Data Studio & Optim Query Tuner
- Advanced query acceleration techniques
 - IBM Smart Analytics Optimizer

SHARE in Boston

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.

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

SHARE in Boston

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>

DB2 10 for z/OS At a Glance




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• IBM Smart Analytics Optimizer


SHARE in Boston

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.

5
8



Questions?



Watch for upcoming white papers and conferences
IOD IDUG Share

SHARE in Boston

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=s_w-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/>

Detailed backup

SHARE in Boston

DB2 10 is an exciting new version for many customers. Here is some more detailed information, as DB2 10 has too much to describe in an hour.

IBM System z, Software & Workload Optimized Systems

Resilience

Scalable

Productive

Managed Secure

Low cost

Tivoli software
 Visibility, control, and automation from System z across your business
 OMEGAMON, Tivoli Service Management Center for z, TSAM, TWSz

Database
 DB2, IMS
 Design, develop, deploy, operate, optimize and govern enterprise data throughout its lifecycle

Applications

WebSphere
 Application infrastructure, connectivity and dynamic business processes
 WAS, CICS, BPM, WMQ, ESB

Rational
 Application Development Tools and Software Delivery Platform
 Compilers (C/C++, PL/I, COBOL), RDz, RTCz

Transactions Analytics

Information Management
 Strong information management platform built for business workloads
 CICS, IMS TM, FileNet, InfoSphere Warehouse, InfoSphere MDM Server, Cognos, SPSS, ISAS 9600

Lotus
 Productivity and Collaboration
 Portal, Connections, Notes Domino

SHARE in Boston

40

IBM Software exploits the capabilities of POWER7 and is integrated with IBM Systems Software and Power System, delivering...


- Huge Capacity
 - Near Linear Scaling upto 128 members
 - 64 members: 95% of linear, 128 members: 84% of linear
 - Oracle RAC typically scales poorly beyond 4 to 8 nodes for a non-partitioned data base
- 73% better performance using WebSphere on POWER7 vs. competitive application server on Nehalem
- Lotus Domino clients can consolidate thousands of users from individual Windows servers to AIX or i and save or Exchange users can move to Domino on Power and gain incredible efficiency and scalability. Lotus Domino on two virtualized POWER7 servers supporting 40,000 users costs **40% less** than Microsoft Exchange on HP Nehalem

DB2 10 for z/OS Summary



- Out-of-the-box savings by improving operational efficiencies
- Unsurpassed resiliency for business-critical information
- Rapid application and warehouse deployment for business growth
- Enhanced query and reporting facilities

SHARE in Boston



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.

Unsurpassed resiliency for business-critical information Business resiliency is a key component of the value proposition of DB2 for z/OS and the System z operating system, supporting your efforts to keep your business running even during unexpected circumstances. 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. Security improvements also contribute to robust business resiliency.

Continuous availability enhancements Version 10 provides online schema enhancements that allow you to make changes to database objects (indexes and table spaces) while maximizing the availability of the altered objects. Through enhancements to ALTER statements, you can now change indexes and table spaces without having to unload the data, drop and re-create the objects, regenerate all of the security authorizations, re-create the views, and reload the data. Now, the changes are materialized when the altered objects are reorganized.

In addition, Version 10 improves the usability and performance of online reorganization in several key ways. This release of DB2 for z/OS supports the reorganization of disjoint partition ranges of a partitioned table space, and improves SWITCH phase performance and diagnostics. Also, Version 10 removes restrictions that are related to the online reorganization of base table spaces that use LOB columns.

Reduced catalog contention In Version 10, the DB2 catalog is restructured to reduce lock contention by removing all links in the catalog and directory. In addition, new functionality improves the lock avoidance techniques of DB2, and improves concurrency by holding acquired locks for less time and preventing writers from blocking the readers of data. In Version 10 new-function mode, you can access currently committed data to dramatically minimize transaction suspension. Now, a read transaction can access the currently committed and consistent image of rows that are incompatibly locked by write transactions without being blocked. Using this type of concurrency control can greatly reduce timeout situations between readers and writers who are accessing the same data row. In addition, this new capability provides flexibility and increased performance to applications that require only available and committed data to be returned from DB2 tables. **Virtual storage relief** Enhancements in Version 10 substantially increase the amount of virtual storage. This release of DB2 for z/OS moves most memory to 64-bit, which provides virtual storage relief and can greatly improve the vertical scalability of your DB2 subsystem. In addition, a 64-bit ODBC driver is now available. The 64-bit ODBC driver runs in 64-bit addressing mode and reduces the virtual storage constraint, because the driver can accept 64-bit user data pointers and access user data above the 2 GB bar in the application address space. Extended support for LOB and XML streaming, improvements to LOB and XML processing, and enhancements to the FETCH statement for LOB and XML data also contribute to reducing virtual storage consumption, and improve the stability of your DB2 subsystem.

Security enhancements This release of DB2 for z/OS provides critical enhancements to security and auditing. These enhancements strengthen DB2 security in the z/OS environment. For example, Version 10 provides increased granularity for DB2 administrative authority, and offers a new a DB2 data security solution that enables you to manage access to a table at the level of a row, a column, or both. In addition, you can define and create different audit policies to address the various security needs of your business.

Fast application and warehouse deployment for business growth Staying competitive in today's global economy is tougher than ever, requiring agility, adaptability, and responsiveness. SQL, pureXML®, and optimization enhancements in Version 10 help extend usability, improve performance, and ease application portability to DB2 for z/OS. These improvements support your efforts to create a sustainable, competitive advantage.

SQL improvements SQL enhancements deliver new function for improved productivity and improve DB2 family consistency and the ability to port to DB2 for z/OS from other platforms and database management systems. Enhancements are provided for SQL scalar functions. SQL table functions are added. Native SQL procedure language enhancements make porting easier from other database management systems. Implicit casting makes porting simpler, as DB2 SQL is more consistent with other products and platforms. Allowing more flexibility in the number of digits for fractions of seconds and allowing timestamps with time zones simplify porting. Moving sums and moving averages help in warehouse queries and in porting.

Temporal tables and versioning In this release of DB2 for z/OS, you have a lot of flexibility in how you can query data based on periods of time. DB2 supports two types of periods, which are the system time (SYSTEM_TIME) period and the business time (BUSINESS_TIME) period. The SYSTEM_TIME period is a system-maintained period in which DB2 maintains the beginning and ending timestamp values for a row. For the BUSINESS_TIME period, you maintain the beginning and ending values for a row. In addition, Version 10 introduces versioning, which is the process of keeping historical versions of rows for a temporal table that is defined with a SYSTEM_TIME period, or both time periods.

4
2

Software > Information Management > DB2 Product Family > DB2 for z/OS >

Announcing DB2 10 for z/OS Beta
The undisputed leader in total system availability, scalability, security and reliability

DB2 for z/OS Family
Library
Technical resources
Success stories
News
How to buy
Events
Training and certification
Services
Support

Introducing DB2 10 for z/OS Beta

In today's business and economic environment, the challenge for IT is clear: improve operational efficiencies, reduce costs, and adapt quickly to support business growth -- all without sacrificing the resiliency required for today's demanding business requirements. DB2 for z/OS is the undisputed leader in total system availability, scalability, security, and reliability at the lowest cost per transaction. DB2 10 builds on the formidable capabilities of DB2 9 for z/OS and continues to set the standard, delivering key innovations and resource savings, including:

Out-of-the-box Savings by improving operational efficiencies
IBM continues to invest in new features to support your efforts to make your business more efficient, and DB2 10 delivers great value in this area. Compared to previous DB2 versions, some customers can achieve a 5% to 10% out-of-the-box CPU savings for traditional workloads and up to 20% out-of-the-box CPU savings for non-traditional workloads. Productivity improvements in DB2 10 for database and system administrators can drive additional operational efficiencies and cost savings. Synergy with other IBM System z platform components reduces CPU use by leveraging the latest processor improvements, larger amounts of memory, solid-state disk and z/OS enhancements.

DB2 10 for z/OS Highlights
DB2 10 for z/OS Beta
Cut Costs & Improve Performance
→ Announcement Letter
→ DB2 for z/OS and Data Warehousing Nordic Seminars
→ Hear what people are saying about DB2

DB2 for z/OS Roadshow - including DB2 10 for z/OS Sneak Preview
Coming to a city near you!
Feb 9 St. Louis, Feb 11 NYC, Feb 18 Toronto, Mar 4 Chicago and many more
→ Register Now - Track 2

DB2 10 for z/OS TECHNICAL PREVIEW

SHARE inBoston

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

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

The announcement itself is here:

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

A version if this presentation is on the web.

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

When should I migrate to DB2 for z/OS Version... DB2 10 and migration

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

Announcing DB2 10 for z/OS Beta

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

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

DB2 10 Roadshow <http://www.ibm.com/software/os/systemz/summit/handouts/>

Share <http://www.share.org>

http://ew.share.org/proceedingmod/project.cfm?project_id=4189&conference_id=22

II14477: DB2 9 MIGRATION/FALLBACK INFOAPAR TO/FROM DB2 10

<http://www.ibm.com/support/docview.wss?uid=isg1114477>

II14474: DB2 V8 MIGRATION/FALLBACK INFOAPAR TO/FROM DB2 10

<http://www.ibm.com/support/docview.wss?uid=isg1114474>



DB2 10 Performance, Scalability Objectives

- Significant scalability and performance improvements
 - Synergy with latest System z hardware & software
 - High n-way scalability
 - Large real memory exploitation
 - Hardware level optimization
 - Improve transaction and response times
 - Lower CPU demands for large & small DB2 subsystems

- Virtual storage is no longer a customer vertical constraint
 - Expanded number of concurrent threads
 - Expect ten times current thread maximum
 - Addressed next tier of constraints: latches & concurrency

SHARE in Boston


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

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

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

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

4



DB2 10 for z/OS What's exciting?

Efficiency	<ul style="list-style-type: none"> • CPU reduced: transactions & queries • Ten times more concurrent users
Resiliency	<ul style="list-style-type: none"> • More online schema changes • Concurrency for catalog & utilities • Improved security controls and audit
Applications	<ul style="list-style-type: none"> • Versioned data or temporal queries • pureXML and SQL enhancements

→ Productivity improved for DBAs, application programmers, & systems

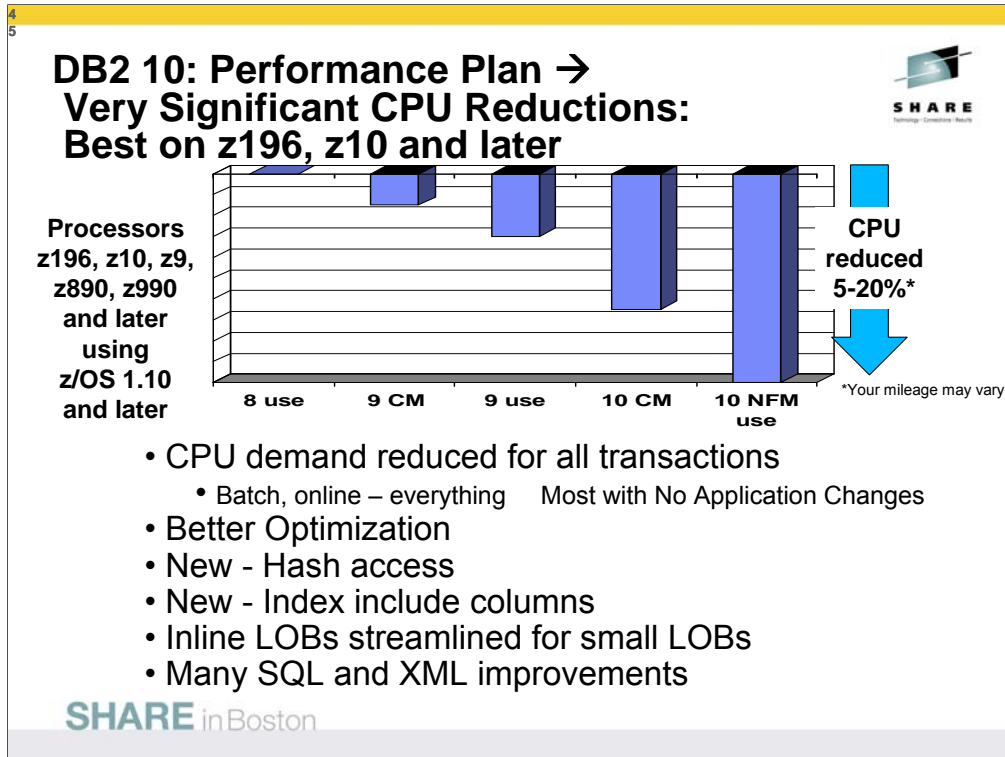
SHARE in Boston

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

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

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

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

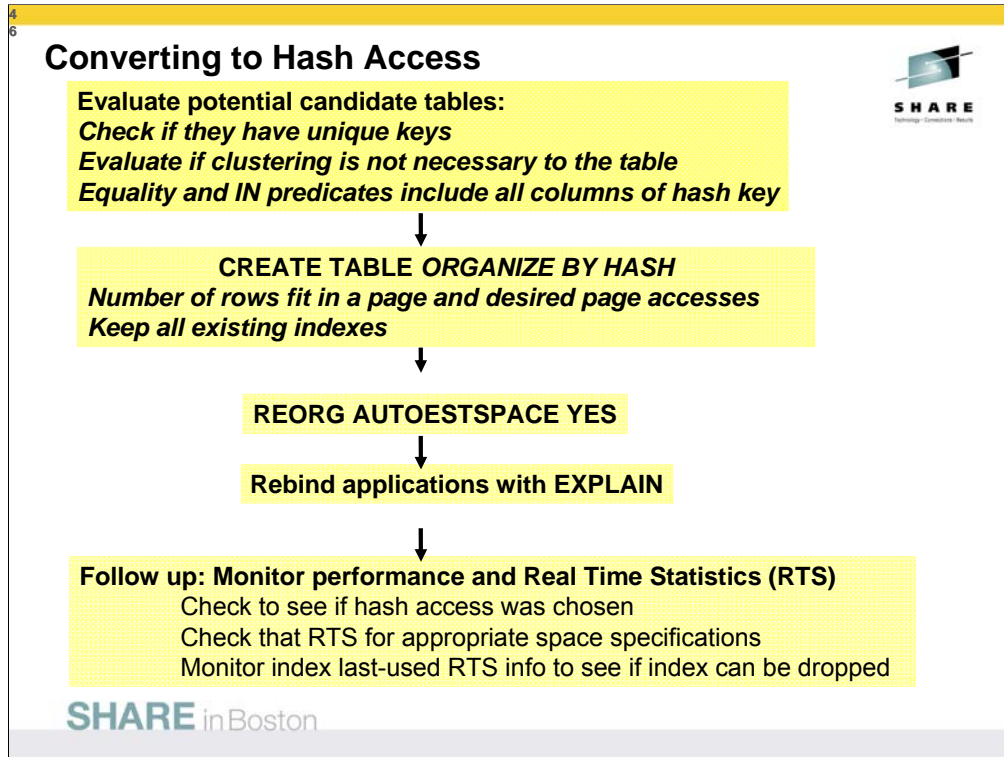


Reducing CPU from DB2 9 to DB2 10 without significant administration or application changes is the primary thrust of the performance work. Most of the changes work with CPU caching and path lengths, so that applications are not changed. We can take advantage of new instructions without needing to have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

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

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

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



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

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

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



DB2 10 for z/OS Resiliency

Scalability	<ul style="list-style-type: none">• Ten times more concurrent users• More concurrent utilities
Availability	<ul style="list-style-type: none">• More online schema changes• Improved concurrency on user & catalog data
Audit and Security	<ul style="list-style-type: none">• Improved audit and control• Increase administrative authority granularity• Masking for data in database

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

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

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

Catalog Restructure for improved concurrency

- Remove links from the catalog and directory
 - DSN1CHKR no longer needed
- Eighty new table spaces; 7 old ones dropped
- Row level locking enabled for catalog and directory
- Conversion during ENFM for migrated systems
- BIND, PREPARE, and DDL will run with better concurrency, fewer timeouts/deadlocks
- Allow online REORG for all catalog and directory table spaces.

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


Other Catalog Changes

- Partition-by-growth (PBG) catalog/directory table spaces
 - Allow packages to grow beyond 64GB using LOBs
 - V8/DB2 9 APAR PK80375 adds zparm for SPT01 compression; no compression in DB2 10
- DB2 managed catalog and directory data sets
 - SMS-controlled, DB2-managed catalog required
 - Eases admin and management burden
- New CLOB & BLOB columns for storing SQL statements
 - Today, SQL statements can be split into several records with sequence numbers
 - CLOBs will make it easier to query SQL statements
- Online REORG enabled for all catalog/directory objects

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

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

5
0



Availability

- Online schema changes for table spaces, tables and indexes – PENDING with ALTER and Online REORG instead of DROP/CREATE or REBUILD INDEX

Alterations occur with REORG, unless noted otherwise

- Page size for table spaces and indexes BUFFERPOOL
- DSSIZE for table spaces
- SEGSIZE
- Convert single table segmented into UTS PBG
- Convert single table simple into UTS PBG
- Convert classic partitioned table space into UTS PBR
- Convert UTS PBR to UTS PBG
- Convert PBG to hash (immediate, but RBDP index)
- MEMBER CLUSTER
- Ability to drop pending changes

SHARE in Boston

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.

5
1

Availability ...



- More ALTERs (not pending)
 - LOB INLINE LENGTH, default
 - VERSIONING
 - ACCESS CONTROL
 - MASK, PERMISSION
 - TRIGGER SECURED
 - FUNCTION SECURED
 - TIMESTAMP precision, time zone
 - INDEX INCLUDE cols
 - BUFFERPOOL PGSTEAL NONE
 - MAXPARTITIONS
- Online REORG for LOBs, other Online REORG / utility improvements & restriction removal
- Online add active log

SHARE in Boston

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.

Data Sharing Improvements



- ACCESS DATABASE command wildcarding support - DB2 9 PK80925
- Sub-group attach (DB2 9 usermod)
- Vertical scalability improvement, ability to reduce members, LPARs
- Buffer pool scan avoidance
- Delete data sharing member – offline for data sharing group
 - Offline utility for “deactivate”, “reactivate”, “destroy”
- MEMBER CLUSTER support for universal table spaces
- DDF Restart Light enhancements: Handle DDF indoubt URs
- Auto rebuild CF lock structure on long IRLM waits during restart
 - Can avoid group-wide shutdowns
- LRSN spin avoidance for inserts to the same page (e.g. Multi Row Insert)
- IFCID 359 for index split
- DB2 group restart avoidance for lock resolution problems
- New dsnzparm to force deletion of CF structures on group start (e.g. disaster recovery testing)

SHARE inBoston

Many improvements help with data sharing. Some delivered in the DB2 9 service stream, such as the ability to use wildcards with the ACCESS command.

New support for MEMBER CLUSTER on universal table spaces can improve insert performance substantially.

Being able to make change to distributed connections dynamically, instead of by stopping and restarting DDF is a major change for increased availability.

5
3

DB2 10 Security Benefits

- More flexible authorization
- Separation of duties
- Do job without access to data
- Policies for audit
- Simpler control
- Tighter security
- Avoid cascade delete
- Avoid views and application security logic
- Allow more tools
- Evolve security policies
- Easier to manage security policy

→ Improved productivity & tighter security

SHARE
Technology. Connected. Ready.

SHARE in Boston

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

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

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

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

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

Eliminating views and security logic in applications

Preventing access from authorities including SYSADM and DBADM

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

Making the evolution of security policies easy to deal with

Relieving the difficulties of managing security policies



DB2 10 more utilities enhancements

- Support of spanned records for UNLOAD of LOB data
 - Currently unload of LOBs >32K must use FRVs
 - Allow inline LOBs with base row in unload data set
 - Provides portability of data
 - Performance enhancement for FRV processing with PDS data sets, also in DB2 9
 - UNLOAD 33% elapsed time reduction
 - LOAD 84% elapsed time reduction
- Autonomic RUNSTATS & table profile

SHARE in Boston

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




DB2 10: More Utility Improvements

- Improved COPY CHANGELIMIT performance
 - Use RTS instead of SM page scans
- Data set level FlashCopy option
- FlashCopy backups with consistency and no application outage
- FlashCopy backups as input to:
 - RECOVER (fast restore phase)
 - COPYTOCOPY, DSN1COPY
- RECOVER “back to” log point
- REPORT RECOVERY for system level backups

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

5
6



Autonomics and DBA Productivity...

- Checkpoint intervals based on both time and log records
- Run 'must complete' backout under pre-emptable SRB
- Identify unused packages
- SQL Statement level monitoring
 - Statement ID introduced
 - Trace records & messages extended to include statement ID
 - New trace class for statement detail
 - GetPages, Locks, I/Os, cpu/elapsed time, etc. at statement level
- Manage max threads, connections, idle thread timeout on an application basis
 - Warning or exceptions issued when threshold is hit

SHARE in Boston

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

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

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

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

DB2 10 Application Enablement, Portability



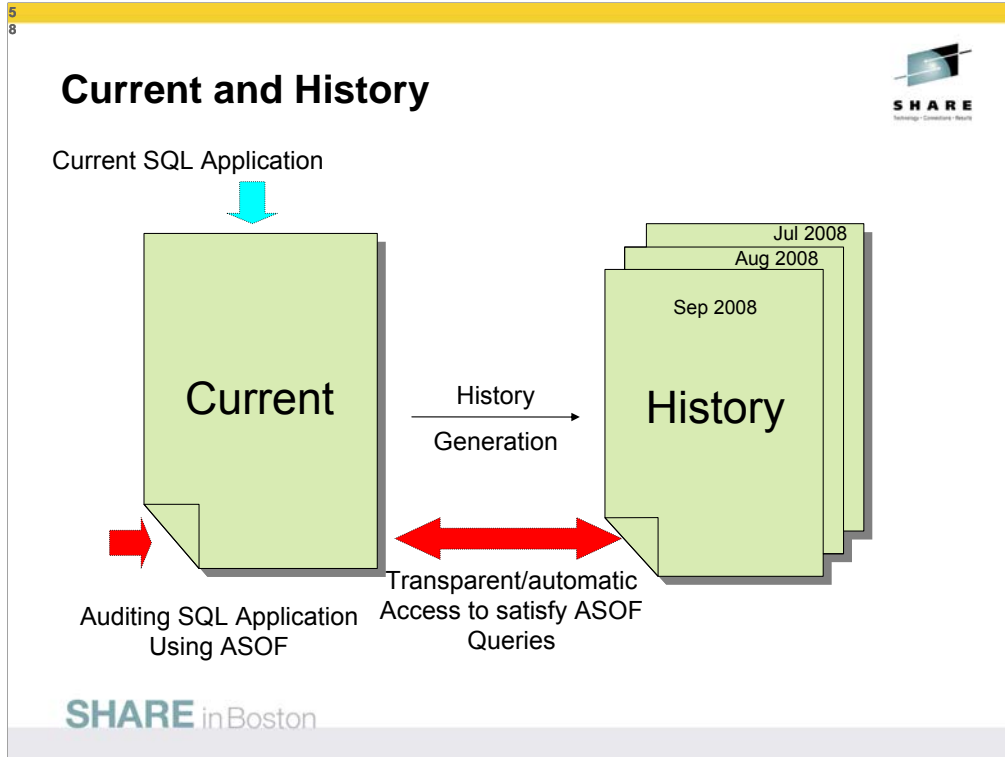
- SQL stored procedure enhancements
 - SQL PL in Scalar UDFs & XML support
- 64-bit ODBC – also DB2 9 PK83072
- Special null indicator to indicate value not supplied or default
- DRDA support of Unicode for system code points
- Instance based statement hints
- Allow caching of dynamic SQL statements with literals
- Data-dependent paging
 - When only a specific part of the result set is needed
 - Efficient access to desired portions of result set, based upon current position

SHARE in Boston

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

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

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



In DB2 10, you can create a versioned data or a temporal table, which is a base table with one or more time periods defined. DB2 supports two types of time periods, system time period and business time period. The system time period is a system-maintained period which reflects the time the data was changed in the system. The business time period is a user-specified period where the business event time is recorded for the row data.

The SYSTEM_TIME period is meaningful because of versioning. Versioning specifies that old rows are archived into another History table. When the temporal table is defined with a History table the logical table contains all the rows but physically and automatically separates the operational data away from the history data. This improves performance of the operations and provides automatic migration of the old data. Many older systems were designed this way and now DB2 provides facilities for referencing all the old and new data as one table through the new SQL FROM clause ASOF.

Using both SYSTEM TIME and BUSINESS TIME 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 and a lot of flexibility for querying periods of time along with calculating moving averages and moving sums.

5
9

DB2 SQL 2001

z z/OS V7
common
luw Linux, Unix & Windows V8.2




z { Range partitioning

c
o
m
m { Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index Support, Self Referencing Updates with Subqueries, Sort Avoidance for ORDER BY, and Row Expressions, Call from trigger, statement isolation

l
u
w { Updateable UNION in Views, ORDER BY/FETCH FIRST in subselects & table expressions, GROUPING SETS, ROLLUP, CUBE, INSTEAD OF TRIGGER, EXCEPT, INTERSECT, 16 Built-in Functions, MERGE, Native SQL Procedure Language, SET CURRENT ISOLATION, BIGINT data type, file reference variables, SELECT FROM INSERT, UPDATE, or DELETE, multi-site join, 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, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT

SHARE in Boston

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

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

DB2 SQL 2004

z/OS V8
common
luw Linux, Unix & Windows V8.2

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, session variables, range partitioning

**c
o
m
m
o
n** { Inner and Outer Joins, Table Expressions, Subqueries, GROUP BY, Complex Correlation, Global Temporary Tables, CASE, 100+ Built-in Functions including SQL/XML, Limited Fetch, Insensitive Scroll Cursors, UNION Everywhere, MIN/MAX Single Index Support, Self Referencing Updates with Subqueries, Sort Avoidance for ORDER BY, and Row Expressions, 2M Statement Length, GROUP BY Expression, Sequences, Scalar Fullselect, Materialized Query Tables, Common Table Expressions, Recursive SQL, CURRENT PACKAGE PATH, VOLATILE Tables, Star Join Sparse Index, Qualified Column names, Multiple DISTINCT clauses, ON COMMIT DROP, Transparent ROWID Column, Call from trigger, statement isolation, FOR READ ONLY KEEP UPDATE LOCKS, SET CURRENT SCHEMA, Client special registers, long SQL object names, SELECT from INSERT

**l
u
w** { Updateable UNION in Views, ORDER BY/FETCH FIRST in subselects & table expressions, GROUPING SETS, ROLLUP, CUBE, INSTEAD OF TRIGGER, EXCEPT, INTERSECT, 16 Built-in Functions, MERGE, Native SQL Procedure Language, SET CURRENT ISOLATION, BIGINT data type, file reference variables, SELECT FROM UPDATE or DELETE, multi-site join, MDC

SHARE in Boston

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

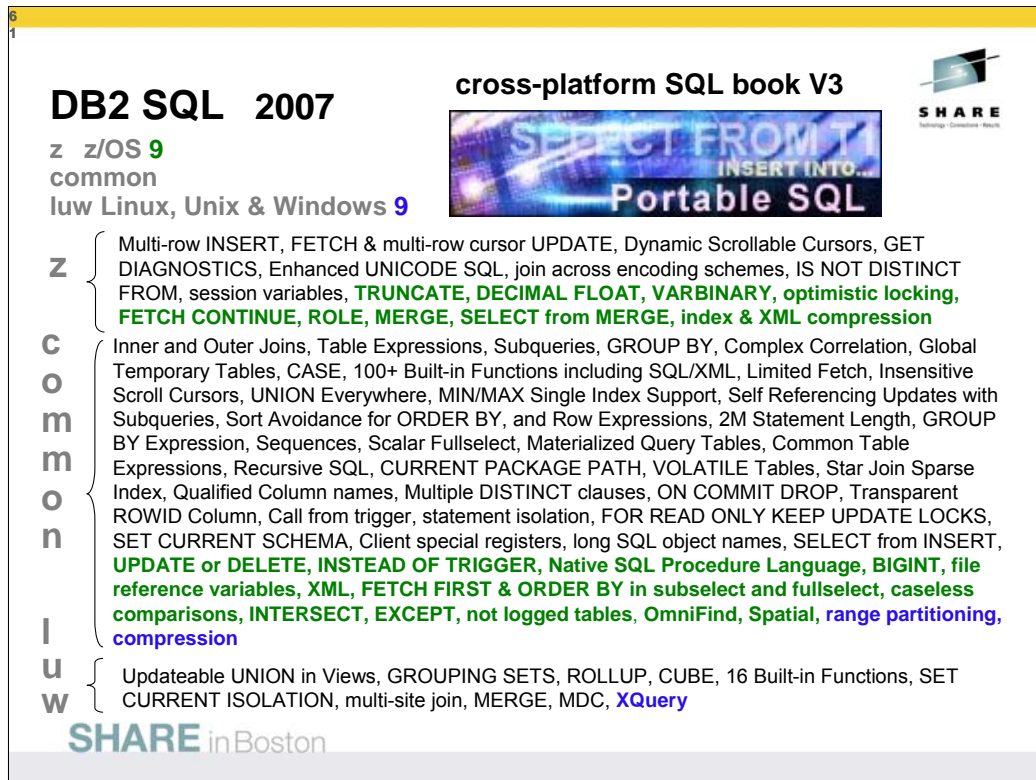
There are three sets of SQL noted above, with some that is unique to DB2 for z/OS in the first group, SQL that is common across DB2 for Linux, Unix, Windows and z/OS in the large group in the middle, then SQL that is unique to DB2 for Linux, Unix and Windows in the bottom group. Sheryl Larsen provided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

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

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

● Cross-Platform Development,

<http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html>



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

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

Cross-Platform Development Version 3,


<http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html>

ftp://ftp.software.ibm.com/ps/products/db2/info/xplatsql/pdf/en_US/cpsqlrv3.pdf

6
2

DB2 SQL 2008 **cross-platform SQL book V3.1**

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



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

**c
o
m
m
o
n** { 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**

**l
u
w** { 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, more vendor syntax**

SHARE in Boston

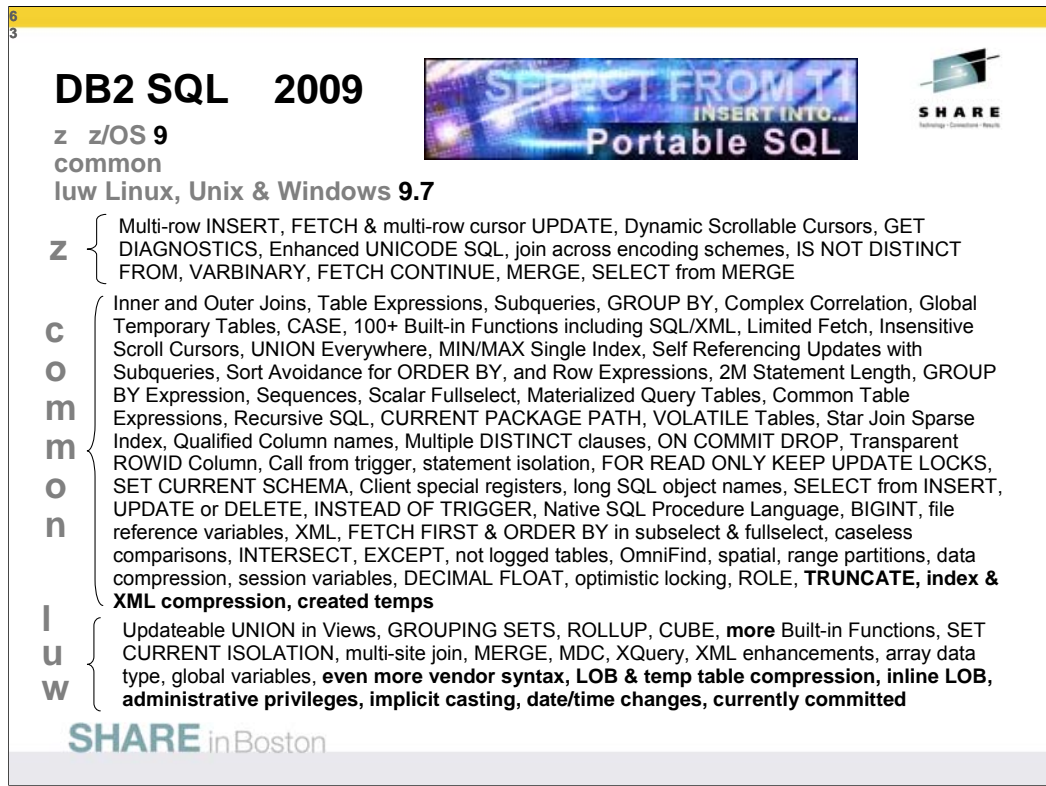
This chart shows the 2008 relationship of DB2 for Linux, Unix & Windows with DB2 for z/OS. This step in the process is DB2 9 for z/OS, (DB2 9). DB2 9 moved about half of the LUW unique items into the common set and adds a little more that is unique to the z platform. DB2 9.5 for LUW, delivered in 2008.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

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

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

Cross-Platform Development Version 3.1,

<http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html>



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

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

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

Cross-Platform Development Version 3.1,

<http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html>

DB2 for z/OS Technical Strategy

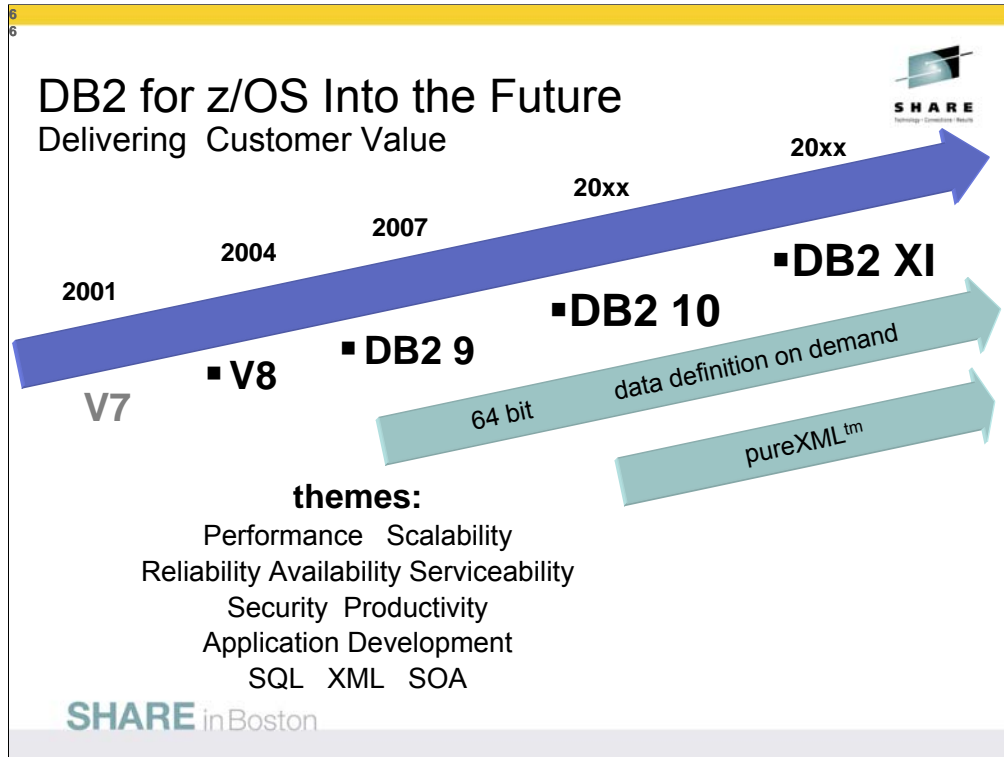


- Application enablement
 - Apps can easily connect to DB2 from anywhere
 - Advanced SQL, XML capability, application portability
- Extend the lead in availability, scalability and performance.
 - Parallel Sysplex: the best scale-out solution in the industry
 - Tight integration between DB2 and the System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



DB2 10 Major Focus Areas

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



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

The themes for future versions will continue to focus on core platform strengths of performance, scalability, reliability, stability, availability, resilience, and security. PureXML and Schema evolution or data definition on demand will be ongoing for a long time. In contrast, most of the 64 bit evolution should be completed in DB2 10.

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

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

DB2 9 and 10 in IBM Redbooks Publications

1. DB2 9 Technical Overview SG24-7330
2. DB2 9 Performance Topics SG24-7473 updated Dec. 2009
3. DB2 9 Stored Procedures SG24-7604
4. Index Compression with DB2 9 for z/OS redp4345
5. SQL Reference for Cross-Platform Development
6. Enterprise Database Warehouse, SG24-7637
7. 50 TB Data Warehouse on System z, SG24-7674
8. New Tools for Query Optimization SG24-7421
9. LOBs with DB2 for z/OS SG24-7270
10. Deploying SOA Solutions SG24-7663
11. Enhancing SAP - DB2 9 SG24-7239
12. SAP Application on Linux z SG24-6847
13. Best practices SAP BI - DB2 9 SG24-6489-01
14. Data Sharing in a Nutshell, SG24-7322
15. Securing DB2 & MLS z/OS SG24-6480-01
16. Data Sharing: Distributed Load Balancing & Fault Tolerance Configuration redp4449
17. Considerations on Small & Large Packages redp4424
18. Backup and Recovery Considerations redp4452
19. Powering SOA with IBM Data Servers SG24-7259
20. Packages Revisited, SG24-7688
21. Data Studio V2.1 Web Services redp4510
22. Ready to Access Solid-State Drives redp4537
23. Distributed Functions SG24-6952
24. Buffer Pool Monitoring & Tuning redp4604
25. Securing & Auditing Data SG24-7720
26. Serialization and Concurrency SG24-4725-01 new
27. Utilities SG24-6289-01 newest
28. Integrated Data Management Service Level Objectives, SG24-7769

SHARE in Boston

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

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

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

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



DB2 10 for z/OS

Planning to migrate




DB2 10 is an exciting new version for many customers, even though it is still in beta. Many customers will want to know what to do to position themselves for it. Step 1 is wait for general availability and additional information about the new version.

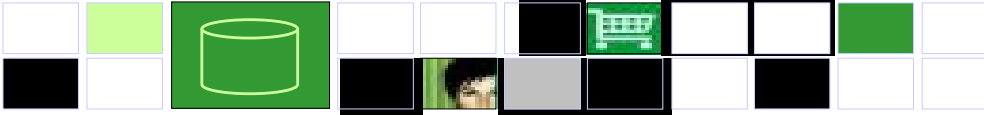
Here are some of the early planning options and work that customers can do to get ready for DB2 10.

6
9

Why Migrate to DB2 10 for z/OS?



- Business needs to save money
 - **Reduce CPU demand**
 - **Improve performance**
- Application developers need improved productivity and integration
 - pureXML for a faster, more capable interface to XML data
 - Powerful new SQL temporal enhancements & portability
- Database Administrators need
 - Improved performance
 - Availability, scalability & memory management
 - Simpler security and regulatory compliance
 - More productive database administration
- On DB2 V8?, then 9 too



SHARE in Boston

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

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


Management needs to cut costs and DB2 10 offers CPU demand reduction and improved performance. Management will be focused on these advantages as the information about DB2 10 is quickly embraced throughout the industry.

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

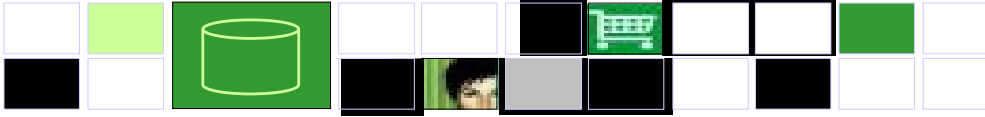
Application developers are most excited by the second release of pureXML, which improves the SQL and XML interface to access XML data stored in a native format. Application developers need powerful new SQL enhancements. Improved SQL and data definition compatibility with other DB2 platforms makes porting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

7
0

Why Migrate to DB2 9 for z/OS?



- Business needs to save money
 - Reduce CPU time & disk space
 - Improve business agility
 - Service Oriented Architecture
- Application developers need
 - PureXML for a powerful SQL and XML interface to XML data
 - Powerful new SQL enhancements
 - Portability with SQL and data definition compatibility
- Database Administrators need
 - Improve availability and performance
 - More flexible security and easier regulatory compliance
 - Better web application & data warehouse function and performance
 - LOB function, performance, usability

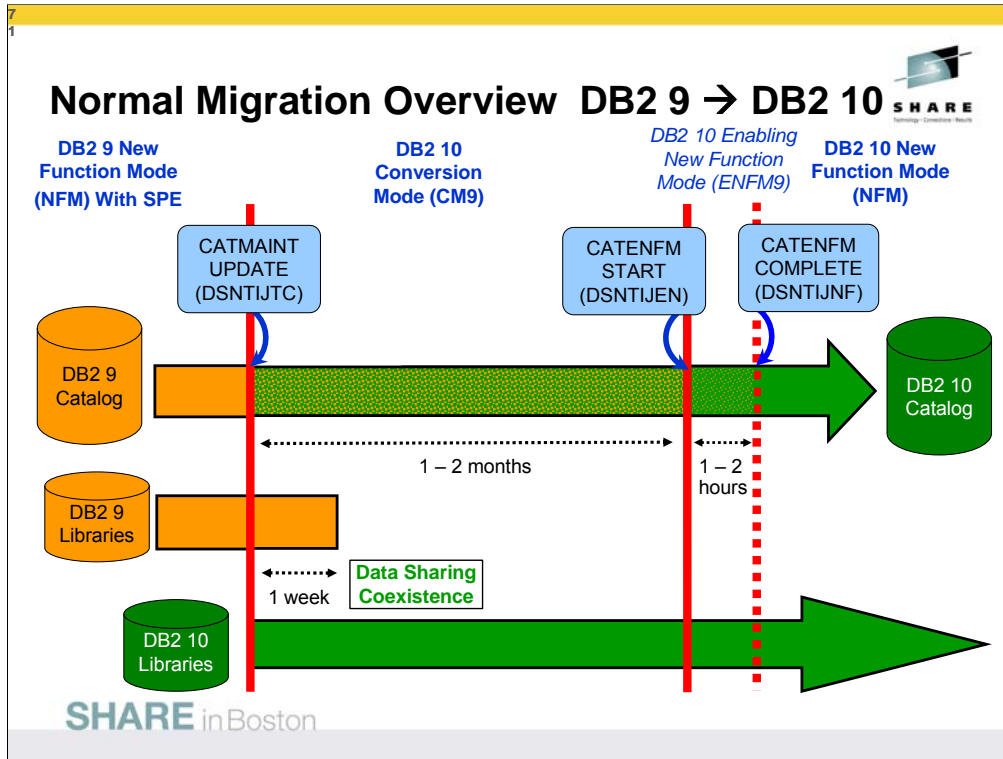


SHARE in Boston

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

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

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



Migration from DB2 9 to DB2 10

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

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

- CM9 Conversion Mode entered when migrating from DB2 9 NFM to DB2 10
- ENFM9 Enabling New Function Mode on a system that migrated from DB2 9 NFM to DB2 10. Once this mode has been entered the system can not fallback to DB2 9 and a DB2 9 member can not be started in a data sharing group.
- NFM This is the NFM mode when all system changes are made on a system that migrated from DB2 9 NFM to DB2 10 and the system is ready for DB2 10 new function
- CM9* The system migrated from DB2 9 NFM to DB2 10 and at one point was in either ENFM9 or NFM on DB2 10.
- ENFM9* The system was migrated from DB2 9 NFM to DB2 10 and at one point was in NFM on DB2 10.

Some migration considerations are:

- A DB2 9 system that has started the migration to DB2 10 can only fallback to DB2 9.
- A data sharing group that migrated from DB2 9 NFM to DB2 10 can not have any V8 members.

Note – this is not necessarily to scale!

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

Overview of Modes when migrating 9 → 10



CM9 Conversion Mode – The mode DB2 is in when DB2 10 is started for the first time after migrating direct from DB2 9. It will still be in CM9 when migration job DSNTIJTC has completed. Very little new function can be executed in CM9. Data sharing systems can have DB2 9 and DB2 10 members in this mode. DB2 can only migrate to CM9 from DB2 9 NFM.

ENFM9 Enabling New Function Mode - This mode is entered when CATENFM START is executed (the first step of job DSNTIJEN). DB2 remains in this mode until all the enabling functions are completed. Data sharing systems can only have DB2 10 members in this mode.

NFM New Function Mode - This mode is entered when CATENFM COMPLETE is executed (the only step of job DSNTIJNF). This mode indicates that all catalog changes are complete and new function can be used.

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

CM9* This is the same as CM9 but the * indicates that at one time DB2 was at a higher level. Objects that were created at the higher level can still be accessed. When DB2 is in CM9* it can not fallback to DB2 9 or coexist with a DB2 9 system.

SHARE in Boston

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

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

- Place the DB2 subsystem in enabling-new-function mode 9 (ENFM9).
- The first time that you run job DSNTIJEN, DB2 saves the RBA or LRSN of the system log in the BSDS.

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

- Migration job DSNTIJCS takes the system to CM9*.
- Migration job DSNTIJNF takes the system to DB2 10 NFM.

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

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

Migration job DSNTIJEN

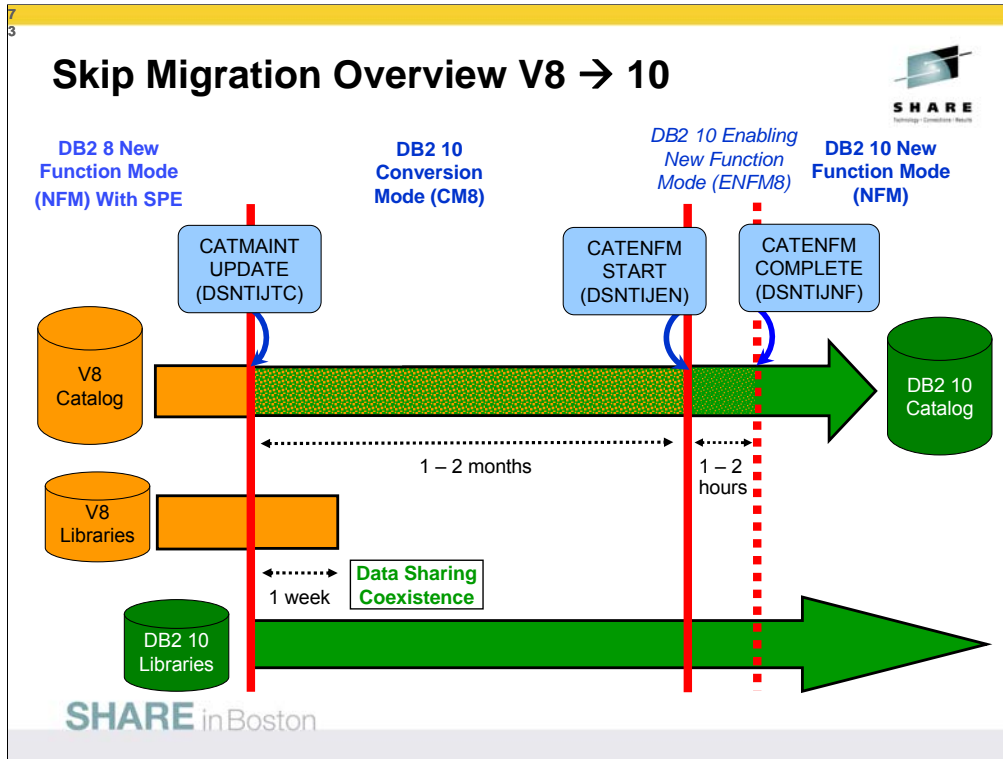
- If DB2 went from NFM or ENFM9* to CM9*, DSNTIJEN takes it to ENFM9*.
- If DB2 went from ENFM9 to CM9*, it completes ENFM9 processing and leaves DB2 in ENFM9.

Migration job DSNTIJNF

- If DB2 went from NFM or ENFM9* to CM9*, DSNTIJNF takes it to NFM.
- If DB2 went from ENFM9 to CM9*, DSNTIJNF does one of two things:
 - If ENFM9 processing had completed it takes DB2 to NFM.
 - If not, DSNTIJNF indicates that DSNTIJEN must be run.

Migration job DSNTIJES

- If DB2 went from NFM or ENFM9* to CM9*, DSNTIJES takes it to ENFM9*.
- If DB2 went from ENFM9 to CM9*, DSNTIJES indicates that DSNTIJEN must be run.



Migration from V8 direct to 10

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

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

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

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

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

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

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

Some migration considerations are:

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

Some differences in the C8 and C9 catalog are:

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

Note – this is not necessarily to scale!

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

Overview of Modes when migrating V8 → 10



CM8 Conversion Mode - This is the mode DB2 is in when DB2 10 is started for the first time after migrating direct from DB2 V8. It will still be in CM8 when migration job DSNTIJTC has completed. Very little new function can be executed in CM8. Data sharing systems can have DB2 V8 and DB2 10 members in this mode. DB2 can only migrate to CM8 from DB2 V8 NFM.

ENFM8 Enabling New Function Mode - This mode is entered when CATENFM START is executed (the first step of job DSNTIJEN). DB2 remains in this mode until all the enabling functions are completed. Data sharing systems can only have DB2 10 members in this mode.

NFM New Function Mode - This mode is entered when CATENFM COMPLETE is executed (the only step of job DSNTIJNF). This mode indicates that all catalog changes are complete and new function can be used.

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

CM8* This is the same as CM8 but the * indicates that at one time DB2 was at a higher level. Objects that were created at the higher level can still be accessed. When DB2 is in CM8* it can not fallback to DB2 V8 or coexist with a DB2 V8 system.

SHARE in Boston

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

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

- Place the DB2 subsystem in enabling-new-function mode 8 (ENFM8).
- The first time that you run job DSNTIJEN, DB2 saves the RBA or LRSN of the system log in the BSDS.

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

- Migration job DSNTIJCS takes the system to CM8*.
- Migration job DSNTIJNF takes the system to DB2 10 NFM.

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

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

Migration job DSNTIJEN

- If DB2 went from NFM or ENFM8* to CM8*, DSNTIJEN takes it to ENFM8*.
- If DB2 went from ENFM8 to CM8*, it completes ENFM8 processing and leaves DB2 in ENFM8.


Migration job DSNTIJNF

- If DB2 went from NFM or ENFM8* to CM8*, DSNTIJNF takes it to NFM.
- If DB2 went from ENFM8 to CM8*, DSNTIJNF does one of two things:
 - If ENFM8 processing had completed it takes DB2 to NFM.
 - If not, DSNTIJNF indicates that DSNTIJEN must be run.

Migration job DSNTIJES

- If DB2 went from NFM or ENFM8* to CM8*, DSNTIJES takes it to ENFM8*.
- If DB2 went from ENFM8 to CM8*, DSNTIJES indicates that DSNTIJEN must be run.

Customer migration paths:



Skip DB2 9 is possible, but most customers will go to DB2 9

Estimate save 20% to 25% vs 2 steps

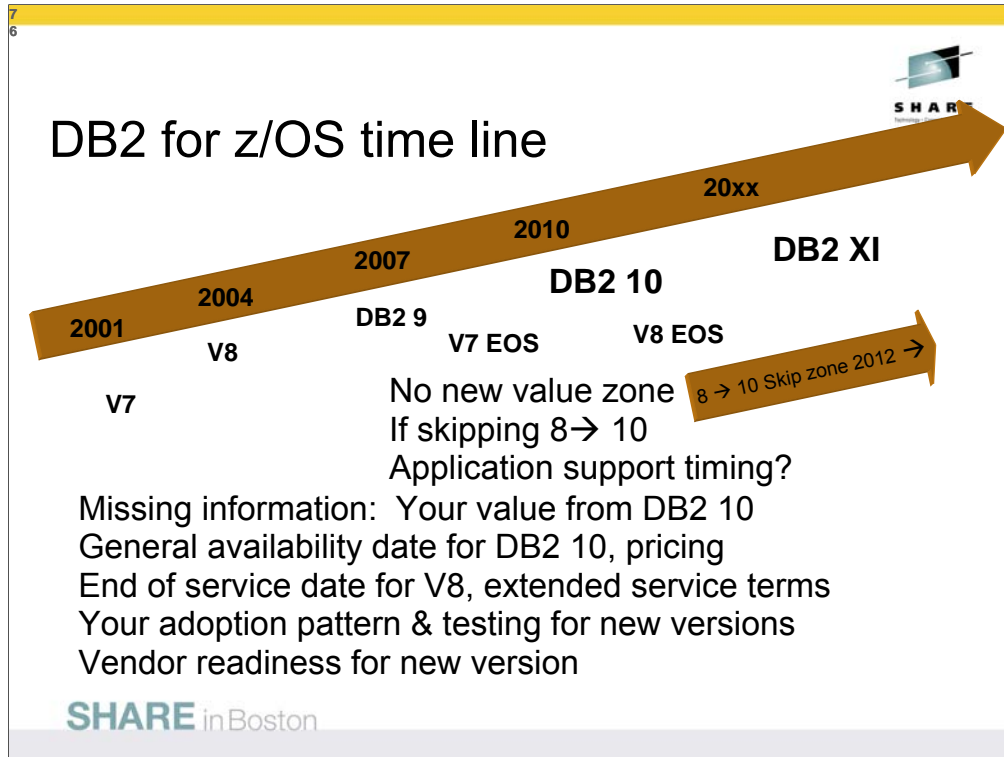
Reasonable timing to skip

Value in DB2 9 deferred until then

When will applications add support for 10? Drop 8?

Skip versions V5 → 7, V8 → 10, not every version

Do you want to migrate from V8 to DB2 10? The first question is the objective in terms of time you expect to save. The next questions are timing and service. Do you want support? DB2 10 will probably become generally available in 2010. When can you migrate to it? DB2 V8 end of marketing was Sept 2009. Migrating to DB2 9 is faster, easier and smoother now. We have experience with skip release migration. We did provide skip level migration from V5 to V7, but not for any other recent releases. The ability to skip helps some customers who have not installed recent versions, but this work tends to help a minority of the DB2 customer base. The key question for skipping DB2 9 is, "How long can you wait for that function?" If you are on DB2 V8 today, how long after GA do you wait before moving to a new version? Support for the new version in applications and tools can delay the process. Waiting for others to deliver success stories adds more time. If you wait three years after GA, then you could start moving to DB2 10 in 2014, and probably get there about a year later. I expect DB2 V8 to be out of support before that, so you would probably have additional expense for extended service. The savings for skipping a version migration are less than 50%, since the education and needed application and administration changes are about the same. Customers who do skip migration reported that the project takes longer, taking about 50% more work than a normal migration path. Changing from DB2 V8 or before to DB2 10 would require a cultural shift that some describe as culture shock. If customers spend the bulk of their migration project time in testing, then savings could be up to 40%, but most customer plans should show 20% to 25% reduction over two migrations. The cost for skipping is primarily later delivery of DB2 9 improvements by two to four years. Do you need the improvements in DB2 9 before that? CPU savings, especially in utilities. Disk savings via compression for indexes. Improved insert and update rates. Improved SQL and pureXML for developer productivity. Better availability. If customers are on V7 or earlier, that version is out of service. Recommend migrating to V8 to be on a supported version and to allow a supported skip version migration.



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

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

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

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

If you are not on DB2 9, look what you are missing.



Availability, Scale & Resiliency	More Online Schema Changes	➔	Business Resiliency
	Volume Level Backup & Recovery	➔	Faster, Cheaper, Granular Recovery
Compliance & Security	Database Roles & Trusted Context	➔	Efficient Auditing & Compliance
Easier Application Development	PureXML	➔	Streamlined Data Integration
OLTP & Warehousing	Query Optimization	➔	Highly Available, Secure Data
Reduce Cost of Ownership & Skills	Index Compression Utility CPU	➔	Save 50% Disk Save 20% - 30% utility CPU

SHARE in Boston

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

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

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

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

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

Key details about DB2 10: getting ready



Prerequisites: migrate from DB2 9 for z/OS or DB2 for z/OS V8

- z/OS V1.10 SMS-controlled DB2-managed DB2 catalog
- System z10, z9, z890, z990, and above (no z800, z900)
- DB2 Connect 9 FP1, 9.7 FP3 for many 10 functions, FP2 beta
- IMS 10 & 11 (not 9) CICS compilers (See announcement)
- SPE PK56922 PK69411 PK61766 PK85956 PM04680 PK87280
PK87281 PM08102 PM08105 PM10227
- Premigration check DSNTIIPA PM04968
- Info APARs II14477 (DB2 9) II14474 (V8)

Items deprecated in earlier versions eliminated: more for V8 mig.

- Private protocol → DRDA (DSNTP2DP, PK92339, PK64045)
- Old plans and packages V5 or before → REBIND
- Plans containing DBRMs → packages PK62876 (9), PK79925 (8)
PK85833 (9), PM01821 (All)
- ACQUIRE(ALLOCATE) → ACQUIRE(USE)
- Old plan table formats → DB2 V8 or 9, Unicode, 59 cols PK85068
- BookManager use for DB2 publications → Info Center, pdf

SHARE in Boston

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

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


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

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

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

7
9

No longer supported in DB2 10 from DB2 9:



- Private protocol: Convert to DRDA.
- EXPLAIN tables prior to Version 8 format: Alter add cols.
- Plans containing DBRMs. Acquire allocate. Old packages.
- DB2 catalog tables are DB2-managed & SMS-controlled. No links. More LOBs & table spaces. No SPT01 compress.
- REORG TABLESPACE SHRLEVEL NONE on LOB table spaces. Use SHRLEVEL CHANGE or REFERENCE.
- DB2 MQ XML functions: Use pureXML functions.
- DB2 XML Extender: Use pureXML.
- DB2 Management Clients feature, Control Center
- msys for Setup DB2 Customization Center
- Some subsystem parameters removed, many changed
- XDBDECOMPXML and XDBDECOMPXML100MB
- Accessories Suite parts Optimization Service Center, Developer Workbench, Visual Explain for DB2 → Data Studio

SHARE in Boston

No longer supported, moving from DB2 9

Private protocol: Convert to DRDA.

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

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

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

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

DB2 MQ XML functions: Use pureXML functions.

DB2 XML Extender: Use pureXML.

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


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

Several subsystem parameters

- Annotated XML schema decomposition using XDBDECOMPXML and XDBDECOMPXML100MB
- Accessories Suite will no longer include Optimization Service Center component and the Data Studio Workbench feature. Visual Explain for DB2 for z/OS is not available for DB2 10. The recommended query optimization and service tools for DB2 for z/OS are Optim Query Tuner and Optim Query Workload Tuner. The recommended no-charge query optimization and service tool for DB2 for z/OS is Data Studio (<http://www.ibm.com/software/data/studio/>). These tools are based and built on the foundation of Optimization Service Center and Optimization Expert.

5
0

No longer supported in DB2 10 from V8:



- Private protocol: Convert to DRDA.
- EXPLAIN tables prior to Version 8 format: Alter add columns
- Plans containing DBRMs; Acquire allocate; Old packages
- DB2 catalog tables are DB2-managed & SMS-controlled
No links. More LOBs & table spaces. No SPT01 compress
- REORG TABLESPACE SHRLEVEL NONE on LOB table spaces. Use SHRLEVEL CHANGE or REFERENCE
- DB2 MQ XML functions: Use pureXML functions
- DB2 XML Extender: Use pureXML
- DB2 Management Clients feature, Control Center
- msys for Setup DB2 Customization Center
- Some subsystem parameters removed, many changed
- ...

SHARE in Boston

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

Private protocol: Convert to DRDA.

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

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

DB2 catalog tables are DB2-managed and SMS-controlled. Catalog and directory tables do not have links, but have more LOBs and more table spaces.

Compression for table space SPT01 is not supported.

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


DB2 MQ XML functions: Use pureXML functions.

DB2 XML Extender: Use pureXML.

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

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

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



No longer supported in DB2 10 from V8: ...


- Net.Data
- DB2-established stored procedures
- Old JDBC driver
- Pascal L string data type from VAX
- Creating simple table spaces
- QMF Visionary Studio
- DB2 Estimator
- BookManager help
- DB2 Extenders: AIV, text, Net Search
- Java stored procedures in resettable JVMs
- ...

SHARE in Boston

No longer supported ...:

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

5
2



Deprecated: may be removed from later versions:

- Simple table spaces
- Partitioned table spaces other than universal table spaces
 - Need to specify SEGSIZE(0) to get classic partitioned
- Default for new table spaces is now universal table space
- DSNHDECP load
- NEWFUN(YES) and NEWFUN(NO)
- DSNHPC7 precompiler
- EXPLAIN stored procedures
- ...

All of these lists are much longer, see Install book or Information Center for details.

SHARE inBoston

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

Simple and partitioned table spaces other than universal table spaces are deprecated. The default table space type for new tables is now the partition-by-growth universal table space. Use alter in new function mode to convert single-table table space to universal.

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

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

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

SMS now a prerequisite for DB2 catalog & directory



Important – All new catalog indexes and new table spaces WILL be SMS controlled.
“SMS now a prerequisite of DB2 10”. So be prepared!

- During migration from (V8 or DB2 9) to DB2 10 conversion mode, all new indexes and new table spaces in the catalog and directory will be created as SMS-controlled with extended addressability (EA) attributes.
- DSNTIJSS provides SMS classes for customers without SMS in use.
 - The environment created by DSNTIJSS is ONLY for DB2 Catalog and Directory data sets.
 - Other DB2 data sets such as logs and BSDS not covered.
- Useful SMS reference material includes:
 - z/OS DFSMS Introduction.
 - z/OS DFSMS Implementing System Managed Storage.
 - z/OS DFSMS Storage Administration Reference.
 - IBM Redbooks publication: Maintaining your SMS environment



SHARE in Boston

SMS is now a prerequisite of DB2 10.


During the process of migrating from V8 / DB2 9 to DB2 10 conversion mode, all the new indexes and new table spaces in the catalog and directory will be created as SMS-controlled.

The environment created by the DSNTIJSS is only for DB2 Catalog and Directory data sets, which MUST be SMS controlled in DB2 10. Other DB2 subsystem data sets such as logs and the BSDS are not accounted for in this environment.

Useful SMS reference material includes:


- z/OS DFSMS Introduction.
- z/OS DFSMS Implementing System Managed Storage.
- z/OS DFSMS Storage Administration Reference.
- The Redbook: Maintaining your SMS environment.

5
4



Convert plans with DBRMs to packages

- Convert from DBRMs in PLAN to packages
- Specify the collection id for the packages
- REBIND plan option: COLLID
- Default collection id DSN_DEFAULT_COLLID_planname
- Able to specify plan name, list of plans, *
- DB2 9 APARs PK62876, PK85833, PM01821
- V8 APARs PK79925, PM01821
- See Packages Revisited book SG24-7688 chapter 4




SHARE in Boston


- PK62876 PK79925 (V8) BIND changes for eliminating acquire(allocate), plans with DBRMs
- New Rebind plan option: COLLID(collection_id|*): Trigger DBRM to Package conversion by specifying collection ID for the packages, which are converted from DBRMs.
 - REBIND PLAN(X) COLLID(*). "*" is the default value as: DSN_DEFAULT_COLLID_X (X is the plan name)
This command will allow customer to convert DBRMs bound with plan X into packages under collection ID: DSN_DEFAULT_COLLID_X automatically.
 - REBIND PLAN(x) COLLID(collection_id): CollID option is used if user intends to convert DBRMs bound with plan x into packages. As shown in this example: users can specify collection id as y for the new packages converted from DBRMs. If there are both DBRMs and package list in plan x, newly converted package entries will be inserted into the front of the existing packagelist. Same behavior applies to *, but using DSN_DEFAULT_COLLID_X instead of letting user specify.
 - REBIND PLAN (X1, X2, X3) COLLID (collection_id|*) Rebind multiple plans which may contain DBRMs. Bind all Plans specified here in SYSPLAN table, and do conversion, until no plans specified here have DBRMs bound within.
 - REBIND PLAN (*) COLLID (collection_id|*) All plans in SYSIBM.SYSPLAN table.
 - REBIND PLAN (x) COLLID (collection_id|*) PKLIST(Z) In this case: If plan has been bound with both DBRMs and package list, the DBRMs will be converted into packages, and the new entries will be inserted into the front of the existing package list.
 - REBIND PLAN (x) COLLID (collection_id|*) NOPKLIST. Three cases :
 - 1. Plan x does not have DBRMs, then the existing pklist will be deleted. No other change required.
 - 2. Plan x only has DBRMs, DBRMs will be converted into packages accordingly, and NOPKLIST will be ignored.
 - 3. Plan x has both DBRMs and PKLIST, the existing pklist will be deleted, while the new pklist converted from the DBRMS will be bound into plan x.

5


Convert private protocol to DRDA

- Trace to find use of private protocol
- Private protocol trace analysis program DSN1PPTP
- Use DRDA protocol REXX tool DSNTP2DP
- Private protocol statement checker DSNTPPCK
- Set PRIVATE_PROTOCOL parameter to NO
- Changes provided in base, PK92339, PK64045
- See Distributed book SG24-6952-01 chapter 5.2

 SHARE
Technology. Connected. Ready.

 IBM

**DB2 9 for z/OS:
Distributed Functions**

 Establish connectivity to and from DB2 systems
Balance transaction workload across data sharing members

SHARE in Boston

The changes to help remove private protocol deliver in the base for DB2 9 and in APARs PK92339 and PK64045 (closed).

DB2 trace for IFCID 0157 provides the detailed information that you need. Programs are provided to analyze the traces, help convert and check.

When private protocol is removed, set the subsystem parameter PRIVATE_PROTOCOL to NO to prevent any future use of private protocol.

See the documentation updates in the APARs and the base as well as the new Distributed book, chapter 5.2 for information on using these facilities.

For more information about DB2 10



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

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/>

Name	Size	Type	Modified
availability-idug-na-2010-roberts.pdf	700 KB	Adobe Acrobat Doc...	5/28/2010 1:02 PM
db2-10-questions-answered.pdf	37.1 KB	Adobe Acrobat Doc...	5/28/2010 3:02 PM
db2-10-webcast-2010-miller.pdf	1.35 MB	Adobe Acrobat Doc...	5/28/2010 3:02 PM
dba-idug-na-2010-miller.pdf	934 KB	Adobe Acrobat Doc...	5/28/2010 1:02 PM
optimizer-idug-na-2010-purcell.pdf	253 KB	Adobe Acrobat Doc...	5/28/2010 1:02 PM
optimizer-iod-2010-purcell.pdf	643 KB	Adobe Acrobat Doc...	5/28/2010 1:02 PM
performance-preview-idug-na-2010-hoshikawa.pdf	413 KB	Adobe Acrobat Doc...	5/28/2010 1:02 PM
SAP-iod-2010-schuetzner.pdf	303 KB	Adobe Acrobat Doc...	5/28/2010 1:02 PM
security-iod-2010-pickel.pdf	845 KB	Adobe Acrobat Doc...	5/28/2010 1:02 PM
xml-iod-2010-zhang.pdf	736 KB	Adobe Acrobat Doc...	5/28/2010 1:02 PM

See the web for presentations from recent conferences. Topics include overview, optimization, performance, advantages for SAP, security and XML.

SHARE in Boston

For more information, see the web for presentations from recent conferences. Topics include overview, optimization, performance, advantages for SAP, security and XML. This is the DB2 10 main page on the IBM web. See some early pointers and quotes from customers and consultants. <http://www.ibm.com/software/data/db2/zos/db2-10/>

The announcement itself is here:

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

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/>

Here are the URLs for presentations that are on the web. You can see the topics in the file names.

Recent conference presentations on DB2 10 performance, availability, optimization, XML, security, DBA improvements, SAP benefits, overview and synergy are on the web.

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/>

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/performance-preview-idug-na-2010-hoshikawa.pdf>

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/availability-idug-na-2010-roberts.pdf>

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/optimizer-idug-na-2010-purcell.pdf>

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/optimizer-iod-2010-purcell.pdf>

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/xml-iod-2010-zhang.pdf>

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/security-iod-2010-pickel.pdf>

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/dba-idug-na-2010-miller.pdf>

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/SAP-iod-2010-schuetzner.pdf>

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/db2-10-webcast-2010-miller.pdf>

<ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/overview/db2-10-technical-overview-iod-2010-campbell.pdf>


<ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/overview/db2-10-overview-idug-na-2010-josten.pdf>

We have a new white paper on DB2 10 from Dave Beulke now

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

Wille Favero just published an article in zJournal. <http://www.mainframezone.com/z-journal>

6
7



Questions and Answers

Thank You

Watch for upcoming white papers and conferences
IOD IDUG Share

SHARE in Boston

Here are some questions and answers to questions from conferences and webcasts. As of June 2010, DB2 10 for z/OS beta has started, but there are no books yet. One white paper is published. Some conference presentations are on the web.

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/>

More details will become public during the fall 2010 conferences and much more at general availability.

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

All of these lists are much longer, see Install book or Information Center for details.

Questions and Answers about DB2 10 for z/OS

August 3, 2010 Share Boston version

One key caveat for all of these answers is that the release is currently in beta, so that changes are to be expected. Some of the answers will probably change. Prediction is hard, especially for the future.

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.

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, 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. When this information is known, we can answer the question.

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

DB2 10 for z/OS beta began in March 2010. We don't really know when DB2 10 will be generally available. The beta duration in recent versions ranged from 6 months to 14 months. We expect a similar timing. If the beta goes as expected and we get positive feedback from the customers, consultants and vendors, then the release is ready to be generally available. Please watch the web for notice.

<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

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

DB2 10 for z/OS is in beta. High level information about the beta product can be found at <http://www.ibm.com/software/data/db2/zos/db2-10/>

You can also learn more at upcoming conferences, including Share in Boston, the IDUG conference in Vienna and IOD conference in Las Vegas.

<http://www.share.org>

<http://www.idug.org/european-conference/registration-information.html>

<http://www-304.ibm.com/jct03001c/services/learning/ites.wss?pageType=page&c=a0008651>

We expect to have another white paper. We anticipate an extended beta program during the summer. If you are interested in participating, talk to your IBM Representative. DB2 books appear at or after general availability, including IBM Redbooks publications.

What information is available now? When can we get more?

DB2 10 for z/OS is in beta. High level information about the beta product can be found at <http://www.ibm.com/software/data/db2/zos/db2-10/>

Some information has been provided at conferences and webcasts. A paper from Dave Beulke provides an overview of the function. Here is the webcast presentation with answers to some frequently asked questions. Get the pdf or register for the webcast with audio.

<ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/overview/db2-10-webcast-2010-miller.pdf>

http://www.ibm.com/software/os/systemz/webcast/13apr/index.html?S_TACT=100GU00M&S_CMP=SP

Sessions at IDUG NA in Tampa and at IOD EMEA 2010 included these topics, which are on the web in this folder.

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/>

S01 Overview Jeff Josten

A01: DBA improvements Roger Miller

B02: What's new from the optimizer Terry Purcell

A03: Performance Preview Akiko Hoshikawa

A06: DB2 and System z Synergy Chris Crone

A08: Availability Enhancements Haakon Roberts

Sessions including DB2 10 for z/OS at IOD EMEA 2010

2911 DB2 10 for z/OS technical overview John Campbell

2908 DB2 10 for z/OS security features help satisfy your auditors Jim Pickel

2894 What's coming from the optimizer in DB2 10 for z/OS? Terry Purcell

3075 DB2 for z/OS trends and directions

3010 New pureXML® Features in DB2® for z/OS: Breaking Relational Limit
Guogen (Gene) Zhang

2969 DB2 Family Design, Implementation and Convergence Roger Miller

2971 Planning your DB2 for z/OS migration Roger Miller

Here are URLs for presentations on the web. You can see the topics in the file names. Recent conference presentations on DB2 10 performance, availability, optimization, XML, security, DBA improvements, SAP benefits, overview and synergy are on the web.

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/>

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/performance-preview-idug-na-2010-hoshikawa.pdf>

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/availability-idug-na-2010-roberts.pdf>

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/optimizer-idug-na-2010-purcell.pdf>

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/optimizer-iod-2010-purcell.pdf>

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/xml-iod-2010-zhang.pdf>

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/security-iod-2010-pickel.pdf>

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/dba-idug-na-2010-miller.pdf>

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/SAP-iod-2010-schuetzner.pdf>

<ftp://public.dhe.ibm.com/software/data/db2/zos/presentations/v10-new-function/db2-10-webcast-2010-miller.pdf>

<ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/overview/db2-10-technical-overview-iod-2010-campbell.pdf>

<ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/overview/db2-10-overview-idug-na-2010-josten.pdf>

New white paper on DB2 10 from Dave Beulke

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

Wille Favero top 10 article in zJournal.

<http://www.mainframezone.com/it-management/top-10-features-of-db2-10-for-z-os-a-preview>

<http://db2onlinehandbook.com/?p=126>

<http://www.mainframezone.com/z-journal>

<http://www.mainframezone.com/it-management/top-10-features-of-db2-10-for-z-os-a-preview>

Blogs, short articles, lots more if you Google for "DB2 10 for z/OS"

<http://davebeulke.com/?p=1030>

<http://www.databasejournal.com/features/db2/article.php/3869096/IBM-DB2-10-for-zOS-Justifying-the-Upgrade.htm>

<http://www.triton.co.uk/blog/category/db2/db2-10/>

<http://db2onlinehandbook.com/?p=163>

<http://ibmsystemsmag.blogs.com/db2utor/2010/06/ten-reasons-to-upgrade-to-db2-10.html>

<http://ibmsystemsmag.blogs.com/db2utor/2010/07/db2-10-purexml.html>

<http://ibmsystemsmag.blogs.com/db2utor/2010/06/data-versioning-in-db2-10.html>

Rick Bowers <http://www.gse-nordic.org/Working%20Groups/GNRC/Conferences/2010/DB2/S48.%20DB2%20for%20z%20Reducing%20Business%20Costs%20and%20Increasing%20Business%20Value>

Security: <http://www.ibm.com/software/sw-events/teleconference/A556672U34009F89.html>
http://www.ibm.com/software/os/systemz/telecon/6jul/index.html?S_TACT=100GU00M&S_CMP=SP

We'll be getting more information out at the fall conferences: Share, GSE, IOD, and IDUG.

This is the IDUG EMEA 2010 track:

A01 - DB2 10 for z/OS Overview	Roger Miller
A02 - DB2 10 for z/OS Migration Planning	Roger Miller
A03 - DB2 for z/OS Trends and Future Direction	Rick Bowers
A04 - DB2 for z/OS V9 and V10 Enhancements for System Programmers	John Campbell, John Tobler
A05 - DB2 10 for z/OS Security Enhancements	James Pickel
A06 - DB2 10 Availability Enhancements	Haakon Roberts
A07 - DB2 10 - What's in it for SAP and Enterprise Applications	Namik Hrle
A08 - What's new from the optimizer in DB2 10 for z/OS?	Terry Purcell
A09 - DB2 for z/OS and Websphere Update	Maryela Weihrauch
A10 - Hash Access to DB2 Data - Faster, Better, Cheaper	Kalpana Shyam
A11 - DB2 and System z Synergy	Chris Crone
A12 - DB2 10 Application topics – A sneak preview	Suresh Sane
A13 - DB2 10 for z/OS Performance Preview	John Tobler
A14 - Utilities Update - DB2 10 for z/OS Preview	Haakon Roberts
A15 - How I Learned to Stop Worrying and Love Plan Management and Access Path Stability	Gerald Hodge
A16 - The Next Level: Data Warehousing with DB2 V10 and zEnterprise	Patric Becker
A17 - Introducing IBM Smart Analytics Optimizer - DB2 Performance Revolution	Namik Hrle

Does the statement “10 times more users” include data sharing systems? Or are additional changes needed in data sharing in order to achieve this?

See slide 22 from the Share DB2 10 overview, which does show data sharing. REBIND and you can run many more threads on each member. Reducing the number of members in this situation would take some changes, especially if you connect to specific members, rather than to the group. Also check slide 52 for data sharing enhancements in DB2 10.

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.

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 only constraint for most customers. REBINDs are required for the EDMPOOL improvements and some working storage. Some ECSA is 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 2010 or 2011 should generally do so. The Share slides 68 – 85 discuss migration options and the session on migration will have more.

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 will show hash access.

In the Share DB2 10 overview slide 18 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.

How is the hash access path established?

See slide 46 for an example of the process. CREATE the table with the hash, populate it and online REORG. Then BIND, check access, and see if indexes are still needed.

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 a UTS, then REORG to change the table space to MEMBER CLUSTER, using a pending ALTER.

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

PC=YES has been the only option since V8.

During PENDING ALTER, is the object is accessible?

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

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 of 4KB – and now add 1MB. DB2 page sizes remain the same: 4KB, 8KB, 16KB, and 32KB.

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

This overhead was addressed in DB2 10.

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/Os are needed to store or retrieve the LOB data. Typically these small LOBs can be anywhere from 2k to 30k so that they fit on the same page.

Does DB2 10 support private protocol?

No. See slides 79, 80, and 85.

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?

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 should not need to change as their timestamp precision will not change. If you have a general application which will process any DB2 data, then it needs to change to handle this variation and timestamps with time zones.

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. 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, allows REBALANCE and ALTER LIMITKEY even 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.

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

Hash access is a new direct access type 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.

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

DB2 10 has even more DB2 family and other DBMS compatibility functions, as noted on slides 31 and 33. 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. The primary utility changes are noted on slides 28, 54 and 55.

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 the SYSLGRNX?

SYSLGRNX has been needed for some recovery situations.

Will R10 restrict BINDING a plan to a single DBRM?

See slides 79, 80 and 84. 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.

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 slide 84 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 have any performance features as dramatic as V8 multi row fetch and insert, but do have many improvements in SQL and XML in DB2 10. See Share slides 31 - 34 for the application improvements, ranging from temporal or versioned data, new timestamp options, improved ability to port from other platforms, and many XML enhancements.

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 some of the improvements.

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

See slide 78 in the Share presentation and the beta announcement. Watch for the general availability announcement.

http://www.ibm.com/common/ssi/rep_ca/5/877/ENUSZP10-0015/ENUSZP10-0015.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 will be out of service soon and is not supported with DB2 10. Migrate to IMS 10 or 11 soon. See the announcement for prerequisites.

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 slides 16 through 19 to understand the feature requirements. REBINDs and adjusting for the function removed are the primary application programming tasks. See slides 68 – 85 in the Share presentation.

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 part of DB2 for LUW, so you could look there.

Will statement level BIND/REBIND be available for static SQL?

REBIND query is for dynamic SQL statements. Static SQL BIND and REBIND are at the package level. DB2 10 builds on the access path stability improvements that were offered in Version 9 by introducing a more comprehensive framework for the management of access paths for both static and dynamic SQL statements. With these new features for access path stability you can capture information for storage in an access path repository, save multiple copies of access paths, and switch between different copies of access paths of the same query. For both bind and rebind processes, you can regenerate runtime structures without changing access paths. In addition, you can compare the new and old access paths at bind or rebind processing and indicate that a warning or error is issued when an access path changes. These access path stability features also allow you to control when dynamic SQL statements are re-optimized, as you can for static SQL statements.

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 is working on increasing that number to about 100K.

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

Yes, the DB2 catalog changes from using manual definition and extension to DB2 managed data sets under SMS control. See Share slide 83 for a start. The DB2 catalog and directory have some table spaces defined as DSSIZE 64G, which requires SMS managed storage.

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

The utilities have been working with large LOBs since DB2 V7 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.

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, allowing more customers to use

this option. Compression will use an SMF compression option. 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 labels, similar to the needs for military confidentiality. Row and column access control is more flexible, with capabilities like 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?

Detailed measurements on DB2 10 are not ready to release, but you can expect improvements. Native SQL procedures came very close to COBOL and C in IBM measurements, within about 15%. Remote native SQL procedures used less CPU than COBOL or C by using zIIP engines.

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

The primary change in DB2 10 is removing some parallel restrictions. The primary change for customers is to design for and use parallelism to increase use of zIIP. Use of remote native SQL procedures can also help.

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.

You seem to be very cautious about skip version migration, suggesting that V8 customers migrate to DB2 9. Why?

My caution is a reflection of the current timing and cycle. DB2 10 is in a beta, not generally available(GA). Some key information is missing. Please see notes under **When should I migrate to DB2 10?** above.

Experience with DB2 migration and with skip version migration suggests caution about skipping during the first few years of a version. Customer misconceptions caused many problems in the past, and I want to avoid a repeat of those problems. One issue is underestimating the time and length of a skip version plan. The second is underestimating the impact of the standard software maturity cycle. Some customers seem to underestimate the value in DB2 9. The misconceptions can lead to decisions which result in failed migrations and costly issues.

Customers who are running a release that is about to go out of service have different experiences than customers running a release that is just GA. The need for testing and handling problems are very different. Can you handle skip version migration? The plan for education and for remediation of incompatibilities needs

to allow more time and resources. When do your key tools and applications handle the new version?

The primary objective for skip version migration is to allow customers who are behind, such as those who just migrated to V8 or are still on V7 to catch up. Early migrations to DB2 10 are best for customers who generally migrate early, but most such customers are on DB2 9 now. As we get more experience and DB2 10 matures in the usual cycle, skip will work for more customers.

What are the key migration considerations for DB2 10 for z/OS?

IBM added many data server capabilities in DB2 10 for z/OS and reduced or removed support for some functions. As you plan and prepare for your subsystem migration to DB2 10, you should be aware of the following changes:

For both DB2 for z/OS V8 and DB2 9 for z/OS:

- DB2 XML Extender no longer supported
- DB2 Management Clients feature is no longer available. Control Center does not support connections to DB2 10.
- EXPLAIN tables prior to Version 8 format are no longer supported. Change tables to Unicode and add any needed columns. Conversion to Unicode can mean different results for joins or sorts.
- Private protocol no longer supported. Convert to DRDA.
- Plans containing DBRMs are no longer supported. Plans will contain packages, which have the DBRMs. Use the COLLID option of REBIND to make the needed changes. Acquire allocate is not a supported BIND option any longer. Use acquire use.
- DB2 catalog tables are DB2-managed and SMS-managed. Catalog and directory tables do not have links, but have more LOBs and more table spaces. Compression for table space SPT01 is not supported.
- DB2 MQ XML functions are no longer supported. Use pureXML functions.
- msys for Setup DB2 Customization Center no longer supported. Use installation panels instead.
- REORG TABLESPACE SHRLEVEL NONE on LOB table spaces is removed. Use SHRLEVEL CHANGE or REFERENCE.
- Several subsystem parameters are removed or no longer supported.

For DB2 for z/OS V8 only:

- Net.Data is removed. WebSphere is the strategic IBM solution for delivering DB2 data to Web applications.
- DB2-established stored procedure address spaces are no longer supported. Workload Manager (WLM) managed stored procedure address spaces is the strategic solution for stored procedure support, and

migration to WLM managed stored procedure spaces is required for use of stored procedures in DB2 10.

- JDBC/SQLJ Driver for OS/390 and z/OS is no longer supported. All Java application programs and Java routines that are currently written to work with the JDBC/SQLJ Driver for OS/390 and z/OS need to be modified to work with the IBM DB2 Driver for JDBC and SQLJ (formerly known as the DB2 Universal JDBC Driver). The steps for migrating JDBC and SQLJ applications from the legacy JDBC/SQLJ Driver for OS/390 and z/OS to the IBM DB2 Driver for JDBC and SQLJ can be found in the Application Programming Guide and Reference for Java (SES1-3023). In addition, all WLM-managed stored procedures address spaces that are set up to execute Java routines must be modified to reference the IBM DB2 Driver for JDBC and SQLJ.
- Connections from VAX machines and the PASCAL L string data type are no longer supported.
- Creation of simple table spaces is no longer supported. DB2 10 for z/OS no longer implicitly creates simple table spaces nor allows customers to create simple table spaces. However, DB2 10 for z/OS continues to support simple table spaces created in previous versions.
- DB2 QMF Visionary Studio program is removed from DB2 QMF Enterprise Edition.
- DB2 Estimator is not available for DB2 9 or DB2 10.
- BookManager-based online help has been removed. The prior help support has been replaced by the Information Management Software for z/OS Solutions Information Center (Information Center). The web-based Information Center is updated periodically during the life of each DB2 version, thus ensuring reader access to the most up-to-date information.
- AIV Extender, Text Extender, and Net Search Extender are removed.
- Java stored procedures no longer run in resettable JVMs.

For DB2 9 for z/OS only:

- Annotated XML schema decomposition using XDBDECOMPXML and XDBDECOMPXML100MB is no longer supported
- Accessories Suite will no longer include Optimization Service Center component and the Data Studio Workbench feature. Visual Explain for DB2 for z/OS is not available for DB2 10. The recommended query optimization and service tools for DB2 for z/OS are Optim Query Tuner and Optim Query Workload Tuner. The recommended no-charge query optimization and service tool for DB2 for z/OS is Data Studio (<http://www.ibm.com/software/data/studio/>). These tools are based and built on the foundation of Optimization Service Center and Optimization Expert.

Plans and packages should be rebound on current releases to benefit from optimizer enhancements. For the DB2 10 beta release, packages from DB2 V5 or earlier will be automatically rebound when accessed by DB2 10.

For more detailed information on migrating to DB2 10, refer to the DB2 10 for z/OS Installation and Migration Guide once it becomes available.

As part of ongoing efforts to deliver the most current technology and to remove those features that no longer provide our customers with strategic benefits, the following features are deprecated in DB2 10 for z/OS and may be dropped from future versions:

- Simple and partitioned table spaces other than universal table spaces are deprecated. Use ALTER in new function mode to convert single-table table space to universal. The default table space type for new tables is now a universal table space.
- Some current use of DSNHDECP is deprecated. If you have code that loads DSNHDECP and maps it with macros, you should plan to change that code by using the new techniques. Customers who want to have one library for multiple DSNHDECP modules need to make this change.
- SQL processing options NEWFUN(YES) and NEWFUN(NO) options are deprecated. Use NEWFUN(V10) rather than NEWFUN(YES). Use NEWFUN(V9) or NEWFUN(V8) rather than NEWFUN(YES).
- The DSNHPC7 precompiler is deprecated. Use the current precompiler or coprocessor.