

DB2 10 for z/OS



Helping you improve operational efficiencies
and gain competitive advantage



© 2009 IBM Corporation


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.



Disclaimer/Trademarks

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

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

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

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

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

2 © 2009 IBM Corporation

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

See the web for current information, the announcement, ...

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

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

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

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

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

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

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

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



DB2 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 	

Beta Announced:
Feb 9, 2010

3

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

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

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

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

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

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

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

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

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

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

4

© 2009 IBM Corporation

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>

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=isg1I14477>

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

<http://www.ibm.com/support/docview.wss?uid=isg1I14474>



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

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

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


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

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




Top 10 in DB2 10 for z/OS





Top 10 in DB2 10 for z/OS

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



7


© 2009 IBM Corporation

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

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


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

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



Why Migrate to DB2 10 for z/OS?

- Business needs to save money
 - **Reduce CPU time**
 - Service Oriented Architecture
- 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



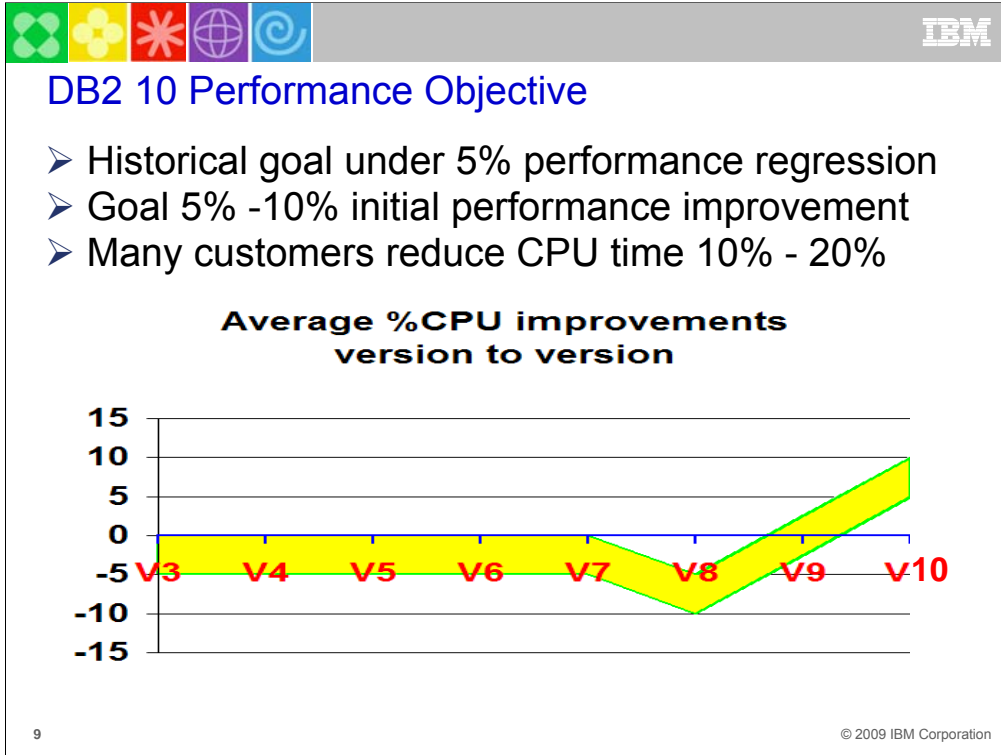
8 © 2009 IBM Corporation

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

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

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

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



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

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


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

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

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

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

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



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

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

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

Scales with less complexity and cost


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

Improved operational efficiencies and lower administration cost

- Automatic diagnostics, tuning, and compression

Even better performance

- Elapsed time improvement for small LOBS and Complex Queries



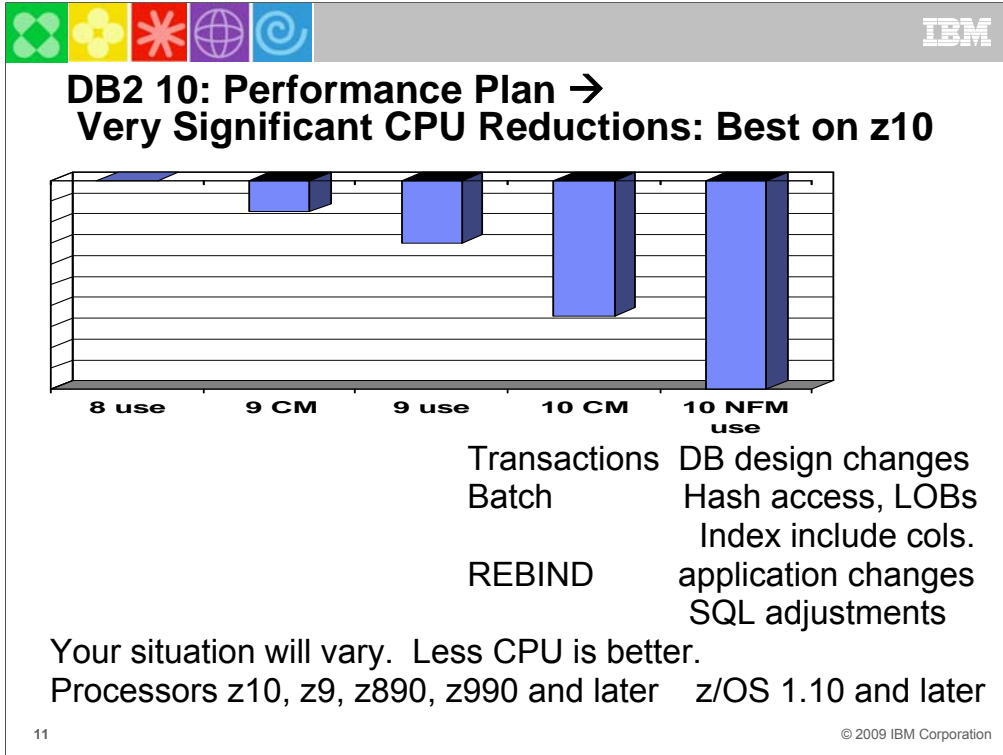
10 © 2009 IBM Corporation

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

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

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

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





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

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

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

We expect DB2 10 to run only on z10, z9, z890, z990, and later processors, and to provide CPU reductions from the beginning, with improvements in CM, but more dramatic reductions for applications that can take advantage of the improvements in 11

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 times
 - Lower CPU usage for large & small DB2 subsystems
- Virtual storage is most common vertical scale constraint for large customers
 - Limited number of concurrent threads for a single member / subsystem
 - Address next tier of constraints: latches, concurrency


12 © 2009 IBM Corporation

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.




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 z10 1MB page size
 - “Fully in memory” option (ALTER BUFFERPOOL)

13 © 2009 IBM Corporation

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

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




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
- 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

14 © 2009 IBM Corporation

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

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



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, PBG
- Insert improvement for UTS


15

© 2009 IBM Corporation

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

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

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

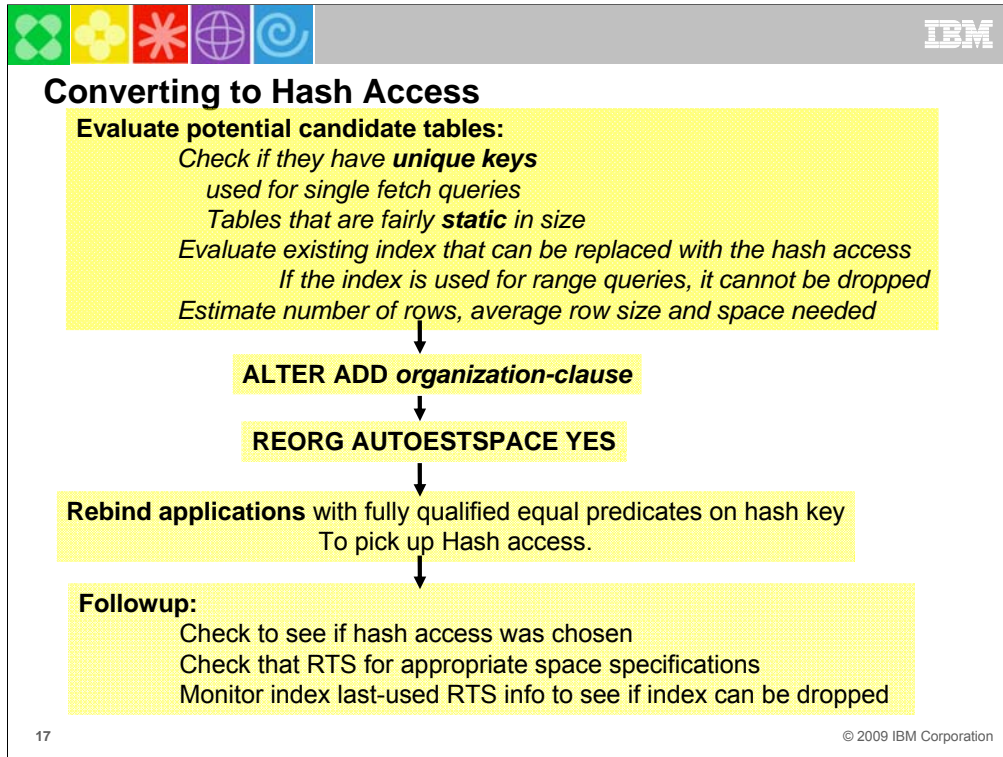


Performance Enhancements requiring NFM + DBA work

- Hash access path Alter + Reorg + rebind to activate
- Index include columns Alter + Rebuild + rebind to activate
- Inline LOBs Alter (need UTS and RRF)
- MEMBER CLUSTER for UTS Alter + Reorg
- DEFINE NO for LOB and XML columns
- Alter to UTS, page size, DSSIZE, SEGSIZE Alter + Reorg

16 © 2009 IBM Corporation


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



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

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

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



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

18 © 2009 IBM Corporation

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.

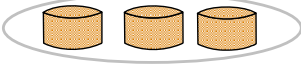


Running a Large Number of Active Threads

Today

Coupling Technology

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

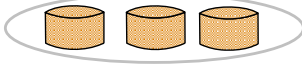


- Data sharing and sysplex allows for efficient scale-out of DB2 images
- Sometimes multiple DB2s per LPAR

DB2 10

Coupling Technology

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




- More threads per DB2 image
- More efficient use of large n-ways
- Easier growth, lower costs, easier management
- Data sharing and Parallel Sysplex still required for very high availability and scale
- Rule of thumb: save 1/2% CPU for each member reduced, more on memory

20 © 2009 IBM Corporation

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 planned to 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 can contend with one another
 - 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


21 © 2009 IBM Corporation

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
 - CLOB and BLOB columns for long strings
 - Online reorganization and check
 - More automatic: DB2-managed SMS-controlled

22 © 2009 IBM Corporation

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.




Catalog Restructure for improved concurrency

- Remove links from the catalog and directory
 - DSN1CHKR no longer needed
- Sixty 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.

23 © 2009 IBM Corporation

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.



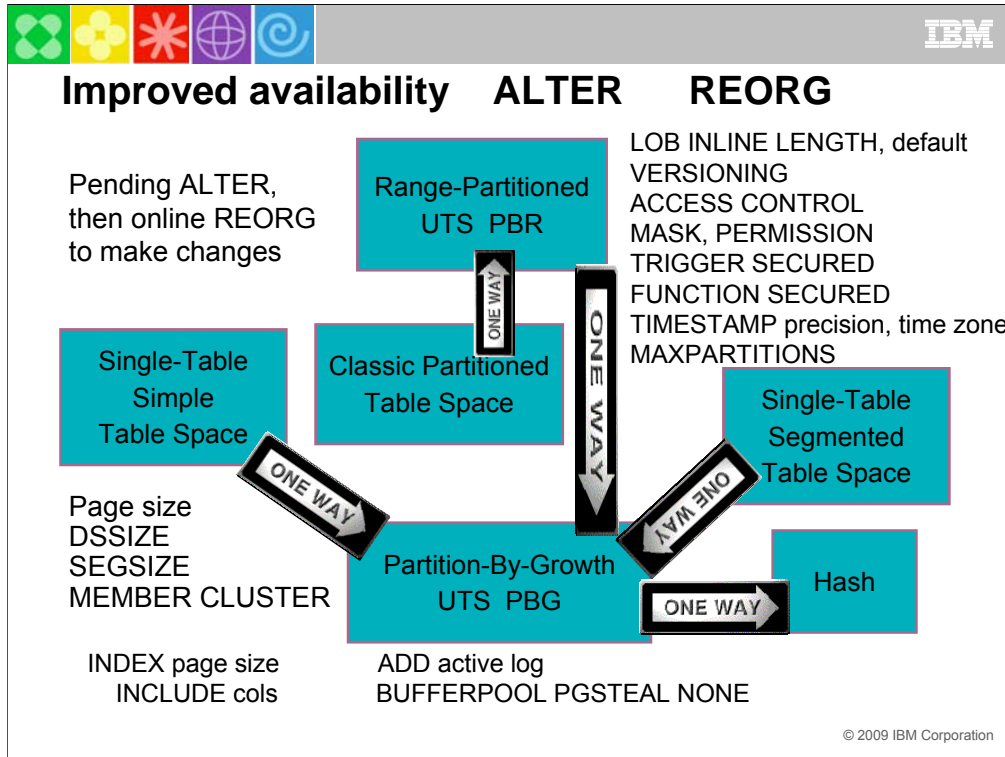
Other Catalog Changes

- 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-managed, 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

24 © 2009 IBM Corporation

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


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



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

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

What is not done? Change from multi-table segmented table space. Change back to classic simple, segmented and partitioned. The strategic choice for table space type is the universal table space. Simple table spaces are deprecated, and this version provides a migration path. The ability to add a new active log data set is included. Many online REORG restrictions are removed, to allow more online operations. If you need more improvements in table spaces, then universal table spaces – either partition by range or partition by growth should be your choice.



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

26 © 2009 IBM Corporation


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.



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

27 © 2009 IBM Corporation

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 - V9 PK80925
- Sub-group attach (v9 usermod)
- BP 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
- Online DDF changes
- 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



DB2 10: 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



Use disk encryption

29 © 2009 IBM Corporation


Customers are being pressed for a wide range of improved security and compliance. Data retention is a growing need. Protecting sensitive data from the privileged users and administrators is required. Separation of authority for security, access, and some common tasks, like EXPLAIN will help. Auditing for privileged users can also make compliance simpler. Access control is refined in several ways with better granularity for the administrative privileges and with finer grained access control at the row and column level, including the ability to mask access to some fields. Auditing is also enhanced.



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



Use disk encryption

30 © 2009 IBM Corporation

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

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

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

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

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


Eliminating views and security logic in applications

Preventing access from authorities including SYSADM and DBADM

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


Making the evolution of security policies easy to deal with

Relieving the difficulties of managing security policies



DB2 10: Productivity – Doing More with Less!

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




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

31 © 2009 IBM Corporation

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

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

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



DB2 10 Utilities Enhancements

- REORG SHRLEVEL(CHANGE) for LOBs
- Online REORG enhancements
 - SHRLEVEL(CHANGE) support for all catalog/directory objects
 - Option to cancel blocking threads
 - Improved 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

32 © 2009 IBM Corporation

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


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

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

Consistent image copies can be provided without a quiesce.

Inline copies to allow for dataset-level FlashCopy.

Online REORG usability and performance enhancements are provided.



DB2 10 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

33 © 2009 IBM Corporation


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 support 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.



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


35 © 2009 IBM Corporation

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.



Optimization Stability and Control

Provide unprecedented level of stability for query performance by stabilizing access paths for

- Static SQL - Relief from REBIND regressions
- Dynamic SQL
 - Remove the unpredictability of PREPARE
 - Extend Static SQL benefits to Dynamic SQL


- Access path repository
- Versioning
- “Fallback”
- “Lockdown”
- Manual overrides. Hints: easily influence access paths without changing apps
- Per-statement BIND options
- Safe query optimization: assess “reliability” of access path choices
- RID pool overflow to workfiles

36 © 2009 IBM Corporation

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

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

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



Many improvements for SAP & web applications


- **Autonomics**
- Compress on the fly on INSERT
- Auto-statistics
- Access path stability and 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

37 © 2009 IBM Corporation

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


38 © 2009 IBM Corporation

The pureXML improvements are noted two slides later.

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

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

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



Versioned data or Temporal Data

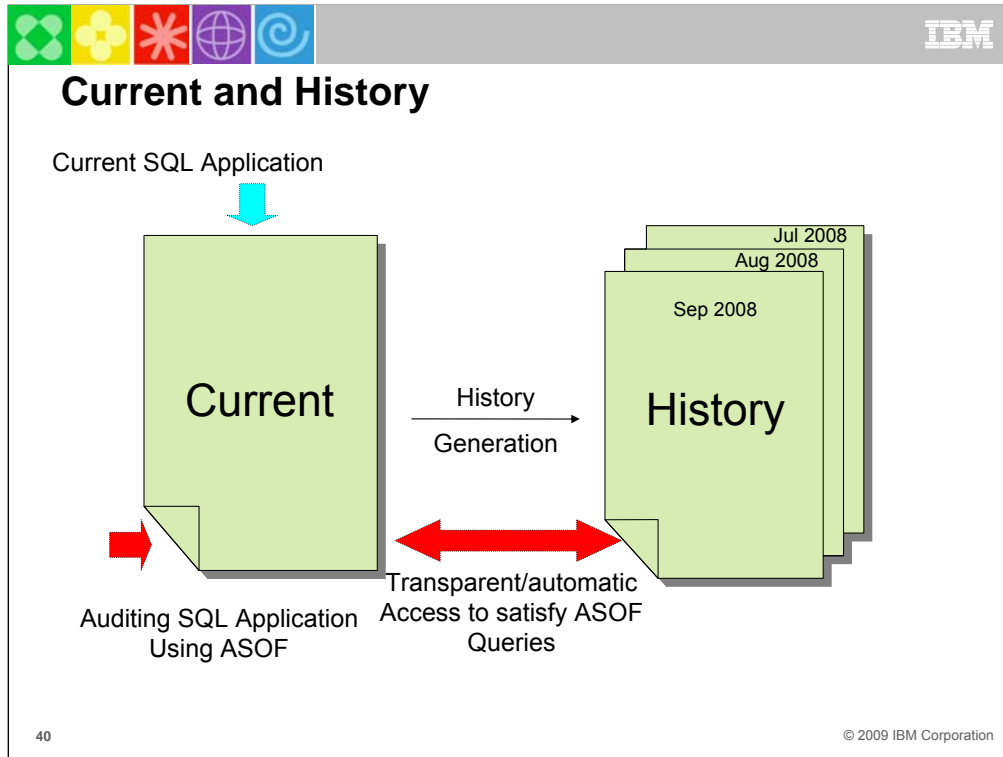
- Table-level specification to control data management based upon time
- Two notions of time:
 - System time: notes the occurrence of a data base change
 - “row xyz was deleted at 10:05 pm”
 - Query at current or any prior period of time
 - Useful for auditing, compliance
 - Business time: notes the 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

39 © 2009 IBM Corporation

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

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


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



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

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

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



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

41 © 2009 IBM Corporation

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.



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

42 © 2009 IBM Corporation


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 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

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

luw { 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

43 © 2009 IBM Corporation

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


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

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

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

● Cross-Platform Development,


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



DB2 SQL 2007

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

cross-platform SQL book V3



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, **TRUNCATE, DECIMAL FLOAT, VARBINARY, optimistic locking, FETCH CONTINUE, ROLE, MERGE, SELECT from MERGE, index & XML compression**

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

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

U

W

44 © 2009 IBM Corporation


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

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

Cross-Platform Development Version 3,

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


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



DB2 SQL 2008

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

cross-platform SQL book V3.1



Portable SQL

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

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**

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, more vendor syntax**

45 © 2009 IBM Corporation


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>



DB2 SQL 2009

z z/OS 9
common
luw Linux, Unix & Windows 9.7



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

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, **TRUNCATE, index & XML compression, created temps**

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, **even more vendor syntax, LOB & temp table compression, inline LOB, administrative privileges, implicit casting, date/time changes, currently committed**

46
© 2009 IBM Corporation


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 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**

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, TRUNCATE, index & XML compression, created temps, **inline LOB, administrative privileges, implicit casting, date/time changes, currently committed, moving sum & avg.**

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, even more vendor syntax, LOB & temp table compression,

47 © 2009 IBM Corporation


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>



pureXML improved performance and 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 sub-parts of XML document
- Stored proc, UDF, Trigger enhanced support
- XML index matching with date/timestamp
- CHECK DATA utility checks XML

48 © 2009 IBM Corporation

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.



DB2 10 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 and more even parallel distribution
 - increased zIIP use
 - In-memory techniques for faster query performance
 - Access path stability and control
- Analysis: instrumentation, Data Studio & Optim Query Tuner
- Advanced query acceleration techniques
 - IBM Smart Analytics Optimizer

Query enhancements in DB2 build on the improvements in DB2 V8 and 9. The CPU reductions can make a differences to queries. Improved SQL with better ability to query, temporal understanding, and XML improvements make the queries simpler.

The key new access techniques are index include columns and hash access. Improvements in access techniques provide more parallel access by reducing restrictions. More parallel means more ability to redirect the work to zIIP and reduce costs. In memory techniques provide improved performance. Access path stability allows better control for queries. Improvements in the instrumentation help all performance monitors. The Data Studio and Optim Query Tuner have replaced some older function. Advanced query acceleration is being previewed in the IBM Smart Analytics Optimizer.



Technology Preview: IBM Smart Analytics Optimizer

What is it?

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




Business Value

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

Targeted Uses for DB2 for z/OS customers:

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

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



DB2 10 for z/OS 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 • More online schema changes
Availability Security Productivity	<ul style="list-style-type: none"> • Improved concurrency: catalog, data, & utilities • Row and column access control, masking • Administrator privileges with finer granularity • Administration productivity enhancements • Versioned data or temporal queries
Application Enablement	<ul style="list-style-type: none"> • pureXML enhancements • Last committed reads • SQL improvements that simplify porting • Moving sum, moving average
Dynamic Warehousing	<ul style="list-style-type: none"> • Many query optimization improvements • Query parallelism improvements • IBM Smart Analytics Optimizer

51 © 2009 IBM Corporation

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



DB2 10 for z/OS

Planning to migrate



© 2009 IBM Corporation

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

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



DB2 10 for z/OS: Skip-Level Migration

May move from V8 to DB2 10,
but just because you can, doesn't mean you always should....

Key considerations:

- Risk/reward analysis
 - What's the 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?
 - Am I missing out on DB2 9 value in the meantime?
- May not see large migration cost savings
 - Expect 20% to 25% cost savings versus two migrations
 - Larger migration project, longer migration timeline, more risk
 - Applications and ISVs may not be ready



If you are on V7 or earlier, go to V8

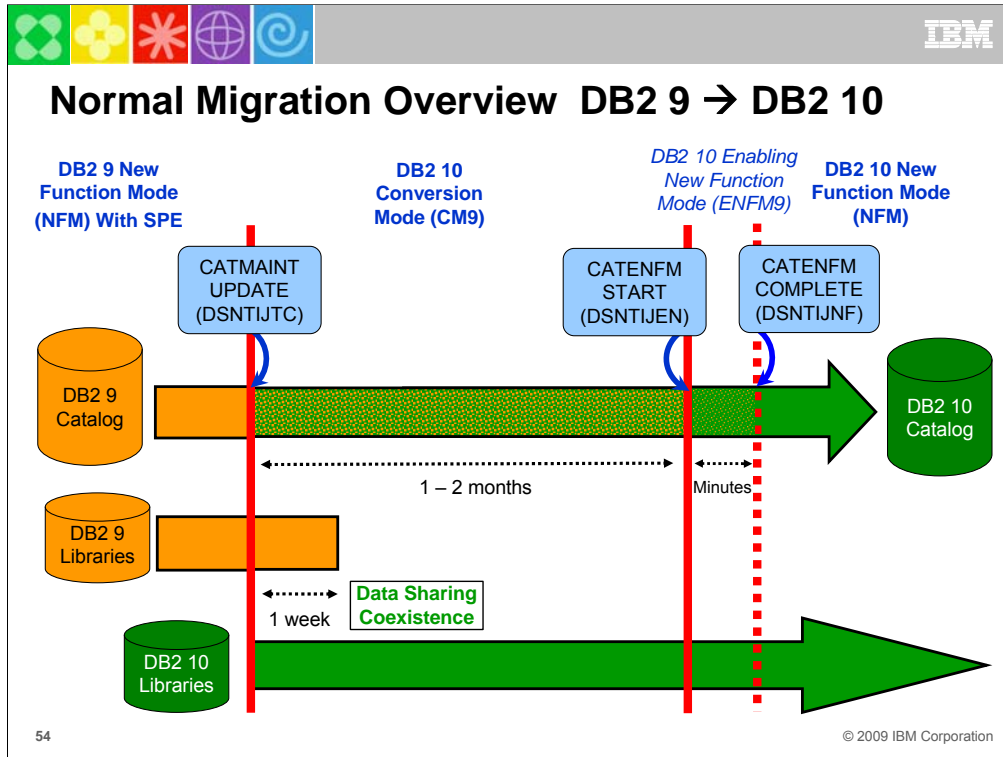
If you plan to migrate in 2010 or 2011, go to DB2 9

If you are on V8 for over 2 years, go to DB2 9

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

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

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



Migration from DB2 9 to DB2 10

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

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

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

Some migration considerations are:

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

Note – this is not necessarily to scale!

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



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.

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

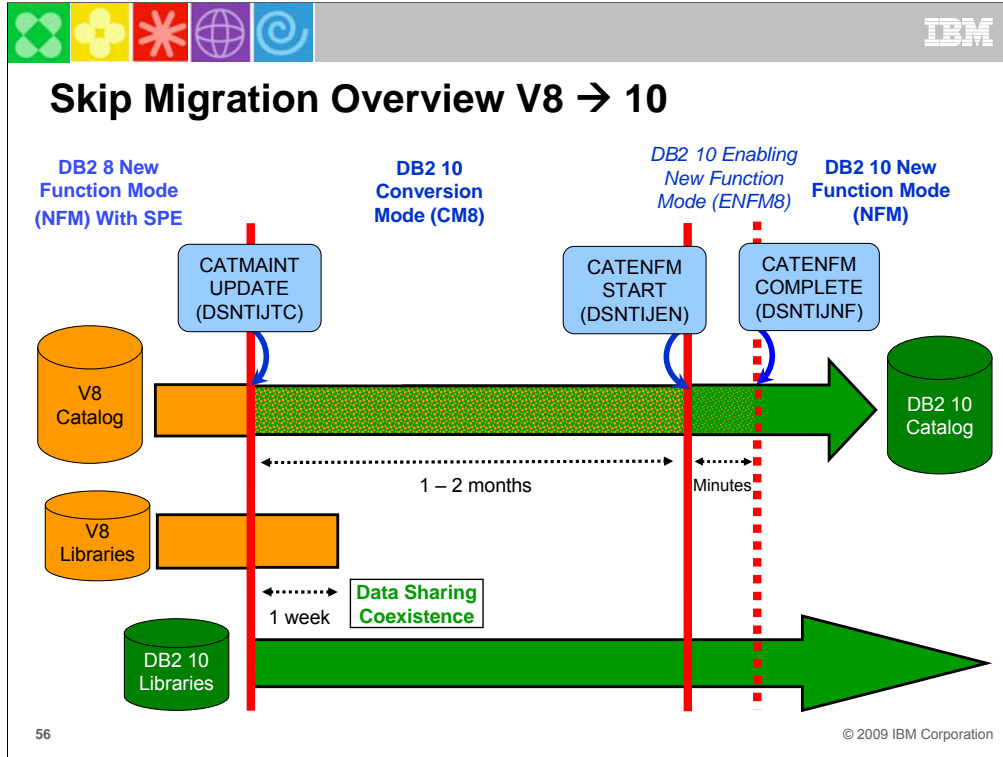
- If DB2 went from NFM or ENFM9* to CM9*, DSNTIJEN takes it to ENFM9*.
- If DB2 went from ENFM9 to CM9*, it completes ENFM9 processing and leaves DB2 in ENFM9.

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.

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

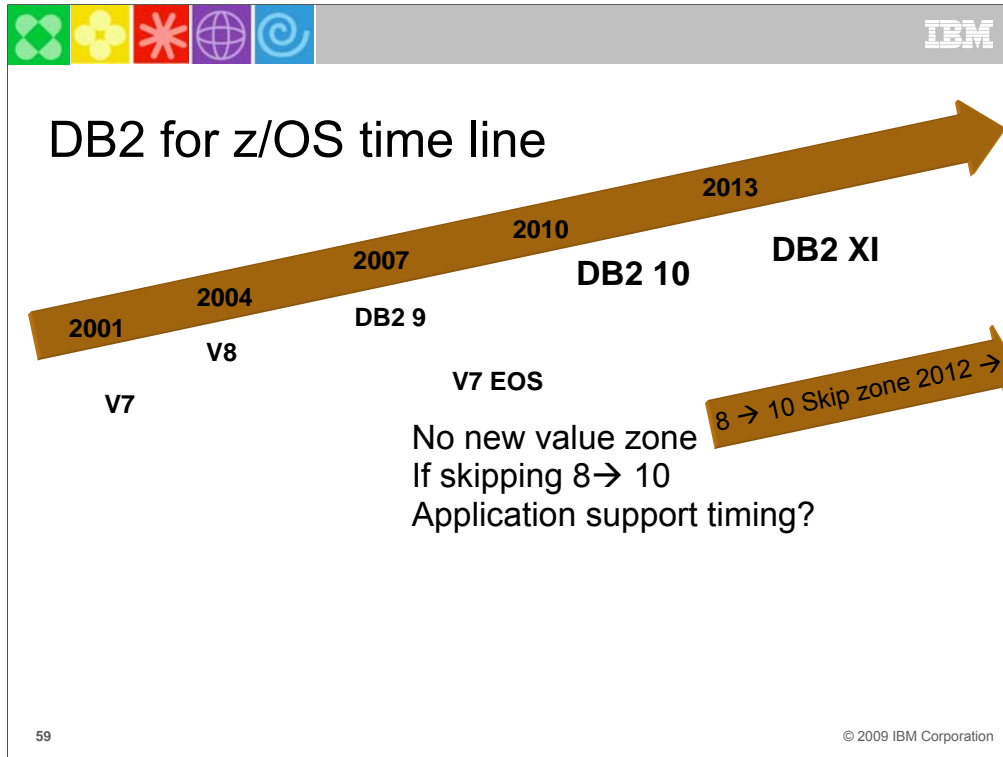
When will applications add support for 10? Drop 8?

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 is Sept 2009. Migrating to DB2 9 is faster, easier and smoother now. We have experience with skip release migration. We did provide skip level migration from V5 to V7, but not for any other recent releases. The ability to skip helps some customers who have not installed recent versions, but this work tends to help a minority of the DB2 customer base. The key question for skipping DB2 9 is, "How long can you wait for that function?" If you are on DB2 V8 today, how long after GA do you wait before moving to a new version? Support for the new version in applications and tools can delay the process. Waiting for others to deliver success stories adds more time. If you wait three years after GA, then you could start moving to DB2 10 in 2014, and probably get there about a year later. I expect DB2 V8 to be out of support before that, so you would probably have additional expense for extended service.

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

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




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

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

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

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



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

Availability, Scale & Resiliency	More Online Schema Changes	➔	Business Flexibility
	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	➔	50% Disk Savings

60 © 2009 IBM Corporation


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.

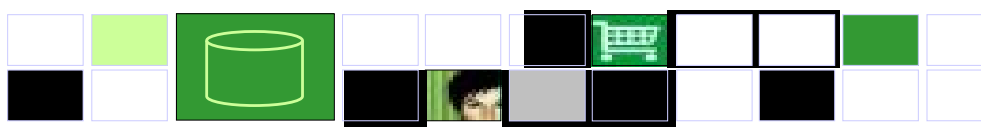
DB2 9 is the best warehouse delivery in DB2 for z/OS ever, with improved SQL, optimization and surrounding products.

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



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




61 © 2009 IBM Corporation

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 61



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-managed 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
- 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


62 © 2009 IBM Corporation

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

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

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

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



No longer supported 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-managed. 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

63 © 2009 IBM Corporation

No longer supported, moving from DB2 9

Private protocol: Convert to DRDA.

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

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

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

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

DB2 MQ XML functions: Use pureXML functions.


DB2 XML Extender: Use pureXML.

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

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

Several subsystem parameters

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



No longer supported 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-managed. 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
- ...

64 © 2009 IBM Corporation

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

Private protocol: Convert to DRDA.

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

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

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

Compression for table space SPT01 is not supported.

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


DB2 MQ XML functions: Use pureXML functions.

DB2 XML Extender: Use pureXML.

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

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

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




No longer supported 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
- ...

65 © 2009 IBM Corporation

No longer supported ...:

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



Deprecated, may be removed from later versions:

- Simple table spaces
- Partitioned table spaces other than universal table spaces
- DSNHDECP load
- NEWFUN(YES) and NEWFUN(NO)
- DSNHPC7 precompiler
- EXPLAIN stored procedures
- ...

All of these lists are much longer, see book for details.

66 © 2009 IBM Corporation


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

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

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

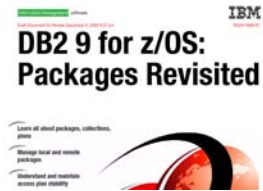
SQL processing options NEWFUN(YES) and NEWFUN(NO) options are deprecated. Use NEWFUN(V10) rather than NEWFUN(YES). Use NEWFUN(V9) or NEWFUN(V8) rather than NEWFUN(YES).

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



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





67

© 2009 IBM Corporation



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.




Convert private protocol to DRDA

- Trace to find use of private protocol
- Private protocol trace analysis program DSN1PPTP
- Use DRDA protocol REXX tool DSNTDP2DP
- 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



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

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



68

© 2009 IBM Corporation

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.



Questions?

Thank
YOU

69

© 2009 IBM Corporation

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



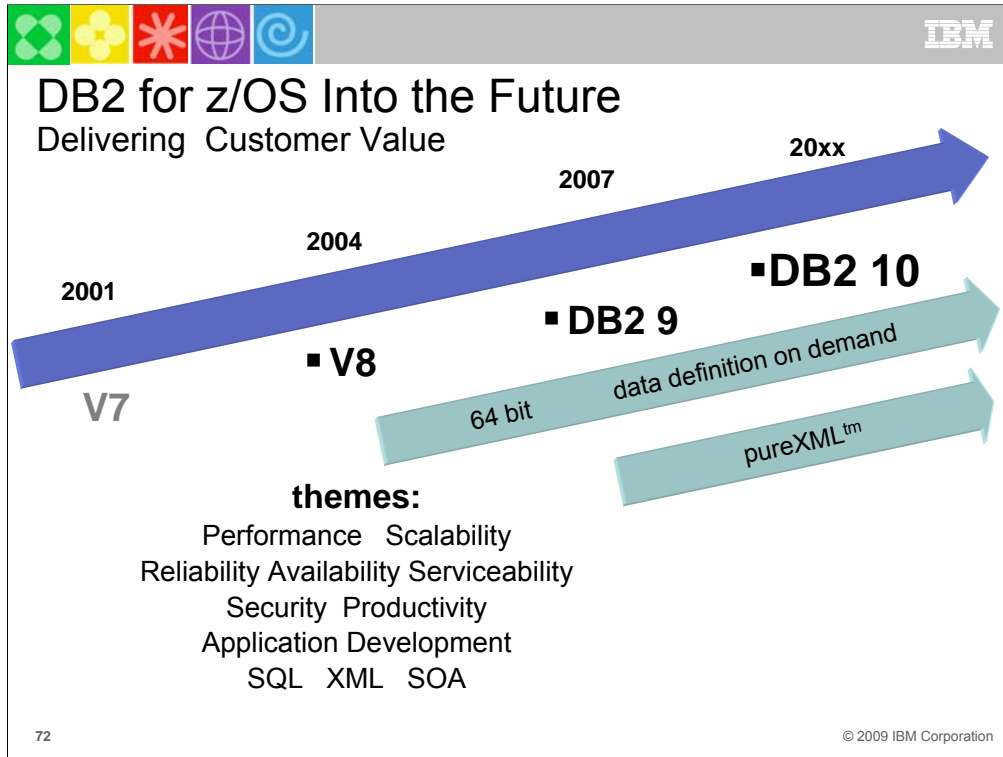
DB2 for z/OS Technical Strategy

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



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 whatever the name becomes.

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


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

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



DB2 9 and 10 in IBM Redbooks Publications

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



73

© 2009 IBM Corporation

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

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

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

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