



DB2 Aktuell
DB2 for z/OS Migration Delivering the Value
or Migration Step 2

Roger Miller, DB2 for z/OS technical evangelist
Dienstag, 23. September 2008



**INNOVATE.OPTIMIZE.
PERFORM.**

IBM INFORMATION ON DEMAND 2008
October 26 - 31, 2008
Mandalay Bay
Las Vegas, Nevada

DB2 for z/OS Migration Step 2: Delivering the Value

We have gone through all the steps to the new version. Are we done yet? What now? Even though some improvements come immediately as you migrate, even more are delivered when you start using the new function in New Value Mode. This session helps you understand and plan to deliver the value.

This is migration step 2, the big one. Understanding which changes deliver the most value and how to make the changes easily is the challenge that we address. Migration step 1, which we'll skip over, discusses migrating to DB2 9, preparing to use the new functions. Process, recommendations from many other customers, and pointers to more information are included. This session will focus on the tasks of the person who installs DB2. Systems programmers and database administrators generally perform this task. The speaker has worked in DB2 development for the past 29 years, and continues working with the team building the install & migration process.

DB2 9 for z/OS has many new features. This presentation will discuss planning to migrate to DB2 9, preparing to use the new functions. We assume that you know DB2 9. We'll emphasize process and recommendations from many other customers, with pointers to more information. The content of this session will focus on the tasks of the person who plans or installs DB2. Systems programmers and database administrators generally perform this task, but there is work for all DB2 users.

Migration Staging

- Planning:
 - Education for planning, Plan to educate
 - Detailed project plan and communications
 - Install prerequisites
 - Find and fix incompatibilities
 - Plan virtual and real storage
 - Plan performance
- Migration: monitoring and tuning
- Stage 2: implementation
 - Using new improvements to deliver value



2



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

Why Migrate to DB2 9 for z/OS?

- Business needs
 - 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



DB2 9 has a lot for everyone. Here are just a few of the strategic highlights.

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

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

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

Common Challenges: DB2 for z/OS Migration

- Code page check and correction (V8)
- Current product and tool support
 - Old COBOL or PL/I or tools
- Cleaning up incompatible change
- Real and virtual storage
 - Commit intervals, Calls, ...
- Plan, monitor and manage performance



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

4

These are the key items that every customer should check first, to see the most common items that slow down a migration to V8. We will discuss each of these points in detail on later slides:

- Checking and cleaning up code page issues, if present. DSNTIJP8 provides the information. Run it on every subsystem.
- Related product and tool support requires looking at the documentation, migrating from out-of-service products to the supported levels, and getting the correct releases and service installed. Talk with your vendors soon.
 - Old COBOL or PL/I is a common example of this situation. Compiles need to use current products. V8 does add some old compiler capability, but DB2 9 is not expected to handle obsolete products.
 - Omegamon V4 is needed for DB2 9. See detail later.
- Cleaning up incompatible changes takes time and a detailed understanding of your shop. Testing is needed.
- Real and virtual storage may improve a little or a lot. Planning is required.
- Plan and monitor performance, as yours can improve substantially or degrade on migration. Be ready to do the needed work.

Migration to DB2 9 is easier than V7 to V8

- Migration process enhancements: ENFM shorter, CM*
- Much less performance regression:
 - Earlier improvements
 - Bind stability & tools for avoiding access path issues
- CCSIDs and old product issues resolved in V8
- Simpler virtual storage considerations
- Less impact from incompatible changes
- Earlier deliveries from vendors



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

Migration process enhancements

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

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

Much less performance regression

More consumable performance improvements

Utility CPU improvements in CM

Bind stability to reduce concern about access path regression.

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

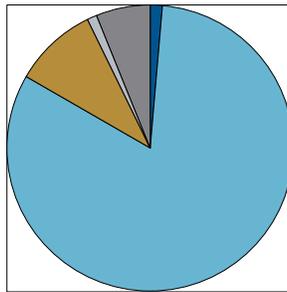
Simpler virtual storage considerations – incremental improvements.

Less impact from incompatible changes

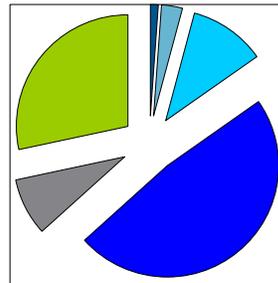
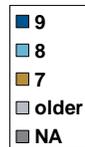
Many vendors are ready today.

Customer reactions to DB2 9

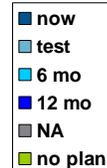
IDUG webcast , 154 people responding (not NA)
100 or 65% moving to DB2 9 in next 12 months



Now on



DB2 9
plans



6



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

Notes from an IDUG webcast January 9, 2008. Two questions were about DB2 9 migration. 154 people on the call answered the second question. Some consultants with no DB2 answered NA or not applicable. When the question was asked about which version they are on,

- 2 answered 9.
- 123 answered V8.
- 14 answered V7,
- 2 answered older, and
- 9 answered NA (e.g. consultants, no subsystems)

When asked about plans for moving to DB2 9, removing NA (null) responses from the percentages:

- 1 said 9 production now;
- 5 said DB2 9 in test now (3%);
- 19 (12%) indicated migration to DB2 9 within 6 months;
- 81 (53%) indicated migration to DB2 9 within 6 - 12 months;
- 48 (31%) indicated no plans for DB2 9 migration;
-
- 14 indicated NA.

DB2 9 for z/OS field update

- **Largest Customers are migrating (12% of top 100)**
- **How's the quality (comparison to V8)?**
 - Lower overall PMR volume
 - Less Sev 1 APARs
 - Lower PE rate
- **V8:**
 - 100% of the top 100 Customers
 - More than 98% of the top 200 Customers
 - Better quality than V7
- **V7 out of service June 30, 2008**



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

7

Now that DB2 9 has been generally available in the field for over a year, let's compare with V8 at the same point. For almost all of the key quality measures, DB2 9 is better. The most important issue for many customers is PTFs in Error (PEs), and the rate is lower on DB2 9. The PMR volume is lower and the number of field severity 1 problems is lower. V7 was widely regarded as a high quality version, and DB2 V8 total field experience is better than V7, with DB2 9 as the best of the three.

Some customers are still migrating to DB2 for z/OS Version 8. The largest DB2 customers have moved to Version 8. All of the largest 100 customers are running Version 8, and large customers are implementing DB2 9 development.

Part of the improvement in overall quality measures is due to more rigorous testing. Some of the improvement is due to the reduced amount of new function added in the service stream. Still these numbers represent the averages, and we know that our customers are not average. Individual customer results are highly skewed, so the average numbers might not reflect your individual experience.

DB2 V7 is now out of service after seven years of general availability. If you are still there, it is time to move now.



DB2 for z/OS V8 field update

- 100% of the top 100 Customers
- More than 98% of the top 200 Customers
- How's the quality (comparison to V7)?
 - Lower overall PMR volume
 - Less Sev 1 APARs
 - Quicker APAR closure time
 - Lower PE rate
 - Less multisystem outages
- **V8 is better than V7 (out of service June 30, 2008)**

Common Criteria
 security certification: EAL3+
 DB2 for z/OS V8



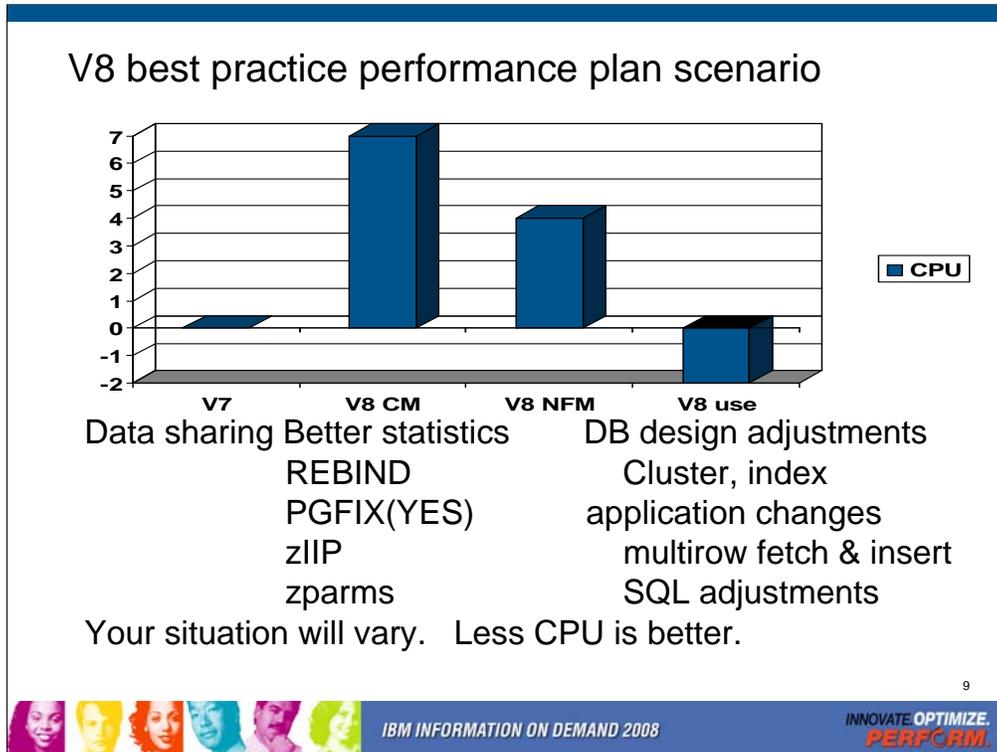
IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
 PERFORM.

Some customers are still migrating to DB2 for z/OS Version 8. The largest DB2 customers have moved to Version 8. All of the largest 100 customers are running Version 8, and large customers are implementing DB2 9 development. Now that DB2 V8 has been generally available in the field for over four years, let's compare with V7 at the same point. For almost all of the key quality measures, V8 is better. The most important issue for many customers is PTFs in Error (PEs), and the rate is lower on V8. The PMR volume is lower and the number of field severity 1 problems is lower. V7 was widely regarded as a high quality version, and DB2 V8 total field experience is better than V7.

Part of the improvement in overall quality measures is due to more rigorous testing. Some of the improvement is due to the reduced amount of new function added in the service stream. Still these numbers represent the averages, and we know that our customers are not average. Individual customer results are highly skewed, so the average numbers might not reflect your individual experience.

DB2 V7 is now out of service after seven years of general availability. If you are still there, it is time to move now.



Your situation and mileage will vary, but this is a common shape for a V8 performance plan, starting with zero for the V7 base line. When you move to V8, CPU time generally increases from 5% to 10%, shown here as 7. Start with long term page fix for buffer pools with high numbers of pages read and written. Reorg and collect improved statistics for non-uniform distribution of data on non-indexed columns. The V8 CM performance plan REBINDs the primary packages, and adjusts DSNZPARMS. The CM REBIND process provides most of the improved access paths. Data sharing batching helps in CM. During CM, a zIIP is added if your peak work load includes DRDA SQL, parallel query or LOAD, REORG and REBUILD.

In moving to NFM, some additional DSNZPARMS are adjusted and REBIND all plans and packages. Database designs start taking advantage of new clustering & indexing options, such as NOT PADDED for large varchar indexed columns. After making the design changes, REORG the data; REORG or REBUILD the indexes; get improved statistics & REBIND. The data sharing group is quiesced, and protocol 2 locking is used.

V8 use takes more advantage of the V8 performance improvements: MQTs, DPSI, more not-padded indexes, multi-row Fetch, cursor Update, cursor Delete, & Insert. Use other SQL improvements to reduce V8 CPU, less than V7. The work may grow, but some of the growth uses the zIIP.

Most consumable DB2 9 improvements

- CM very little to no action:
 - Utility CPU reductions
 - Logging improvements
 - Improved index page split
 - Larger prefetch, write & preformat quantities
 - LOB performance
 - DDF VSCR
 - Optimization Service Center, Opt. Expert, & Data Studio
- Changed online REORG
- Improved RUNSTATS
- Optimization improvements, EDMPOOL VSCR
- Increased parallelism, use of zIIP
- NFM
 - Reordered row format
 - Index: larger page sizes, compression, index on expression
 - LOB lock avoidance



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

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

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

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

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

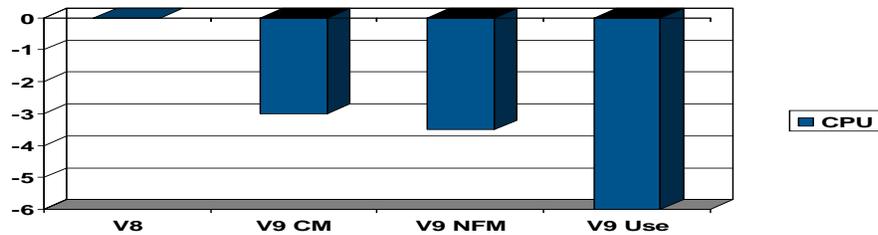
Optimization Service Center takes some learning, but should be fast for those who have used Visual Explain in the past. See the new redbook, SG24-7421, DB2 9 for z/OS: New Tools for Query Optimization.

Reordered row format requires a REORG in NFM and varying length columns.

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

LOB lock avoidance requires a quiesce of all subsystems in NFM, until APAR PK62027 is applied.

DB2 9 z10, z9, z890 & z990 performance plan scenario



Utilities	DB design adjustments
Histogram statistics	Index improvements
REBIND	application changes
DSNZPARMS	native SQL procedures
	SQL adjustments

Your situation will vary. Less CPU is better.
z800 and z900 expect +5% to +10% CPU

11

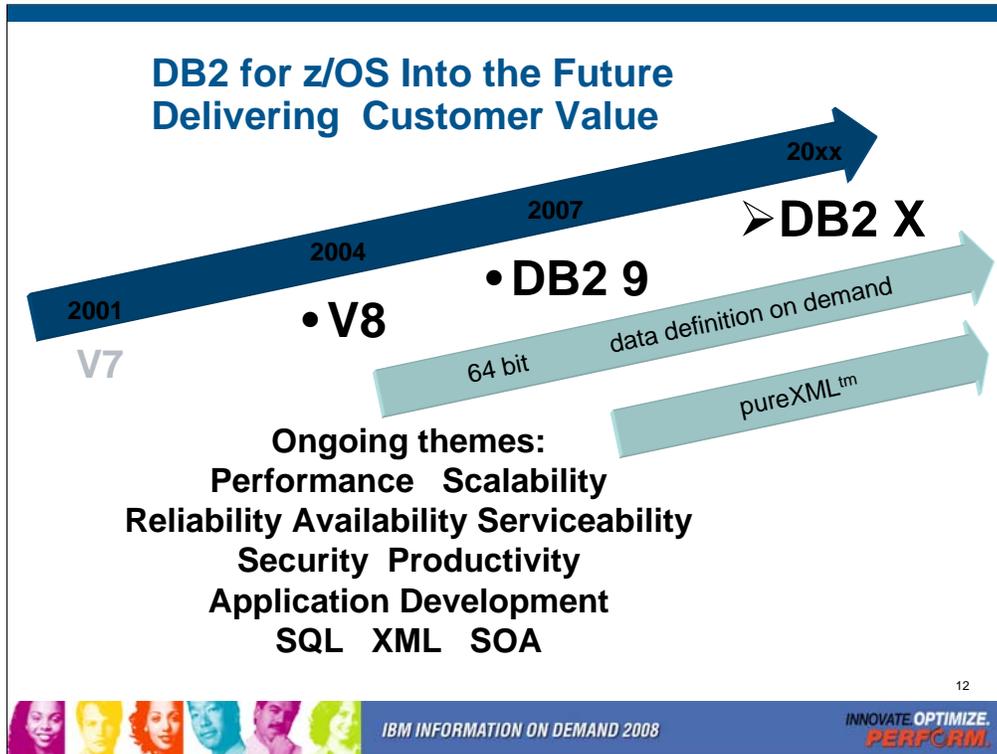


IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

If you have a z9, z990 or z890, this is expected to be a common shape for a DB2 9 performance plan, starting with zero for the V8 baseline. When you first move to DB2 9, total DB2 CPU time generally decreases from 0% to 5% for z9, z890 and z990 customers, shown here as a first step -3%. Utility CPU reductions help immediately. Some work will be about the same (+/-3%). Start with reorgs and collect improved histogram statistics when useful. The DB2 9 CM performance plan REBINDs the primary packages and adjusts DSNZPARMs. The REBINDs provide most of the improved access paths. On z800 or z900 the initial cpu expectation is +5 to +10% regression, more if there are many columns, so making adjustments is more important. In moving to NFM, some additional DSNZPARMs are adjusted and all plans and packages are rebound. The DB2 9 use line takes wider advantage of DB2 9 performance improvements. Database designs start taking advantage of new indexing options, such as compression, index on expression and larger pages. After making the design changes, REORG the data and REORG or REBUILD the indexes, get the improved statistics and REBIND. Native SQL procedures, added use of zIIP, and improved SQL continue the improvements in this phase.

Scenario: Customer mix of DB2 CPU time is 30% in utilities, 70% in SQL access. With 10% improvement for the utilities, we get a -3% net, assuming that SQL is the same as before. With optimization improvements, another -½% improvement shows up in DB2 9 NFM. Then as design adjustments, reorgs and rebinds are performed, we get improvements from varchar improvements, native SQL procedures and improved SQL, another -3%.



DB2 for z/OS V7 became generally available (GA) March 2001, and V8 delivered three years later, March 2004. DB2 9 became generally available in March 2007, three more years. We expect the next version will be roughly 3 years from DB2 9 GA to DB2 10 or DB2 X or whatever the name becomes, but the timing and packaging may change too.

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

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

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

DB2 for z/OS Version Availability Summary

Version	PID	Generally Available	Marketing Withdrawal	End of Service
V3	5685-DB2	Dec. 1993	February 2000	March 2001
V4	5695-DB2	Nov. 1995	Dec. 2000	Dec. 2001
V5	5655-DB2	June 1997	Dec. 2001	Dec. 2002
V6	5645-DB2	June 1999	June 2002	June 2005
V7	5675-DB2	Mar. 2001	Mar. 2007	June 2008
V8	5625-DB2	Mar. 2004		
DB2 9	5635-DB2	March 2007		

13



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

All statements regarding the future direction and intent of IBM are subject to change or withdrawal without notice and represent goals and objectives only. This table summarizes the latest versions of DB2. As of August 2007, DB2 V7 was generally available for more than six years and has been out of service for a few months. If you need service, a contract is available. DB2 V8 has been generally available for over four years. If you are currently running V7, then it's time to finish your migration to V8. I generally think of five years after general availability as the end of practical life for a DB2 release. Migration to DB2 9 is only from DB2 for z/OS V8. There is no standard ability to skip a version. A few customers are still running older versions, but the end of service is past. In general, you can expect versions to be in service for five years and end of service to be announced a year in advance. Extended service at extra cost is available for three more years. If you want to search for these products by version number, the PID number is useful. Almost all of the largest DB2 customers are running V8 or DB2 9 now. End of service of Version 7 was June 2008, after end of service for all of the operating systems that support processors other than System z or zSeries.

To see current DB2 availability notices, start at ibm.com/software/db2zos
Click on Support, then on Product Life Cycle.
ibm.com/software/data/db2/zos/support/plc/

z/OS Support Summary

z/OS									
R6	x	x	x	x	x	x	9/07	1.8	
R7	x	x	x	x	x	x	9/08	1.9	
R8	x	x	x	x	x	x	9/09*	1.10*	
R9	<u>x</u>	<u>x</u>	<u>x</u>	<u>x</u>	<u>x</u>	<u>x</u>	<u>9/10*</u>	<u>1.11*</u>	<u>9/07</u>
R10*	x	x	x	x**	x	x	9/11*	1.12*	9/08*
R11*	x	x	x	x	x	x	9/12*	1.13*	9/09*

z/OS 1.9 Coexistence-supported releases

Release	Coexistence-supported
z/OS 1.9	z/OS 1.7, z/OS 1.8, z/OS 1.9
z/OS 1.10*	z/OS 1.8, z/OS 1.9, z/OS 1.10*
z/OS 1.11*	z/OS 1.9, z/OS 1.10*, z/OS 1.11*

z/OS.e 1.7, 1.8 supported on z800, z890, and z9 BC only. There is no z/OS.e 1.9.

* Statements regarding IBM future direction and intent are subject to change or withdrawal, and represents goals and objectives only.

** zIIP Web Deliverable required for HiperDispatch support on System z10

14

z/OS 1.9 Coexistence-supported releases*

Release	Coexistence-supported
z/OS 1.9	z/OS 1.7, z/OS 1.8, z/OS 1.9
z/OS 1.10*	z/OS 1.8, z/OS 1.9, z/OS 1.10*
z/OS 1.11*	z/OS 1.9, z/OS 1.10*, z/OS 1.11*

z/OS.e 1.3 - 1.8 supported on z800, z890, and z9 BC only. There is no z/OS.e 1.9

x1 – IBM eServer™ zSeries® 990 (z990) compatibility or exploitation feature required (no longer orderable)

* Statements regarding IBM future direction and intent represent goals and objectives only, and are subject to change or withdrawal without notice.

This chart summarizes much of the information about z/OS that you need to know. Server and DASD hardware support, end of service dates, coexistence, and planned availability dates are shown here for existing and planned releases through z/OS R9.

z/OS 1.8 was announced and became available in September 2006. 1.8 is required for some DB2 9 function. z/OS 1.9 became available in September 2007.

z/OS 1.6 is out of service as of September 2007. 1.4 and 1.5 are out of service as of March 2007. z/OS (and z/OS.e) 1.7 planned end of service is September 2008, with extended service available for additional cost.

Only service-supported releases can coexist in the same sysplex

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

This chart notes the support for z/OS with various processors.

Note the z/OS end of service dates as well. End of service for z/OS 1.3 was March 2005. Check the latest information by going to the z/OS web page, then to Support and then to the z/OS marketing and service announce, availability, and withdrawal dates

http://www.ibm.com/servers/eserver/zseries/zos/support/zos_eos_dates.html

Staged Migration process: best practice

- **Test with New Install & Migrate (sandbox)**
- **Migrate to new release in Conversion Mode (1 – 2 months)**
- **New Function Mode (NFM)**
 - Enable New Function Mode (minutes)
 - Move into NFM
- **Deliver value. Take advantage of new function.**



15



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

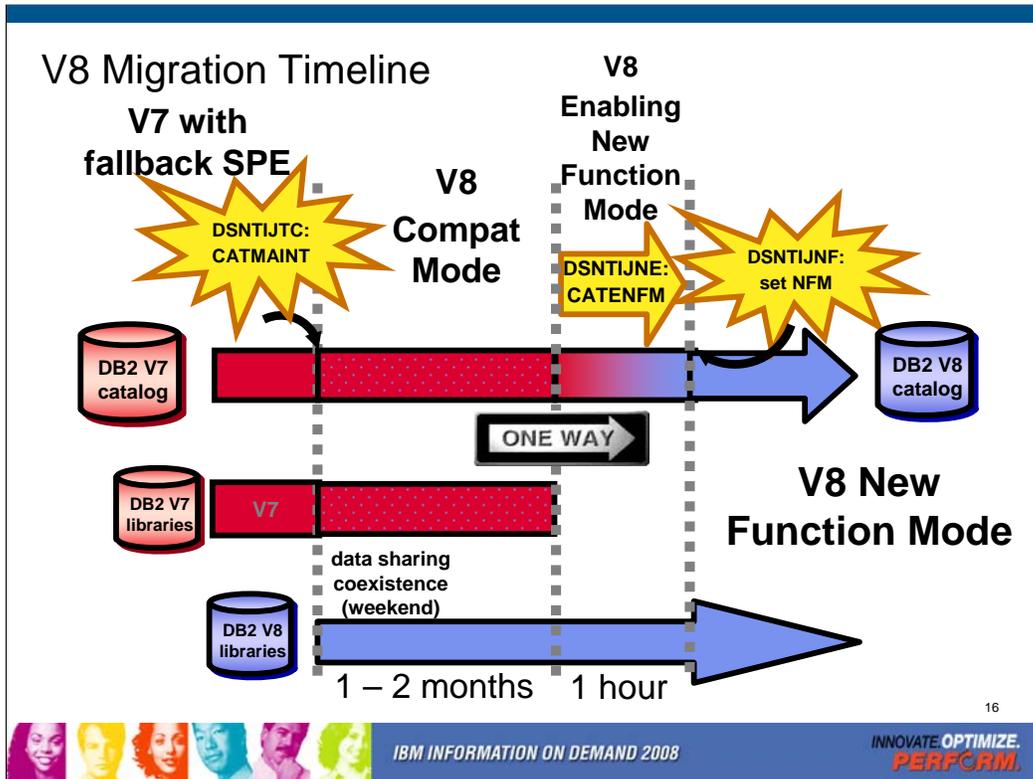
The process of migration will take four major steps.

Use a new install or a migration on a separate subsystem and practice migration there. Test the function you want to use. This will ensure that you are ready to move forward.

Migrate to the new release in Conversion Mode (CM). You are using the new code but almost no new function. Conversion mode is NOT no new function mode, but does have very limited new function. The primary objective is making the full migration smooth, so we prefer to avoid problems, but if they can't be avoided, then CM is much better than NFM

The third step is a few minutes for Enabling New Function Mode (ENFM), typically a one hour process, moving directly to the next, New Function Mode or NFM.

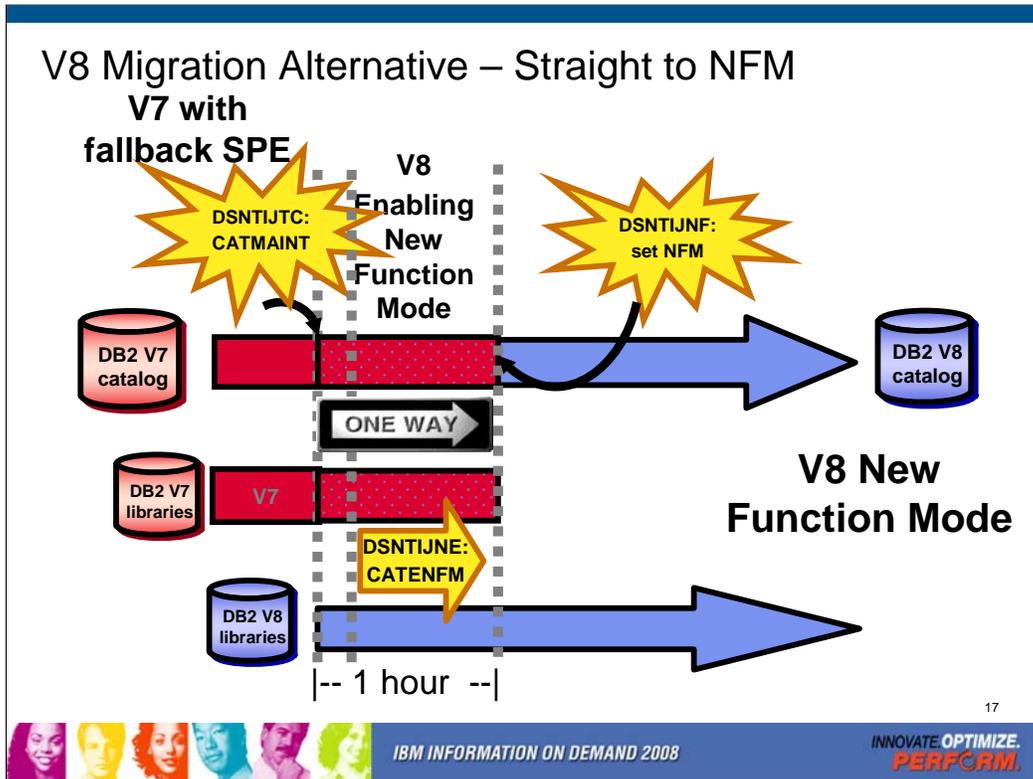
After you are running New Function Mode (NFM), then one or more projects are used to deliver the value of V8 or 9, improving performance, availability and new capabilities in SQL, XML and utilities.



This is an overview of the V8 migration process. There are two steps which modify the catalog, CATMAINT to change into Compatibility Mode (CM), and also CATENFM for Enabling New Function Mode (ENFM), which contains the major catalog changes to long names and Unicode, using a technique of an online (read only or shrlevel reference) reorg. The DB2 9 migration process is similar.

Keep the time for data sharing coexistence fairly short, such as over a weekend or one week. Compatibility mode will often be one or two months, but is not used for some vendors, such as SAP. Once you move to ENFM, you don't go back to CM or to V7. ENFM will generally be on the order of an hour (5 GB catalog size). Then you turn the switch and are taking advantage of New Function Mode, NFM.

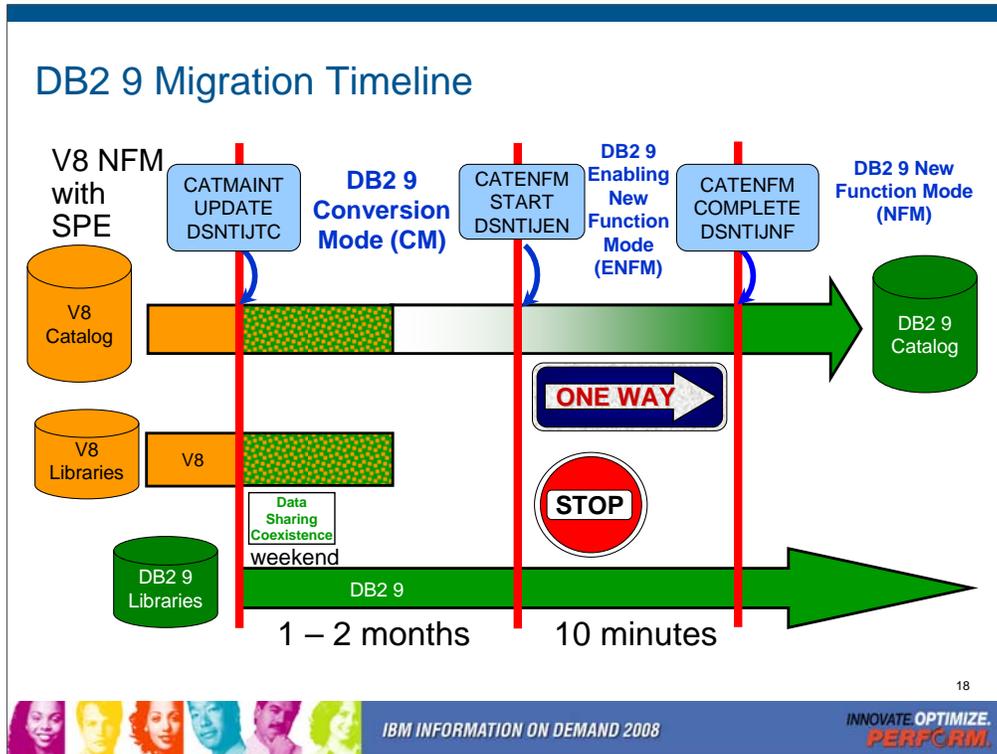
How long should you stay in CM? Generally after more than one or two months, you have adjusted to the improved access paths, storage changes and improved utilities, so that fallback is not feasible. At that point, moving forward to NFM can eliminate some of the restrictions and operational problems. You are running almost exactly the same code, so problems in NFM are generally in new v8 function, found in test.



Some customers will use an alternative migration process, migrating immediately from V7 to V8 NFM. This process means that fallback to V7 is not possible, so thorough testing is needed. The advantages are

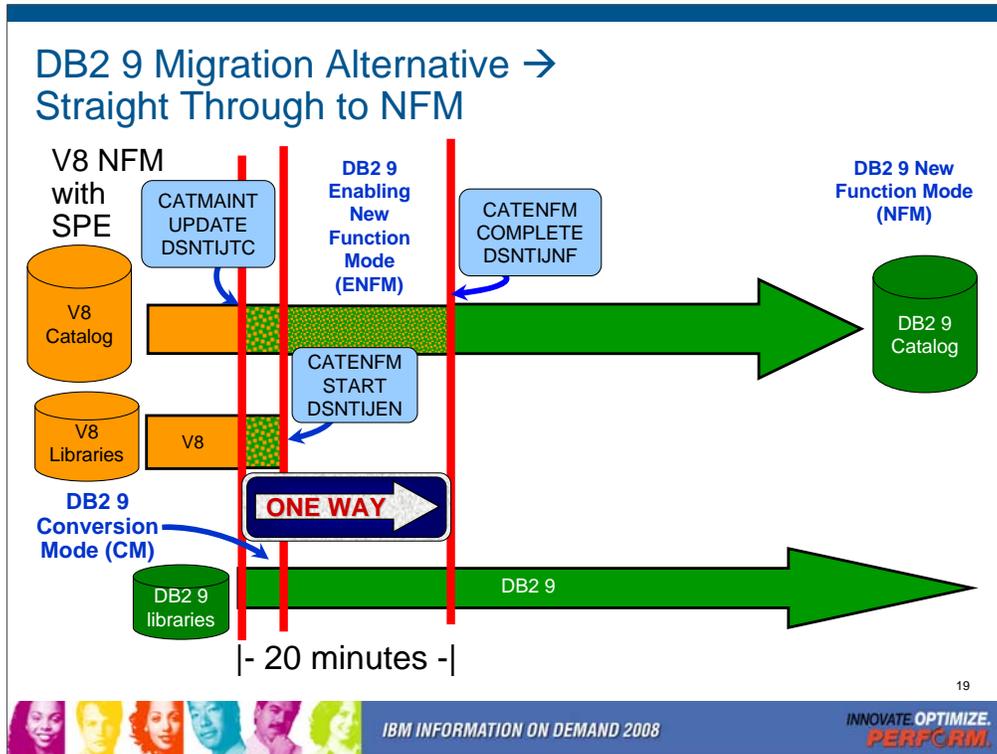
- single window for the migration
- test in only one mode, NFM
- faster delivery of improved value in performance, function and availability

Some vendors, such as SAP, use this process. Customers who cannot allow multiple windows for migration need to consider this alternative and the work and controls needed to migrate while processing continues.



The time line is not to scale. Customers should spend only a few minutes in ENFM and then move immediately to NFM.

The ONE WAY situation is a bit more flexible, with more opportunity to drop back. There are two new modes to denote drop back processes. You can drop back from NFM to ENFM as before, but the label for the drop back mode is now ENFM*. It is also possible to drop back from ENFM or NFM to CM*, but not to V8 or to coexistence.



Some customers will use an alternative migration process, migrating immediately from V8 to DB2 9 NFM. This process means that fallback to V8 is not possible, so thorough testing is needed. The advantages are

- single window for the migration
- test in only one mode, NFM
- faster delivery of improved value in performance, function and availability

Some vendors, such as SAP, use this process. Customers who cannot allow multiple windows for migration need to consider this alternative and the work and controls needed to migrate while processing continues.

ENFM and NFM Considerations



STOP!
THINK!

Attention: You cannot return to the prior version once you enter ENFM. Do NOT move to ENFM until you are certain that you will not need to return.

HOWEVER
The code base for DB2 9 ENFM and NFM is the same as for CM.
You can return to CM* from ENFM or NFM if necessary.

20



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

In DB2 V8, you cannot return to CM when you get to ENFM. In DB2 9, you can return to CM* from ENFM. We have heard a lot of concern about this step which does not have an easy reversal, but have never experienced a problem that could be resolved by this return. It still makes the process easier to explain as always being able to get back to the prior state or a reasonable facsimile thereof.

DB2 9 migration continues the best practice migration technique, with small, safe steps. Every step can be backed off. Customers move first to conversion mode, and can move back to V8 from there. The next step, ENFM, changes the catalog, and customers can go back to CM* from there. The * is our note that the subsystem has moved before. The customer can move to NFM and then can go back to ENFM* or even to CM*. Most customers move directly to NFM at the end of the 10 minute ENFM, so there are two steps to the process.

DB2 9 Modes – An Overview

CM Conversion Mode (previously called compatibility mode) - This is the DB2 mode DB2 9 migrated from V8 and is started for the first time. It will still be in CM when migration job DSNTIJTC has completed. Very little new function can be executed in CM. Data sharing systems can have V8 and DB2 9 members in this mode, called coexistence. Coexistence should be very short, such as a weekend or a week. DB2 can only migrate to CM from V8 NFM. This mode is not compatibility, but the ability to fall back. We try to move most (but not all) problems for migration from NFM to CM, so that fallback can be used, if necessary.

ENFM 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 9 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.

ENFM* This is the same as ENFM but the * indicates that at one time DB2 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 ENFM* it can not fallback to V8 or coexist with a V8 system.

CM* This is the same as CM but the * indicates that at one time DB2 was at ENFM and possibly NFM. Objects that were created in NFM can still be accessed. When DB2 is in CM* it can not fallback to V8 or coexist with a V8 system. Functionally, CM* is very similar to ENFM.

21



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

DB2 9 builds upon the structure from V8. The modes are almost the same, but DB2 9 clarifies the situation if you drop back after moving to ENFM or NFM. There are two more modes with names CM* and ENFM*, and the restrictions are clarified a little more. The process is very similar to the best practice that successful customers have used for years. It is possible to move from NFM or ENFM to CM*.

While the first mode was labeled as “Compatibility Mode” or CM, it is not compatible. CM does help with the conversion, but the incompatibilities have been moved to this mode, where possible, since fallback is simple from this stage. This change removed problems in moving to ENFM and NFM.

I think of NFM as New Value Mode, the opportunity to deliver the value of the new version.

What works in Conversion Mode?

- Logging improvements, improved index page split
- DSN1COMP index compression estimate, NOT index compression
- Larger prefetch, write and preformat quantities
- LOB performance other than LOB lock avoidance
- Command enhancements: REFRESH, DISPLAY, START, STOP, & MODIFY TRACE qualifications and ACCESS DATABASE
- Virtual storage use above the 2 GB bar DDF & EDMPOOL (when bound)
- Package BIND stability (PLANMGMT & EXCHANGE)
- Optimization improvements (when bound)
- Increased parallelism and use of zIIP
- Optimization Service Center & Optimization Expert (work with V8 if PTFs are applied)
- Most utility enhancements and CPU reductions (see next slide)

22



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

New functions are available only in new-function mode, unless explicitly stated otherwise in the product documentation. A few general exceptions exist for utilities and for optimization. The following specific improvements are available in conversion mode:

- Logging improvements
- Improved index page split
- Larger prefetch, write and preformat quantities
- LOB performance other than LOB lock avoidance
- Virtual storage use above the 2 GB bar DDF & EDMPOOL (when bound)
- Package BIND stability (PLANMGMT & EXCHANGE)
- Optimization improvements (when bound)
- Optimization Service Center & Optimization Expert (work with V8 if PTFs are applied)
- Utility CPU reductions and most other utility improvements that do not depend on SQL changes. See the following slide and web for detail.
- <http://www.ibm.com/support/docview.wss?rs=64&uid=swg21307783>

DB2 9 Conversion Mode versus New Function Mode

Utility enhancements available in both CM and NFM are:

- CPU reductions
- Parallel unload, reload and log apply processing by the REORG utility
- Removal of the BUILD2 phase for REORG at the partition level
- Log Phase addition for REORG SHRLEVEL REFERENCE PART
- REORG SHRLEVEL REFERENCE for LOB table spaces
- CHECK INDEX SHRLEVEL REFERENCE with Parallelism
- Histogram Statistics for RUNSTATS
- CHECK LOB SHRLEVEL CHANGE
- CHECK DATA SHRLEVEL CHANGE
- REPAIR LOCATE SHRLEVEL CHANGE
- REBUILD INDEX SHRLEVEL CHANGE
- MODIFY RECOVERY support for the RETAIN option
- Object level recovery from a system-level backup by the RECOVER Utility
- BACKUP SYSTEM & RESTORE SYSTEM system-level backups to tape
- Skipping locked rows for transaction updates by the UNLOAD Utility
- TEMPLATE switching support based on size for image copy datasets

Other utility enhancements are only available for DB2 9 NFM, such as...

- LBI (large Block Interface) support for sequential datasets on tape
- Recover to Point in Time (PIT) with consistency (defaults as new behavior)
- Clones effectively provide LOAD REPLACE SHRLEVEL CHANGE
- Online Load (LOAD RESUME YES SHRLEVEL CHANGE) without logging using NOT LOGGED

23



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

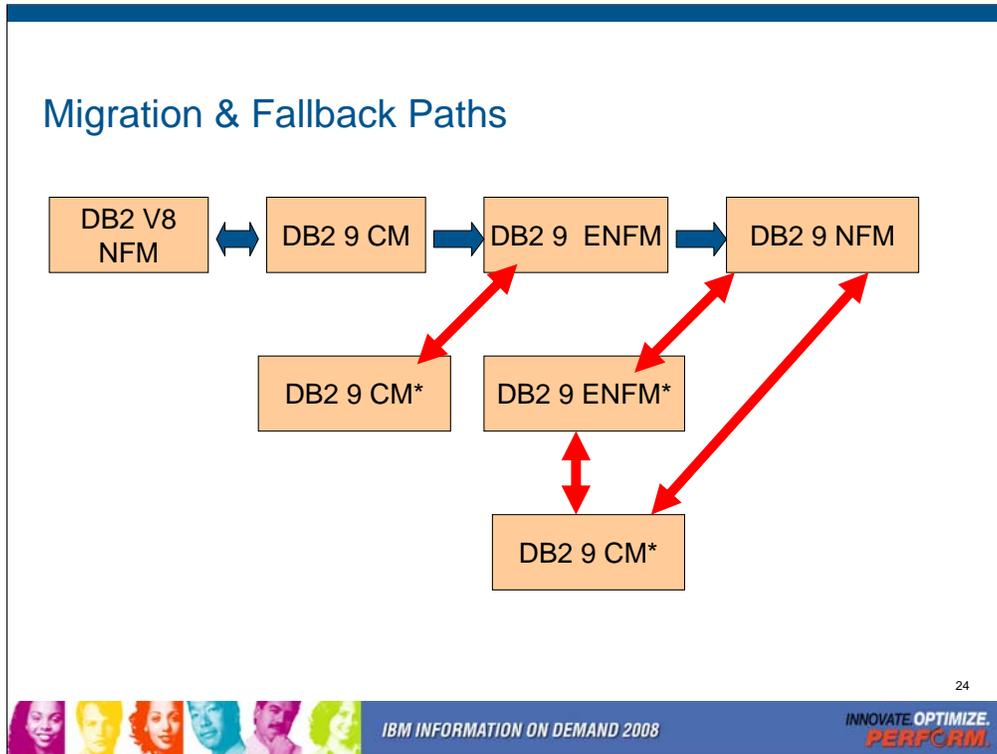
DB2 9 has many utility enhancements available in both Conversion Mode (CM, previously called Compatibility Mode) and New Function Mode (NFM), Some exceptions are only available for NFM, and not for CM.

Utility enhancements available in both CM and NFM are:

- Parallel unload, reload and log apply processing by the REORG utility
- Removal of the BUILD2 phase for REORG at the partition level
- Log Phase addition for REORG SHRLEVEL REFERENCE PART
- REORG SHRLEVEL REFERENCE for LOB table spaces
- CHECK INDEX SHRLEVEL REFERENCE with Parallelism
- Histogram Statistics for RUNSTATS
- CHECK LOB SHRLEVEL CHANGE
- CHECK DATA SHRLEVEL CHANGE
- REPAIR LOCATE SHRLEVEL CHANGE
- REBUILD INDEX SHRLEVEL CHANGE
- MODIFY RECOVERY support for the RETAIN option
- Object level recovery from a system-level backup by the RECOVER Utility
- BACKUP SYSTEM & RESTORE SYSTEM system-level backups to tape
- Skipping locked rows for transaction updates by the UNLOAD Utility
- TEMPLATE switching support based on size for image copy datasets

Other Utility enhancements are only available for DB2 9 NFM, such as...

- LBI (large Block Interface) support for sequential datasets on tape
- Recover to Point in Time (PIT) with consistency (defaults as new behavior)
- Clones effectively provide LOAD REPLACE SHRLEVEL CHANGE
- Online Load (LOAD RESUME YES SHRLEVEL CHANGE) without logging can be accomplished by altering the table space to NOT LOGGED, running the Online Load, altering the table space back to LOGGED, and taking an image copy.



CM is the mode when DB2 9 is started for the first time from DB2 V8 New Function Mode (NFM). Limited new function can be executed in CM. This is the only mode that permits fallback to DB2 V8. CM does help with the conversion, but most incompatibilities occur in this mode, where possible, since fallback is simple from this stage.

Enabling New Function Mode (ENFM) is the mode that is entered when CATENFM START is executed which is the first step of DSNTIJEN. ENFM can fall back to a Conversion Mode which is identified as CM*.

New Function Mode (NFM) is the mode that DB2 enters once CATENFM COMPLETE is run. This is the only step of job DSNTIJNF. New functions can be used at this time. Fallback can occur to Enabling New Function Mode which will be identified as ENFM* or to Conversion Mode which will be identified as CM*.

Enabling New Function Mode Star (ENFM*) is the same as ENFM but indicates that at one time DB2 was in NFM.

Conversion Mode Star (CM*) is the same as CM but indicates that at one time DB2 was in either ENFM or NFM. You can not fallback to DB2 V8 from CM*.

DB2 9 migration modes are C, C*, E, E*, or N in message DSN7100I

DB2 Catalog Evolution

The DB2 catalog continues to grow with every DB2 release.

DB2 Version	Table Spaces	Tables	Indexes	Columns	Table Check Constraints
V1	11	25	27	269	N/A
V3	11	43	44	584	N/A
V5	12	54	62	731	46
V6	15	65	93	987	59
V7	20	84	117	1212	105
V8	21	87	128	1286	105
DB2 9	28	106	151	1668	119

Doesn't include objects for XML Schema Repository

25



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

The DB2 catalog continues to add new catalog table spaces, tables, columns, indexes and check constraints in each new version.

The newer table spaces are segmented, while the older ones are simple. Some use techniques that are unique to the catalog.

DB2 9 for z/OS Migration Planning Workshop

- Understand breadth of features in DB2 for z/OS
- Bring together a toolbox of resources for your migration planning
- Explain the current migration process
- Bring a project focus to migration
- Remain relevant through GA life of the product
 - Updated with field experiences



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

26

What is the Migration Planning Workshop (MPW)? • This offering enables customers to understand the breadth of features delivered in DB2 for z/OS versions. • Many resources are brought together in a single offering. • The migration process is explained. • Leave the session with materials that you can use to start your installation / migration immediately, or in the future. • Your questions are given attention, typically in a small group setting. While most MPWs are small, multi-company events, they can be delivered for individual companies or even user groups.

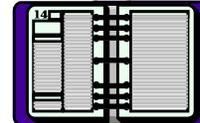
Who should attend? An MPW contains content that is appropriate to: • Application Developers • Database Administrators • System Administrators • Architects • IT decision makers • Project Managers

The morning session provides information for all parties. The afternoon is more migration focused, and therefore a better fit for Database Administrators, System Administrators, and Project Managers. A typical agenda provides a DB2 9 for z/OS Overview in the morning and DB2 9 Migration (• Preparations • Planning • Process) in the afternoon.

What you can expect? • An understanding of the features delivered with DB2 9 and how they can benefit your enterprise. • Clarity of the migration process. • References for many subjects, including: Migration, Fallback, Prerequisites & preparations. You will leave with Presentation materials, Checklists, Project plan framework, Related documentation, Networking and Contacts.

Migration Strategy: Test and Production

- Short time for data sharing coexistence
- One to two months for conversion mode
 - ▶ Not able to use most new performance, function
- Minutes for enabling new function mode
- Move to new function mode
 - ▶ Stage in new function
- Test to QA to Production
 - ▶ Better controls for not using new function
 - ▶ Still don't want long gap test to production
 - ▶ Vendor support may affect strategy, e.g. NFM certified
 - ▶ Improved granularity, stage by modes



27



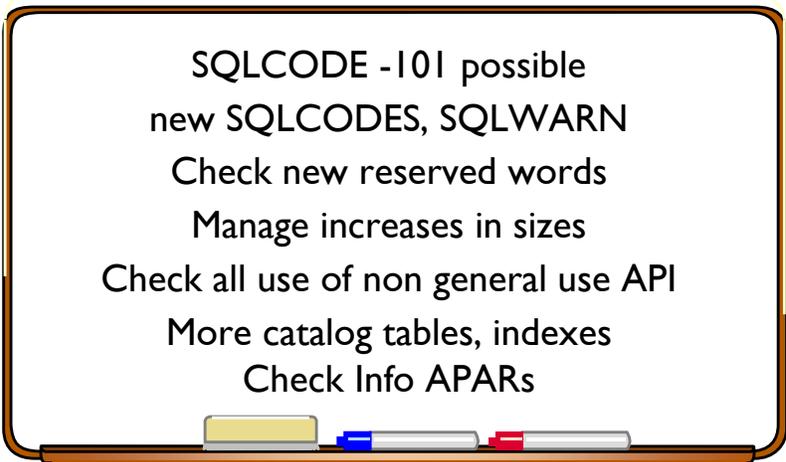
IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

In the past we have recommended a short time for data sharing coexistence, and that is still true. A short period for ENFM is also recommended highly. Support from vendors may affect the migration staging. One concern for conversion mode is that most new performance improvements cannot be used, and some regression is possible, as always.

The timing for moving from test to QA to production has more options to consider. There are better controls for not using new function, but a long gap between test and production levels is not advisable. Now you have more granularity in the migration and can move mode by mode. Some customers are migrating both test and production to CM, then changing to NFM in a short time.

DB2 considerations for any migration



- SQLCODE -101 possible
- new SQLCODES, SQLWARN
- Check new reserved words
- Manage increases in sizes
- Check all use of non general use API
- More catalog tables, indexes
- Check Info APARs

28



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

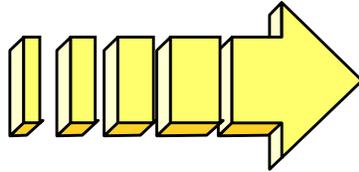
Whenever you move from one release to another, there are some changes to manage. It is possible, but rare to find SQL which processed in the prior release that will not run on the new one. If you have users who test the boundaries, instead of reading the limits, then this is more common. Reserved words are handled in context when possible, but new reserved words can cause problems in SQL that worked on the prior release. Check for APARs. The Info APAR for DB2 9 migration is I112423.

Some catalog tables grow. There are many more catalog tables and indexes. The concern for V8 was often the long names and Unicode, but the catalog growth for that was generally 10% or less. As you REBIND, the SPT01 and SYSPKAGE table spaces often grow 50%.

Error information changes more frequently, note any dependencies upon error information. There are always new SQLCODEs, and we still find a few applications that do not have tests for other cases. Additional complexity means that storage sizes increase in DB2 & application address spaces.

If you use an interface identified as a Product-Sensitive Programming Interface, it must be checked. If it's not in the books, then it's not an interface.

Premigration Activity ...



- Find problems with incompatibilities
 - ▶ It is important to look at these now because many have actions that you must address before migrating.
 - ▶ Get the complete list from the Installation Guide (Get the latest from the web).
 - ▶ Use V8 job DSNTIJP9 APAR PK31841

29



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

Read these pages of the DB2 Installation Guide carefully. Many items need to be addressed before migration. Check the items several times, including just before running CATMAINT.

General REBIND advice: When to REBIND depends upon complexity of the SQL, how well tuned the statements are, and your experiences with improvements & regression. For simple cases, a rule of thumb is that the statistics need to change by 10% or 20% or cross a key threshold to make a difference in the access path.

Real Time Statistics address when to run REORG, RUNSTATS & COPY. Read Administration Guide. Some rules of thumb & threshold techniques discussed in the book may be useful for REBINDs. One of the best times for RUNSTATS & REBIND is after REORG. If you update statistics & the differences are large or cross a boundary, then REBIND is more likely to make a difference. Statistics history tables may help in the decision. After REBIND, check access path & processing estimates.

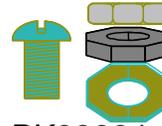
Install Prerequisite Software

- Prerequisites and Functional Needs

- ✓ All products need to be at proper maintenance level (prior DB2 release too)

Migration & Fallback SPE APARs PK11129, PK33664, PK39850, PK38799, PK44884

- ◆ Ensure that you apply all hiper APARS, PE fixes
- ◆ Use the DB2 Program Directories, the Installation Guide, PSP (preventative service planning), web & Info APAR II12423



- ◆ Other software ...Talk to your Vendors

- ◆ z/OS R7 R8 R9



- ◆ Job DSNTIJPM or DSNTIJP9



30



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

You need to know what functions will be used first. Many functions have additional prerequisites. Read the section of Release Guide carefully to see if you need to upgrade other software. The most successful policy is to keep software at similar levels. z/OS should be current & in service. Some functions need later releases. z/OS V1R7 is minimum for DB2 9 and the minimum in service today.

IBM & other vendor software must be checked for levels required & needed for the function you want. Some APARs require vendor changes. Contact your vendors.

Program Directories:

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

<http://www.ibm.com/software/data/db2/zos/v9books.html>

DB2 9 incompatibilities

- **Installation Guide for details**
 - **Chapter 2. Preparing your system to install or migrate DB2**
 - **Changes to DB2 that might impact your migration to DB2 9 conversion mode**
 - Application and SQL release incompatibilities
 - Utility release incompatibilities
 - Command release incompatibilities
 - Storage release incompatibilities
 - Other release incompatibilities
 - Functions that are no longer supported
 - Functions that are deprecated
 - Preparing a data sharing group or distributed environment for migration
 - **Chapter 5. Installing or migrating DB2**
 - **Migrating your DB2 subsystem to new-function mode**
 - **Implications of converting to DB2 9 new-function mode**
- **Application Programming and SQL front**
 - **Application and SQL release incompatibilities**

31



IBM INFORMATION ON DEMAND 2008

INNOVATE.OPTIMIZE.
PERFORM.

For detailed product dependencies and incompatibility descriptions, see

- Application Programming and SQL Guide (APSG)

- Installation Guide

- Program Directories

- Web page for DB2 9 books:

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

- Release incompatibilities are divided into incompatibilities in conversion mode and incompatibilities in new-function mode and are documented in lists in two separate sections in the Installation Guide. Here are links to the two topics that contain these lists:

Parent topic for conversion mode incompatibilities:

Changes to DB2® that might impact your migration to DB2 9 conversion mode:

<http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp?topic=/com.ibm.db29.doc.inst/dmigcn2.htm>

Parent topic for new-function mode incompatibilities:

<http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp?topic=/com.ibm.db29.doc.inst/dmigcn2.htm>

Conversion considerations:

The Installation Guide section is about 18 pages long. Some of the items expected to be the most common are noted here, but your database administrators and application programmers should review the full list, as some practice might be common at your shop.

Application programmers should review the entire section in the APSG, about 9 pages long, with the primary changes that are anticipated to require changes. The following slides have the highlights.

DB2 9 removed function

- DB2-managed stored procedures → WLM-managed SPs
- Create simple table space → create segmented or partition by growth
- Temp databases & table spaces → work files
- Very old plans and packages (V3 and before) → REBIND
- AIV Extender → Content Management
- Text, Net Search Extenders → new text search
- Net.Data → WebSphere, REXX, command line proc., ...
- DB2 Estimator → capacity planning products
- Old JDBC / SQLJ driver → Current driver
- Java stored procedures → non-resettable JVMs
- Online help → Information Center
- Parameters removed RELCURHL, DBPROTCL, STORPROC, ...

32



IBM INFORMATION ON DEMAND 2008

INNOVATE.OPTIMIZE.
PERFORM.

DB2 9 for z/OS supports migration only from V8 NFM. Customers not yet running V8 should plan to migrate to DB2 for z/OS, V8 as preparation for migration to DB2 9. IBM added many data server capabilities in DB2 9 for z/OS and reduced or removed support for some functions: AIV Extender is removed. Text Extender is removed. Net Search Extender is removed. Net.data, which was stabilized at the V7 level, is removed. WebSphere is the strategic IBM solution for delivering DB2 data to Web applications. For some situations a scripting language, such as REXX or the Command Line Processor which comes with DB2 9 can provide the needed function. DB2 Estimator, available via Web download for DB2 V8, will not be provided for DB2 9 for z/OS.

DB2-managed stored procedure support is removed. Workload Manager (WLM) managed stored procedure address spaces is the strategic solution for stored procedure support. Simple table space creation support is removed. DB2 9 for z/OS no longer implicitly creates simple table spaces nor allows customers to create simple table spaces. However, DB2 9 for z/OS continues to support simple table spaces created in previous versions.

The Online help facility is replaced, replaced by the Information Management Software for z/OS Solutions Information Center (Information Center). JDBC/SQLJ Driver for OS/390 and z/OS support is removed. 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). Java stored procedures no longer run in resettable JVMs.

Plans and packages must be rebound on DB2 9 to benefit from optimizer enhancements. Plans & packages DB2 V3 & before, will be automatically rebound when accessed by DB2 9. See the Installation Guide section "Functions that are no longer supported" for more information and additional items. Temporary databases move to work files, authorization needs to be permitted on the work files.

DB2 9 deprecated function

- Private protocol → DRDA (new help in DSNTDP2DP)
- Plans containing DBRMs → packages
- ACQUIRE(ALLOCATE) → ACQUIRE(USE)
- XML Extender → XML type
- DB2 MQ XML user-defined functions and stored procedures → XML functions
- DB2 Management Clients feature (DB2 Administration Server, Control Center, & Development Center) → IBM Data Studio application & administration services
- msys for Setup DB2 Customization Center → install panels
- BookManager use for DB2 publications → Info Center, pdf
- Simple table spaces → segmented or partitioned by growth

33



IBM INFORMATION ON DEMAND 2008

INNOVATE.OPTIMIZE.
PERFORM.

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

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

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

DB2 9 migration process changes, APARs

- Enterprise PL/I change JCL error DSNTEP2 PK63659
- Rebind required for V3 and earlier plans & packages
- Best process is rebind all in CM or NFM, top 10% in CM
 - Highly recommended to rebind all V8 and below
- PDSE is required for SDSNLOAD. PDSE service needed.
- Premigration job DSNTIJP9 APAR PK31841
- SPE PK11129, PK33664, PK39850, PK38799, PK44884
- Info APARs: Install II12423, Unicode II13048 & II13049, DFSORT II14047, II14213 & II13495, DDF II14203, zIIP II14219, Performance II13579, Storage II10817, Parallel II12836, LOBs II13767, DSNUTILS sec. id. II13133, Coprocessor II13782, PDSE II14067, II14255, II14403 ...

34



IBM INFORMATION ON DEMAND 2008

INNOVATE.OPTIMIZE.
PERFORM.

All DSNTEP2 programs on any version can be affected by a PL/I run time change, if the JCL does not reflect LRECL=133. This was being overridden by the DCB, but will be an error. A temporary bypass is available in PK63659, so the JCL problems need to be fixed.

Plans and packages must be rebound on DB2 9 to benefit from optimizer enhancements. DB2 V3 and earlier plans & packages will be automatically rebound when accessed by DB2 9. A much better process is to rebind all plans on the new version. If there is not adequate time to rebind in DB2 9 CM, see if you can rebind the most frequently accessed packages. If any plans have not been bound on V8, rebind them.

PDSE is documented as required on V8, but some customers noted that a PDS could be used (as long as the largest load module remained under 16MB). The PDSE is the only option for DB2 9, as there are program objects and larger load modules.

The premigration job DSNTIJPM on DB2 9 is shipped as DSNTIJP9 in APAR PK31841 for V8. Running this job on every subsystem helps find the work needed before migration. The SPE for DB2 9 is PK11129 and PK33664 on V8. Get all the service in this chain.

There are many useful information APARs. These are the ones I use the most. For secondary ids used with DSNUTILS, read II13133.

DB2 9 default and process changes

- Changes to defaults
 - BIND ISOLATION defaults to CS. Query RR
 - Not changed for distributed, REBIND
 - BIND CURRENTDATA defaults to NO. Query Yes.
 - Not changed for REBIND
 - DB2-managed data set secondary extent allocation default YES, even for migration
 - Varying length data changes row format
 - Need to find and complete or delete incomplete definitions
 - Online REORG processes change
 - Configure shared virtual storage (128 GB for each DB2)
 - Work files: more 32K space, avoid small files for DGTT

35



The default value for bind option CURRENTDATA changes from YES to NO for the BIND PLAN and the BIND PACKAGE subcommands, as well as the CREATE TRIGGER for trigger packages, and the CREATE PROCEDURE and the ALTER PROCEDURE ADD VERSION SQL statements for SQL PL procedure packages. CURRENTDATA NO is the best option for performance. The default value for bind option ISOLATION is changed from RR to CS for the BIND PLAN and the remote BIND PACKAGE subcommands. For the BIND PACKAGE subcommand, the current default (plan value) stays. The default change does not apply to distributed. Although you can specify DBPROTOCOL(PRIVATE) for the DBPROTOCOL parameter of the BIND option, DB2 issues a new warning. All BIND statements for plans and packages that are bound during the installation or migration process specify the ISOLATION parameter explicitly, except for routines that do not fetch data. The current settings are maintained for compatibility.

DB2-managed data set secondary extent allocation default setting is YES and is changed by installation CLIST from NO to YES. In migration mode, the input migration value for this field is overridden to YES, so you need to reset it to NO if you want to continue to disable secondary extent optimization.

REORG and LOAD REPLACE utilities automatically convert table spaces to a new reordered row format (RRF) if there is no VALIDPROC or EDITPROC. If you use the REORG SHRLEVEL CHANGE utility, you might need to change your operating procedures if you run concurrent REORG TABLESPACE SHRLEVEL CHANGE PART x on the same table space. Configure 128 GB of shared virtual storage for each DB2 subsystem on the MVS image for memory shared between DB2 DDF and DBM1.

DB2 9 Prerequisites: See Program Directories

- **Processor requirements:**
 - Processor that supports z/Architecture, including IBM System z10, z9, zSeries z990, z890, z900, z800 or later
 - DB2 9 will generally require increased real storage for a workload compared to DB2 V8
- **Base Software Requirements:**
 - z/OS V1.7
 - DFSMS Version 1 Release 7
 - Language Environment Base Services
 - z/OS Version 1 Release 7 Security Server (RACF)
 - IRLM V2.2
 - z/OS Unicode Services & appropriate conversion definitions
 - DB2 Utilities Suite or equivalent function

36



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

These prerequisites apply whether you are migrating or installing. The best source for current information is the Program Directory. Be sure that the products are in support by checking on the web.

This is the first slide of several noting the primary prerequisites for DB2 9. The System z processors are required, with no changes from those needed by V8, although performance is better with the later processors. These are the base needs for DB2 9. The following slides note the optional and functional requirements.

IRLM V2.2 comes with DB2 9. Some of the basic operation of a DBMS is provided by utility functions, such as backup, recovery, reorganization, loading and unloading data, gathering statistics and checking data, indexes and large objects. You should ensure that these functions are provided either by ordering DB2 Utilities Suite for z/OS, Version 9, (5655-N97), or by obtaining equivalent function elsewhere.

DB2 9 Prerequisites: Current software

- Migrate only from DB2 V8 NFM, expand BSDS
- DB2 9 for Linux, UNIX, Windows FP1 or V8.1 FP13 (V8.2 FP6 is the same as V8.1 FP13)
- DB2 9 Connect FP1 or V8.1 FP13
- IMS V10, V9, V8
- CICS V3.1, V2.3
- Enterprise COBOL V3.4 Move to V4 recommended
- Enterprise PL/I V3.4 Move to V3.7 recommended
- IBM Tivoli OMEGAMON XE for DB2 Performance Expert on z/OS, V4 or Performance Monitor, PK36297

37



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

These are the primary changes in levels for DB2 9 over those in DB2 V8. Current products are required, and some of the product levels needed by DB2 V8 are now out of service. The general rule is to stay within the service interval for all products if you want service.

The DB2 for LUW and DB2 Connect levels need fixpaks to function correctly. Fixpak 2 has just delivered, and some customers will need those fixes.

If you are not on current levels of COBOL, move to V3.4, as all others are out of service as of April 2007. If you are not on current levels of PL/I, move to V3.6, as all levels before V3.4 go out of service this year or earlier.

The Omegamon level needed is V4, with APAR PK36297 to support new function. The older V3, DB2 PM and DB2 PE products do not support DB2 9.

DB2 9 Prerequisites: z/OS 1.8 for:

- RACF Access Control Module and roles
- RACF user-mapping plug-in Enterprise Identity Mapping
- Coprocessor with C or C++
- Volume-based BACKUP, RESTORE SYSTEM & RECOVER enhancements

38



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

- RACF SAF User-mapping plug-in Enterprise Identity Mapping (EIM) requires z/OS V1.8 with trusted context. See redbook SG24-6480-01 appendix C for more detail on EIM.

- For full RACF support of roles (not limited to ownership checks for roles only), RACF Access Control Module (DSNXRXAC) requires z/OS V1.8.

- The C and C++ compilers which come with z/OS 1.8 are required to use the DB2 coprocessor, instead of the precompiler for application development.

- See the next page for volume-based BACKUP and RESTORE.

DB2 9 Requirements for Dependent Functions

- Volume based utilities require:
 - z/OS 1.8
 - DFSMSHsm 1.8
 - DFSMSdss 1.8
 - FlashCopy Version 2
- Group Buffer Pool Batching:
 - CF level 15 or 14 recommended when using System Managed Duplexing for lock or SCA structures
 - CFLEVEL 15 structure sizes grow
 - CF level 13 recommended for castout and lock table cleanup enhancements
 - CF level 12 for improved V8 and 9 performance GBP batching

These are highlights only, not a complete list!

See Program Directories.

39



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

z/OS V1.8, DFSMSHsm, DFSMSdss and FlashCopy V2 is required for DB2 9 function in System level Utilities: BACKUP SYSTEM, RESTORE SYSTEM, and RECOVER. z/OS V1.7, DFSMSHsm, DFSMSdss and FlashCopy V2 is needed for System level utilities: BACKUP SYSTEM and RESTORE SYSTEM (if you cannot currently satisfy one or more of the preceding program requirements you can use V8 DB2 functions in the BACKUP SYSTEM and RESTORE SYSTEM utilities).

z/OS V1.7, DFSMSdss (and FlashCopy V2 is highly recommended, although not technically required, for performance reasons) for Utilities: CHECK INDEX, CHECK DATA, and CHECK LOB.

Use the latest Coupling Facility (CF) level for the best performance.

CFLEVEL 15 is the current level and is needed for some of the z/OS 1.9 improvements, but it does cause growth in structure sizes. See WSC FLASH01572 for more.

CF level 14 is recommended when using System Managed Duplexing for lock or SCA structures. Coupling Facility (CF) level 13 is recommended for castout enhancement and lock table clean up enhancement. Coupling Facility (CF) level 12 is required for Group Buffer Pool (GBP) Batching.

DB2 9 Functional Software Requirements

- WebSphere Application Server for z/OS V6.1 for trusted context
- WebSphere MQ for z/OS V6.1
- z/OS XML System Services (z/OS XML) & SDK5
- Java SDK 1.4.2
- DB2 QMF components of DB2 9 for z/OS
- ... details in Program Directories on the web

40



IBM INFORMATION ON DEMAND 2008

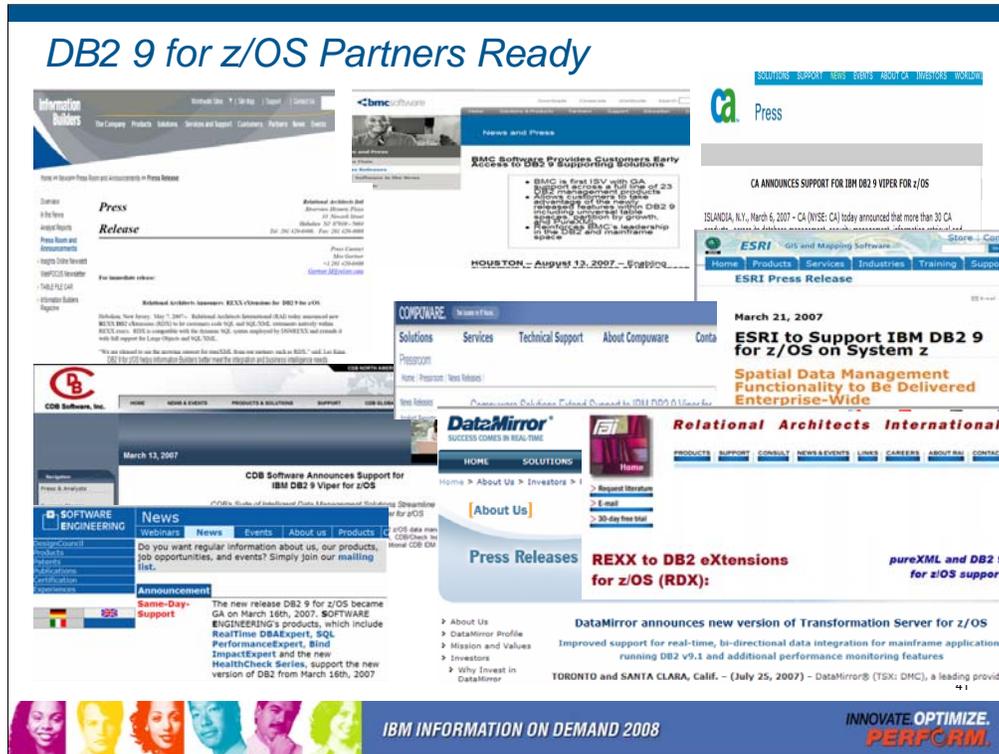
INNOVATE OPTIMIZE.
PERFORM.

If you use WebSphere Application Server for z/OS, trusted contexts require V6.1.

IBM WebSphere MQ V5.3.1 or later is required for MQListener.

XML System Services (XMLSS). XMLSS requires either z/OS V1.8 or z/OS Version V1.7 with APAR OA16303. In addition, use of XML schemas requires 31-bit SDK for z/OS, Java 2 Technology Edition, Version 5 (5655-N98) (SDK5)

IBM SDK for z/OS, Java 2 Technology Edition, Version 1.4 (SDK1.4.2) is required for Application Execution: Applications or stored procedures written in Java, such as those using the JDBC or SQLJ interfaces to DB2.



Here are just a few of the DB2 partners that have announced that they are ready to run with DB2 9 for z/OS. These are images of the web pages from IBI, BMC, CA, ESRI, Compuware, CDB, Software Engineering, SoftBase, DataMirror and Relational Architects. Talk to your vendors about DB2 9.

Of course, if your tool vendor is IBM, you can check on the web for the precise release and PTF levels of the Information Management Tools releases and service needed to work with DB2 9 for z/OS. Note that the matrix lists some levels of tools with an X in the NS or not supported column. See the later version, with an X under T for toleration or under E for exploit or the ability to use those new improvements.

<http://www.ibm.com/support/docview.wss?rs=434&uid=swg21256800>

DB2 Tools Portfolio Day 1 DB2 9 support

<p><u>Application Management</u></p> <ul style="list-style-type: none"> ▪ DB2 Path Checker ▪ DB2 Bind Manager ▪ DB2 Table Editor 	<p><u>Database Administration</u></p> <ul style="list-style-type: none"> ▪ DB2 Administration Tool ▪ DB2 Object Comparison Tool ▪ DB2 Storage Management Utility ▪ DB2 Change Management Expert 	<p><u>Backup and Recovery</u></p> <ul style="list-style-type: none"> ▪ Application Recovery Tool for IMS and DB2 Databases ▪ DB2 Archive Log Accelerator ▪ DB2 Change Accumulation Tool ▪ DB2 Cloning Tool ▪ DB2 Log Analysis Tool ▪ DB2 Object Restore Tool ▪ DB2 Recovery Expert
<p><u>Utilities Management</u></p> <ul style="list-style-type: none"> ▪ DB2 Utilities Suite ▪ DB2 Automation Tool ▪ DB2 Utilities Enhancement Tool ▪ DB2 High Performance Unload 	<p><u>Performance Management</u></p> <ul style="list-style-type: none"> ▪ OMEGAMON XE DB2 Performance Expert ▪ OMEGAMON XE DB2 Performance Monitor ▪ DB2 SQL Performance Analyzer ▪ DB2 Buffer Pool Analyzer ▪ DB2 Optimization Expert ▪ DB2 Query Monitor ▪ DB2 Performance Expert 	<p><u>Data Governance</u></p> <ul style="list-style-type: none"> ▪ IBM Optim Data Growth ▪ IBM Optim Data Privacy ▪ IBM Optim Test Data Management ▪ DB2 Audit Management Expert ▪ IBM Database Encryption Expert ▪ Data Encryption for DB2 and IMS
<p><u>Business Intelligence</u></p> <ul style="list-style-type: none"> ▪ IBM DataQuant ▪ IBM QMF ▪ DB2 Web Query Tool 	<p><u>Information Integration</u></p> <ul style="list-style-type: none"> ▪ WebSphere Classic Data Event Publisher ▪ WebSphere Classic Federation Server ▪ WebSphere Classic Replication Server ▪ WebSphere DataStage ▪ WebSphere Replication Server ▪ DataMirror 	

* Available only on System z
 * Available on System z and LUW
 * Available only on LUW

IBM INFORMATION ON DEMAND 2008

IBM offers a broad portfolio of DB2 Tools for z/OS across numerous functional categories. We continue to fulfill our commitment to provide new tools and enhanced function to meet your requirements. new Tools-- DB2 Recovery Expert for z/OS and DB2 Cloning Tool for z/OS.. We enhanced DB2 SQL Performance Analyzer to include Easy Explain functionality and to integrate with the OMEGAMON XE DB2 performance tools that were announced in 2005 along another new tool, DB2 Thread Expert . In addition we continue to provide new releases for tools to increase the value they bring to your DB2 environment. We also continue to grow our portfolio of multiplatform tools that support DB2 in the Linux UNIX and Windows environments. All DB2 Tools exploit Version 8. IBM tools make the same commitment to support DB2 9. See slides 82 and 83 for the tools web site for the release levels and APARs needed for DB2 9.

Old COBOL?

- ✓ Understand the options DB2 CICS V3
 - ➔ Compiler and Run-Time Migration Guide
 - ➔ No compiles OS/VS COBOL or VS COBOL II CICS V3 also
 - ➔ Run Old COBOL modules
- ✓ Set strategy & plan for Enterprise COBOL V4
 - ➔ Use V8 precompiler?, NEWFUN(NO)?
- ✓ Plan staging and Migrate

43



IBM INFORMATION ON DEMAND 2008

INNOVATE.OPTIMIZE.
PERFORM.

DB2 9 supports current Enterprise COBOL V3.4 and V4. Older compilers are no longer supported. OS/VS COBOL has been out of support since 1994. VS COBOL II support ended in 2001. Only LE runtime is supported. End of support for Enterprise COBOL V3.2 was October 2005, V3.3 was April 2007 so your target should be V4.

<http://www.ibm.com/servers/eserver/zseries/zos/le/history/cobmvs.html>

The key resource is the COBOL Migration Guide, GC27-1409. Here is where to get it from the web:

<http://www.ibm.com/software/awdtools/cobol/zos/library/>

<http://publibfp.boulder.ibm.com/epubs/pdf/igy3mg40.pdf>

Key questions and guidance are in the first chapter. Most customers will stage migration. Do you have Enterprise COBOL V4 or V3.4 now? Is that the release developers use? Old load modules are not a problem. Precompiling is the challenge, so if new development and changes to old programs switch to the new compiler, then that is a possible migration strategy.

DB2 generates COBOL code that we expect to fail on the old compilers, although we have not tested. Programs that use the DB2 9 precompiler or coprocessor require current Enterprise COBOL V3.4 or V4.

<http://www.ibm.com/support/docview.wss?rs=64&uid=swg21166881>

Old PL/I?

- ✓ Understand the options
 - ➔ Compiler and Run-Time Migration Guide
 - ➔ Run Old PL/I modules
- ✓ Set strategy & plan for Enterprise PL/I V3.7
 - ✓ Support requires at least V3.4 now
 - ➔ Use V8 precompiler?, NEWFUN(NO)?
- ✓ Plan staging and Migrate

44



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

DB2 for z/OS supports only current PL/I compilers. Every compiler prior to V3.4 is out of service in 2007, including OS PL/I V2.3, IBM PL/I for MVS™ & VM V1R1.1, V2, and V3R1 to R3.

<http://www.ibm.com/servers/eserver/zseries/zos/le/history/plimvs.html>

The Program Directory lists compilers you can use for **PL/I**:

The key resource is the Enterprise PL/I for z/OS Compiler and Run-Time Migration Guide *Version 3 Release 3*, GC27-1458-02. Get it from the web:

<http://www.ibm.com/software/awdtools/pli/plimvs/support/>

<http://publibfp.boulder.ibm.com/epubs/pdf/ibm3m101.pdf>

There are some basic choices to make. Key questions and guidance are in the first chapter. Most customers will stage migration to new compilers to improve programmer productivity and avoid migration costs. Do you have Enterprise PL/I V3.2 or V3.3 now? Is that the release developers use? Old load modules are not a problem. Precompiling is the challenge, so if new development and changes to old programs switch to the new compiler, this migration strategy improves costs and productivity.

<http://www.ibm.com/support/docview.wss?rs=64&uid=swg21191789>

DFSORT and DB2 for z/OS Utilities Suite

- V8 removed prerequisite for external sort product
- V8 only uses DFSORT: SORT, MERGE functions only
- Sort 32K page & improve reliability & performance
- DFSORT is part of standard z/OS install
- DB2 for z/OS Utilities Suite has license to use DFSORT when invoked by Utilities
- Must have access to DFSORT R14 plus APAR PQ68263 applied. See Info APARs II14047, II14213 & II13495
- DB2 9 requires APAR PQ92594
- z/OS 1.10 DFSORT PARMLIB changes requires PK59399
- Must update DB2 space statistics for sort space

45



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

The requirement for an external sort package was removed. You do not need a DFSORT license for use by DB2 for z/OS Utilities Suite V8. The base DB2 for z/OS has never used a sort product. If you use DFSORT for other purposes, then you do need a DFSORT license. You must update the DB2 space statistics to accurately reflect the data volumes in order for DFSORT to accurately allocate space. This support was delivered by [APAR PQ68263 \(PTF UQ90054\)](#) on RSU 0312. With [APAR PK04076 \(UK03983\)](#), MSGDSNU1640I will be issued when the utility is unable to locate DFSORT code at APAR PQ68263 maintenance level. So DFSORT is mandatory and you have to use it for the utilities. However, there is good news. You already have DFSORT installed. You may not be using it, but it is there. DFSORT is always shipped and installed with z/OS. However, no license is required for the DB2 utilities to use DFSORT. The only action a customer must take is to add the DFSORT load libraries to link list after the OEM sort libraries or add a steplib/joblib to the DB2 utilities batch JCL. DB2 for z/OS Version 8 has a license to use DFSORT so the customer does not need a license for any external sort product. However, no non-DB2 task can use DFSORT without being licensed. Use of DFSORT other than by DB2 would require a license. Maintenance for DFSORT is shipped with z/OS. If a customer has problems with DB2's use of DFSORT, they open an ETR with DB2 Level 2 support, just as they do today for any other DB2 related issue. The DB2 V8 requirement for DFSORT R14 or above is described in the following places: [URL for informational APAR II14047 describing DB2 use of DFSORT](#)

DB2 Installation Guide (All manuals can be found at

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

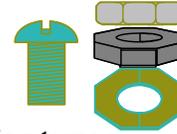
DB2 for z/OS Version 8: Everything You Wanted to Know... and More (SG24-6079)

([Download from IBM RedBook web site](#))

[DB2 v8 Program Directory](#)

<http://www.elink.ibm.com/public/applications/publications/cgibin/pbi.cgi?CTY=US&FNC=SRX&PBL=GI10-8566-03#>

Install Prerequisite Software



- **Functional Requirements for identified features and functions e.g. DB2 Connect, clients**

- ✓ See What's New in DB2?
- ✓ Program Directories for Details

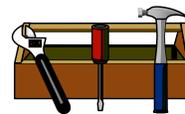
- **Service Level: IBM Consolidated Service Test**



- **Plan for more frequent service**

- ✓ More current releases = more frequent
- ✓ Refresh service level quarterly

- **Prepare to manage performance**



46



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

IBM and other vendor software must be checked for levels required and needed for the function you want. Some APARs require vendor changes. DB2 Connect: See next page and the web.

ibm.com/software/data/db2/db2connect/

Use IBM Consolidated Service Test levels to help provide a well tested set of products at a recommended service level.

ibm.com/servers/eserver/zseries/zos/servicetst/

Customers who are relatively current on service encounter fewer problems overall. Refreshing the service level each quarter makes it easier to avoid encountering old problems.

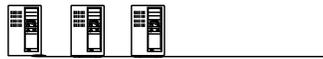
DB2 9 requires DB2 Connect 9 FP1 or V8 FP13
Recommended level is latest fixpacks.

9Performance & Scalability

9Improved Application Availability

9Innovative manageability

9Programmer productivity



V8 News and APARs PKnnnnn

- SOAP UDF enhancements for authentication & security PK48773
- SORTNUM elimination PK45916
- Relief for 254 compressed part limit in LOAD/REORG PK51853
- “V7 COBOL precompiler” precompile with old COBOL PK46170
 - Note: V8 Co-processor does not work with old COBOL
- Ignore KEEPDICTIONARY first time with RRS, NFM PK41156
- BatchPipes input to LOAD for fast loading PK34251
- IDENTITYOVERRIDE keyword on LOAD utility PK27287
- Allow RESTART(LIGHT) to not wait for INDOUBT URs PK29791
- MQListener support for logical message grouping PK51290
- APARs stored procedures for Data Studio, OSC, OE
- Large volumes EAV in z/OS 1.10 PK58291, PK58292, PK61105
- LOAD & REORG performance PK61759
 - 5-10% CPU improvement for LOAD, REORG RELOAD
- SORTBLD with small SECQTY PK60956



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

SOAP UDFs. 4 new SOAP UDFs to allow for a) optional SOAP parms for std security hdr for passing auth info, b) support for HTTPS, c) detailed error messages when UDF is not successful. Long-running readers threshold parameter default. IFCID 313, Ship “v7 COBOL precompiler” with v8. Allow easier migration to v8, allow usage of old COBOL or PL/I compilers

Add new MQ UDFs to support traditional Request/Reply scenarios

MQListener will support logical message grouping (available with MQ v6.0). This is a customer requirement that will allow MQ messages to be processed in a logical order and will prevent application issues arising from processing messages in a undefined order

allow ALTER SET DATA TYPE to increase a varchar length in CM

SORTNUM elimination. Difficult to determine SORTNUM (number of sort data sets), V8 migration consideration for some customers who were using non IBM sort. Enabled by 2 Zparms:

UTSORTAL to enable predictable sort space allocation, IGNSORTN to ignore SORTNUM

A limit was placed on LOAD/REORG to process a max of 254 parts on a compressed table space. This limit was put in place to prevent very large amounts of real storage consumption which can cause an LPAR wide outage. A new ZPARM is being added to raise this limit. By default DB2 will enforce the 254 part limit but customers with sufficient real can exploit the new zparm

Add support for BATCHPIPES SUBSYS to TEMPLATE U

Allow RESTART(LIGHT) to not wait for INDOUBT UR

IDENTITYOVERRIDE keyword to LOAD utility. LOAD utility does not allow data unloaded from a table with identity columns with GENERATED ALWAYS to be reloaded and preserve the original identity column values. Additionally, a LOAD REPLACE with an empty input data set (used as a fast way to empty a table space) against a table space that has identity columns defined with GENERATED ALWAYS requires each table and column to be specified which is error-prone.

OSC Fixpack 4 is available on the web, providing support for DB2 V8.

<http://www.ibm.com/software/data/db2/zos/downloads/osc.html>

OE V1.2 was announced on March 4, GA on March 14th. OE V1.2 Fixpack 1 provides coexistence mode and fixes. <http://www.ibm.com/software/data/db2imstools/db2tools/opti-expert-zos/>

DB2 9 Post GA Deliveries APAR PKxxxxx

- Text search server Accessories Suite 1.2
- New XMLTABLE and XMLCAST functions (next page) PK51573
- Incremental FlashCopy PK41001, z/OS 1.8 APAR OA17314
- Trusted context enhancements PK44617, PK47579
- New storage class Zparm for online CHECK utilities PK41711
(Needed when PPRC is used)
- ALTER TABLE ALTER COLUMN DROP DEFAULT PK56392
- Allow RESTORE SYSTEM without log truncation PK51979
- Spatial phase II PK51020, PK54451, ...
- BIND package stability changes PK52522, PK52523
- BIND change to convert plans with DBRMs to packages PK62876
- LOAD COPYDICTIONARY PK63324, PK63325
- Protocol(3) without data sharing group quiesce PK62027⁴⁹



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

DB2 9 has been generally available for over a year now, and some new function has delivered in the service stream. Delivery in APARs is used when the need is urgent and the function can be delivered with a very small risk of disruption to existing customers. The delivery criteria become stronger as more customers move to the version. A text search server delivered in Accessories Suite 1.2 at the end of 2007. XML performance improvements and new XMLTABLE and XMLCAST functions are very important for early XML applications and consistency. Incremental FlashCopy makes use of the new z/OS capability. Trusted context enhancements are needed by early users. A new storage class parameter is needed when PPRC is used and is provided for online CHECK utilities.

ALTER TABLE ALTER COLUMN DROP DEFAULT extends the capability to ALTER a default. DB2 allows RESTORE SYSTEM recovery without requiring log truncation. Many new spatial functions deliver compatibility with the DB2 for LUW function. Changes come in DB2 to work with new fixpacks of OSC, OE and Data Studio. BIND stability changes deliver the capability to keep two or three versions of packages and the ability to switch from the current access paths to the older ones. This function can help you with regression from a REBIND.

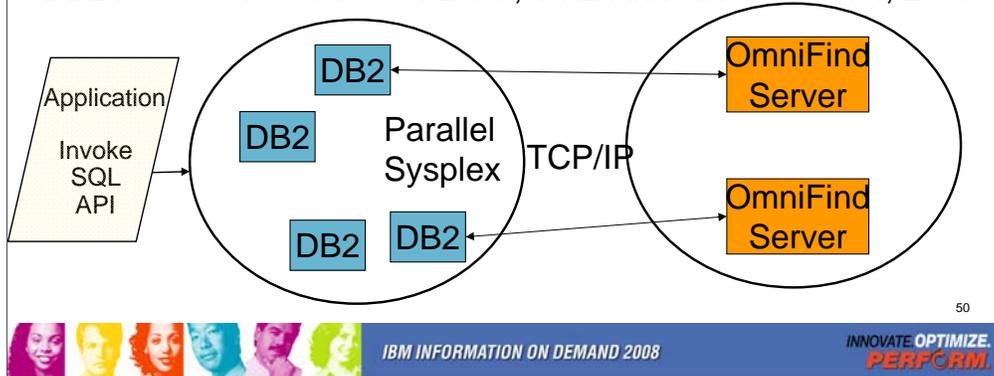
Plans containing DBRMs, rather than using packages, are deprecated. That is to say we expect plans containing DBRMs to disappear some time in the future. To help with migration from DBRMs to packages, we expect to deliver a change to BIND which can perform the conversion to packages.

The newest capability is to copy a dictionary from one partition to another using the LOAD utility. Watch for APAR PK62027 to avoid the need to quiesce the data sharing group to get the new LOB lock avoidance.

OmniFind Text Search Support

- CONTAINS() built in function for text search
- Search CHAR, VARCHAR, LOB, & XML columns
- OmniFind provides a text index server
- Efficient communication interactions with DB2 for z/OS
- OmniFind text indexes are persisted into DB2 tables for backup/recovery purposes

DB2 Accessories Suite for z/OS, V1.2 Announced Dec. 4, 2007



We have added a new built-in function, CONTAINS() to allow better text searching. This is the basic architecture of our solution. A specialized text search engine on a separate server is expected to come in DB2 9, but will deliver after general availability. While the server is separate, the text indexes are saved into DB2 tables to provide improved backup and recovery and consistency with DB2 data. See Peggy Zagelow's IOD session 1696 for the details.

For the announcement, see DB2 Accessories Suite for z/OS, V1.2, part of Announcement 207-319, dated December 4, 2007.

http://www.ibm.com/common/ssi/rep_ca/9/897/ENUS207-319/ENUS207319.PDF

More XML support in DB2 9 in APARs XMLTABLE: return XML in tabular format

```

SELECT X.* FROM dept,
XMLTABLE ('$d/dept/employee' passing deptdoc as "d")
COLUMNS
  "empID"      INTEGER PATH '@id',
  "firstname"  VARCHAR(30)  PATH 'name/first',
  "lastname"   VARCHAR(30)  PATH 'name/last',
  "office"     INTEGER PATH 'office') AS "X"
  
```

XML Input:

```

<dept bldg=101>
  <employee id=901>
    <name>
      <first>John</first>
      <last>Doe</last>
    </name>
    <office>344</office>
  </employee>
  <employee id=902>
    <name>
      <first>Peter</first>
      <last>Pan</last>
    </name>
    <office>216</office>
  </employee>
</dept>
  
```

empID	firstname	lastname	office
901	John	Doe	344
902	Peter	Pan	216

51



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

XML is having strong acceptance among early customers, but we are seeing many needed changes. We are expecting to see several come in DB2 9 as APARs. Check marks indicate that PTFs are delivered.

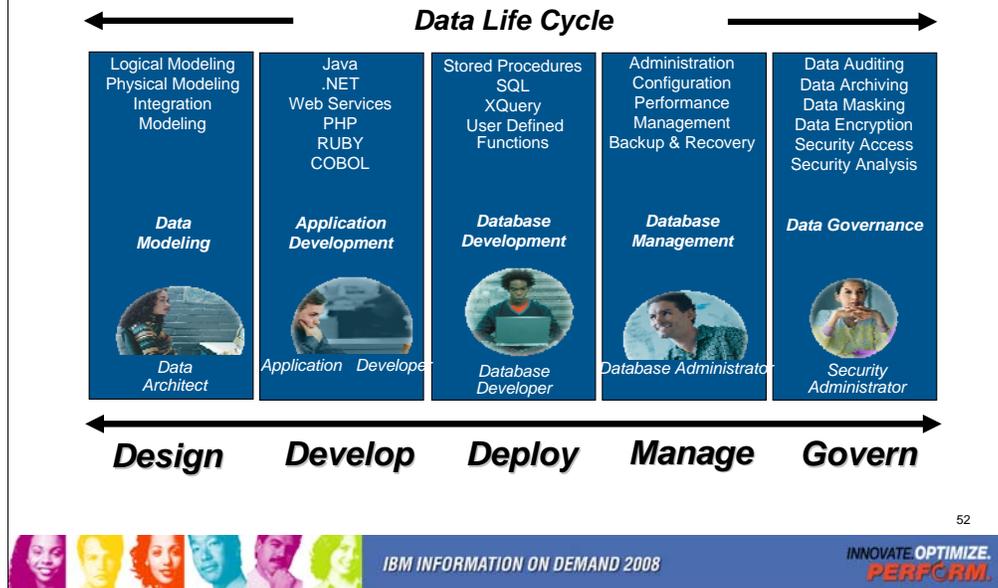
- ✓PK57786 XML query performance XMLAGG
- ✓PK55585 XPath function, phase 1
- ✓PK55831 XPath function, phase 2
- ✓PK51573 XMLTABLE, XMLCAST
- PK56337 XPATH performance
- PK57409 XMLTABLE enhancement

The ability to return data from XML in the form of a table is one. The ability to cast explicitly is another. Some additional XPath functions are a third.

- PK58914 XMLTABLE performance enhancement

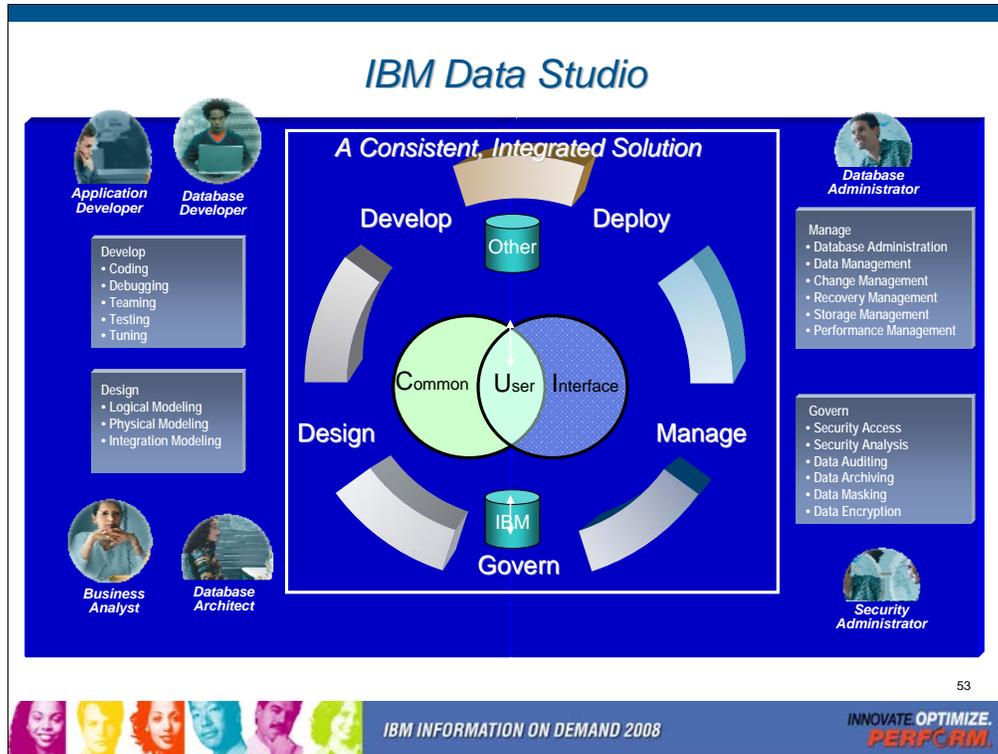
Some customers and vendors are using XML extensively and we want them to be aware of the recent and upcoming changes.

Install IBM Data Studio for many reasons



IBM Data Studio is for designers, developers, and administrators who work across the entire data life cycle and the entire IBM relational DBMS portfolio: DB2 for LUW, DB2 for z/OS, DB2 for i5/OS and IDS. Almost every customer will need the tools included with DB2, and many will need the additional ones. See the IBM Data Studio web page for all the changes in IBM Data Studio. Watch closely, as this area is changing fast.

<http://www.ibm.com/software/data/studio/>



IBM Data Studio is for designers, developers, and administrators who work across the entire data life cycle and across the full set of IBM relational DBMS: DB2 for LUW, DB2 for z/OS, DB2 for i5/OS and IDS. Almost every customer will need the tools included with DB2, and many will need the additional ones..

DB2 is changing to deliver the needed interfaces, usually new stored procedures.

See the IBM Data Studio web page for all the changes in IBM Data Studio. Watch closely, as this area is changing fast.

<http://ibm.com/software/data/studio>

IBM Data Studio v1.1

- Empowering developers and database administrators
- Complimentary and available in October of 2007
- Support for DB2 on all platforms and IDS

DB2 for LUW	DB2 for z/OS	DB2 for i5/OS	IDS
<ul style="list-style-type: none"> • Physical Data Modeling • Data Distribution Viewer • Integrated Query Editor • SQL Builder • SQL Routine Debugger • Java Routine Debugger • XML Editor • XML Schema Editor • pureQuery for Java • Data Web Services • Object Management • Data Management • Update Statistics • Health Monitoring * • Visual Explain • Security Access Controls • Project Management <p>* Technical Preview</p>	<ul style="list-style-type: none"> • Physical Data Modeling • Data Distribution Viewer • Integrated Query Editor • SQL Builder • SQL Routine Debugger • Java Routine Debugger • XML Editor • XML Schema Editor • pureQuery for Java • Data Web Services • Object Management • Data Management • Update Statistics • Visual Explain • Security Access Controls • Project Management 	<ul style="list-style-type: none"> • Physical Data Modeling • Data Distribution Viewer • Integrated Query Editor • SQL Builder • SQL Routine Debugger • Java Routine Debugger • XML Editor • XML Schema Editor • pureQuery for Java • Data Web Services • Object Management • Data Management • Security Access Controls • Project Management 	<ul style="list-style-type: none"> • Physical Data Modeling • Data Distribution Viewer • Integrated Query Editor • SQL Builder • XML Editor • XML Schema Editor • pureQuery for Java • Data Web Services • Object Management • Data Management • Security Access Controls • Project Management

54



IBM INFORMATION ON DEMAND 2008

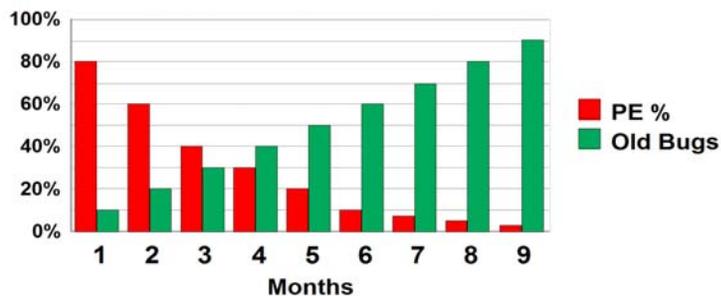
INNOVATE OPTIMIZE.
PERFORM.

IBM Data Studio is for designers, developers, and administrators who work across the entire data life cycle and across the full set of IBM relational DBMS: DB2 for LUW, DB2 for z/OS, DB2 for i5/OS and IDS. Almost every customer will need the tools included with DB2, and many will need the additional ones. See the IBM Data Studio web page for all the changes in IBM Data Studio. Watch closely, as this area is changing fast.

<http://www.ibm.com/software/data/studio/>

How current should we be?

- Stay reasonably current with DB2 releases & fixes
- Balance for severities. Use Consolidated Service Test.
 - ▶ Potential PTF in Error or PE: Enhanced HoldData
 - ▶ Problems encountered instead of avoided
 - ▶ Business constraints and work load



55



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

Choosing service timing requires judgement, not simple rules. The best practice for service means an understanding of your business drivers and constraints. These scales are not the same. Timing for moving preventive service into production is a careful balance between potential problems caused by PTFs containing an error (PEs) & potential problems avoided by applying the maintenance sooner. About 70% of PEs are found within four months after the PTF is closed. Currently, about 6 PEs are found every month across all versions. Green bars reflect the percentage of PEs yet to be found. Systems programming work load & availability of windows for installing service can affect timing.

I recommend upgrading service levels two or three times a year. Use current CST recommended service level, all service plus a two to four month later level for hipers & PE fixes. Stage that level through development systems & production. Match your levels with CST, across operating system & key subsystems. Enhanced HoldData provides a much faster way to discover missing hipers and PEs without fixes.

<http://www.ibm.com/servers/eserver/zseries/zos/servicetst/mission.html>

<http://service.boulder.ibm.com/390holddata/390holddata.html>

Service Best Practice

Use Consolidated Service Test
Recommended Service
Cross product tested levels
Best practice for service:
two or three per year

<http://www.ibm.com/servers/eserver/zseries/zos/servicetst/mission.html>

Consolidated Service Test and the RSU

The Consolidated Service Test Mission

Our mission is to provide a consolidated, tested, and recommended set of service for z/OS and key subsystems. In a quarterly, RSU with published results all IBM maps that we have tested your current service levels. We have a service package with a tested level of service for all these products.

How we do our testing

We test in a customer-like production sysplex environment in an IBM test lab and are continuously improving our CST test environment. We run batch and data-sharing applications that exploit and stress the latest functions with up to two levels of subsystems on three levels of z/OS systems. As we test, we observe how each product runs and how they interact in this environment. We report problems to the IBM Support Center. We have the capability to generate reports on systems. For more information about our environment, see our [Consolidated Service Test](#).

We have a quarterly, CST, test cycle!

- 1st month: - Install service from prior quarter, - Run existing workload, - Plan new test cases
- 2nd month
- 3rd month
- 4th month: - Provide recommended service package, - Publish quarterly report, - Run in high stress environment, - Run recovery tests

56



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

Consolidated Service Test provides a way to leverage more IBM testing and experience from many customers across a much broader set of products, including z/OS, IMS, CICS, DB2, WebSphere and many tools. This technique is working for very large, very small and in between customers. If you don't have an extensive test suite - like most customers, then this technique can help a lot in reducing research and finding a stable service level.

Effective June 15, 2007, the DB2 for z/OS Version 8 SUP tape is available world-wide for new customer orders. This SUP tape was built differently than prior SUPs, such that its build in May integrated PTFs COR-closed through December 2006, which had also completed a Consolidated Service Test (CST) cycle. Once PTFs complete a CST cycle, they are assigned a Recommended Service Upgrade (RSU) status. Our DB2 for z/OS Version 8 SUP tape contains CST tested PTFs which were marked "RSU 0703" (they completed CST testing in March 2007). This SUP build integrated a total of 964 PTFs (the delta since our December 2005 SUP), which is nearly 1,000 fewer PTFs our new customers need to deal with during their DB2 V8 installation experience. For additional information on CST and RSU, please see: <http://www.ibm.com/servers/eserver/zseries/zos/servicetst/mission.html>

Installing two or three CST levels a year means more current service, but not too current. If you want a process that has proven track record for success and can reduce your work, then look at the CST. Use current CST recommended service level: all service plus a two to four month later level for hipers & PE fixes. Stage that level through development systems & production. Match your levels with CST, across operating system & key subsystems. Enhanced HoldData provides a much faster way to discover missing hipers and PEs without fixes.

Premigration Activity ... Manage Performance



Prepare to manage performance

- Save accounting reports
- Save critical access paths
 - ▶ Be sure you have **PLAN_TABLE** access paths & **DSN_STATEMNT_TABLE** & **DSN_FUNCTION_TABLE** for critical applications (or Visual Explain reports).
 - ▶ Run **EXPLAIN** on queries before migrating.
 - ▶ Add new columns to **PLAN_TABLE**
- Differentiate by plans & packages

57



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

Save critical access paths & accounting data! Changes can cause unwanted access path changes. Identify important queries, plans & packages. Be sure plan tables contain access paths & costs. ALTER current plan tables to add new columns. EXPLAIN REBIND may change access paths, so extract plans & run EXPLAIN under a dummy collection, a different application or program name.

Keep accounting reports for crucial queries & applications. If you have a problem & send in long accounting reports & explain data, we find the issue quickly. If you don't have data, then we guess.

Monitoring & managing are more difficult with few plans. Never ending threads & plan reuse accumulate EDMPOOL & working storage. Each COMMIT needs package accounting. Recycle threads. Finding unused packages & cleaning up mistakes are difficult unless ownership is known. Performance monitoring may require analyzing everything in a CICS region instead of just one plan. Plan is a required short name. Package name is optional, long, multipart and not always easy to find. Best practice is to differentiate by plan within a CICS region, having at least 5 to 20 plans within a region. You can work faster & make changes on a more granular basis.

Practice Three Rs of DB2

- **REORG**
 - ▶ Real time statistics
 - ▶ Performance, space
- **RUNSTATS**
 - ▶ After significant change
- **REBIND**
- **Recheck**
 - ▶ Explain and run time performance
 - ▶ **MODIFY, COPY**



58



IBM INFORMATION ON DEMAND 2008

INNOVATE.OPTIMIZE.
PERFORM.

This practice starts with the basics, but performing the three Rs effectively can avoid a wide range of problems. In other cases, they are the steps needed to resolve a problem. These steps are needed to keep clustering effective, maintain free space, avoid index splits, keep sequential prefetch performance high, ... Using real time statistics to drive the process is being used more and more.

After the REORG is completed or during the process, gather the updated statistics. Balance the cost of gathering the information with the improved information for optimization. Improved statistics can mean improved access paths – ones that were not available before.

Using the improved information means using dynamic SQL or rebinding. After the rebind, checking the explains and the run time performance for regression can avoid problems. Do you want to rebind only in NFM or do you rebind some packages in CM too. How long will you be in CM? The longer the period, the more you will want or need to rebind in CM.

To REBIND or not to REBIND

- **WHEN? is the question**
 - ▶ When to take the chance of access path regression
 - ▶ For the potential benefits of improved performance & memory
- **When migrating? Test Before**
 - ▶ Parallel individual binds
 - ▶ More disruptive, longer
- **Incrementally or later**
 - ▶ More potential for problem calls
 - ▶ Resolution could be more difficult



59



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

To rebind or not to rebind, “When?” is the question: Whether 'tis nobler in the mind to suffer the slings and arrows of outrageous access paths, Or to take arms against a sea of troubles, And by rebinding, end them? To die: to sleep; No more; and by a sleep to say we end the heart-ache and the thousand natural shocks that old plans are heir to. 'tis a consummation devoutly to be wished To die, to sleep. To sleep: perchance to dream: ay, there's the rub; For in that sleep of death what dreams may come when plans have shuffled off this mortal coil, must give us pause: There's the respect that makes calamity of so long life for old plans and packages; [Apologies to the Bard and to Hamlet 3.1.]

The question is not whether to REBIND, but rather when to REBIND. Planned rebinds avoid surprises. A fix may need to rebind everything. Performance gains in a new version often occur at REBIND. When you rebind, performance is

1. Same: Most common case is plus or minus 5%, which is usually ignored
2. Worse: Least common, but usual regression causes are new access path or increased parallelism. If you compare to the prior access plan, then you can find most problems easily. If you wait until a BIND occurs, analysis is more difficult; you are less likely to have all of the performance & plan data. Use a tool or compare access paths to find problems.
3. Better: Many of the V8 performance improvements only occur if you bind. Some SQL is likely to benefit from optimization improvements, such as index use, joins, subquery transformations or UNION improvements. When you find a problem - errors or performance - the fix is often to REBIND. Best practice for migration is to test binds before migration & bind at migration time (comparing before and after). Use a tool for comparison and the task is simpler, but simple comparisons of cost estimates catch many problems.

<http://www.ibm.com/support/docview.wss?rs=64&&uid=swg21078662>

Plan to REBIND packages

Old plans and packages have more to gain, less to lose
BIND changes coming for Bind stability, to eliminate private
protocol and DBRMs bound into plans

CM – how long? Very little new value

- ▶ Most optimizations included
- ▶ Rebind key packages at least

New function mode

- ▶ Adjust database design, applications
 - ▶ Improved index options: compression, larger page sizes,
on expression, New SQL
 - ▶ Native SQL procedures, more zIIP
 - ▶ Rebind after changes
- ▶ Rebind all plans and packages once CM or NFM (may stage)



60



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

Best practice for migration is to test binds before migration & bind at migration time (comparing before and after). Use a tool for comparison and the task is simpler, but simple comparisons of access path differences catch most problems. There are a number of good tools on the market. For the best chance of improvements with the lowest risk of regression, rebind the oldest plans and packages. If there was a problem that resulted in a plan or package error, then the only way to resolve the issue is to rebind (or free and bind). Use a dummy collection for the test bind and save the old package to revert back if needed. Since most of the improvements in optimization are already functional in CM, most customers bind the key packages and plans during conversion mode. If CM is a very short period, or you do not need the improvements, then you might wait for new function mode to rebind. If you stay in CM for a long time, then you will want to rebind more of the packages and plans.

In new function mode, all of the optimization improvements can be made. Often you will want to look at performance challenges for your customer and see how they can be resolved. such as compression, index on expression and larger pages. After making the design changes, REORG the data and REORG or REBUILD the indexes, get the improved statistics and REBIND. Native SQL procedures, added use of zIIP, and improved SQL continue the improvements in this phase.

<http://www.ibm.com/support/docview.wss?rs=64&uid=swg21078662>

Optimization

Addressing

SQL – static or dynamic - performance regression for new BIND or REBIND
Who?
Administrators, performance
Impacts: CPU, elapsed time, productivity
Substantial impact when access path is bad

Solution

Aspects of success:
Large reduction in regression
If regression occurs, easy switch to old access paths
Reduce, almost eliminate effort
Save CPU, IO, elapsed time

Less fear of BIND, REBIND
Allow optimizations to be more aggressive, automation, parallelism,
Easier migrations to new versions, ...

61



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

Unexpected changes in access paths lead to performance regressions

Regressions typically seen during migration to a new release

Also seen during the lifetime of a release after maintenance fixes are applied

changes in data and/or schema: growth, need for REORG, getting current statistics, ...

changes in the applications require REBIND or BIND REPLACE

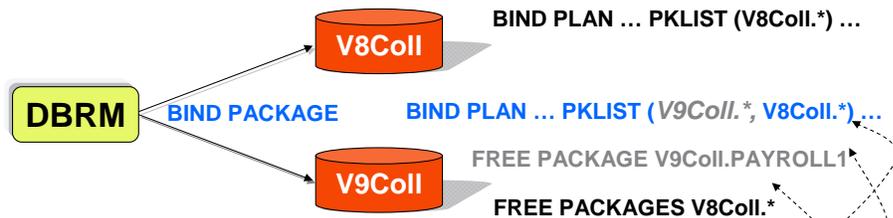
Both Static and Dynamic SQL affected

Static SQL regressions seen after BINDs and REBINDs

Dynamic SQL even more exposed

Package Stability – Using a dummy collection

- REBIND packages into a new collection ID



- BIND PLAN with new collection ID added
- Free regressed packages from the new collection
- Free packages from the old collection when new packages demonstrate expected performance

62



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

This part is a technique that can be used today with a dummy collection used to REBIND packages for comparisons.

Package or Bind Stability APAR PK52523

- Old access paths for a package retained at REBIND
REBIND PACKAGE (HRCOLL.SALPKG)
PLANMGMT(BASIC | EXTENDED)
 - BASIC retains 2 copies: Current and Previous
 - EXTENDED retains 3 copies: Current, Previous, Original
 - PLANMGMT new subsystem parameter
 - Can provide default for command options
- Fallback to an older access path
REBIND PACKAGE (HRCOLL.SALPKG)
SWITCH(PREVIOUS | ORIGINAL)
 - Conceptually, this can do a “ROLLBACK” of a prior REBIND

63



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

These APARs (PK52522 for preconditioning and PK52523) support the following function:

1. A new REBIND PACKAGE option called PLANMGMT that can be used to control whether REBIND PACKAGE saves old package copies. This option has three possible settings - OFF, BASIC and EXTENDED.

PLANMGMT(OFF) - No change to existing behavior. A package continues to have one active copy

PLANMGMT(BASIC) - A package has one active copy, and one additional old copy (PREVIOUS) is preserved.

PLANMGMT(EXTENDED) - A package has one active copy, and two additional old copies (PREVIOUS and ORIGINAL) are preserved.

Preserving old copies has no impact on query performance. However, using the PLANMGMT(BASIC) option can double the disk consumption in table space SPT01, and using the PLANMGMT(EXTENDED) option can triple it. The extra space is needed to maintain old copies. In addition, using the BASIC or EXTENDED option adds a 10-40% CPU overhead to the performance

Package or BIND Stability in DB2 9 PK52523

- At REBIND, save old copies of packages
 - Catalog tables
 - Directory (SPT01 table space)
- Two flavors
 - BASIC and EXTENDED
 - REBIND options
 - Default new DSNZPARM
- REBIND PACKAGE ...
 - PLANMGMT (BASIC) 2 copies: Current & Previous
 - PLANMGMT (EXTENDED) 3: Current, Previous, & Original
- Most bind options can be changed at REBIND.
- REBIND PACKAGE ...
 - SWITCH(PREVIOUS)
Switch between current & previous
 - SWITCH(ORIGINAL)
Switch between current & original
- FREE PACKAGE ...
 - SCOPE(ALL) – Free package completely
 - SCOPE(INACTIVE) – Free old copies
- Each copy invalidated separately
- Cost of stability
 - 2-3 X size in SPT01
 - 30% CPU overhead for REBIND
 - No impact on statement runtime processing, fast switch

64



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

1. PK52523 has the main function. PK52522 is the pre-conditioning APAR. Both of these have PTFs now. The DB2 books and the online Information Center document this functionality so customers can already read about it online.
2. The old copies are primarily saved in the SPT01 table space and it will grow in size. For BASIC, the overhead is 2x. For extended, it is 3x.
3. We have a few solutions to mitigate this disk space increase:
 - (a) Users could use FREE PACKAGE ... PLANMGMTSCOPE to selectively free old copies of packages that they think are stable
 - (b) In DB2 X, we expect to implement an 'access path comparison' function that will not save an old copy if it's the same as the new copy.

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, *
- APAR PK62876 is open, so watch for changes

65



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

PK62876 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 pklst 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 pklst will be deleted, while the new pklst converted from the DBRMS will be bound into plan x.

How much room will you get in DBM1?

It depends. For DB2 9,

Estimate 10% to 15%

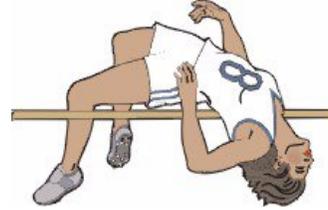
5 EDM pools

REBIND needed for savings

DDF - shared memory

Above the bar shared virtual storage required

You need to calculate, monitor and manage



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

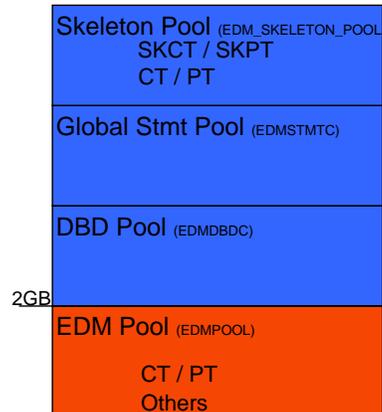
66

You still need to do the calculations, monitor and manage storage. DB2 9 improvements are generally about 10% to 15% more storage moved out from below the bar, with two more parts of the EDMPOOL, some of dynamic statement cache, and use of 64 bit storage for the DDF address space, with some above-the-bar storage shared with the DBM1 address space. You need to configure the 128 GB virtual shared area above the 2 GB bar. EDM virtual storage savings only occur if you rebind packages. Note that plans containing DBRMs are deprecated. Watch for some improvements in migration techniques for plans containing DBRMs.

DB2 9 64 bit Evolution (Virtual Storage Relief)

Virtual Storage Constraint is still an important issue for many DB2 customers.

- EDMPOOL Changes:
 - V8 – DBD storage moved above 2GB bar.
 - DB2 9 – SKCT, SKPT, some CT, PT storage moved above 2 GB bar.
 - DB2 9 approx. 60% reduction in EDMPOOL size observed for lab workloads
- Other changes:
 - Some storage acquired for distributed applications moved above 2GB bar.
 - Control blocks for table spaces and RTS move above the bar.
 - DSC statement text moved above the bar
 - SAP tests have shown almost 300 MB reduction in virtual storage below 2 GB bar



IFCID 217: detailed DBM1 virtual storage health
IFCID 225: consolidated DBM1 virtual storage health

67



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

Objective was 10-15% relief. EDMPOOL – you can estimate DB2 9 EDM pool size from V8 stats as follows: $(\#pgs \text{ for SKCT/SKPT}) * 0 + (\#pgs \text{ PT/CT}) * 70\%$. Rough ROT is V8 EDM pool size can be reduced by 60%. Extensive use of SQL SPs can drive up EDM pool usage since these pkgs are larger. General recommendation, keep this the same in DB2 9. EDM_SKELETON_POOL – added in DB2 9.

For SAP SD transaction with 360 user threads, the total DBM1 virtual storage below 2GB was 1091MB in V8 and 819MB in DB2 9 for almost 300MB reduction. Most of that comes from local dynamic statement cache storage going down from 466MB to 172MB. Net of 1% real storage increase overall. For TPCE for example, 532 pages below 2GB were used for PT in V8. In DB2 9, 763 pages are used below 2GB and 316 pages are used above 2GB. So there is an increase in below 2GB virtual storage usage for PT in DB2 9. 1268 pages used for SKPT in V8 are all moved above 2GB in DB2 9 with a total of 1329 pages. EDM pool pages used below 2GB therefore dropped from 532+1268 to 763, or 58%.

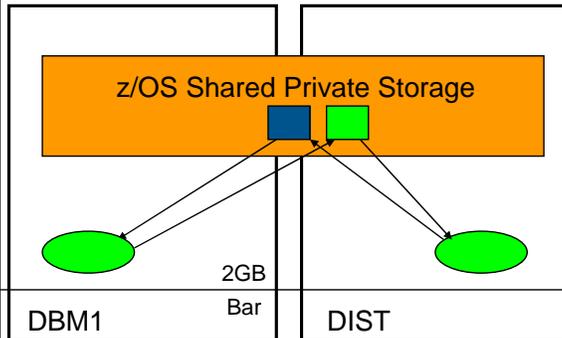
EDM LRU latch contention relief (LC24). Previously one LRU chain in EDM pool with only one corresponding latch. Now split into three pools each with 3 latches: 1) SKCT/SKPT above 2GB 2) CT/PT above 2GB 3) CT/PT below 2GB

DBM1, the following are moved above the bar in DB2 9

- Parse trees peak below-the-bar storage for full prepare reduced 10%
- EDM fixed pools: V8 customer dumps show as much as 50m will be moved. Allows larger above the bar EDM pools
- SKPTs / SKCTs (primarily static SQL). Also part of the CTs/PTs; New EDM pool for skeletons; Savings in below the bar 10m to 300m
- Page set blocks, RTS blocks: up to 10's of MB savings
- Local SQL statement cache: rough ROT: about 60% moves above bar
- Thread-related storage: Certain Run Time Structures, space block DMTR 10's of MB or more in savings
- Average V8 customer may see another 200 MB of relief: RID Pool, Sort Pool, compression dictionaries, castout buffers and others already above the bar (V8).

64-bit DDF – Shared Private with DBM1

- DDF address space runs in 64-bit addressing mode
 - Shared 64-bit memory object avoids cross memory moves between DBM1 and DDF and improves performance



- Shared memory: new virtual storage type allowing multiple address spaces to share storage.
- Similar to ECSA – always addressable, avoids AR and XM moves.
- Different from ECSA – only available to those address spaces registering with z/OS to share this storage.
- Reduces data formatting and data movement
- Reduces virtual storage
 - It exists once, instead of in each address space

68



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

Skeletons above the bar: For customers that use heavy package and plan activity such as banks, this is the most significant DBM1 below the bar storage relief in DB2 9. For customers that use very few or small packages, such as SAP environments, the savings is smaller.

LI702 – move spaceblk (SPA) above the bar. SPA to be split into 2, 1 above, 1 below. Only a few, non-complex RTs are being considered for DB2 9. Simple insert, delete Expected results will vary by SQL mix. (-5 to 30%)?

Has storage key & fetch protection

Defaults to 2TB size

DB2 requires a minimum of 128GB configured or DB2 9 will not run

Even if not running DIST
Set by HVSHARE in Parmlib

Migration - data sharing group

- ▶ V8 and 9 CM members in group for short time (coexistence)
- ▶ **MUST** apply fallback SPE to all members first
- ▶ First member migrated must be successful before migrating other members
 - Run CATMAINT only for the first member
- ▶ Use Data Sharing book for migration
- ▶ Migrate all members before using new function
- ▶ Restrict execution of packages & plans bound on new version to migrated members
- ~~▶ Group quiesce in DB2 9 NFM for locking protocol 3 (LOBs)~~
- ▶ CFLEVEL 15, 14, 13, 12 for best performance



69



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

Remember to work with new DB2 9 function only on a group that has migrated to 9 NFM. Therefore, it is best to migrate all of your members and test with existing applications before implementing any new features or functions.

If you have multiple releases in a group for more than a week, then you need to have much better controls in place and be able to test your processes more thoroughly.

Persistent structure size changes, group attach improvements, restart light and enhanced constraint management are more complex if multiple releases are involved.

The minimum CFLEVEL is to have a level that is in service. Most data sharing customers will want to be at CFLEVEL 15, 14, 13 or 12. There is quite a bit of information about CFLEVEL on the web:

<http://www.ibm.com/servers/eserver/zseries/psocftable.html>

The need for a group quiesce is removed in APAR PQ62027.

Get the detailed steps from the Installation Guide & Data Sharing: Planning & Administration.

Customer Experiences



- **Customer experiences say pay attention to:**
 - **Read the web, books, Info APARs, PSP**
 - **Have a plan that includes performance**
 - **Run premigration job DSNTIJPM**
 - **Use V8 DSNTIJP9 job**
 - **On every subsystem**
 - **Current maintenance for DB2, z/OS & CF**
 - **Get the latest books: library, red books, ...**



70



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

Most of these were mentioned before. They are mentioned again because other customers missed them. Many customer problems dealt with items on the foil titled Premigration Action. Run the premigration job. Check CATMAINT messages.

Programmers must add colons for all host variables. Make sure DBAs understand the new authorizations and meaning of schema names. Check for release dependent objects on fallback or coexistence.

Having current service & books saves time & avoids problems.

Check other early customer experiences here.

<ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/migration/>

<ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/customer-experience/>

Step 2: Deliver the new value

- Reduce CPU time
- Save disk space
- Service Oriented Architecture
- Productivity for applications & administrators
 - PureXML & Data Studio
 - SQL & LOBs
 - Portability via SQL & data definition compatibility
- Improve availability
- Make security more flexible
- Ease regulatory compliance
- Web applications
- Data warehouse function and performance



This is the largest and most important phase of migration. You will spend more time in this phase than in all of the others combined. This is the phase where you deliver the largest improvements to your customers.

What are your most strategic and most pressing needs in the areas of business, applications and administration? That question and the efforts needed should determine the schedule and order for using the new function. Some of the simpler changes for performance are noted on slide 9.

Most consumable DB2 9 improvements

- CM very little to no action:
 - Changed online REORG
 - Improved RUNSTATS
 - Optimization improvements, EDMPOOL VSCR
 - Increased parallelism, use of zIIP

NFM

- LOB lock avoidance
- Reordered row format
- Index: larger page sizes, compression, index on expression



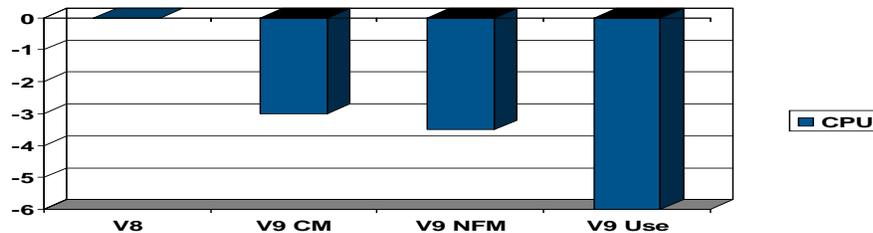
72



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

DB2 9 z10, z9, z890 & z990 performance plan scenario



Utilities	DB design adjustments
Histogram statistics	Index improvements
REBIND	application changes
DSNZPARMS	native SQL procedures
	SQL adjustments

Your situation will vary. Less CPU is better.
z800 and z900 expect +5% to +10% CPU

73



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

If you have a z9, z990 or z890, this is expected to be a common shape for a DB2 9 performance plan, starting with zero for the V8 baseline. When you first move to DB2 9, total DB2 CPU time generally decreases from 0% to 5% for z9, z890 and z990 customers, shown here as a first step -3%. Utility CPU reductions help immediately. Some work will be about the same (+/-3%). Start with reorgs and collect improved histogram statistics when useful. The DB2 9 CM performance plan REBINDs the primary packages and adjusts DSNZPARMs. The REBINDs provide most of the improved access paths. On z800 or z900 the initial cpu expectation is +5 to +10% regression, more if there are many columns, so making adjustments is more important. In moving to NFM, some additional DSNZPARMs are adjusted and all plans and packages are rebound. The DB2 9 use line takes wider advantage of DB2 9 performance improvements. Database designs start taking advantage of new indexing options, such as compression, index on expression and larger pages. After making the design changes, REORG the data and REORG or REBUILD the indexes, get the improved statistics and REBIND. Native SQL procedures, added use of zIIP, and improved SQL continue the improvements in this phase.

Scenario: Customer mix of DB2 CPU time is 30% in utilities, 70% in SQL access. With 10% improvement for the utilities, we get a -3% net, assuming that SQL is the same as before. With optimization improvements, another -½% improvement shows up in DB2 9 NFM. Then as design adjustments, reorgs and rebinds are performed, we get improvements from varchar improvements, native SQL procedures and improved SQL, another -3%.

What works in Conversion Mode?

- Logging improvements, improved index page split
- DSN1COMP index compression estimate, NOT index compression
- Larger prefetch, write and preformat quantities
- LOB performance other than LOB lock avoidance
- Command enhancements: REFRESH, DISPLAY, START, STOP, & MODIFY TRACE qualifications and ACCESS DATABASE
- Virtual storage use above the 2 GB bar DDF & EDMPOOL (when bound)
- Package BIND stability (PLANMGMT & EXCHANGE)
- Optimization improvements (when bound)
- Increased parallelism and use of zIIP
- Optimization Service Center & Optimization Expert (work with V8 if PTFs are applied)
- Most utility enhancements and CPU reductions (see next slide)

74



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

New functions are available only in new-function mode, unless explicitly stated otherwise in the product documentation. A few general exceptions exist for utilities and for optimization. The following specific improvements are available in conversion mode:

- Logging improvements
- Improved index page split
- Larger prefetch, write and preformat quantities
- LOB performance other than LOB lock avoidance
- Virtual storage use above the 2 GB bar DDF & EDMPOOL (when bound)
- Package BIND stability (PLANMGMT & EXCHANGE)
- Optimization improvements (when bound)
- Optimization Service Center & Optimization Expert (work with V8 if PTFs are applied)
- Utility CPU reductions and most other utility improvements that do not depend on SQL changes. See the following slide and web for detail.
- <http://www.ibm.com/support/docview.wss?rs=64&uid=swg21307783>

DB2 9 Conversion Mode versus New Function Mode

Utility enhancements available in both CM and NFM are:

- CPU reductions
- Parallel unload, reload and log apply processing by the REORG utility
- Removal of the BUILD2 phase for REORG at the partition level
- Log Phase addition for REORG SHRLEVEL REFERENCE PART
- REORG SHRLEVEL REFERENCE for LOB table spaces
- CHECK INDEX SHRLEVEL REFERENCE with Parallelism
- Histogram Statistics for RUNSTATS
- CHECK LOB SHRLEVEL CHANGE
- CHECK DATA SHRLEVEL CHANGE
- REPAIR LOCATE SHRLEVEL CHANGE
- REBUILD INDEX SHRLEVEL CHANGE
- MODIFY RECOVERY support for the RETAIN option
- Object level recovery from a system-level backup by the RECOVER Utility
- BACKUP SYSTEM & RESTORE SYSTEM system-level backups to tape
- Skipping locked rows for transaction updates by the UNLOAD Utility
- TEMPLATE switching support based on size for image copy datasets

Other utility enhancements are only available for DB2 9 NFM, such as...

- LBI (large Block Interface) support for sequential datasets on tape
- Recover to Point in Time (PIT) with consistency (defaults as new behavior)
- Clones effectively provide LOAD REPLACE SHRLEVEL CHANGE
- Online Load (LOAD RESUME YES SHRLEVEL CHANGE) without logging using NOT LOGGED

75



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

DB2 9 has many utility enhancements available in both Conversion Mode (CM, previously called Compatibility Mode) and New Function Mode (NFM), Some exceptions are only available for NFM, and not for CM.

Utility enhancements available in both CM and NFM are:

- Parallel unload, reload and log apply processing by the REORG utility
- Removal of the BUILD2 phase for REORG at the partition level
- Log Phase addition for REORG SHRLEVEL REFERENCE PART
- REORG SHRLEVEL REFERENCE for LOB table spaces
- CHECK INDEX SHRLEVEL REFERENCE with Parallelism
- Histogram Statistics for RUNSTATS
- CHECK LOB SHRLEVEL CHANGE
- CHECK DATA SHRLEVEL CHANGE
- REPAIR LOCATE SHRLEVEL CHANGE
- REBUILD INDEX SHRLEVEL CHANGE
- MODIFY RECOVERY support for the RETAIN option
- Object level recovery from a system-level backup by the RECOVER Utility
- BACKUP SYSTEM & RESTORE SYSTEM system-level backups to tape
- Skipping locked rows for transaction updates by the UNLOAD Utility
- TEMPLATE switching support based on size for image copy datasets

Other Utility enhancements are only available for DB2 9 NFM, such as...

- LBI (large Block Interface) support for sequential datasets on tape
- Recover to Point in Time (PIT) with consistency (defaults as new behavior)
- Clones effectively provide LOAD REPLACE SHRLEVEL CHANGE
- Online Load (LOAD RESUME YES SHRLEVEL CHANGE) without logging can be accomplished by altering the table space to NOT LOGGED, running the Online Load, altering the table space back to LOGGED, and taking an image copy.

Package or BIND Stability in DB2 9 PK52523

- At REBIND, save old copies of packages
 - Catalog tables
 - Directory (SPT01 table space)
- Two flavors
 - BASIC and EXTENDED
 - REBIND options
 - Default new DSNZPARM
- REBIND PACKAGE ...
 - PLANMGMT (BASIC) 2 copies: Current & Previous
 - PLANMGMT (EXTENDED) 3: Current, Previous, & Original
- Most bind options can be changed at REBIND.
- REBIND PACKAGE ...
 - SWITCH(PREVIOUS)
Switch between current & previous
 - SWITCH(ORIGINAL)
Switch between current & original
- FREE PACKAGE ...
 - SCOPE(ALL) – Free package completely
 - SCOPE(INACTIVE) – Free old copies
- Each copy invalidated separately
- Cost of stability
 - 2-3 X size in SPT01
 - 30% CPU overhead for REBIND
 - No impact on statement runtime processing, fast switch

76



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

1. PK52523 has the main function. PK52522 is the pre-conditioning APAR. Both of these have PTFs now. The DB2 books and the online Information Center document this functionality so customers can already read about it online.
2. The old copies are primarily saved in the SPT01 table space and it will grow in size. For BASIC, the overhead is 2x. For extended, it is 3x.
3. We have a few solutions to mitigate this disk space increase:
 - (a) Users could use FREE PACKAGE ... PLANMGMTSCOPE to selectively free old copies of packages that they think are stable
 - (b) In DB2 X, we expect to implement an 'access path comparison' function that will not save an old copy if it's the same as the new copy.

To REBIND or not to REBIND

- **WHEN? is the question**
 - ▶ When to take the chance of access path regression
 - ▶ For the potential benefits of improved performance & memory
- **When migrating? Test Before**
 - ▶ Parallel individual binds
 - ▶ More disruptive, longer
- **Incrementally or later**
 - ▶ More potential for problem calls
 - ▶ Resolution could be more difficult



77



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

To rebind or not to rebind, “When?” is the question: Whether 'tis nobler in the mind to suffer the slings and arrows of outrageous access paths, Or to take arms against a sea of troubles, And by rebinding, end them? To die: to sleep; No more; and by a sleep to say we end the heart-ache and the thousand natural shocks that old plans are heir to. 'tis a consummation devoutly to be wished To die, to sleep. To sleep: perchance to dream: ay, there's the rub; For in that sleep of death what dreams may come when plans have shuffled off this mortal coil, must give us pause: There's the respect that makes calamity of so long life for old plans and packages; [Apologies to the Bard and to Hamlet 3.1.]

The question is not whether to REBIND, but rather when to REBIND. Planned rebinds avoid surprises. A fix may need to rebind everything. Performance gains in a new version often occur at REBIND. When you rebind, performance is

1. Same: Most common case is plus or minus 5%, which is usually ignored
2. Worse: Least common, but usual regression causes are new access path or increased parallelism. If you compare to the prior access plan, then you can find most problems easily. If you wait until a BIND occurs, analysis is more difficult; you are less likely to have all of the performance & plan data. Use a tool or compare access paths to find problems.
3. Better: Many of the V8 performance improvements only occur if you bind. Some SQL is likely to benefit from optimization improvements, such as index use, joins, subquery transformations or UNION improvements. When you find a problem - errors or performance - the fix is often to REBIND. Best practice for migration is to test binds before migration & bind at migration time (comparing before and after). Use a tool for comparison and the task is simpler, but simple comparisons of cost estimates catch many problems.

<http://www.ibm.com/support/docview.wss?rs=64&&uid=swg21078662>

Reduce CPU time and save disk space

- Native SQL procedures
- DB design adjustments
 - Index improvements
 - Universal table spaces – new table space
- application changes – mostly in new applications
 - SQL
 - XML
- See the selected performance topics presentation

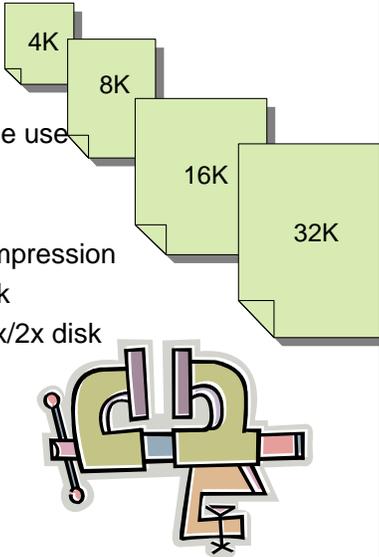


Many of the performance improvements in DB2 9 do not require changes or need only small changes. Review the items on slide 72 and check off the ones that are already done.

Then we'll look at the next level of improvements in this section for saving CPU time and disk space. During the session on selected performance topics, we'll address more.

Indexing Enhancements

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



79

IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE. PERFORM.

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

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

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

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

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

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

Data Studio pureQuery

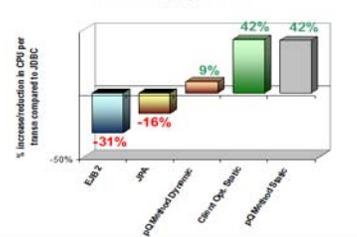
- **Reduce time to market for new Java applications**
 - Bridge Java and data
 - Balance productivity and control
- **Improve quality of service for new and existing Java applications**
 - Improve performance
 - Lock in access plans
 - Increase accuracy of capacity planning
- **Get more value from existing investments**
 - Increase system throughput
 - Speed up problem resolution
- **Enhance security**
 - Limit user access – principle of least privilege
 - Minimize SQL injection risk
 - Improve audit readiness



"IBM Data Studio enables us to bridge the gap between object-oriented design and relational database technology. As a result, we can speed the development of high quality applications and improve developer productivity by between 25 and 50 percent"

In-house tests show double-digit CPU reduction

% increase/reduction in CPU per transaction compared to JDBC using Type 2 driver



<http://www.ibmdatabasemag.com/story/showArticle.jhtml?articleID=208802229>



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

One key technology that we've introduced over the last year and the center of the Data Studio portfolio is pureQuery. pureQuery is about.

- productivity for the developer without sacrificing control over the SQL
- giving you a higher level of quality of service (whether you wrote the code or it was generated by a framework) where you get better performance and and better stability by locking-down access paths.
- better utilization of your existing resources, both system and staff
- better security minimizing privileges and minimizing opportunity for attacks

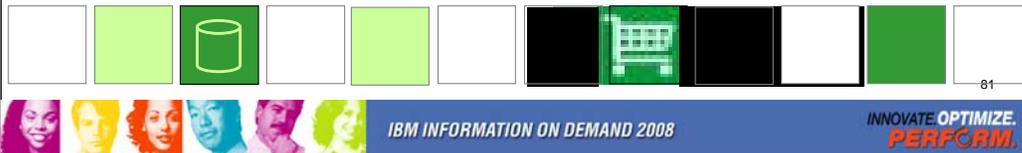
Univar and other customers corroborate the productivity advantages.

Customers and IBM have measured the performance advantages of static SQL many times (as shown here). And now we've made it much easier to get those benefits from your Java environment.

<http://www.ibmdatabasemag.com/story/showArticle.jhtml?articleID=208802229>

Online Schema Evolution → Database Definition On Demand

- Online reorganization with no BUILD2 phase
- Fast replacement of one table with another
- Table space that can add partitions, for growth
- Ability to rebuild an index online
- Rename column and index
- Modify early code without requiring an IPL
- Alter table space and index logging
- Create & alter STOGROUP SMS constructs
- Alter column set default



One of the key initiatives of V8 was online schema evolution. Online schema evolution is expanding and changing to be data definition on demand.

Online table space reorganization for a few partitions is improved a lot, removing the BUILD2 phase for all types of secondary indexes. 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. Table space and index logging can be altered. A new ability to change the DB2 early code does not require an IPL. SMS constructs MGMTCLAS, DATACLASS and STORCLAS can be defined on a STOGROUP and can be altered.

Availability: DB2 9 Utilities Highlights

- Online REORG BUILD2 phase elimination
- Rebuild Index SHRLEVEL CHANGE
- Reorg LOB now supports SHRLEVEL REFERENCE (space reclamation)
- Check data, LOB and repair locate ... SHRLEVEL CHANGE
- Check index SHRLEVEL REFERENCE supports parallel for > 1 index
- Recover to point on log with consistency
- Improved BACKUP SYSTEM / RESTORE SYSTEM
- Clones for “online LOAD REPLACE”
- UNLOAD SKIP LOCKED DATA option

82



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

If you look at all the utility offerings,

Universal Table Spaces

- **Universal Table Space**
 - Combination of segmented with partitioning options
 - Better space management
 - Support of mass deletes / TRUNCATE
 - If partitioned
 - Must be one table per table space
 - Can choose **Range Based partitioning**
 - Can choose **Partitioned By Growth (PBG)**
 - DROP / CREATE to migrate existing table spaces
 - Simple table spaces can not be created
 - Default table space is now Segmented

83



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

We need a partitioning table space with some of the segmented table space features.
The advantages of segmented space maps for partitioned table spaces ?

Universal Table space (Partition By Range) A partitioned segmented table space.
Partitioning column required. One table per table space.

Universal Table space (Partition By Growth) A partitioned segmented table space.
No partitioning column required. One table per table space.

Partition By Growth (PBG)

Single-table table space, where each partition contains a segmented page set
(allows segmented to increase from 64GB to 16TB or 128 TB with 32K pages)

Eliminates need to define partitioning key and assign key ranges

Partitions are added on demand. A new partition is created when a given partition
reaches DSSIZE. DSSIZE defaults to 64G. Up to MAXPARTITIONS

Retains benefits of Utilities and SQL parallelism optimizations for partitioned tables
SEGSIZE defaults to 4 & LOCKSIZE defaults to ROW

Considerations: Single-table table space Always defines as LARGE

Need PBR for query partition elimination

No LOAD PART, ALTER ADD PART, or ROTATE PART All indexes are
NPSIs

Automatic repositioning of Variable columns to end of row

Length attributes replaced with Indicators positioned after fixed length columns

Any table space created in DB2 9 NFM

To Convert: REORG or LOAD REPLACE a table space or partition ADD PARTITION

No EDITPROCs or VALIDPROCs

PIT RECOVER will set the table space to the row format of the PIT

Catalog / Directory remains in Basic Row Format (BRF)

CLONE Tables

- Allows fast replacing production data without renames and rebinds
 - An alternative to online load replace
- ALTER TABLE to create a Clone Table
 - All indexes, LOB and XML objects are also cloned
 - Structures ONLY – not data
 - Base and Clone tables share the same table space and index names
 - Underlying data sets are differentiated by a data set instance number
- On single-table table spaces (partitioned or non-partitioned)
- Use insert or load to populate clone tables
- Utilities (except RUNSTATS) can operate on clone tables with new CLONE keyword
- EXCHANGE DATA switches logical names with underlying data

84



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

Fast replacement of a table with another, also known as clone tables:

This function will allow you to generate a copy of a current table with the same attributes and same data, in the same table space. It appears to an application or user as a very fast replacement of data within a table. Web-based applications, striving for maximum availability will benefit from the option of implementing a pair of tables that are clones of each others' structure. Copies for application testing and auditing can be easily created. These *clone* tables will have the unique ability to change names quickly. Applications can quickly and almost transparently switch between dual mirror tables.

On single-table table spaces (partitioned or non-partitioned)

Use insert or load to populate clone tables

Utilities (except RUNSTATS) can operate on clone tables with a new CLONE keyword

Use EXCHANGE to switch logical names with underlying data

FLIGHTS now references the data that was in
FLIGHTS_CLONE

Indicated in SYSTABLESPACE .CLONE

Improve flexibility for security & ease regulatory compliance

- DB2 9 provides more flexible authorization, control and audit
 - Database Roles
 - Trusted Context
 - Audit: Improved audit selectivity improves security performance, problem isolation and performance monitoring
- Extend encryption to tape controller (now) and disk storage (future direction)



Common Criteria security certification: DB2 for z/OS V8 evaluated at EAL3+ level.

85



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

DB2 for z/OS V8 provides many enhancements for security, DB2 9 builds on that base and improves control, flexibility and ability to audit. DB2 9 roles are used to provide a more flexible technique than groups or users in assigning and controlling authorization, while improving consistency with the industry. A network trusted context provides a technique to work with other environments more easily, improving flexibility. The instead of trigger is an SQL technique that allows a trigger to be used in place of a view update, consistent with DB2 for LUW. Improved audit selectivity is needed for being able to see that security is functioning. Secure Socket Layer or SSL implementation provides encryption of data on the wire. Some additional techniques for data encryption will help protect data at rest and in backups – now for tapes and in the future for disks.

Common Criteria: z/OS V1.7 with the RACF optional feature has achieved EAL4+ for Controlled Access Protection Profile (CAPP) and Labeled Security Protection Profile (LSPP). DB2 for z/OS Version 8 is evaluated under the Common Criteria for CAPP and LSPP with a conformance claim of EAL3+, earning the service mark above.

See: <http://www.ibm.com/systems/z/security/>

<http://www.ibm.com/systems/z/security/mls.html>

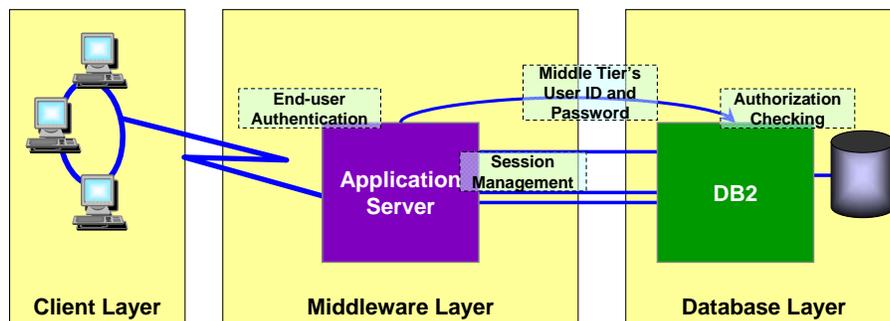
http://www.ibm.com/systems/z/security/ccs_certification.html

http://www.ibm.com/security/standards/st_evaluations.shtml

<http://www.ibm.com/software/tivoli/governance/action/08162007.html>

<http://www.ibm.com/systems/storage/product/showcase/index.html>

Current Authentication in a Multi-tier Architecture



- Three-tiered application model with DB2 as the database server:
 - Middle layer authenticates users running client applications
 - It also manages interactions with the database server
 - Middle layer's user ID & password used for database authentication
 - Privileges of associated authorization id are checked when accessing the database, including all access on behalf of all end-users

86



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

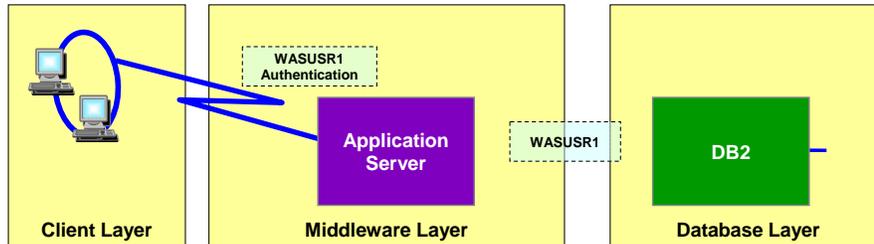
In a three-tier architecture model all interactions with the database server occur under middle-tier authorization ID. This results in loss of end-user identity, over granting of privileges to the middle-tier authorization ID, and weakened security. Hence, need better access control from application servers with end to end auditing of users.

When a user creates an object, the user becomes the owner. So, when the user leaves the company, in order to remove privileges from the user the object has to be dropped and recreated.

In some customer shops, a generic TSO id is created and granted SYSADM authority. This is done in order to avoid the cascading effect of revoking SYSADM when individual team members leave the organization. Unfortunately the use of the generic TSO user id does not provide the individual accountability that is now required for Sarbanes-Oxley (SOX) compliance. A function similar to shared DBADM/SYSADM is needed with auditing capability.

Also, DBADM can create view/alias for another ID, but has no DROP/ALTER privilege

Trusted Authentication in a Three-tier Architecture



- Application server's user ID and password are used to establish the trusted connection.
- User is switched in the trusted connection and client user ID is propagated to the server
- Client authorization ID's privileges are checked for database access

87



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

Trusted Contexts

- **TRUSTED CONTEXT** establishes a trusted relationship between DB2 and an external entity such as a middleware server. For example:
 - WebSphere Application Server
 - Lotus Domino
 - SAP NetWeaver
 - PeopleSoft V7
- A set of *trust attributes* is evaluated to determine if a specific context is to be trusted
- A trusted context allows the external entity to use a database connection under a different user ID without the database server authenticating that ID
- It also allows an AUTHID to acquire database privileges associated with that trusted context, and not available outside it, via a **ROLE**

WebSphere software

Lotus software

SAP

PEOPLE
Soft

88



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

A **TRUSTED CONTEXT** establishes a trusted relationship between DB2 and an external entity such as a middleware server. For example:

WebSphere Application Server

Lotus Domino

SAP NetWeaver

PeopleSoft V7

A set of *trust attributes* is evaluated to determine if a specific context is to be trusted

A trusted context allows the external entity to use a database connection under a different user ID without the database server authenticating that ID

It also allows an AUTHID to acquire database privileges associated with that trusted context, and not available outside it, via a **ROLE**.

Secure DBA Activities

- Security administrator can control use of special DBA privileges by:
 - a. Revoking DBA privileges from individual IDs
 - b. Granting special privileges to a DBA role
 - c. Creating trusted context and assign the DBA role to the DBA IDs
 - d. Start DB2 traces for audit (always on)
- When a DBA needs to perform a database change, the security administrator then
 - Enables trusted context to allow access to sensitive objects
 - DBA can now connect and performs the database change
 - Disables trusted context to protect sensitive objects
- An auditor can review audit trace to ensure compliance

89



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

Using trusted context and role, DBA privileges can easily be disconnected and reconnected to individual employees. This provides function similar to shared SYSADM or DBADM user ids, but avoids the audit compliance problems associated with shared user ids.

Role has the ability to “own” DB2 objects, so that revoking a person’s ROLE does not cause the objects to be cascade deleted.

With these capabilities, customers are able to create DBA procedures that can be audited and protected so that one individual cannot violate the established rules without being detected during the audit review.

Allow DBA to run as another USER

For example, a DBADM who created view for other IDs can DROP / ALTER view owned by the other IDs

```
CREATE TRUSTED CONTEXT CTXLOCAL
  BASED UPON CONNECTION USING SYSTEM AUTHID PRODDBA1
  ATTRIBUTES (JOBNAME 'DBAJOB*')
  WITH USE FOR PRODOWNR
  ENABLE;
```

```
//DBAJOBA JOB USER='PRODDBA1'
//IKJEFT1B EXEC PGM=IKJEFT1B
//SYSTSPRT DD SYSOUT=*
//SYSPRINT DD SYSOUT=*
//SYSTSIN DD *
  DSN SYSTEM(DB1P) ASUSER(PRODOWNR)
  END
//SYSIN DD *
  ALTER VIEW PROVIEW REGENERATE;
  COMMIT ;
//
```



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

DBADM under certain circumstances can create view for others. However, this same DBADM is not able to drop or grant privileges on these views to others. Since views do not have an underlying database, DBADM can not be recognized for views.

A trusted context can allow DBADM or any other permitted user to assume the identity of another user like view owner and then perform the desired actions.

To accomplish view maintenance on behalf of other user, create a trusted context at local DB2 with DBADM id as SYSTEM AUTHID and switch to view owner id.

SQL/XML – work in both worlds

- Full power of SQL to address structured fields
- Join XML documents and tables
- SQL and XML Predicates
- Create XML from structured fields
- Materialize tables from XML documents



```
select d.deptID , u.headcount, xmlquery('$deptdoc/dept/name' passing d.deptdoc as
"deptdoc")
  from dept d, unit u
  where d.deptID = u.unitID and u.headcount > 200
     and xmlexists('$deptdoc/dept[@bldg = $b]' passing d.deptdoc as "deptdoc",
     u.bldg as "b")
     and xmlexists('$deptdoc/dept/employee/name' passing d.deptdoc as "deptdoc")
```

91



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

Now let's discuss the wide range of new function provided inside the DB2 for z/OS engine. This work is being done in parallel with similar changes in DB2 for Linux, Unix and Windows.

DB2 for z/OS extensibility was initially implemented via *extenders*. The XML extender provides for the storage and indexing of an XML document as a character large object (CLOB), or for the shredding of the XML document into relational columns for storage and query. DB2 Version 8 expanded on XML support by implementing several XML publishing operations as built-in DB2 functions. This allows you to perform XML document composition from relational data with improved performance without the XML Extender.

XML in DB2 9 integrates much more function into the engine. This includes an XML data type, native storage of XML documents, integration of the XPath language, and catalog extensions to support definitions of XML schemas. Utilities support creation and maintenance of XML data. Here are some examples in an article on the web:

<http://www.ibmdatabasemag.com/dbadmin/showArticle.jhtml?articleID=207801115>

DB2 9 – Summary of pureXML Support

- XML as a native data type
- Pure XML storage and indexing
- SQL/XML and XPath support
- Integration with traditional relational data
- XML Schema Repository
- Schema validation
- Application Support (Java, C/C++, .NET, PHP, COBOL, PL/1 etc.)
- Visual Tooling, Control Center Enhancements
- DB2 Utilities: Load, Unload, Reorg, etc.
- ...and more

DB2 9

Secure and Resilient Infrastructure for a New Breed of Agile Applications

92



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

Until DB2 9, managing XML data records with a relational data servers meant decomposing the data into columns - a process known as shredding. Or by storing the entire data record in a single cell as a character large object - known as a CLOB. The CLOB approach does not cost overhead as the data records go in. But when you query these records you pay the overhead of parsing each one at runtime which can be a significant performance impact to the application. With shredding, overhead is paid up front to turn the data into a relational record that can be queried efficiently. But overhead is also paid later if the record needs to be recreated for delivery in XML format. This process also affects the fidelity of the record itself - leading to an approach that uses both shredding and CLOB methods for applications that require both performance and fidelity. This results in even more overhead to ensure the records remain in sync.

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

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

DB2 SQL

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 for SQL, join across encoding schemes, IS NOT DISTINCT FROM, Session variables, range partitioning

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

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

93



IBM INFORMATION ON DEMAND 2008

INNOVATE.OPTIMIZE.
PERFORM.

This text just shows the relationship of DB2 for Linux, Unix & Windows with DB2 for z/OS, comparing 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 probably 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.

New cross-platform SQL Reference on web now

DB2 SQL

z z/OS 9

common

luw Linux, Unix & Windows 9



z	{	Multi-row INSERT, FETCH & multi-row cursor UPDATE, Dynamic Scrollable Cursors, GET DIAGNOSTICS, Enhanced UNICODE for SQL, join across encoding schemes, IS NOT DISTINCT FROM, Session variables, TRUNCATE, DECIMAL FLOAT, VARBINARY, optimistic locking, FETCH CONTINUE, ROLE, MERGE, SELECT from MERGE
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, range partitioning, compression
l	{	Updateable UNION in Views, GROUPING SETS, ROLLUP, CUBE, 16 Built-in Functions, SET CURRENT ISOLATION, multi-site join, MERGE, MDC, XQuery
u		
w		

94



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

This text just shows the 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 moves about half of the LUW unique items into the common set and adds a little more that is unique to the z platform. At about this time we'll also have a new release of DB2 9 for LUW, code named Viper. 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.

DB2 Family SQL Reference



**New cross-
platform SQL
Reference
March 28, 2008**

IBM DB2

SQL Reference for Cross-Platform Development
z/OS i5/OS AIX HP-UX Solaris Linux Windows

SQL Reference for DB2 family

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

Version 3 DB2 9 for Linux, UNIX, Windows & z/OS
DB2 for iSeries V5R4

95



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

Do you want to design or write applications for the entire DB2 family, rather than for just one of the platforms? If so, you need the IBM DB2 Universal Database SQL Reference for Cross-Platform Development.

Version 3 covers these product versions:

- DB2 9 for Linux, UNIX & Windows

- DB2 9 for z/OS

- DB2 for iSeries V5 R4

SQL consistency across the DB2 family has improved substantially in the past few versions, while significant new common function has been added. SQL that is common to the DB2 UDB relational database products and the SQL 1999 Core standard is much more comprehensive. The 934 page book describes the rules and limits for preparing portable programs for these versions.

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

DB2 SQL

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



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

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, range partitions, data compression, **session variables**, **DECIMAL FLOAT**, **optimistic locking**, **ROLE**

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

96



IBM INFORMATION ON DEMAND 2008

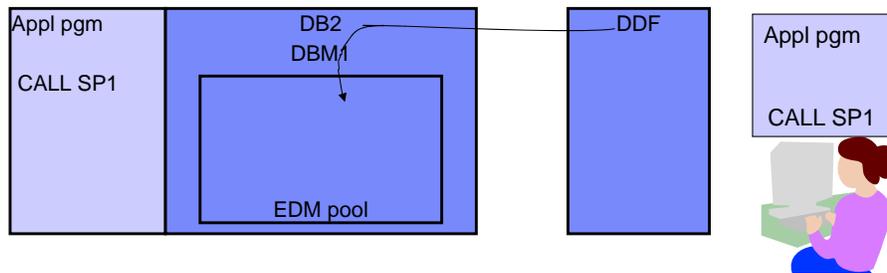
INNOVATE.OPTIMIZE.
PERFORM.

This text just shows the newest delivered 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 moves 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, code named Viper 2 is already generally available. We are able to move more from the unique z list to the common list with DB2 9 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.

Native SQL Procedural Language

- Eliminates generated C code and compilation
- Fully integrated into the DB2 engine
 - An SQL procedure created without FENCED or EXTERNAL is a native SQL procedure



97



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

LOB Improvements

- Progressive Streaming for LOB Locator Values
 - DB2 uses LOB size to determine whether to send LOB data to Java or DB2 CLI clients in one go (<32KB), in chunks (<1MB) or as LOB locator (>=1MB)
 - Transparent to application using LOB locators
- FETCH CONTINUE
 - Allows applications to retrieve LOB/XML data in pieces without the use of locators
- File reference variables
 - A file reference variable allows direct transfer of LOB data between DB2 and the file named in the variable
- Utility Changes
 - LOAD / Cross load LOB column lengths > 32KB supported
 - Logging for > 1GB LOBs
 - REORG LOB reclaim space
 - Online CHECK LOB and DATA
- Elimination of LOB locks for improved availability and performance



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

98

For Java and DB2 CLI programs that use locators with LOBs. Improves performance and less network traffic for LOBs that are less than 1MB. Default behavior if using DB2 9 for z/OS. Requires DB2 Connect 9.1 FP 1. No changes required to programs using locator values. DB2 Client and Type-4 driver manage progressive streaming of data to program. DB2 for z/OS determines whether to flow LOB values or Locators to client based on size thresholds for JDBC, SQLJ, and CLI.

For small LOBs, (Default <= 32KB) the performance should approximate that of retrieving a VARCHAR column of comparable size.

Medium size LOBs (Defaults > 32KB and <= 1MB)

For large LOBs (Default over 1MB) locators are still used.

Specific FETCH that contains LOB or XML columns

Used with programs that materialize LOBs

Application uses a buffer that might not be large enough to hold the entire LOB or XML value.

If any of the fetched LOB or XML columns do not fit, DB2 returns information about truncated columns and the actual length.

Retrieve LOB or XML data in multiple pieces without use of locators

Must specify WITH CONTINUE on initial FETCH

Subsequent fetches use FETCH CURRENT CONTINUE

Application must manage buffers & reassemble data

Not required to fetch entire object before moving to next

SQLCA indicates whether data is truncated

LOAD / Cross load LOB column lengths > 32KB supported

Logging for > 1GB LOBs

REORG LOB reclaim space

SHRLEVEL(REFERENCE)

Allows LOG NO

SHRLEVEL(NONE) is still an option.

Online CHECK LOB and DATA

Elimination of LOB locks

Now using LRSN & page latching for consistency checks

Prior to DB2 9, LOB locks were held until commit

Even for UR

Space search for LOB allocation

No LOB locks acquired for space search

Read LSN granularity improved to page level in LOB table space

Improved availability & performance

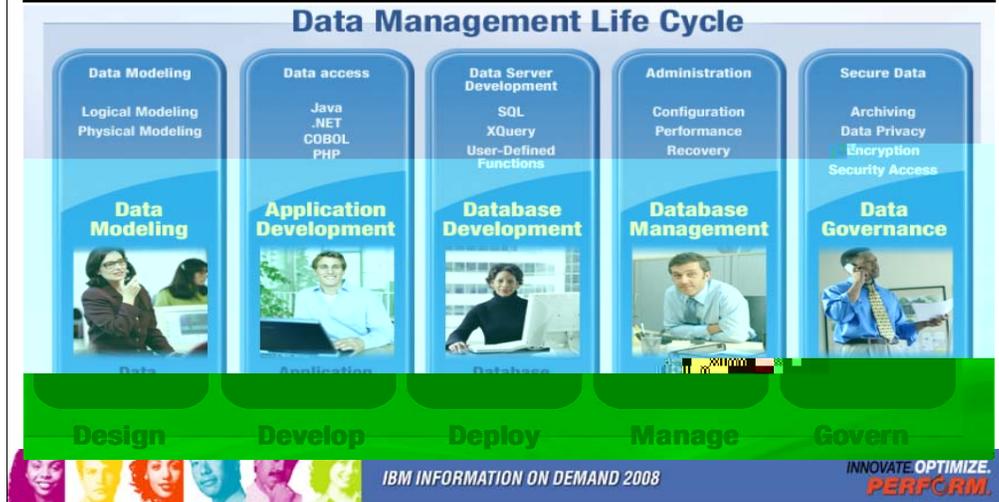
Particularly for UR readers

Requirements: NFM "Locking protocol 3" GBP changes Automatic in non-data sharing
Clean group-wide shutdown in data sharing once NFM enabled until PK62027

IBM Data Studio

A Foundation for Innovation

*IBM Data Studio is a comprehensive data management solution.
With IBM Data Studio you can design, develop, deploy and manage your data,
databases and database applications throughout the data management life cycle.*



Read from the slide directly.

In this initial release, most features are geared for the development cycle

Improved Java data access



pureQuery is a high-performance Java data access platform focused on simplifying the tasks of developing and managing applications that access data

Optimizing Java database development and deployment

- Integrated tools, APIs, and runtime
- Flexible deployment options with Static SQL support
- Unified API to query databases and in-memory Java objects with SQL



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

100

The benefits of using pureQuery extend throughout the development, deployment, management, and governance stages of the application life cycle.

pureQuery provides access to data in databases and in-memory Java objects via its tools, APIs, and runtime environment.

Embraces SQL as the common query language

Improves the Java data access life cycle

Development, deployment, management, and governance stages

Key component of IBM's end-to-end problem determination strategy

Before pureQuery

JDBC or SQLJ to access database

Tedious to develop, customize, maintain

Developer spends a lot of time on lower level data access instead of business logic

Technologies based on proprietary query languages

No visibility to the efficiency of the generated SQL

Problem determination difficult

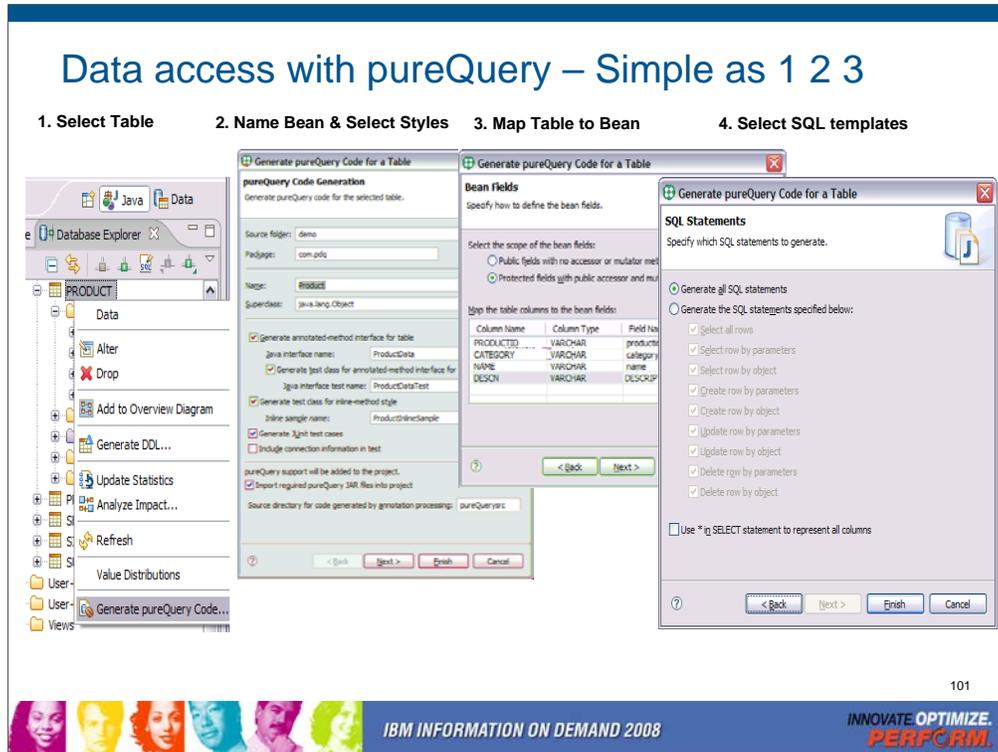
Tied to single vendor

Not sophisticated enough to handle complex scenarios

Performance boost from Static SQL requires new skills

No simple integrated way to work with database and in-memory Java objects

With pureQuery you can use standard SQL to access data from databases or in-memory Java objects



IBM® Data Studio Developer provides an integrated database development environment.

Data Studio Developer provides a complete development and testing environment for building database objects, queries, database logic, and pureQuery applications.

Develops database applications faster with integrated query editor for SQL and XQuery

Optimizes your applications and queries with ease

Builds and tests your stored procedures (Java and SQL) with the interactive routine debugger

Rapidly develops or customizes the SQL inside Java applications

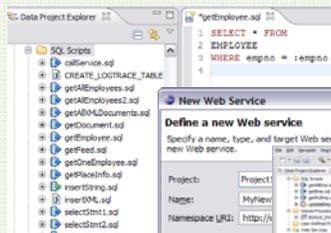
Auto-generates test applications for your pureQuery applications

Provides heterogeneous data server support for IBM DB2 Data Server and Informix Dynamic Server across the platforms.

Data Web Services and Data Studio

Data Web Services without programming

1. Create and Test Queries or Stored Procedures

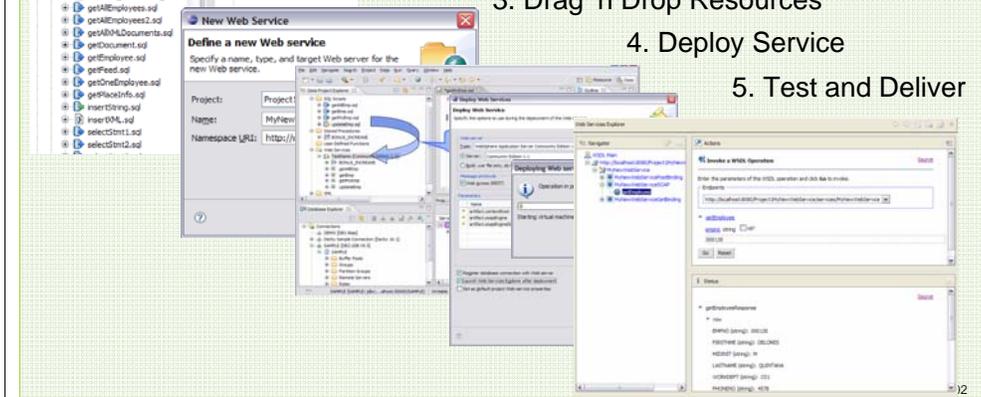


2. Create Service

3. Drag 'n Drop Resources

4. Deploy Service

5. Test and Deliver



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

The IBM® Data Studio solution includes many new features, among these is the Data Web Services support. Data Web Services is the next generation Web service provider for IBM database servers. Data Web Services provides a full Web service interface, that supports both REST-style (HTTP GET/POST binding) as well as SOAP over HTTP and automatic WSDL generation. Data Web Services supports the IBM DB2® family of databases, IBM Informix® Dynamic Server (IDS), and many popular Web servers, such as Apache Tomcat v5.5, and IBM WebSphere® Application Server V6.1 and Community Edition V1.1.

<http://www.ibm.com/developerworks/edu/dm-dw-dm-0711pauser-i.html>

http://www.ibm.com/developerworks/db2/library/techarticle/dm-0712bommireddipalli/?S_TACT=105AGY82&S_CMP=GENSITE

DB2 9 deprecated function

- Private protocol → DRDA (new help in DSNTDP2DP)
- Plans containing DBRMs → packages
- ACQUIRE(ALLOCATE) → ACQUIRE(USE)
- XML Extender → XML type
- DB2 MQ XML user-defined functions and stored procedures → XML functions
- DB2 Management Clients feature (DB2 Administration Server, Control Center, & Development Center) → IBM Data Studio application & administration services
- msys for Setup DB2 Customization Center → install panels
- BookManager use for DB2 publications → Info Center, pdf

103



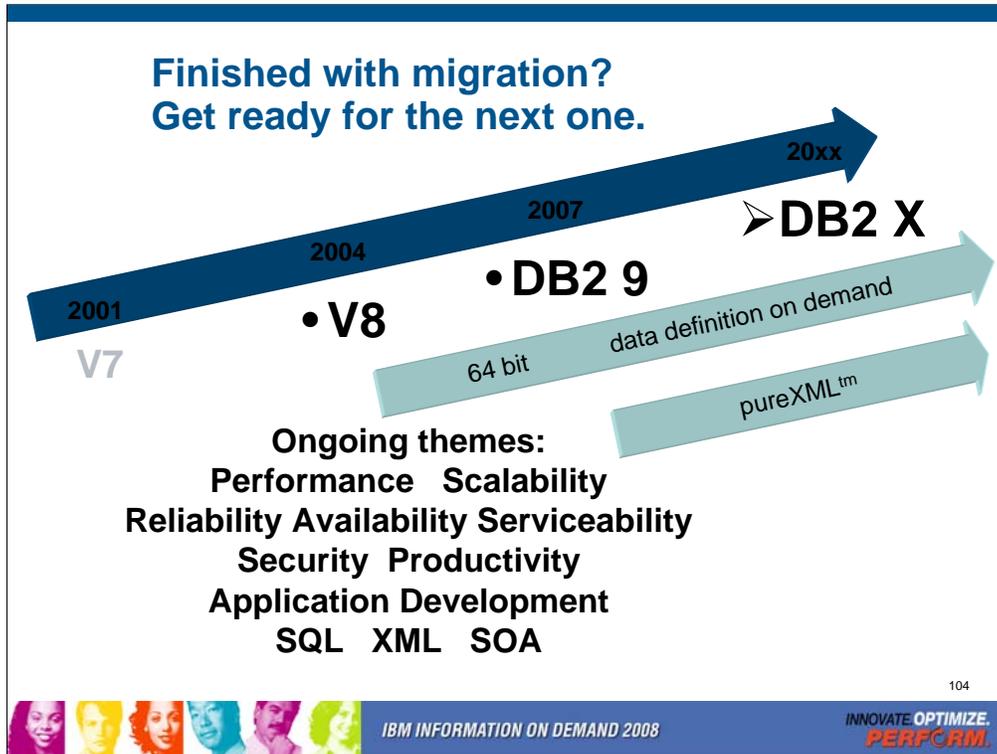
IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

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

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

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

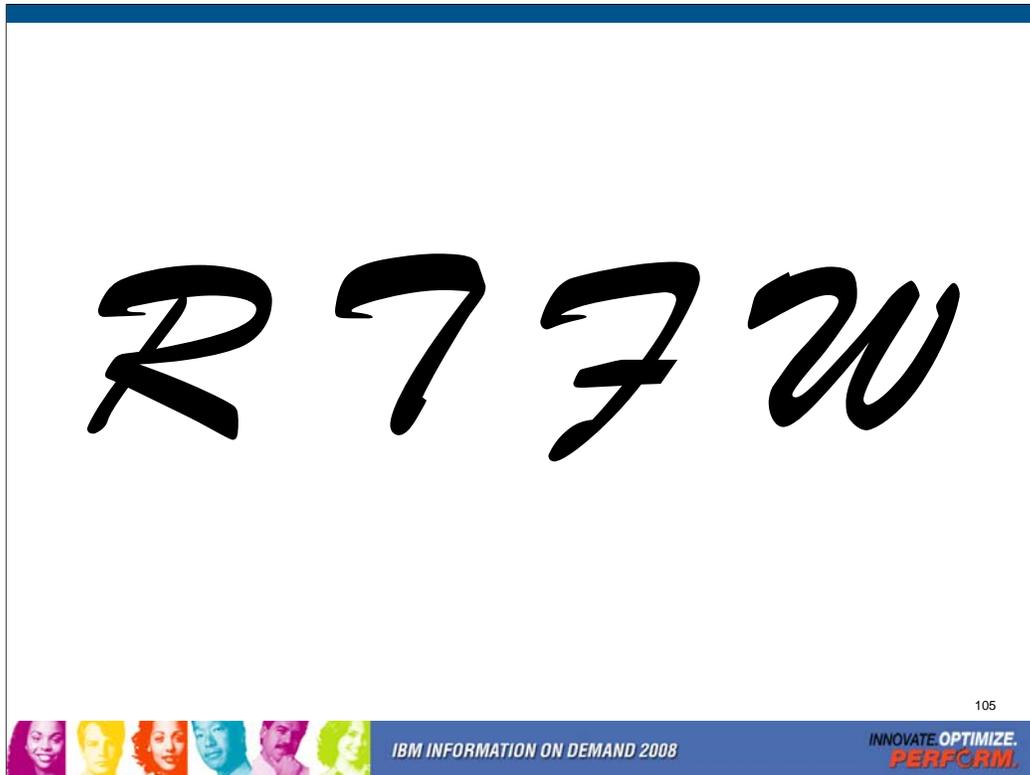


DB2 for z/OS V7 became generally available (GA) March 2001, and V8 delivered three years later, March 2004. DB2 9 became generally available in March 2007, three more years. We expect the next version will be roughly 3 years from DB2 9 GA to DB2 10 or DB2 X or whatever the name becomes, but the timing and packaging may change too.

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

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

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



RTFW is the acronym for Read the Friendly Web. Let's take a short walk on the wild, wild web. The problem with the web is not too little information, but rather too much information. The experience is a bit like trying to take a drink from a fire hose. So I'd like to help a bit by narrowing the search with the twenty five cent tour of a few of my favorite DB2 web sites. A lot more information has been added in the past month or two, with many new books and web pages.

Let's start with the DB2 family. I'll generally show the short form or alias of the URL, omitting `http://`

Here are some tips for avoiding the 404. You don't need `www` in front of `ibm.com` in most situations. There is often something after `www`, such as the `-306` in `www-306` when you get the URL back from the browser. Remove the hyphen and number when you save the URL, since that number changes more quickly than the rest of the URL.

Get updated books latest September 2008
<http://publib.boulder.ibm.com/infocenter/imzic/>

- 9Program Directories
- 9Installation Guide
- 9Redbooks
- %Data Sharing: Planning and Administration
- %Application Programming Guide & Reference for JAVA™
- %RACF Access Control Module Guide
- %ODBC, Spatial, ...
- 9Messages
- 9Codes
- 9Administration Guide
- 9Performance Guide
- 9Utility Guide and Reference
- 9Internationalization Guide (Unicode)

106

For installation, you need many books. Some are optional, for example the data sharing book is not needed if you don't use data sharing. If you use data sharing, then part of the install process is in that book. Part of the Java install process is in the Java book. Part of the ODBC install process is in the ODBC book. If you use RACF access control, then you need the RACF Access Control Module Guide. You can get most of the books from the DB2 Library web page. The books were updated in December 2007 and February, March, June, and August 2008, with some coming later, so get the latest ones. Some of the Redbooks may be helpful. You may need books from the z/OS Library as well.

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

<http://www.ibm.com/support/docview.wss?rs=64&uid=swg27011658>

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

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

DB2 9 in IBM Redbooks Publications

- DB2 9 Technical Overview SG24-7330
- DB2 9 Performance Topics SG24-7473
- DB2 9 Stored Procedures SG24-7604
- Index Compression with DB2 9 for z/OS redp4345
- SQL Reference for Cross-Platform Development
- Enterprise Database Warehouse, SG24-7637
- DB2 9 Optimization Service Center SG24-7421
- LOBs with DB2 for z/OS SG24-7270
- Powering SOA with IBM Data Servers SG24-7259
- Enhancing SAP - DB2 9 SG24-7239
- Best practices SAP BI - DB2 9 SG24-6489-01
- Data Sharing in a Nutshell, SG24-7322
- Securing DB2 & MLS z/OS SG24-6480-01
- Data Sharing: Distributed Load Balancing & Fault Tolerant Configuration redp4449
- Considerations on Small & Large Packages redp4424
- Backup and Recovery Considerations



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

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.

- DB2 9 Technical Overview, SG24-7330 <http://www.redbooks.ibm.com/abstracts/SG247330.html>
- DB2 9 Performance Topics, SG24-7473, <http://www.redbooks.ibm.com/abstracts/SG247473.html>
- DB2 9 Stored Procedures, SG24-7604, <http://www.redbooks.ibm.com/abstracts/SG247604.html>
- Index Compression DB2 9, REDP4345, <http://www.redbooks.ibm.com/abstracts/redp4345.html>
- 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
- Enterprise Data Warehousing, SG24-7637, <http://www.redbooks.ibm.com/abstracts/sg247637.html>
- Powering SOA IBM Data Servers, SG24-7259
<http://www.redbooks.ibm.com/abstracts/SG247259.html>
- LOBs: Stronger & Faster SG24-7270, <http://www.redbooks.ibm.com/abstracts/SG247270.html>
- Securing DB2 & MLS z/OS, SG24-6480-01, <http://www.redbooks.ibm.com/abstracts/sg246480.html>
- Enhancing SAP, SG24-7239, <http://www.redbooks.ibm.com/abstracts/SG247239.html>
- Best practices SAP BI, SG24-6489-01, <http://www.redbooks.ibm.com/abstracts/sg246489.html>
- Optimization Service Center, SG24-7421, <http://www.redbooks.ibm.com/abstracts/sg247421.html>
- Data Sharing in a Nutshell, [SG24-7322](http://www.redbooks.ibm.com/abstracts/sg247322.html), <http://www.redbooks.ibm.com/abstracts/sg247421.html>
- DB2 9 for z/OS Data Sharing: Distributed Load Balancing and Fault Tolerant Configuration
<http://www.redbooks.ibm.com/abstracts/redp4449.html>
- DB2 for z/OS: Considerations on Small and Large Packages
<http://www.redbooks.ibm.com/abstracts/redp4424.html>

ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/

Name	Size	Type	Modified
application-development		File Folder	2008-05-08 17:02
customer-experience		File Folder	2008-06-09 12:02
migration		File Folder	2008-05-19 13:02
overview		File Folder	2008-05-27 13:02
performance		File Folder	2008-05-19 15:02
security		File Folder	2008-01-21 18:02
v8-new-function		File Folder	2008-05-29 11:02
v9-new-function		File Folder	2008-05-29 12:02

Name	Size	Type	Modified
administrators-idug-na-2007-teng.pdf	119 KB	Adobe Acrobat 7.0 ...	2008-04-17 13:02
availability-beta-works-2007.pdf	1.19 MB	Adobe Acrobat 7.0 ...	2008-04-21 19:02
backup-recovery-id-2006-teng.pdf	422 KB	Adobe Acrobat 7.0 ...	2008-04-17 13:02
buffer-pool-management-id-2007-teng.pdf	204 KB	Adobe Acrobat 7.0 ...	2008-04-17 13:02
clone-tables-idug-na-2007-hye.pdf	461 KB	Adobe Acrobat 7.0 ...	2008-04-17 13:02
data-sharing-beta-works-2007.pdf	605 KB	Adobe Acrobat 7.0 ...	2008-04-21 19:02
leverage-native-xml-id-2006-zhang.pdf	444 KB	Adobe Acrobat 7.0 ...	2008-05-29 12:02
labs-id-2006-weltrauch.pdf	314 KB	Adobe Acrobat 7.0 ...	2008-05-29 12:02
labs-id-2007-weltrauch.pdf	314 KB	Adobe Acrobat 7.0 ...	2008-04-15 19:02
new-sql-id-2007-favero.pdf	730 KB	Adobe Acrobat 7.0 ...	2008-04-15 19:02
optimizer-id-2006-purcell.pdf	332 KB	Adobe Acrobat 7.0 ...	2008-04-17 15:02
partitioning-clone-tables-beta-works-2008.pdf	1.24 MB	Adobe Acrobat 7.0 ...	2008-04-21 17:02
pure-xml-idug-na-2007-zhang.pdf	1.14 MB	Adobe Acrobat 7.0 ...	2008-04-17 13:02
reordered-row-format-not-logged-beta-works-2007.pdf	825 KB	Adobe Acrobat 7.0 ...	2008-04-21 17:02
reordered-row-format-read-log-beta-works-2007.pdf	1.79 MB	Adobe Acrobat 7.0 ...	2008-04-21 17:02
scalability-beta-works-2008.pdf	0.98 MB	Adobe Acrobat 7.0 ...	2008-04-21 19:02
sql-enhancements-beta-works-2007.pdf	1.77 MB	Adobe Acrobat 7.0 ...	2008-04-21 18:02
xml-db2-9-roadshow-2007.pdf	1.52 MB	Adobe Acrobat 7.0 ...	2008-05-08 16:02
xml-part1-foundation-dba-beta-works-2007.pdf	1.05 MB	Adobe Acrobat 7.0 ...	2008-04-21 18:02
xml-part2-application-development-beta-works-2007.pdf	799 KB	Adobe Acrobat 7.0 ...	2008-04-21 19:02
availability-beta-works-2007.wmv	47.8 MB	Windows Media Aud...	2008-04-21 20:02

108



The ftp server has moved and has been reorganized. Now, if you use the url above, you'll see categories for new function in V8, new function in DB2 9, application development, customer experience, migration, performance, security and overview (none of the above).

If you move to one of the topics (DB2 9 new function in this example), you can check the individual presentations. Many have slides in pdfs and audio with the slides in .wmv files.

Here are slides from the 2008 road show:

<http://www.ibm.com/support/docview.wss?rs=865&uid=swg27013016>

More information on the web: hours of education

COBOL and DB2 V8 multirow fetch, insert

<http://www.ibm.com/support/docview.wss?&uid=swg27012149>

<http://www.ibm.com/support/docview.wss?&uid=swg27012160>

DB2 9 education from BetaWorks

XML <http://www.ibm.com/support/docview.wss?&uid=swg27012156>

SQL <http://www.ibm.com/support/docview.wss?&uid=swg27012157>

Clones <http://www.ibm.com/support/docview.wss?&uid=swg27012150>

Scalability <http://www.ibm.com/support/docview.wss?&uid=swg27012151>

Data sharing <http://www.ibm.com/support/docview.wss?&uid=swg27012152>

Availability <http://www.ibm.com/support/docview.wss?&uid=swg27012153>

Security <http://www.ibm.com/support/docview.wss?&uid=swg27012155>

Data <http://www.ibm.com/support/docview.wss?&uid=swg27012158>

Install <http://www.ibm.com/support/docview.wss?&uid=swg27012181>

Log detail <http://www.ibm.com/support/docview.wss?&uid=swg27012154>

109



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

There is a lot of new information on the web, and the information has been reorganized.

The easiest way to see the information is from the events page, then clicking on presentations from previous conferences. See the performance presentations, a course on DB2 9 with slides and audio (very big downloads).

Here are the topics for the items, and the suffix of URL

XML 156 SQL 157 scalability 151 data sharing 152

Partitioning and clones 150 availability 153 log detail 154

Security 155 reordered row format & no log 158

Migration suffix 181, also

<ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/migration>

The ftp site (next slide) can also be used to access these

<ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/>

<ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/v9-new-function/reordered-row-format-read-log-betaworks-2007.pdf>

Main DB2 for z/OS Web Page
ibm.com/software/db2zos

- Library (books)
- Events
- Education
- Services
- Support
- V8, DB2 9
- Developer Domain
- DB2 Magazine
- z9 zIIP
- Tools

Home Business solutions IT services Products Support & downloads

Welcome Mr. Roger Miller

Software > Information Management > DB2 Product Family >

DB2 for z/OS

Features and benefits

System requirements

Library

Success stories

News

How to buy

Events

Training and certification

Services

Support

Related software

DB2 for z/OS

[+ Add to My interests](#)

Overview	Editions
<ul style="list-style-type: none"> • Data serving with IBM's mainframe and DB2 for z/OS helps you reduce cost and complexity in your IT infrastructure, simplify compliance, and leverage your core asset - your data. Learn more about industry-leading innovations: <ul style="list-style-type: none"> • zIIP enablement is ready for DB2 for z/OS Version 8 • System z9 • System z9 Integrated 	<ul style="list-style-type: none"> • DB2 Version 9.1 for z/OS Available March 16, 2007 • DB2 for z/OS Version 8 • DB2 for z/OS and OS/390 Version 7 - Service ends June 30, 2008 • DB2 for z/OS Product Lifecycle - Updated Marketing and support service availability by version

IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE. **PERFORM.**

This is the main DB2 for z/OS web page. You can get to the other DB2 for z/OS pages from here, so I often call this my home page. This page changes frequently, so look at the highlighted NEW items. Do you want to look in a DB2 book? Click on Library to see books on DB2 and QMF Version 8 (about 40), Version 7, 6 or even 5. V6 and V5 are out of service. You can check the latest changes by looking at the Information Updates or go to the Information Center. From this page, you can look for conferences (Events), specific classes (Education), or services. If you want to see the latest on DB2 9 or DB2 Version 8, click on the DB2 9 or the V8 link. If your primary concern is application development, the Developer Domain is for you. DB2 Magazine covers a broad range of topics about DB2. The latest machines System z9, z990 and z890 are on the System z page. Click DB2 and IMS Tools to see the wide range of help we provide.

DB2 for z/OS Lifecycle

ibm.com/software/data/db2/zos/

IBM Software Support Lifecycle

Product support lifecycle

The announcement letter dates below are US only. Information for other country announcements is located [here](#).

Product name/comments	Version/Rel./Mod.	PID/License type	General availability*/End of marketing/End of support	Last updated
DB2 for z/OS	8.1	5625-DB2 ICA	GA: 26-Mar-2004,204-029 EOM: EOS:	06-Mar-2007
DB2 for OS/390 and z/OS Comments: Replacement: DB2 for z/OS v8 (5625-DB2)	7.1	5675-DB2 ICA	GA: 30-Mar-2001,201-054 EOM: 05-Mar-2007,906-254 EOS: 30-Jun-2008,907-023	07-Feb-2007
DB2 for OS/390	6.1	5645-DB2 ICA	GA: 30-Jun-1998,299-154 EOM: 04-Dec-2001,901-307 EOS: 30-Jun-2005,904-021	30-Jan-2007
DB2 for OS/390	5.1	5655-DB2 ICA	GA: 27-Jun-1997,297-146 EOM: 31-Dec-2001,901-184 EOS: 31-Dec-2002,901-202	30-Jan-2007

Related solutions

- DB2® and IMS Tools
- IBM Announces DB2 9 for z/OS - General availability

Editions

- DB2 Version 9.1 for z/OS Available March 16, 2007
- DB2 for z/OS Version 8
- DB2 for z/OS and OS/390 Version 7 - Service ends June 30, 2008
- DB2 for z/OS Product LIFECYCLE - Updated**
Marketing and support service availability by version

Tools and components

- DB2 and IMS Tools
IBM DB2 and IMS Tools are specifically designed to enhance the performance of DB2 and IMS. Affordable database power tools at your

We're here to help

Easy ways to get the answers you need.

Call me now

E-mail us

or call us at 877-426-3774
Priority code: 104CBW60

DB2 9 for z/OS
Secure Information Services You Can Trust
Learn more

111

IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE. PERFORM.

Customers can track product availability on their own from this site. Start on the main DB2 for z/OS web page, then click on the Lifecycle.

Data Studio
<http://www.ibm.com/software/data/studio/>

United States [change]

Home Solutions - Services - Products - Support & downloads - My IBM -

Welcome Mr. Roger Miller [Not you?] [IBM Sign in]

Software > Information Management >

IBM Data Studio

The comprehensive data management solution to design, develop, deploy, and manage data-driven applications.

"With IBM Data Studio, we can discover faulty assumptions and get resolution from business sponsors early in the design process. This will help us reduce development costs as it's 50 to 75 percent more expensive to fix a design problem after we implement the changes"

Kevin Campbell, Application Architect, Univar USA

Why IBM Data Studio

IBM Data Studio is an Integrated Data Management Environment. Learn how IBM Data Studio can increase productivity and reduce development cost throughout the data lifecycle.

Related links

- DB2 9
- Informix
- U2
- IMS
- IBM Balanced Warehouse
- DB2 for z/OS

We're here to help

Easy ways to get the answers you need.

Call me now

E-mail us

or call us at
877-426-3774
 Priority code:
104CBW60

Finalist International

112

IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE. PERFORM.

See this page for all the changes in IBM Data Studio. Watch closely, as this area is changing fast. IBM Data Studio is an Integrated Data Management Environment. Learn how IBM Data Studio can increase productivity and reduce development cost throughout the data lifecycle.

IBM Data Studio is an Integrated Data Management Environment. Learn how IBM Data Studio can increase productivity and reduce development cost throughout the data lifecycle.

IBM Data Studio Developer: An Integrated Development Environment for creating and testing database and pureQuery applications.

IBM Data Studio pureQuery Runtime: A high-performance Java data access platform -- improves security and manageability of Java application connections to databases.

<http://www.ibm.com/software/data/studio/>

zIIP ibm.com/systems/z/ziip/ II14219 II12836 II10817

Home Products Services & solutions Support & downloads My account

Servers > Mainframe servers > Hardware >

Mainframe servers
About IBM mainframe
Hardware
Software
Operating systems
Solutions
Customer successes
Literature
FAQs
Support
Education
News and events

zIIP – Overview

zIIP Getting started
Overview About Resources

Today's corporate data
Data is at the core of every business. But data management can be a big project. It needs to be available globally; it must be kept up in real time, and accessible 24x7. At the same time, it needs to be protected from unauthorized access and to comply with new regulations. The IBM System z™ mainframe family has a strong heritage in data serving, and its architecture is designed for massive data access, whether across the Internet, to storage devices, or to remote backup sites. Scalability, availability, security—these are all core competencies of the mainframe.

Learn More
→ What they are saying about zIIP
→ Data Serving Press release
→ Data Serving Roadmap z990
→ z/OS website
→ DB2 V9 for z/OS preview
→ System z security

Exciting new features
Earlier this year, IBM extended its mainframes data serving capabilities, delivering a new roadmap for the future of data serving and information on demand, previewing new DB2 function, and introducing a new specialty engine directed toward data serving workloads.
The new specialty engine, the IBM System z9 Integrated Information Processor (IBM zIIP), is now available on the System z9 Enterprise Class (EC) and System z9 Business Class (BC) servers.

We're here to help
IBM Mainframes just got even more cost competitive. Let us show you how with a customized quote.
Request a quote
or call us at 1-888-SHOP-IBM

Related links
• Resource Link
• Resources for IBM Business Partners
• Resources for developers
• ShopSystemZ
• Printing solutions
• ISV software support
• IBM Training
• IBM Design Centers

113

IBM INFORMATION ON DEMAND 2008 INNOVATE. OPTIMIZE. PERFORM.

See zIIP information on the web. This web site has the most current information, & pointers to more resources. The zIIP is for customers who are concerned about costs for growth. One big cost reduction is hardware cost, which is much less than a standard processor. The biggest cost reductions are in software, as IBM does not charge for software running on the specialty processors. The zIIP fits some customers very well, but will not apply for all. As a specialty processor, not all work can use the zIIP, which will only process work running under an enclave SRB. Most applications cannot run in SRB mode. The specifics of the software charging need to be considered. Customers must be current on hardware (System z9), current on software (z/OS 1.6 or later, DB2 V8 or later) and have a work load peak using the types of work eligible for zIIP:

Remote SQL processing of DRDA network-connected applications over TCP/IP: These DRDA applications include ERP (e.g. SAP or PeopleSoft), CRM (Siebel), and business

DB2 for z/OS Roadmap Page
ibm.com/software/db2zos/roadmap.html

Software > Information Management > Data Servers > DB2 Product Family

DB2 for z/OS Information Management software

DB2 for z/OS
 DB2 for z/OS Information Roadmap

Other roadmaps

- Content Manager
- IMS™
- SQL Replication Version 8.2

DB2 for z/OS

- Features and benefits
- System requirements
- Library
- Success stories
- News
- How to buy
- Events
- Training and certification
- Services
- Support

- Product overview
- Planning
- Installing
- Migrating
- Configuring
- Administering
- Developing
- Scenarios
- Samples
- Debugging
- Testing
- Security
- Monitoring and tuning
- Troubleshooting and Support
- Reference

This roadmap outlines the information resources that are available for installation, administration, and programming for DB2 for z/OS.

PDF files require Adobe® Reader®

Product overview

114

IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE. PERFORM.

Don't you wish you could find all of the different kinds of information about DB2 for z/OS in one place? I think you'll be happy to learn about a new Web page called the DB2 for z/OS Information Roadmap. Information roadmaps are being developed by many IBM product teams in an effort to simplify the lives of users by providing a single place to find all sorts of useful information resources, such as: Product information from the official library IBM Redbooks Articles in a variety of different online publications, including the DB2 Magazine, presentations, support information like answers to frequently asked questions or technotes, web pages with high-level descriptions of different products that DB2 works with and more. We tried to provide a simple way to get to many of the best resources for various tasks. This will be a great resource to people just starting to use a DB2 capability that they haven't yet used.

The information roadmap is organized according to different tasks that you might perform, such as migrating, configuring, administering, and debugging. Although the information roadmap is available for Version 8, it is available to anyone with Web access, and many of the resources are relevant and useful to users who aren't yet on Version 8.

I encourage you to check out the new DB2 for z/OS Information Roadmap: <http://www.ibm.com/software/data/db2/zos/roadmap.html>

You can also start from the main DB2 for z/OS web page and click on the Information Roadmap news item or <http://www.ibm.com/software/db2zos>

DB2 Support Page
ibm.com/software/db2zos/support.html

Software > Information Management > DB2 Product Family >

DB2 for z/OS
 Product support

Primary support resources

Flash 2 May, 2007: How can I redeploy an SQL stored procedure without rebuilding it?
 Flash 1 Nov, 2006: URGENT: Effects of Daylight Saving Time changes on your systems and applications
 Flash 21 Jul, 2006: IBM Software Support Toolbar
 [View all Flashes]

Solve a problem

Forums and newsgroups
 Featured documents

Search Support (this product)
 Enter terms, error code or APAR #

Technotes | APARs

115



Do you need Answers to Frequently Asked Questions? Product Information? Technical Presentations? Redbooks? White Papers? Support for all of these is on this web page. Click Support on DB2 web pages. From this page, you access roughly 800 Technotes or answers to Frequently Asked Questions, technical presentations from recent conferences, a dozen white papers, and over 50 Redbooks. For more presentations, see the Events page. The Consolidated Service Test (CST) documents IBM recommended service level for a range of z/OS products.

To search for the presentation you want, put a word or two into the Additional search terms or Order by date – newest first.

There are new and updated redbooks on Powering SOA with IBM Data Servers, LOBs with DB2 for z/OS: Stronger and Faster, V8 Performance Topics, Data Sharing in a Nutshell, Data Integrity, Stored Procedures: Through the Call and Beyond, Design Guidelines for High Performance and Availability, Disaster Recovery, and Security Improvements.

ibm.com/software/db2zos/support.html

ibm.com/support/docview.wss?rs=64&uid=swg21307783

Software > Information Management > DB2 Product Family >

Support & downloads

DB2 for z/OS

Features and benefits

System requirements

Library

Technical resources

Success stories

News

How to buy

Events

Training and certification

Services

Support

Functions available in conversion mode vs. new-function mode

Technote (FAQ)

This document applies only to the following language version(s):
English

Question
Which new functions work only in DB2 9 for z/OS new-function mode (NFM), and which new functions also work in conversion mode (CM, previously called compatibility mode)?

Answer
The short answer
New functions are available only in new-function mode, unless explicitly stated otherwise in the product documentation. A few general exceptions exist for utilities and for optimization.

Where is this information documented?
This general guidance is stated in the information center (http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db29.doc.inst/db2z_relincompatassumptions), and in a statement at the beginning of each PDF book, with some specific exceptions. For most information, the guidance is that the function is not supported in conversion mode unless noted. For utilities and optimization, the function is available in conversion mode unless noted. The functions available in conversion mode (CM), conversion mode* (CM*), enabling-new-function mode (ENFM), and enabling-new-function mode* (ENFM*) are

Related software

- WebSphere
- CICS
- IMS
- DB2 Connect
- z/OS
- Linux on IBM System z

116

IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

See this web page for more detail on what works in what mode.

Question: Which new functions work only in DB2 9 for z/OS new-function mode (NFM), and which new functions also work in conversion mode (CM, previously called compatibility mode)?

Short answer: New functions are available only in new-function mode, unless explicitly stated otherwise in the product documentation. A few general exceptions exist for utilities and for optimization.

Where is this information documented?

This general guidance is stated in the information center, and will be updated there.

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

<http://publib.boulder.ibm.com/infocenter/imzic/>

DB2 9 Technical Resources page
ibm.com/support/docview.wss?rs=64&uid=swg27011656

DB2 for z/OS - Technical Resources

Support & downloads
 DB2 for z/OS
 Features and benefits
 System requirements
 Library
 Technical resources
 Success stories
 News
 How to buy
 Events
 Training and certification
 Services
 Support

Download 30 PDF books
Access html books
DB2 new August 2008
Program Directories
Information Center

Related software
 • WebSphere
 • CICS
 • IMS
 • DB2 Connect
 • z/OS
 • Linux on IBM System z

Related hardware
 • IBM TotalStorage
 • IBM System z

Related solutions
 • DB2 and IMS Tools
 • Java on OS390 and z/OS
 • DB2 for Linux

Related services
 • IBM Consolidated Service Test

Product documentation

Abstract
 Here is your complete DB2 for z/OS information resource! Here you can find DB2 for z/OS technical information, information about related tools, and valuable technical tips.

Content
 To support recent enhancements to the DB2 for z/OS Library page, the DB2 for z/OS technical information has a new address on the Web. **Please update your bookmarks!**

Information center topics
 → Information Management Software for z/OS Solutions Information Center
 Search topics across all information for DB2 for z/OS, IMS, and DB2 and IMS Tools, including QMF. The Information Management Software for z/OS Solutions Information Center is the primary source for the information about DB2 for z/OS.

Search the information center topics

As part of IBM's ongoing effort to improve your information experience, we are optimizing the information delivered with the DB2 for z/OS product for viewing in the information center. [Learn more.](#)

Technical information by version
 → DB2 Version 9.1 for z/OS
 → DB2 for z/OS Version 8

Documentation
 Product
 Softw
 Data
 Mana
 Data
 (Data
 Mana
 Syste
 DB2
 Operati
 z/OS
 Softwar
 7.0, 6
 Softwar
 All Et
 Referen
 7011
 IBM Grc
 Softw
 Modific
 2008

Transla
 Select
 Trn

Rate th
 Help
 117

 **IBM INFORMATION ON DEMAND 2008** **INNOVATE.OPTIMIZE.PERFORM.**

This is the Technical Resources page for DB2. I clicked Technical Resources to get here. There are about 40 PDF books and about the same number of books in html or BookManager format. Choose the style you want. I generally download the PDF books to my computer. From this page you can get books on DB2 and QMF, including both the library and program directories. If you need to look for some tools that can be downloaded, such as Estimator or Visual Explain, click on Downloads. You can also get to the Information Center from here. Watch this web page and see the dates next to each book for the latest updates.

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

DB2 Connect
 ibm.com/software/data/db2/db2connect/
 Software > Information Management > DB2 Product Family >

DB2 Connect

INTRODUCING DB2 9.5
 → Learn more

DB2 Connect makes your company's host data directly available to your personal computer and LAN-based workstations. It connects desktop and palm-top applications to your company's mainframe and minicomputer host databases for access to your enterprise information no matter where it is. **DB2 Connect** provides the application enablement and robust, highly scalable communication infrastructure for connecting Web, Windows, UNIX, Linux and mobile applications to z/OS and AS/400 data. **DB2 Connect** is included in many of the DB2 products.

Product Editions
 + Compare editions

Related links
 → What's new in DB2 Connect V9.5
 → Blog: An Expert Guide to DB2 Technology

We're here to help
 Easy ways to get the answers you need.
 Call me now
 E-mail us
 or call us at 877-426-3774
 Priority code:

118

INNOVATE. OPTIMIZE. PERFORM.

DB2 Connect makes your company's host data directly available to your personal computer and LAN-based workstations. It connects desktop and palm-top applications to your company's mainframe and minicomputer host databases for access to your enterprise information no matter where it is. **DB2 Connect** provides the application enablement and robust, highly scalable communication infrastructure for connecting Web, Windows, UNIX, Linux and mobile applications to z/OS and AS/400 data. **DB2 Connect** is included in many of the DB2 products.

<http://www.ibm.com/software/data/db2/db2connect/>

DB2 Cross-Platform SQL Reference

updated for DB2 9 March 28, 2008 another coming

ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html

The SQL Reference for Cross-Platform Development

Level: Introductory

Staff, IBM, Staff, IBM

01 Jun 2002

Updated 28 Mar 2008

Great news for people who want to write portable applications! This reference makes it easy for you to develop applications using SQL that is portable across the [DB2 database family](#), including [DB2 for z/OS®](#) Version 9, [DB2 for iSeries™](#) Version 5 Release 4, and [DB2 for Linux, UNIX®, and Windows®](#) Version 9.

Introduction

The *SQL Reference for Cross-Platform Development* uses the familiar format of the product SQL Reference manuals and includes the following topics:

- Database concepts
- Built-in functions
- Statements and queries
- SQL procedural language control statements
- Limits
- Programming in C, Java™, COBOL, REXX, and external routines in general
- SQL return codes (SQLSTATES)
- CCSIDs

New in Version 3!

119



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

This is the web page for the SQL Reference for Cross-Platform Development. It has three versions of the cross-platform SQL Reference and pointers to the base product SQL Reference books as well. If you want to write programs that work well across the DB2 family, then this is the book to download and use.

These books publish frequently, so you need to choose the one for the proper combination of DB2 platforms and versions. The current book is Version 3 for platforms Version 9 of DB2 for z/OS, Version 5 Release 4 of DB2 for i5/OS, and Version 9 of DB2 for the Linux, UNIX and Windows. The addition of newer versions of the DB2 products means the following great SQL features are now included:

- BIGINT data type
- OLAP specifications
- New built-in functions
- INTERSECT and EXCEPT
- Recursive common table expressions
- Additional common features for ALTER COLUMN on ALTER TABLE
- Table partitions
- INSTEAD OF triggers
- SQL PL extensions including the FOR statement and support for nested ATOMIC compound statements
- Other portability enhancements

Other books are for DB2 for z/OS Version 7 and Version 8.

<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

Events for more info ...

Software > Information Management > Data Servers > DB2 Product Family > DB2 for z/OS Events

IDUG

IOD

SHARE

Presentations

120

IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE. PERFORM.

There are many excellent conferences coming up. Use the Events page to see which one you can attend:

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

If you missed a conference or two, then check for recent presentations from previous conferences. More than 150 presentations are on the web, some with audio as well. 2008 conferences:

<u>conference</u>	<u>dates</u>	<u>location</u>	<u>web location</u>
Share	Aug 10 - 15	San Jose, CA	www.share.org
IDUG Europe	Oct 13 - 17	Warsaw, Poland	www.idug.org
IBM IOD NA	Oct 26 – 31	Las Vegas, NV	

ibm.com/events/InformationOnDemand

IBM Technical Conferences

<http://www.ibm.com/services/learning/ites.wss?pageType=page&c=a0001438>

ibm.com/software/data/db2imstools/

Click
Support
For
Compati-
bility
Matrix
→ → →
&
V8 or
DB2 9

The screenshot shows the IBM DB2 and IMS Tools website. The page has a blue header with navigation links: Home, Products, Services & solutions, Support & downloads, and My account. Below the header, there is a breadcrumb trail: Software > Information Management >. The main heading is "IBM DB2 and IMS Tools". To the left of the main content is a sidebar menu with links: IBM DB2 and IMS Tools, Library, News, How to buy, Events, Training and certification, and Services. The main content area contains a welcome message: "Welcome to the website for IBM® DB2 and IMS Tools. These pages include information on tools for DB2 for z/OS, [DB2 on Linux, UNIX, and Windows](#), IMS, and Information Integration." Below this is a section for news: "Get the latest [news](#) on [announcements](#), [webcasts](#), and [customer case studies](#). Link to the [library](#) to find...". On the right side, there is a "Portfolio of DB2 & IMS Tools" banner. At the bottom of the page, there is a footer with a row of colorful portraits of people, the text "IBM INFORMATION ON DEMAND 2008", and the IBM logo with the slogan "INNOVATE. OPTIMIZE. PERFORM." and the number "121".

This is the primary web page for DB2 and IMS tools. If you want to know exactly which levels of each tool work with DB2 for z/OS Version 8 or with DB2 9 for z/OS, then go to the Support page for the tools. You may find the item directly on the Support page. Or click Technotes (FAQs) and search for items that include DB2 9 or V8 and PTF. Select the item DB2 Data Management Tools and DB2 for z/OS V8.1 or DB2 9 compatibility. This table provides the minimum maintenance required for DB2 Tools to support DB2 for z/OS V8. The Support page has a wide range of other detail about these products.

ibm.com/support/docview.wss?rs=434&uid=swg21256800

DB2 Data Management Tools and DB2 for z/OS V9.1 Compatibility

Flash (Alert)

Abstract
This table provides the minimum maintenance required for DB2® Tools to support DB2 for z/OS V9.1.

Content
Key:
● **NS = Not Supported**
The product does not function properly on DB2 V9.
● **T = Tolerate**
Allows the product to execute as it did on a previous release of DB2 but does not utilize any of the new DB2 V9 functions.
● **E = Exploit**
The product utilizes the new functions and features available in DB2 V9.

NOTE: If the product is not listed in the matrix, it is not supported on DB2 V9. For new products that have been released after this matrix was published contact your IBM Sales Representative for DB2 V9 compatibility.

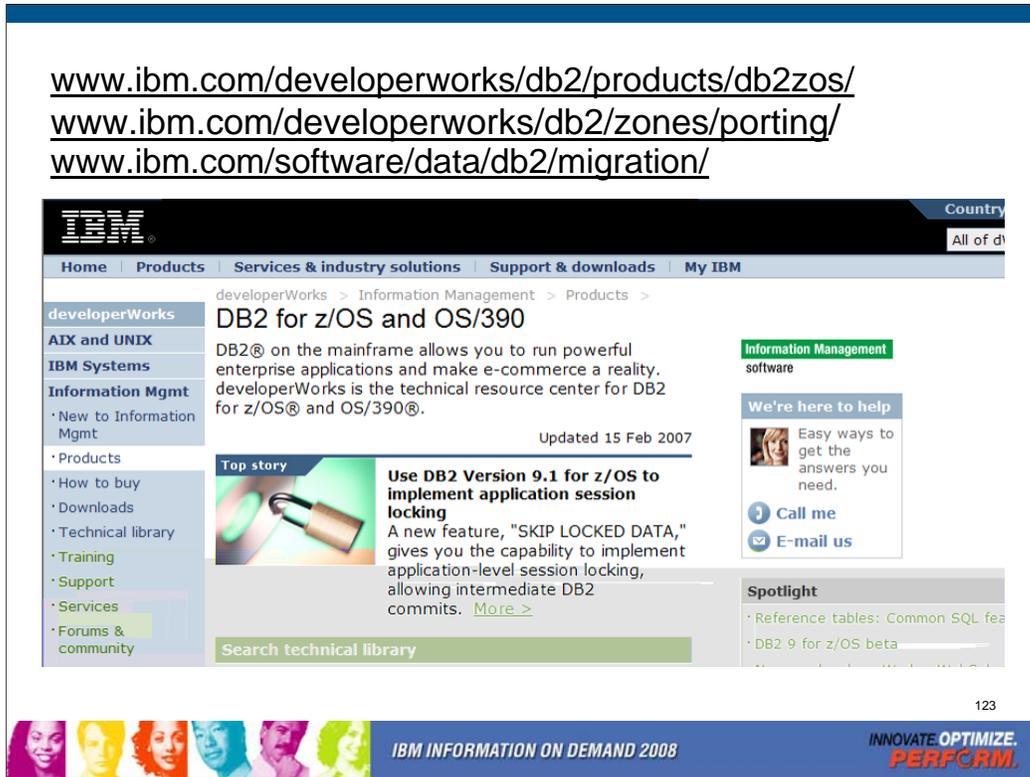
DB2 Tools	DB2 V9 Support				Comments
	NS	T	E		
Application Recovery Tool				X	PK33728
Data Encryption				X	No PTF necessary
DB2 Administration Tool				X	No PTF necessary
DB2 Administration Toolkit for z/OS SAP Edition				X	No PTF necessary
DB2 Archive Log				X	PK38148(See * below)

Software version: **9.1.0**
Reference #: **1256800**
IBM Group: **Software Group**
Modified date: **2007-03-16**

122

This is the resulting page today, looking for the Information Management Tools releases and service needed to work with DB2 9 for z/OS. Note that the matrix lists some levels of tools with an X in the NS or not supported column. See the later version, with an X under T for toleration or under E for exploit or the ability to use those new improvements.

www.ibm.com/developerworks/db2/products/db2zos/
www.ibm.com/developerworks/db2/zones/porting/
www.ibm.com/software/data/db2/migration/



The screenshot shows the IBM DeveloperWorks website for DB2 for z/OS and OS/390. The page has a dark blue header with the IBM logo and a 'Country' dropdown menu. Below the header is a navigation bar with links for 'Home', 'Products', 'Services & industry solutions', 'Support & downloads', and 'My IBM'. The main content area is titled 'DB2 for z/OS and OS/390' and includes a brief description of the product. A 'Top story' section highlights a new feature in DB2 Version 9.1 for z/OS, 'SKIP LOCKED DATA', which allows for application-level session locking. The page also features a sidebar with various links and a 'Spotlight' section.

There is extensive help for applications and porting to DB2 on the web, with suggestions for the process and many resources, including a migration toolkit. The dramatically improved DB2 family compatibility in DB2 for z/OS V7, V8 and DB2 9 makes porting to DB2 for z/OS from other platforms much easier.

<http://www.ibm.com/developerworks/db2/products/db2zos/>

<http://www.ibm.com/developerworks/db2/zones/porting/>

<http://www.ibm.com/software/data/db2/migration/>

DeveloperWorks has many discussions of techniques. Here is the URL for the main DeveloperWorks page:

<http://www.ibm.com/developerworks/db2/>

There is also a DeveloperWorks page on DB2 for z/OS.

<http://www.ibm.com/developerworks/db2/products/db2zos/>

<http://www.ibm.com/developerworks/db2/library/techarticle/dm-0807stolze2/>

For migration information for specific databases, try

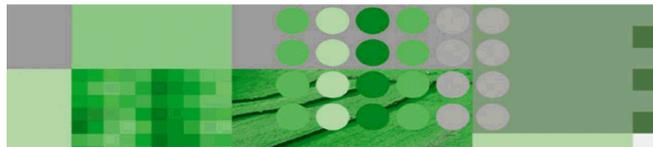
<http://www.ibm.com/developerworks/ondemand/migrate/db.html>

Version 7 end of service & final information refresh!

- **June 30, 2008**
 - End of service for DB2 for z/OS Version 7
- **August 1, 2008**
 - Final information refresh for Version 7 available on the Web and in the IBM Publication Center
- <http://www.ibm.com/support/docview.wss?rs=64&uid=swg27011660>

DB2. DB2 Universal Database for OS/390 and z/OS

Version 7



124



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

Softcopy formats are the information center, PDF, and BookManager.

Incorporates updates for approved DCRs, verified defects, and comments received by db2zinfo up to the cutoff date.

Changed “Library” page

- Focused toward marketing materials
- Contains links to the information center, technical resource pages, and IBM Redbooks

Library

Library items below are for DB2 for z/OS. For other library items, try the [library search](#).

↓ Analyst reports	↓ Data sheets	↓ Product documentation
↓ Books	↓ Demos	↓ White papers
↓ Brochures	↓ Information centers	

SWF files require [Adobe® Flash® Player](#)
PDF files require [Adobe® Reader®](#)

DB2 for z/OS

DB2 for z/OS

Features and benefits

System requirements

Library

Technical resources

Success stories

News

How to buy

Events

Training and certification

Services

Support

Add to My interests

Overview

- [DB2 for z/OS Value Unit Edition](#) The same robust DB2 for z/OS data server at a one-time charge price and available for eligible net new applications workloads on the System z platform.
- [Data serving](#) with IBM's mainframe and DB2 for z/OS helps you reduce cost and complexity in your IT infrastructure, simplify

125



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

Softcopy formats are the information center, PDF, and BookManager.

Incorporates updates for approved DCRs, verified defects, and comments received by db2zinfo up to the cutoff date.

New “Technical Resources” pages

- Contain the traditional content of library pages including:
 - Information Center search
 - Information Center links
 - PDF links
 - Updates links
 - BookManager search
- Changed URLs

DB2 for z/OS					
Title (Order Number)	Published	HTML	PDF	Updates	
Administration Guide SC18-9840-02	Feb 2008	HTML		Updates	
Application Programming & SQL Guide SC18-9841-01	Oct 2007	HTML		Updates	
Application Programming Guide and Reference for JAVA SC18-9842-01	Oct 2007	HTML		Updates	
Codes GC18-9843-02	Feb 2008	HTML		Updates	

126



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

Softcopy formats are the information center, PDF, and BookManager.

Technical Resources pages

Page	URL
DB2 for z/OS	http://www.ibm.com/support/docview.wss?rs=64&uid=swg27011656
DB2 9 for z/OS	http://www.ibm.com/support/docview.wss?rs=64&uid=swg27011658
DB2 for z/OS Version 8	http://www.ibm.com/support/docview.wss?rs=64&uid=swg27011659
DB2 for z/OS Version 7	http://www.ibm.com/support/docview.wss?rs=64&uid=swg27011660

Technical information by version

- ➔ [DB2 Version 9.1 for z/OS](#)
- ➔ [DB2 for z/OS Version 8](#)
- ➔ [DB2 UDB for z/OS & OS/390 Version 7](#)

Additional information resources

- ➔ [Information Roadmap](#)
Links to a wide variety of useful information resources about DB2 for z/OS
- ➔ [developerWorks - DB2 for z/OS Exchange](#)
Sample programs, SQL statements, utility jobs, and more submitted by DB2 for z/OS users

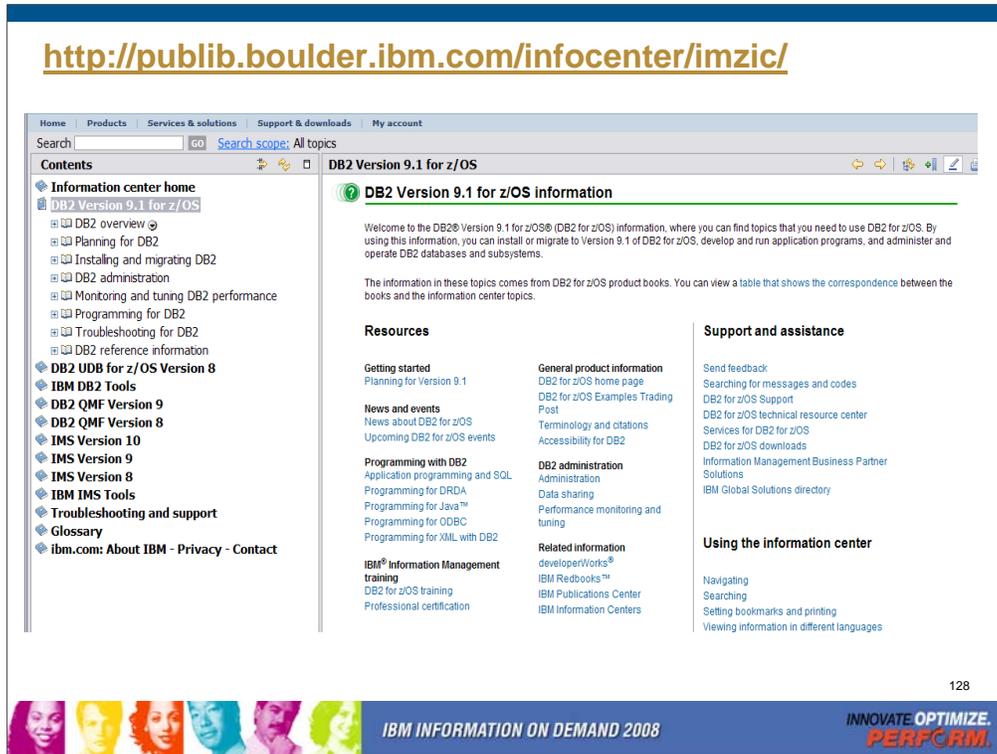
127



IBM INFORMATION ON DEMAND 2008

INNOVATE OPTIMIZE.
PERFORM.

Softcopy formats are the information center, PDF, and BookManager.



This is the Information Center, with a wide spectrum of information and access to books for DB2 for z/OS, DB2 tools, QMF, IMS, IMS tools and more. You can get to this page from the Library page, by clicking Information Center. The Information Center provides information across the books and across multiple products.

If you click “Troubleshooting and Support”, then expand under “Searching knowledge base” and click “Web search:...”, you’ll find a helpful Web search page. From this page, you can search IBM support, DeveloperWorks, or even the whole Internet using Google.

Emphasizing information center topics

- IC search enabled from the technical resource pages
 - Launches complete IC
 - Similar to the BookManager “Bookshelf search”
- IC links for each Publication
 - DB2 9 for z/OS - “HTML” links point to the information center instead of BookManager
 - Version 8 – Links are provided for IC and Bookmanager
- BookManager search moved to the bottom of the page

The screenshot shows the 'Information center topics' page. At the top, it says 'Information Management Software for z/OS Solutions Information Center'. Below this is a search box containing the text 'MSYSTABLESPACESTATS' and a 'Search' button. A mouse cursor is pointing at the search button. Below the search box, there is a paragraph of text: 'As part of IBM's ongoing effort to improve your information experience, we are optimizing the information delivered with the DB2 for z/OS product for viewing in the information center. [Learn more.](#)'

ODBC Guide and Reference SC18-9850-01	Oct 2007	HTML	Updates
Performance Monitoring and Tuning Guide SC18-9851-03	Feb 2008	HTML	Updates
RACF Access Control Module Guide SC18-9852-01	Oct 2007	HTML	Updates
Reference for Remote DRDA Requesters and Servers SC18-9853-01	Oct 2007	HTML	

m/infocenter/dzichelp/v2r2/topic/com.ibm.db29.doc/db2z_perf.htm

129



IBM INFORMATION ON DEMAND 2008

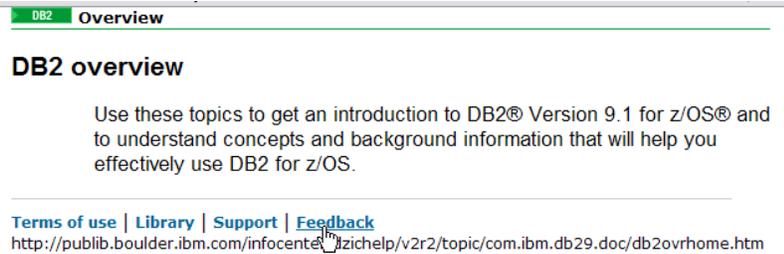
INNOVATE. OPTIMIZE.
PERFORM.

Softcopy formats are the information center, PDF, and BookManager.

Incorporates updates for approved DCRs, verified defects, and comments received by db2zinfo up to the cutoff date.

Submitting feedback

- From each information center topic, click the “Feedback” link in the topic footer, and fill out the online Reader Comment Form <http://publib.boulder.ibm.com/infocenter/imzic/index.jsp>
- Send a detailed email to db2zinfo@us.ibm.com



DB2 Overview

DB2 overview

Use these topics to get an introduction to DB2® Version 9.1 for z/OS® and to understand concepts and background information that will help you effectively use DB2 for z/OS.

[Terms of use](#) | [Library](#) | [Support](#) | [Feedback](#)

<http://publib.boulder.ibm.com/infocenter/imzic/help/v2r2/topic/com.ibm.db29.doc/db2ovrhome.htm>

130



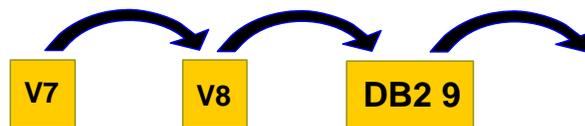
IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

Softcopy formats are the information center, PDF, and BookManager.

Incorporates updates for approved DCRs, verified defects, and comments received by db2zinfo up to the cutoff date.

Migration Path: DB2 V7 to DB2 9



Objectives? Time? Availability? Do you need fallback?
Data Sharing Coexistence? Outage permitted? How long?
When to migrate? Service desired? Version in service?
Option 1: V8 stage 2 – 6 months, stabilize & move
Option 2: Unload application, definitions & data; Load
~~Option 3: V8 on weekend then to DB2 9~~

Fallback only to V8 → no fallback, no coexistence
DO NOT skip steps

131



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

Do you want to migrate from V7 to DB2 9? We have a number of questions to ask. V8 provided a massive reengineering of the DB2 structure, memory and catalog, so it is particularly difficult to skip. The first question is the objective in terms of time you expect to save, need for availability and safety. If you need data sharing coexistence, then V7 to DB2 9 is not an option. In essence, if you need fallback, V7 to DB2 9 is not an option, since the only fallback would be to V8. Fallback to a version you have not run might not be an improvement.

The next questions are timing and service. Do you want support? DB2 9 became generally available March 2007. When can you migrate to it? DB2 V7 went out of service June 2008. Migrating to DB2 V8 is faster, easier and smoother.

The best option, option 1 would be to move to V8, stabilize and adjust, move into NFM for a few months at least, and then migrate to DB2 9. If an application, the definitions and data can be unloaded and reloaded, then you can skip versions. You don't have subsystem coexistence or fallback. So option 2 takes more work and requires more testing, but the migration can be more granular, if your applications permit. An option that we cannot recommend is a fast jump from V7 to DB2 9. You do not have migration or coexistence. Our experience from other customer situations is that skipping migration steps can cause significant problems.

In short, the safest process, by far, is the one tested by years of work in the lab and thousands of customer experiences. You can choose your own path, but the shortcut can be the longest way. Your own path has much more risk. An unsupported path adds risk and can add a lot of cost.

Important Disclaimer

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY.

WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE.

IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION.

NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, OR SHALL HAVE THE EFFECT OF:

- **CREATING ANY WARRANTY OR REPRESENTATION FROM IBM (OR ITS AFFILIATES OR ITS OR THEIR SUPPLIERS AND/OR LICENSORS); OR**
- **ALTERING THE TERMS AND CONDITIONS OF THE APPLICABLE LICENSE AGREEMENT GOVERNING THE USE OF IBM SOFTWARE.**

132



IBM INFORMATION ON DEMAND 2008

INNOVATE. OPTIMIZE.
PERFORM.

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

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

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

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

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only. This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

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

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

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