




DB2 10 for z/OS Migration Planning Webcast February 22, 1011

Roger Miller
IBM Silicon Valley Lab



© 2009 IBM Corporation

This presentation discusses migrating to DB2 10, preparing to use the new functions, understanding process and planning. Process, customer recommendations, 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. This session will help you get ready to migrate to the new version and to be ready for the next one. Whether you are ready to migrate now or are still years away, the planning step should start now, so that you can use your time more effectively.

The speaker has worked in DB2 development for the past 31 years, and continues working with the team building the install & migration process.

Objectives Early planning: deciding to migrate, what can be gained Prerequisite planning Avoiding incompatibilities Planning performance and storage Resources on the web



Disclaimer/Trademarks

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

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

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

Trademarks The following terms are trademarks or registered trademarks of other companies and have been used in at least one of the pages of the presentation:

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

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

General availability began October 22, 2010.

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

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

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

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

Here are some blogs on the topic of DB2 10 for z/OS:

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

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

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

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

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

Agenda points: DB2 10 for z/OS migration

Basic choices, Justifying DB2 10 for z/OS

Existing limits: CPU, virtual storage, latching, DB2 catalog and directory, utilities, ...

Process of getting ready, migrating

Prerequisites and deprecated functions

What to expect from migrating to DB2 10 for z/OS on day 1 and later

This presentation discusses migrating to DB2 10, preparing to use the new functions, understanding process and planning. Process, customer recommendations, 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. This session will help you get ready to migrate to the new version and to be ready for the next one. Whether you are ready to migrate now or are still years away, the planning step should start now, so that you can use your time more effectively.

The session outline will touch lightly on the first two items, which are covered in the opening session, then spend most of the time on the key decisions and process of migration, pointing to other resources.

Existing limits: CPU, virtual storage, latching, DB2 catalog and directory, utilities, ...

What's in DB2 10 to overcome these limitations

Key decisions, process, and best practices of migrating

Prerequisites and deprecated functions

What to expect from migrating to DB2 10 for z/OS on day 1 and later.



DB2 V8 DB2 9 DB2 10

Key Questions are WHEN? and HOW?

4

© 2009 IBM Corporation

If you are running DB2 9 today, then DB2 10 is an opportunity. If you are running DB2 V8 today, then your need is more urgent. DB2 V8 goes out of service in just over a year, April 2012, so your choices are to migrate to DB2 10 or 9 or to purchase extended service. Is your current fish bowl getting constrained? What is limiting you? Is it CPU? Virtual storage? Latching? DB2 catalog and directory? Concurrent utilities? Are you currently running DB2 9? V8? V7?

Should you migrate to DB2 10? To 10, or not to 10, when and how are the questions. Whether 'tis nobler in the mind to suffer the slings and arrows of outrageous limits, Or to take arms against a sea of troubles, And by migrating, 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 versions 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 versions have shuffled off this mortal coil, must give us pause: There's the respect that makes calamity of too long life for old versions; [With abject apologies to the Bard and to Hamlet act 3 scene1]

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

While DB2 10 is expected to be better than prior versions, it will have maturity and service delivery like other software, with more defects at first, then fewer as the software matures. Determining when the software is ready for a specific customer and when the customer is ready for the software depends upon the specific customer resources for testing, prior experience, and the value for the improvements versus the need for stability. Many customers depend upon tools or other software, and having that software work with DB2 is a prerequisite. When this information is known, we can answer the question.

DB2 for z/OS Lowering TCO

Maximum value for dollar investment

- Hardware pricing
 - CPU saving specialty engines (zIIP, zAAP..)
 - Compression of disk space (data, index)
- Software pricing
 - Reduction for tiers
 - Parallel Sysplex aggregation
 - z990, z9, z10 technology dividend
 - 10% reduction in charge units for each
 - zEnterprise pricing, AWLC
 - zNALC, Value Unit Edition, Subcapacity pricing, Solution Editions



**CPU + Memory + I/O and disk + Software + Energy and floor space
+ People = Improved Total Cost of Ownership (TCO)**

Software costs are higher than hardware costs for most customers, generally in a range of three to ten times higher. This is the reason that specialty engines are so attractive, as the software running on them does not incur software charges. Unit software costs continue to decline in many ways. Software charges have many options, but some of the basic tenets are consistent for most. System z software is charged for the processing capacity of the machine, taking into account multiprocessor effects, rather than using a linear scale. Larger amounts of processing power reduce the unit cost by more than 90%. The z990, z9 and z10 generations have each reduced the charging units by 10% over the prior generation. Changing from z900 to z10 processors would reduce the software charge units by about 28% for the same processing power. Additional options for reducing software charges include parallel sysplex charging, subcapacity pricing, zNALC pricing, and the new Value Unit Edition. I/O costs include the disk space and the transfer to the processor. Both costs are addressed by the work to use System z effectively. Compression for data uses z/Architecture instructions to be efficient, with a solid track record since 1993 and improved performance in the latest generations. Index compression in DB2 9 adds to the savings.

The key for pricing is the value. DB2 for z/OS delivers a very high quality of service. The service is shared, so that administration time is reduced. The amounts of processing, memory, disk and people can be reduced by effective sharing, delivering the best value for the money.

Feedback: Top Items Driving DB2 10 Migration Decisions

CPU / Performance improvements

Virtual storage enhancements

- Reduce number of members, save money

Stability & regression


Security enhancements

- Built-in security, trace & audit features, new roles, end-to-end auditing
- Cleaner/safer environment; Better audit/compliance

Temporal

Skip-level migration DB2 V8 → DB2 10

Vast majority of beta customers plan production in 2011



Customers have been enthusiastic about DB2 10 for z/OS performance and scalability. The security and temporal function are seeing stronger than expected early acceptance.

Temporal: ***As much as 80% of our applications can use this, which will drastically save developer time and even more importantly make applications easier to understand to improve business efficiency and effectiveness***

..... will allow us to drastically simplify our data-related queries and reduce our processing cost by having DB2 handle data movement more efficiently than our custom code

Building the Business Case for DB2 10 for z/OS Justifying the Upgrade

- In other sessions, we have covered the highlights and benefits of DB2 10. What follows are customer examples that describe a financial methodology to help you justify the upgrade to DB2 10.

- Examples in a scenario
 - Customer migrating from DB2 9
 - Customer migrating from DB2 V8

In DB2 V8 migration, the key differences were for SAP, Java, and transactions. The primary differences for DB2 10 migration is whether the customer is migrating from DB2 V8 or DB2 9. With the recent price changes, DB2 V8 and DB2 9 costs are similar, while more advantages deliver to customers who skip from V8 to DB2 10. Customer practices need to change more for skipping.

True TCO Should Include Both Direct and Indirect Costs and Benefits

- ❑ Direct: Purchase Price or cost of acquisition (software, hardware, maintenance), labor costs for planning and implementing the upgrade, ongoing labor costs, ongoing maintenance costs, application costs, any consultant fees or contracted support costs.
- ❑ Indirect: service shortcomings, security, system down time or system availability.

Notes:

- According to the Robert Frances Group, cost can be over \$1Million per 1 hour of system down time ("Total Cost of Ownership for Linux in the Enterprise", Robert Frances Group, July 2002).
- See also SAP R/3 ITG report on availability.
- Over 32% of Fortune 500 companies spent over \$500K on resolving IT security issues. As reported in USA Today, November 5, 2003. From "Demonstrate Better TCO with IBM Solutions" IBM Software Group presentation.

Building the Business Case for V8: Justifying the Upgrade: Caveats

While these scenarios use examples of real hardware and software costs, they are by necessity extreme generalizations based on a number of factors and assumptions that vary significantly from customer to customer. However, the models presented should be of use to many customers in providing an outline of the potential financial position.

In other words, your mileage may vary – a lot.

The caveats are hard to overemphasize. We'll show some examples, but customer scenarios are very different, depending upon the current portfolio of hardware and software, current pricing, needs for growth and individual license agreements. Some customers will save more CPU than others. Avoiding an upgrade or capacity limit makes the business case simple. Reducing the CPU time depends upon the licensing technique.



Sample Customer Environment from V8

Infrastructure: eServer z990 (330 MSUs) z/OS 1.4

Annual Ongoing Costs

- Total operating costs: \$5.6 million
- Includes hardware (system and storage), operating system, middleware and ISV software

→ z9 10% more power for same software cost

Added zIIP engines, no increase in software

→ z10 10% more than z9 power, same software cost

On zIIP engines 40% - 80% more than z9 power, same cost

For the purposes of our examples, our customer is running an eServer z990[1] rated at approximately 330 MSUs[2] (e.g. a Model 2064-2C8) running z/OS 1.4. The total operating cost for this system (including hardware, operating system, middleware and ISV software) is \$5,573,000 per year.

This customer is currently paying \$410,000 per annum in PSLC[3] license fees for DB2 Version 7 (included in the OWNERSHIP COST figure above), and will pay an additional \$61,000 per annum following the upgrade to Version 8.

In addition, due to the overheads associated with the move to a 64-bit environment, this customer will experience approximately a 7% increase in CPU requirements, amounting to an additional annual cost of \$387,000 per year[4].

Staff costs for the upgrade are estimated as follows: Planning (per customer) – 20 days systems programmer effort, 20 days DBA effort, 5 days application developer effort. Implementation (per DB2 system) – 5 days systems programmer effort, 10 days DBA effort, 10 days developer effort. This customer uses a nominal internal cost of \$60 per hour for all staff when calculating project costs and benefits.

Notes: [1] Customers intending to move to a z990 platform as part of the migration should consider the additional impact of the “technology dividend” announced by IBM in September 2003. This benefit has **not** been included in the scenarios that follow, as it has been assumed that the sample customers are already on the z990 platform prior to the Version 8 migration.

[2] MSU - Million Service Units. Performance/capacity ratings assigned to IBM S/390 and zSeries systems. [3] PSLC – Parallel Sysplex License Charge.

[4] Calculated by taking 7% of the system total ownership cost (\$5,573,000) with the disk storage costs (\$46,000) subtracted, as they are not impacted by move to Version 8.

Sample Customer Environment

Infrastructure: System z10 (600 MSUs) z/OS 1.10

Annual Ongoing Costs

- Total operating costs: \$7.5 million
- Includes hardware (system and storage), operating system, middleware and ISV software

→ z196 AWLC 10% more than z10 power for same software cost or reduce cost for same power, even more at high end

zIIP 40% more than z10 power, same cost

Move 8 or 9 → 10 & reduce MSUs: under 1% of ongoing costs

Net: Stack price and total ongoing costs reduced 2.5%

A medium sized customer is now about 600 MSUs, where the MSU represents more MIPS than for earlier generations of processors.

The z196 uses a new workload pricing, with a new curve and several new tiers. This will deliver a similar improvement to technology dividends, but more for larger numbers of MSUs. The zIIP engines continue to increase in power, with constant or declining prices. Moving to DB2 10 will result in the ability to reduce the MSUs in the LPAR, and that reduction for the software stack can be more than the increase in the DB2 price.

Sample Improvements

Run time CPU reductions	5% - 10%	
1 MB page size	0% - 5%	z10, z196
Page fix buffers	0% - 8%	V8 & high IO, in use?
Release deallocate	0% - 15%	short trans, batch
Virtual storage constraints	0% - 5%	memory, latches
Data sharing fewer members	1%	for each 2 members
Improved dynamic SQL cache	0% - 20%	literals
Insert	0% - 40%	high volume insert
Predicate evaluation	0% - 60%	complex predicates
Access: hash, index include	0% - 5%	access improved
Increased use of zIIP	0% - 3%	IO, RUNSTATS, parallel
Utilities (from V8)	3% - 20%	about same for 9 → 10
Productivity: memory, temporal, security, admin, ...		priceless

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

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

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

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

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

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

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

Migration path options:



- DB2 9 → DB2 10
- DB2 V8 → DB2 9 → DB2 10
- DB2 V8 → DB2 10

This presentation will include information for different migration paths to DB2 10, with the primary paths as shown.

In an August 2010 web conference, customers were asked how they expected to migrate to DB2 10.

10% responded moving to DB2 10 as soon as possible

46% responded DB2 9 → DB2 10 when convenient

22% responded DB2 V8 → DB2 9

19% responded DB2 V8 → DB2 10

4% responded DB2 V7 or older

Earlier in the webcast, customers were asked their current release

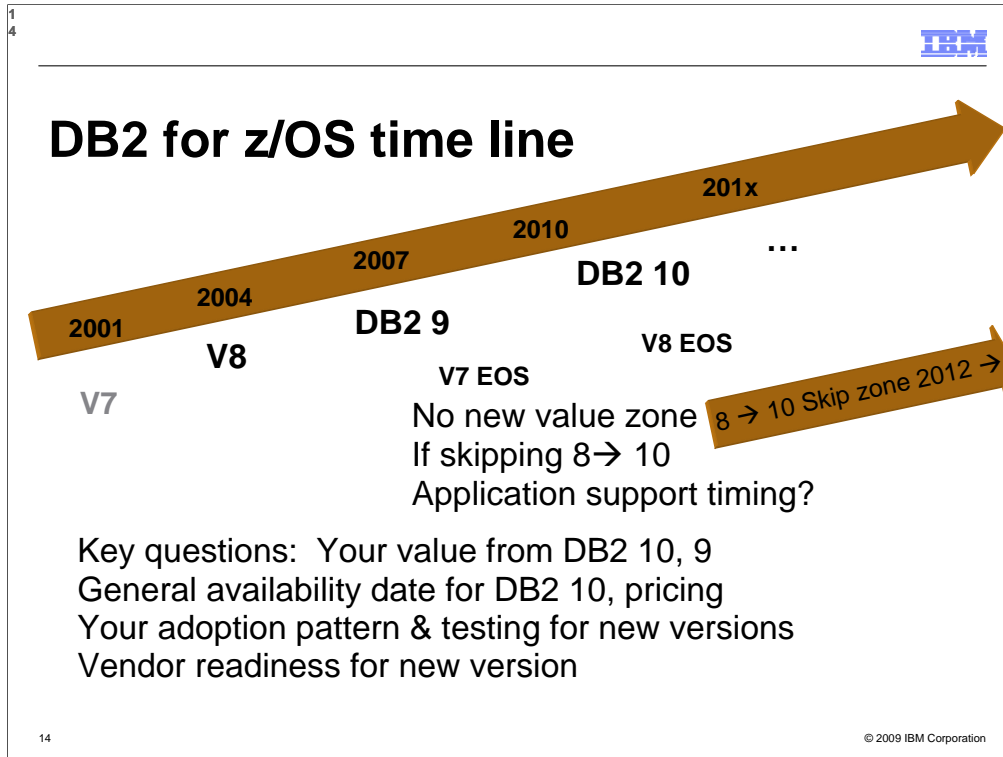
4% indicated moving to DB2 10 soon or early

38% indicated DB2 9 NFM

9% indicated DB2 9 CM or ENFM

47% indicated DB2 V8

2% indicated DB2 V7 or older



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

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

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

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

DB2 10 for z/OS: Skip-Level Migration

May move from V8 to DB2 10,

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

Migration, fallback and data sharing coexistence fully supported

Mix of DB2 9 and 10 or DB2 V8 and 10

Key considerations:



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


DB2 10 supports migration from DB2 9 NFM or from V8 NFM. Customers not yet running V8 or DB2 9 should plan to migrate to DB2 for z/OS V8 NFM as preparation for a migration to DB2 10. We estimate that about one customer in five migrated using a skip version technique for V5 to V7, and we'll see a similar fraction this time. The key value for skip customers on V7 or new on V8. When should I migrate to DB2 10? While DB2 10 is expected to be better than prior versions, it will have maturity and service delivery like other software, with more defects at first fewer as the software matures. Determining when the software is ready for a specific customer and when the customer is ready for the software depends upon the specific customer resources for testing, prior experience, and the value for the improvements versus the need for stability. Many customers depend upon tools or other software, and having that software work with DB2 is a prerequisite. When this information is known, we can answer the question.

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

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

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

Customer migration paths:



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

Estimate save 20% to 25% vs 2 steps

Reasonable timing to skip

Value in DB2 9 deferred until then

When will applications add support for 10? Drop 8?


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

Do you want to migrate from V8 to DB2 10? The first question is the objective in terms of time you expect to save. The next questions are timing and service. Do you want support? DB2 10 will probably become generally available in 2010. When can you migrate to it? DB2 V8 end of marketing is Sept 2009. Migrating to DB2 9 is faster, easier and smoother now.

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

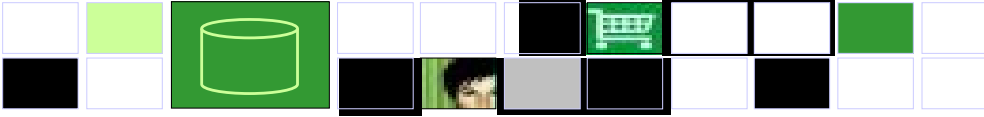
The savings for skipping a version migration are less than 50%, since the education and needed application and administration changes are about the same. Customers who do skip migration reported that the project takes longer, taking about 50% more work than a normal migration path. Changing from DB2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

1
7



Why Migrate to DB2 9 for z/OS?

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



17

© 2009 IBM Corporation

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

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

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

If you are not on DB2 9, look what you are missing.		
Availability, Scale & Resiliency	More Online Schema Changes ➔	Business Resiliency
	Volume Level Backup & Recovery ➔	Faster, Cheaper, Granular Recovery
Compliance & Security	Database Roles & Trusted Context ➔	Efficient Auditing & Compliance
Easier Application Development	PureXML ➔ ➔	Streamlined Data Integration
OLTP & Warehousing	Query Optimization ➔	Highly Available, Secure Data
Reduce Cost of Ownership & Skills	Index Compression ➔ ➔ Utility CPU	Save 50% Disk Save 20% - 30% utility CPU

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

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

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

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

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

Are you ready for DB2 10?

- Check prerequisites
- Migration planning workshop
- Plan gains, testing, memory, and performance
- Build detailed migration plan
- Contact vendors
- Get all the parts out of the box
- Check information APAR
- Apply required service
- Run premigration checks DSNTIJPA (or M) early and often
- Resolve incompatible changes
- Get rid of private protocol
- Convert to packages from DBRMs in plans
- Upgrade plan table formats to Unicode V8 or DB2 9 level
- Get ready for SMS
- Save performance and access path information

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

See the checklist in the DB2 10 Upgrade paper.

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

Free Migration Planning Workshops DB2 10, 9, 8



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. Customers leave the session with materials they can use to start their installation / migration immediately, or in the future. Their questions are given attention, typically in a small group setting. While most MPWs are small, multi-company events, they can be delivered for individual companies or even user groups.

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

Customers should contact their IBM representative for more information about the Migration Planning Workshop.

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

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

No longer supported in DB2 10 from DB2 9:

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

21

© 2009 IBM Corporation

No longer supported, moving from DB2 9

Private protocol: Convert to DRDA.

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

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

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

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

DB2 MQ XML functions: Use pureXML functions.

DB2 XML Extender: Use pureXML.

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

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

Several subsystem parameters

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

No longer supported in DB2 10 from V8:

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

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

Private protocol: Convert to DRDA.

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

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

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

Compression for table space SPT01 is not supported.

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

DB2 MQ XML functions: Use pureXML functions.

DB2 XML Extender: Use pureXML.

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

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

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

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

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

No longer supported ...:

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

Deprecated, may be removed from later versions:

- Simple table spaces
- Partitioned table spaces other than universal table spaces
 - Need to specify SEGSIZE(0) to get classic partitioned
- Default for new table spaces is now universal table space
- DSNHDECP load
- NEWFUN(YES) and NEWFUN(NO)
- DSNHPC7 precompiler
- EXPLAIN stored procedures
- All of these lists are much longer, see Install book or Information Center for details.

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

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

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

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

The DSNHPC7 precompiler is deprecated. Use the current precompiler or coprocessor.

Key details about DB2 10: getting ready

Prerequisites: migrate from DB2 9 for z/OS or DB2 for z/OS V8
z/OS V1.10 SMS-managed DB2-managed DB2 catalog
System zEnterprise 196, z10, z9, z890, z990, above (no z800, z900)
DB2 Connect 9 FP1, 9.7 FP3a for many 10 functions
IMS 10 & 11 (not 9) CICS compilers (See announcement)
SPE PK56922 PK61766 PK85956 PM04680 PK87280 PK87281 PM08102
PM10227 PM12572 PM14377 PM09148 PM19523 PM23243
(9 PK69411 PM08105 PM13631 PM13525 PK31163 PM19221 PM19841
PM19844 PM21557) (8 PM14630)
Premigration check DSNTIJPA PM04968 PM15965
Info APARs II14477 (DB2 9) II14474 (V8)
Items deprecated in earlier versions eliminated: more for V8 mig.
Private protocol → DRDA (DSNTP2DP, PK92339, PK64045)
Old plans and packages V5 or before → REBIND
Plans containing DBRMs → packages PK62876 (9), PK79925 (8)
PK85833 (9), PM01821 (All)
ACQUIRE(ALLOCATE) → ACQUIRE(USE)
Old plan table formats → DB2 V8 or 9, Unicode, 59 cols PK85068
BookManager use for DB2 publications → Info Center, pdf

© 2009 IBM Corporation

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

Note the new premigration job DSNTIJPA provided in APARs PM04968, PM15965, and other service to help discover some of the function removed in DB2 10.

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 10 News book: DB2 10 Performance Topics coming

- MODIFY DDF ALIAS (PM26480, PM26781 ptf)
- APREUSE, APCOMPARE (PM25679, PM24937)
- Delete data sharing member (PM31009, PM31003, PM31004, PM31007)
- Enhancements for new DBA authorities (PM28296, PM26977)
- Compression, BIND performance, Inline LOBs for SPT01 (PM24721, PM27811, PM27073, PM27973)
- Online REORG concurrency materializing deferred ALTERs (PM25648)
- Temporal enhancements
 - TIMESTAMP WITH TIMEZONE support (PM31314)
 - Enhancement for data replication (PM31315)
 - ALTER ADD COLUMN, propagate to history table (PM31313)
- New system profile filters based on “client info” fields (PM28500)
- New zparm force deletion of CF structures on group restart (PM28925)
- Relief for incompatible change in CHAR of decimal data PM29124

DBA authority enhancements: a) prevent users with TRACE privilege to stop audit traces, b) separate implicit system privileges (e.g. stogroup, bufferpool) from System DBADM, c) enhance IFCID 361 to trace all sysadm actions.

Additional APARs

New IFCID 225 real storage values with z/OS support (PM24723)
 1MB page – z/OS (OA31116, OA33702, OA33529)
 Getpage increase using subquery (PM25357 UK63087)
 Improve DB2 catalog perf for GRANT stored procedures (PM28796)
 UTS update on small record goes overflow unnecessary (PM27828)
 Large Notify Support - PM25271, PM25282
 JDBC Type 4 with Java Stored Procs – PM28100
 Built-in Function enhancements – PM29900, PM29901
 Start/maintain minimum # WLM standard procs – PM27097
 DSNACCOX to consider HASH for Reorg - PM25652
 Reorg shrlvl change for increasing inline LOB length – PM29032
 Recover Backout (YES) with mass delete/truncate – PM30991
 SMF decompression job for the SMF compression – PM27872
 Add restrictions for extended implicit cast for set operators – PM28458
 HASH LOAD performance – PM31214
 DB2 9 SPEs (open as of Jan. 28, 2011)
 Implicit auto regeneration for native SQL procedures – PM13525
 Implicit auto regeneration for index on expression – PM13631

PM13525 Implicit auto regeneration for native SQL procedures

Until this change is implemented, probably in March 2011, native SQL procedures created or regenerated

in DB2 10 CM9 can't run in DB2 9 - either coexistence or fallback. Avoid creating or regenerating

procedures until you won't fallback or find and regenerate these procedures when falling back to DB2 9.

PM13631 Implicit auto regeneration for index on expression

Until this change is implemented, probably in April 2011, indexes on expressions created or regenerated

in DB2 10 CM9 can't run in DB2 9 - either coexistence or fallback. Avoid creating or regenerating indexes

on expressions until you won't fallback or find and regenerate these indexes when falling back to DB2 9.

Several open APARs in SPE for DB2 9 migration as of January 18, 2011

PM13525 Implicit auto regeneration for native SQL procedures

Until this change is implemented, probably in March 2011, native SQL procedures created or regenerated in DB2 10 CM9 can't run in DB2 9 - either coexistence or fallback. Avoid creating or regenerating procedures until you won't fallback or find and regenerate these procedures when falling back to DB2 9.

PM13631 Implicit auto regeneration for index on expression

Until this change is implemented, probably in April 2011, indexes on expressions created or regenerated in DB2 10 CM9 can't run in DB2 9 - either coexistence or fallback. Avoid creating or regenerating indexes on expressions until you won't fallback or find and regenerate these indexes when falling back to DB2 9.

Convert plans with DBRMs to packages

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



PK62876 PK79925 (V8) BIND changes for eliminating acquire(allocate), plans with DBRMs

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

Convert private protocol to DRDA

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

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



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

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

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

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

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

SMS now a prerequisite for DB2 catalog & directory



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

All new indexes & new table spaces in the catalog and directory created SMS-controlled with extended addressability (EA & EF)
DSNTIJSS provides SMS classes for customers without SMS in use

- Environment ONLY for DB2 Catalog & Directory data sets
 - Other DB2 data sets such as logs and BSDS not covered.

Useful SMS reference material includes:

- z/OS DFSMS Introduction.
- z/OS DFSMS Implementing System Managed Storage.
- z/OS DFSMS Storage Administration Reference.
- IBM Redbooks publication: Maintaining your SMS environment



SMS is now a prerequisite of DB2 10.

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

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

Useful SMS reference material includes:

- z/OS DFSMS Introduction.
- z/OS DFSMS Implementing System Managed Storage.
- z/OS DFSMS Storage Administration Reference.
- IBM Redbooks publication: Maintaining your SMS environment.

Package or Access Path Stability

PLANMGMT default changed to EXTENDED

- Retains 3 versions: Current, Previous, Original
- Ability to SWITCH to another version or fallback
- Catalog information provided, new structures
- BIND costs in CPU, disk space, and memory
- If cost is too high, change default, use selectively

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 of the REBIND PACKAGE command.

2. A new system parameter called PLANMGMT for specifying the default setting of PLANMGMT option of REBIND PACKAGE. Possible settings are: OFF, BASIC and EXTENDED. The default value of this parameter is OFF. To use a setting other than OFF, update your DB2 9 subsystem parameter (DSNZPxxx) modules, then either use the SET SYSPARM command or stop and start DB2 for the change to take effect.

3. A new REBIND PACKAGE option called SWITCH to allow users to revert a package to an older copy, effectively undoing a prior REBIND invocation.

SWITCH(PREVIOUS) - The PREVIOUS copy is activated

SWITCH(ORIGINAL) - The ORIGINAL copy is activated

4. A new FREE PACKAGE option called PLANMGMTSCOPE to allow users to free older copies that are no longer necessary.

PLANMGMTSCOPE(ALL) - Free the entire package including all copies. This is the default.

PLANMGMTSCOPE(INACTIVE) - Free all old copies

Migration Staging


Migration 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

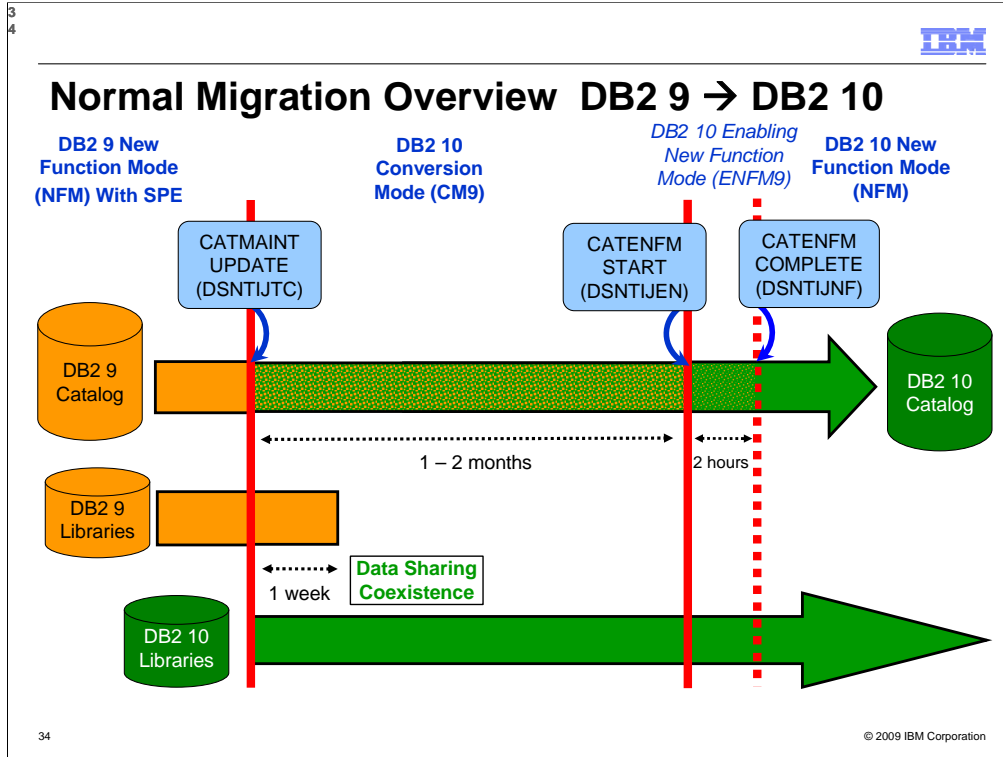
Implementation

- Using new improvements to deliver value



The primary stages of migration to a new version are planning, migration and implementation of the new improvements. This presentation will focus on the planning.

The key items are making sure that everyone is educated, with what they need to succeed. A detailed project plan and communication to all involved are crucial for success. Some of the preparation can occur very early, understanding, obtaining and installing the prerequisites. Applications can make needed changes over a longer period to make the process easier and less costly. Plans for monitoring virtual and real storage as well as performance are necessary. An early health check, communication of the needed changes and staging of the work will make the project more successful.



Migration from DB2 9 to DB2 10

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

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

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

Some migration considerations are:

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

Note – this is not necessarily to scale!

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

Overview of Modes when migrating 9 → 10

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

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

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

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

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

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

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

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

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

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

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

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

Migration job DSNTIJEN

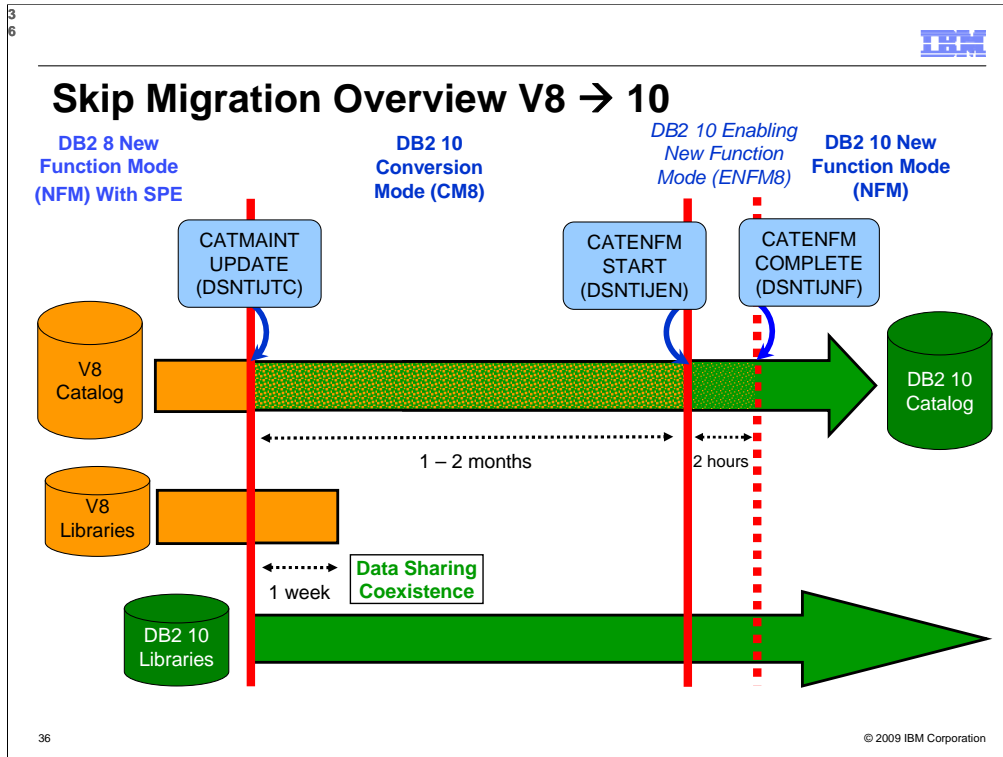
- If DB2 went from NFM or ENFM9* to CM9*, DSNTIJEN takes it to ENFM9*.
- If DB2 went from ENFM9 to CM9*, it completes ENFM9 processing and leaves DB2 in ENFM9.

Migration job DSNTIJNF

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

Migration job DSNTIJES

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



Migration from V8 direct to 10

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

DB2 will support migrating from V8 NFM to DB2 10 without ever starting the system in DB2 9.

When a system is migrating from V8 NFM to DB2 10 conversion mode the possible DB2 10 modes are:

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

ENFM8 Enabling New Function Mode on a system that migrated from V8 NFM to DB2 10.

Once this mode has been entered the system can not fallback to V8 and a V8 member can not be started in a data sharing group.

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

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

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

Some migration considerations are:

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

Some differences in the C8 and C9 catalog are:

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

Note – this is not necessarily to scale!

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

Overview of Modes when migrating V8 → 10

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

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

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

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

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

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

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

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

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

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

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

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

Migration job DSNTIJEN

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

Migration job DSNTIJNF

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

Migration job DSNTIJES

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

Performance enhancements requiring few changes (CM)

SQL runtime improved efficiency

Address space, memory changes to 64 bit, some REBINDs

Faster single row retrievals via open / fetch / close chaining

Distributed thread reuse High Performance DBATs

DB2 9 utility enhancements in CM8

Parallel index update at insert

Workfile in-memory enhancements

Index list prefetch

Solid State Disk use

Buffer pool enhancements

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

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

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

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

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

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

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

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

Performance enhancements requiring REBIND (CM)

Most access path enhancements

Further SQL runtime improvements

Use of RELEASE(DEALLOCATE)

SQL paging performance enhancements

- Single index access for complex OR predicates:

IN list performance

- Optimized Stage1 processing (single or multiple IN lists)
- Matching index scan on multiple IN lists

Safe query optimization

Query parallelism improvements

More stage 2 predicates can be pushed down to stage 1

More aggressive merge of views and table expressions

- Avoid materialization of views

If migrating from V8, get new RUNSTATS before mass rebind

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

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

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

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

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

Performance enhancements requiring NFM

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


Some of the performance improvements require new function mode and some work by database administrators to tune the database design and often to REBIND. Efficient caching for literals needs a rebind in NFM. One of the enhancements that will help everyone with large dynamic SQL applications are the improved processes that handle the Dynamic Statement Cache. They now consolidate SQL statements that are the same but have different literals. This reduces the SQL statement space used in the Dynamic Statement Cache and reuses the security and object verification and access path already developed for the SQL statement. This dramatically improves Dynamic Statement Cache cache hits, reduces the duplicate SQL statement previously held, makes room for more SQL improving the overall workload performance.

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

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

Small WORKFILES are now available for simple predicate evaluation for improving performance

DB2 provides native support for the SQL procedural language eliminating



Performance enhancements which need NFM + DBA work

Hash access path	Create + Reorg + rebind to activate
Index include columns	Alter + Rebuild + rebind to activate
Inline LOBs	Alter (need universal table space and reordered row format)

DEFINE NO for LOB and XML columns

MEMBER CLUSTER for universal table space Alter + Reorg

Alter to universal table space, page size, data set size, segment size Alter + Reorg

Online reorg for all catalog and directory table spaces

41
© 2009 IBM Corporation

Some of the performance improvements require new function mode and some work by database administrators to tune the database design and often to REBIND. The DBAs can improve on their direct keys access and implement the new Hash access table space type reducing a key access to potentially a single I/O. This is great for the customer or product id lookouts that are done millions or billions of times a day.

The Index include non-unique columns within a UNIQUE index definition provides great way to eliminate indexes and consolidate other indexes for better optimization and overall performance. This is great also for better access paths because all the columns are now in the single index for better access paths and fewer indexes to reference.

Inline LOBs use the reorder row format and handle the LOB better for overall streaming and application performance. This along with the DEFINE NO feature allows the row to be used and the data set for the LOB not to be defined. The application can still reference the row and get the proper results. The system doesn't define the LOB data set until the LOB is saved which can 41

Improvements to installation & migration information based on customer feedback

Pre-migration and migration checklists were added to help plan for and keep track of the migration process. The checklists include links to the pre-migration and migration steps. Access the checklists from the following locations:

- Introduction to migration from DB2 Version 8:

http://publib.boulder.ibm.com/infocenter/ecsimzic/v1r0/topic/com.ibm.db2z10.doc.inst/db2z_intro2migfromv8.htm

- Introduction to migration from DB2 Version 9.1:

http://publib.boulder.ibm.com/infocenter/ecsimzic/v1r0/topic/com.ibm.db2z10.doc.inst/db2z_intro2migfromv9.htm

Installation and migration steps clearly define tasks to complete. Related concept and reference information are included as links from the tasks.


Several new jobs simplify the setup and installation of DB2-supplied stored procedures and user-defined functions. The process is documented in [Installation step 20: Set up DB2-supplied routines \(optional\)](#) and [Migration step 25: Set up DB2-supplied routines \(optional\)](#).

For the V10 development cycle, the information development team focused on addressing solutions to top customer pain points. Based on customer feedback, one of the areas we focused on improving is the installation and migration information.

Our team will be asking you specifically about your migration experience with V10 to gather feedback on the success of these improvements.

Update installation CLIST input (DSNTIDxx) to reflect current subsystem parameters & bufferpool settings

DSNTXAZP (job DSNTIJXZ)



CLIST INPUT MEMBER GENERATION REPORT ** 2010-05-21 06:28:18

CLIST PARAMETER REPORT:

```

0001 PARAMETER NAME           = ABEXP
    ZPARAM/BUFFERPOOL PARAMETER = ABEXP
    PARAMETER TYPE              = CHAR
    DATA SHARING SCOPE         = M
    MINIMUM VALUE                = NO
    MAXIMUM VALUE                = YES
    CURRENT CLIST VALUE          = YES
    CURRENT INSTALLED VALUE      = YES
    STATUS                       = RETAINED
-----
0030 PARAMETER NAME           = AUDIT
    ZPARAM/BUFFERPOOL PARAMETER = AUDITST
    PARAMETER TYPE              = CHAR
    DATA SHARING SCOPE         = M
    MINIMUM VALUE                = NONE
    MAXIMUM VALUE                = NONE
    CURRENT CLIST VALUE          = YES
    CURRENT INSTALLED VALUE      = 1
    STATUS                       = UPDATED
  
```

CHANGE SUMMARY REPORT:

CLIST PARAMETER NAME	ZPARAM/BUFFERPOOL NAME	VALUE
ABEXP	-SAME-	YES
ASSIST	-SAME-	NO
AUDIT	AUDITST	1 (YES)

Stored procedure information improvements

Stored procedure information easier to find

- Implementing stored procedures
- DB2 reference information added

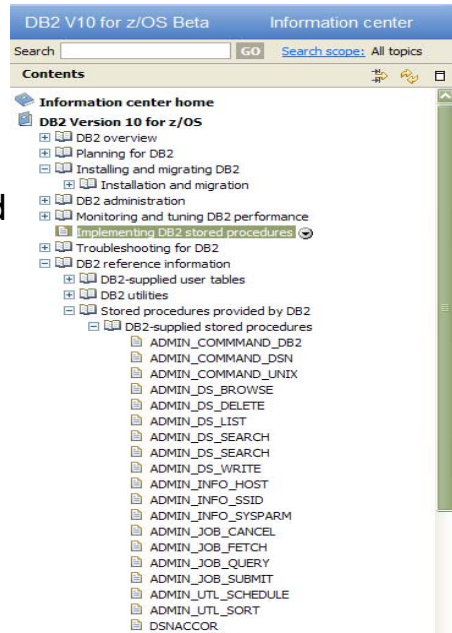
Added more information:

Setting up WLM

Granting privileges

Moving stored procedures from test to production.

Clearer information



For the V10 development cycle, the information development team focused on addressing solutions to top customer pain points. Based on customer feedback, another area we focused on improving is the stored procedure information.

We made the stored procedure information easier to find.

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

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

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

We also rewrote information to make it clearer what you need to do to accomplish a particular goal.

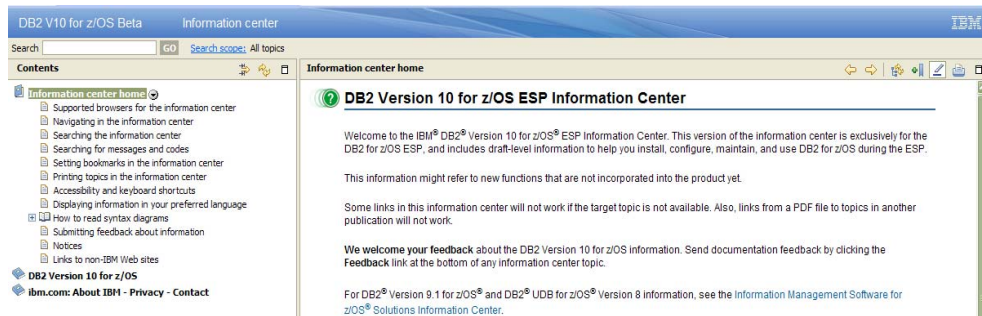
Important features of information center

Find helpful usage instructions in the “Information center home” section.

Easily send feedback by clicking the **Feedback** link at the bottom of any topic.

Conveniently download the PDF version of the information from the link at the bottom of any topic. *(Look for the PDF icon!)*

Efficiently search for the information you need by using the search features.



This information center might refer to new functions that are not incorporated into the product yet.

Some links in this information center will not work if the target topic is not available. Also, links from a PDF file to topics in another publication will not work.

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

Why Migrate to DB2 9 for z/OS?

Business needs to save money


- Reduce CPU time & disk space
- Improve business agility
- Service Oriented Architecture

Application developers need

- PureXML for a powerful SQL and XML interface to XML data
- Powerful new SQL enhancements
- Portability with SQL and data definition compatibility

Database Administrators need

- Improve availability and performance
- More flexible security and easier regulatory compliance
- Better web application & data warehouse function and performance
- LOB function, performance, usability



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

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

Almost every tool from vendors is ready today.

DB2 9 for z/OS field update 2010 December

DB2 9: Continues to climb

- **V8 End of service April 2012, announced August 2010**
- Most large customers
- V8 Withdrawal from marketing Sept. 2009

DB2 V8: Migration finished in most of world

- 100% of Top 100 Over 99% of Top 200
- V7 End of Service: June 30, 2008

How's the Quality (compared to V8)?

- Lower overall PMR volume
- Less Severity 1 APARs
- Lower PE rate



Now that DB2 9 has been generally available in the field for over 3 years, 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. The largest DB2 customers have moved to Version 9. All of the largest 100 customers are running Version 8, and many large customers are implementing DB2 9 development, moving to production. 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 has been out of service for about 2 years after seven years of general availability. If you are still there, it is time to move now. V8 was withdrawn from marketing in 2009.

Best Practices/Recommendations

Start with latest RSU + Identified Hipers

Leverage CST/RSU process: DB2 9 & 10

- Apply 2 to 3 preventative service drops annually
- Exploit Enhanced HOLDDATA to be vigilant on HIPERs and PEs

Use the DB2 9 'Package Stability' function for static SQL

- Offers access path preserving option. Recovers to prior access path if regression is encountered
 - PK52523 - DB2 Access Path Stability

Minimize potential query performance issues

- Use Optimization Service Center to capture SQL statements
- Run Stats Advisor to generate the recommendation for stats collection
- Run RUNSTATS to ensure critical stats are collected as recommended by the advisor

Ensure a PMR is opened prior to migration start

50

© 2009 IBM Corporation

As you migrate to DB2 9, some of the best practices and recommendations can help you have a better experience. The CST process is being used successfully by most customers. The current service level recommendation is to get to RSU0903 and some specific identified HIPER APARs as a minimum level.

The DB2 package stability or access path stability function has been very helpful with customers who have access path regression in getting back to the old access paths. This process is noted in John Campbell's presentation and many others.

The best techniques for minimizing potential access path issues are to

Use the Optimization Service Center (OSC) to capture the SQL statements and related information. You can start using OSC on V8 to get the needed information and to help tune your queries.

Run the statistics adviser to generate the recommended statistics to be collected.

Get current RUNSTATS data, so that the DB2 9 optimizer is using DB2 9 statistics. Get the statistics noted in the item above.

Work with your DB2 advisor to get current information and help. The Migration Project Workshop is a proven technique for transferring skill. Get other presentations from the web and from conferences to help with knowledge. Terry Purcell's presentations are among the best for access path information.

<http://www.ibm.com/support/docview.wss?rs=865&context=SSEPEK&uid=swg27015988>

So where do I go from here?

No need to fear DB2 migration; but be cautious & thoughtful

- Ensure good planning, staying reasonably current on preventative service and application regression testing
- Follow best practices/recommendations
- PMR opened
- Trends are improving!

Procedural upgrade easier

Every customer experience is different Compared to V8?

- Lower overall PMR volume
- Less Severity 1 APARs
- Lower PE rate
- Lower HIPER rate



The majority of customers are current today, and use of the older versions is dropping sharply. Running on supported levels of software is less expensive and less risky. There is a high degree of geographic skew in the number. The majority of customers are running DB2 V8, but customers are moving to DB2 9 in substantial numbers. One third of the largest 100 customers are using DB2 9. About 27% of the largest 200 customers license DB2 9.

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


DB2 V7 is now out of service after seven years of general availability. If you are still there, it is time to move now, and DB2 V8 should be the target.

If customers are interested in seeing experiences from other customers, one of the best opportunities is John Campbell's webcast. The slides and audio are on the web.

<ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/migration/v9-early-migration-experience-iod-2008-campbell.pdf>
[v9-early-migration-experience-iod-2008-campbell.mp3](ftp://ftp.software.ibm.com/software/data/db2/zos/presentations/migration/v9-early-migration-experience-iod-2008-campbell.mp3)

Most consumable DB2 9 improvements

- CM very little to no action:
 - Optimization Service Center, Data Studio and Optim Query Tuner
 - Utility CPU reductions
 - Archive log striping, compression
 - Larger prefetch, write & preformat quantities
 - LOB performance
 - DDF VSCR
 - Index lookaside
 - Changed online REORG
- Package stability
- Improved RUNSTATS
- Optimization improvements, EDMPOOL VSCR, increased parallel & zIIP
- NFM
- LOB lock avoidance
- Logging for data sharing
- Improved index leaf page split
- Reordered row format, native SQL
- Index: larger page sizes, compression, index on expression



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.

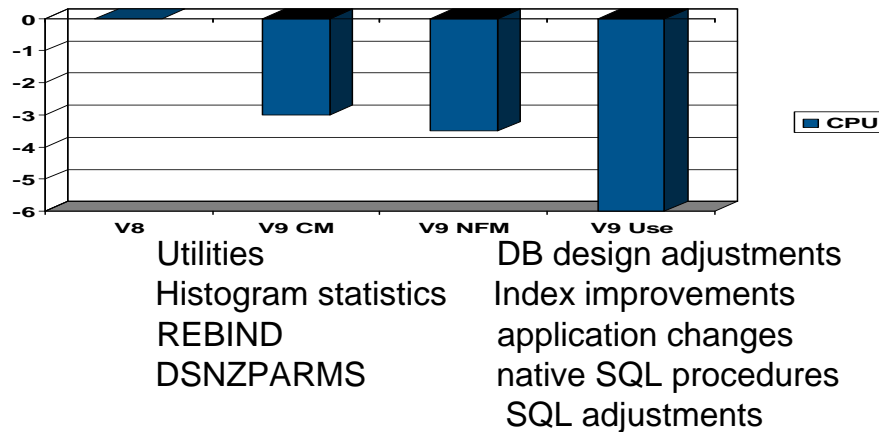
Optimization improvements are automatic for dynamic SQL, but require work to REBIND for static SQL. In both cases, we need baselines to check for regression. REOPT(AUTO) for dynamic SQL needs analysis to be sure the improvement is working. EDMPOOL virtual storage constraint relief also requires a REBIND.

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.

LOB lock avoidance requires a quiesce of all subsystems in NFM until APAR PK62027. Reordered row format requires a REORG in NFM and varying length columns, and use with small columns can reduce compression. See APARs.

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.

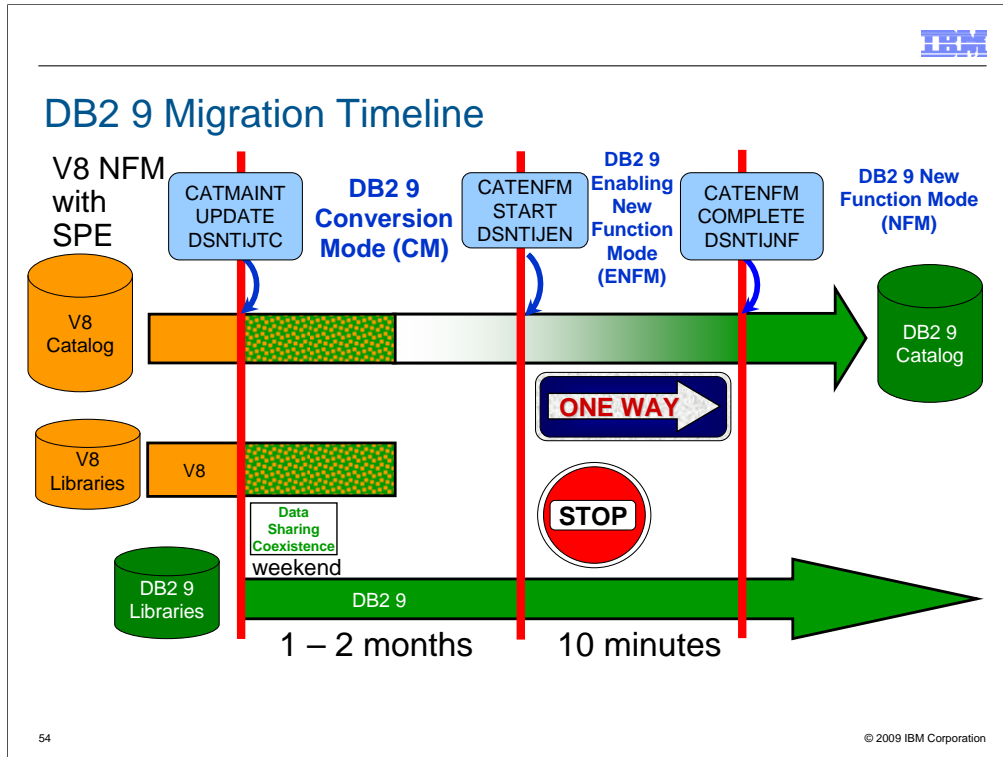
DB2 9 zEnterprise, z10, z9, z890 & z990 performance scenario



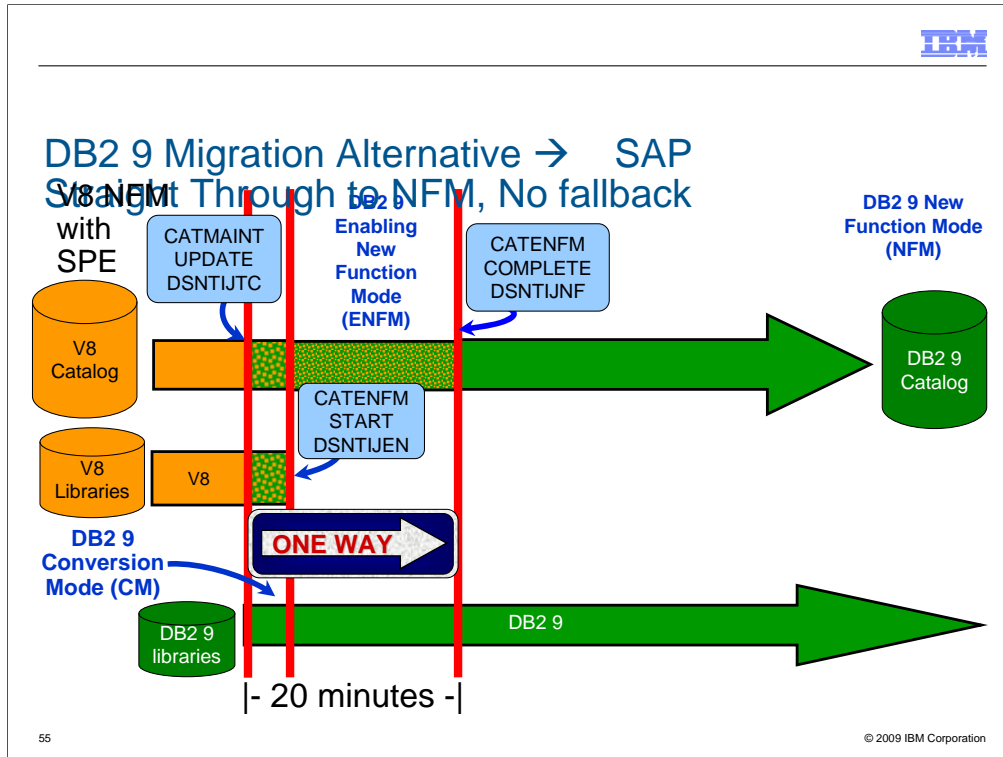
Your situation will vary. Less CPU is better.
z800 and z900 expect +5% to +10% CPU

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%.



The time line is not to scale. Customers should spend only a few minutes in ENFM and then move immediately to NFM. Most customers spent less than three months in CM on DB2 9. 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, such as those who use SAP, 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. This process is like cooking Chinese cooking, three minutes to cook, after three hours of preparation. If you do months of planning and testing, ensure that you don't need to fall back, then the steps can be compressed from two to one.

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.

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.

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 improvement BSAM instead of BDAM
- Index lookaside
- DSN1COMP index compression estimate
- 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)
- Optimization Service Center & Optim Query Tuner & Optim Query Workload Tuner (work with V8 if PTFs are applied)
- Most utility enhancements and CPU reductions (see next slide)

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
- DSN1COMP index compression estimate
- Larger prefetch, write and preformat quantities
- LOB performance other than LOB lock avoidance
- Commands: 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)
- 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 Utility Modes: Conversion vs New Function

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

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.

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




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.

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 **II14401 II14464 II14477** 9→10, **II14474** 8→ 10



- ◆ Other software ...Talk to your Vendors

- ◆ z/OS R7 R8 R9 R10 R11 R12



- ◆ Job DSNTIJPM or DSNTIJP9 or DSNTIJPA

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**

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
- Visual Explain → Optimization Service Center, Data Studio
- Old JDBC / SQLJ driver → Current driver
- Java stored procedures → non-resettable JVMs
- Online help → Information Center
- Parameters removed RELCURHL, DBPROTCL, STORPROC, ...

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 (DSNTP2DP, PK92339, PK64045)
- Plans containing DBRMs → packages PK62876 PK79925
- ACQUIRE(ALLOCATE) → ACQUIRE(USE)
- DB2 catalog → DB2 managed, SMS managed, no links, more LOBs, more table spaces, no SPT01 compression
- XML Extender → XML type
- Old plan tables → V8 | 9 format (59 cols, Unicode) PK85068
- 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

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. For plan table format upgrades, use the Optimization Service Center or see the ALTERs in edited sample member DSNTESC.

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 ..PK61841
- SPE PK11129, PK33664, PK39850, PK38799, PK44884
- Info APARs: Install **II14401**, **II14464**, **II14477**, **II14474** DDF **II14203** XML **II14426** Unicode II13048 & II13049, zIIP II14219, DFSORT II14047, II14213 & II13495, Performance II13579, Storage II10817, Parallel II12836, LOBs **II14334**, DSNUTILS sec. id. II13133, Coprocessor II13782, PDSE II14067, II14255, II14403, Inccorout **II14441**, send dumps II11945, ...

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

- BIND ISOLATION defaults to CS. Query catalog 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
- Work files → much more 32K space needed, avoid small files
- Varying length data change row format PK78958, PK78959
- DB2 disables DATA CAPTURE CHANGES for SYSROUTINES during migration (DSNTIJTC) and SYSPACKSTMT during ENFM (DSNTIJEN) PK72214
- Need to find and complete or delete incomplete definitions
- Online REORG processes change
- Configure shared virtual storage (128 GB for each DB2)

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. Define much more 32K space, moving over from 4K.

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.

Old COBOL? Old programs? DSNHPC7

- ✓ 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 V4.2
 - ➔ Use old precompiler?
 - ➔ NEWFUN(V9 or V8)?
- ✓ Plan staging and Migrate

DB2 10 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 V4.2. DB2 10 still includes DSNHPC7 for obsolete compilers, which can help in some cases.

<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/igy3mg50.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? Old programs? DSNHPC7

- ✓ Understand the options
 - ➔ Compiler and Run-Time Migration Guide
 - ➔ Run Old PL/I modules
- ✓ Set strategy & plan for Enterprise PL/I V3.9
 - ✓ Support requires at least V3.7 now
 - ➔ Use V8 precompiler?, NEWFUN(NO)?
- ✓ Plan staging and Migrate

DB2 for z/OS supports only current PL/I compilers. Every compiler prior to V3.7 is out of service in 2010, including OS PL/I V2.3, IBM PL/I for MVS™ & VM V1R1.1, V2, and V3R1 to R6. DB2 10 still includes DSNHPC7 for obsolete compilers, which can help in some cases.

<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>

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 V3.8 or V3.9 now? Do developers use new ones? 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
- DB2 DFSORT → zIIP needs z/OS 1.10 PK85856, V8 or 9 APAR PK85889
- New option for DB2 Sort

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 or PMR 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**

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 and 10 require DB2 client & Connect 9 FP1 Recommend latest level

9.7 fixpack 3a needed for DB2 10 new functions:

- ✓Dynamic statement cache enhancements.
- ✓Timestamp with timezone.
- ✓Greater timestamp precision.
- ✓Client/Server Binary XML Format
- ✓Extended indicator variables
- ✓Explain mode special register
- ✓Unicode collection and package names

<http://www.ibm.com/support/docview.wss?uid=swg24028306>

DB2 Connect V8 out of
service April 30, 2009



72

© 2009 IBM Corporation

The recommended level for client & DB2 Connect is as current as possible to improve availability, resolve problems and deliver the changes needed in Java, stored procedure and web services environments. DB2 10 new function comes in 9.7 fixpack 3a. Client and DB2 Connect V9.7 fixpack 3a are needed for a number of DB2 10 new functions: dynamic statement cache enhancements, timestamp with timezone, greater timestamp precision, client/server binary XML format, extended indicator variables, explain mode special register, and Unicode collection and package names. Get it from <http://www.ibm.com/support/docview.wss?uid=swg24028306>

If you are moving to DB2 9 or DB2 10 for z/OS, then the minimum supported level is DB2 Connect 9 FP1. If you are not at the minimum levels, then it's time to migrate to the most current level you can, DB2 9.7 fixpak 3a. If you get current now, then migration to DB2 9 for z/OS will not encounter the same problem. So if a client is down level, then get it migrated to current.

If you have a specific set of problems, then you can check the fixpak levels to see where those problems are resolved. But then you still need to test to make sure that the problem is resolved without introducing any new ones. Later fixpaks generally resolve a lot more problems. If you are migrating, find a good one and move to that level. Running with code that is out of support or missing a few thousand fixes is problematical. DB2 Connect V7 and V8 are not supported, and generally do not work. We don't generally separate CM from NFM on our discussions of fixpaks. Some new function comes in CM. See the web for fixpak information.

<http://www.ibm.com/support/docview.wss?rs=56&uid=swg27007053>

<http://www.ibm.com/support/docview.wss?rs=71&uid=swg21255572>

<http://www.ibm.com/support/docview.wss?rs=71&uid=swg21256235>



DB2 9 Post GA Deliveries APAR PK.....

PK52522 PK52523 BIND package stability changes
PK62876 PK79925 BIND convert plans with DBRMs to packages
Accessories Suite 1.2 Text search server
PK41001, z/OS 1.8 APAR OA17314 Incremental FlashCopy
PK44617, PK47579 trusted context enhancements
PK41711 New storage class for online CHECK utilities (PPRC)
PK56392 ALTER TABLE ALTER COLUMN DROP DEFAULT
PK51979 Allow RESTORE SYSTEM without log truncation
PK51020, PK54451, ... Spatial phase II
PK63324, PK63325 LOAD COPYDICTIONARY
PK62027 Protocol(3) without data sharing group quiesce
PK62161 Add rows to accounting & statistics
PK75643 parameter OPTIOWGT default to enable
PK62178 default 10,000 maximum implicit DB
PK70269 USS pipes for LOAD
PK49348 REOPT(AUTO) change

73

© 2009 IBM Corporation

DB2 9 has been generally available for over 3 years 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, a change to BIND is shipped which can perform the conversion to packages.

The capability to copy a dictionary from one partition to another using the LOAD utility as added. APAR PK62027 avoids the need to quiesce the data sharing group to get the new LOB lock avoidance.

DB2 9 Post GA Deliveries APAR PK....., PM.....

PK70060 PM02528 PM17336 work file separation for allocation
 PK92339 PK64045 BIND private protocol help
 PK87348 PK85881 ... RRF parameter, utility keyword
 PK80925 -ACCESS command wildcards
 PK79228 PK79327 parameter group attach like V8
 PK72214 install stop change data capture
 PK85068 migrate plan tables to Unicode, V8, 9, format
 PK80375 compress package directory table space SPT01
 PK81151 EAV for sequential data, requires z/OS V1.11
 PK85856 PK85889 utility sort use zIIP
 PK83072 PK99362 64 bit ODBC
 PK75626 enable WLM bufferpool management
 PK91610 precompiler similar to V7
 PK83735 UTS performance
 PM12256 increase zIIP redirect for DRDA to 60%
 PK87762 REORG multiple part ranges

DB2 9 has been generally available for about 3 years 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.

Rather than creating 60,000 databases implicitly, changes in PK62178 allow default to 10,000 and allow you to set the number.

If you want to pipe information to the LOAD utility, look at PK70269.

PK72214 changes install to stop changed data capture for catalog tables as needed, asking you to restart them after switching to DB2 9 and ENFM.

Changes are in process for the combination of compressed table spaces and RRF. Some customers with many short varchars, e.g. varchar(1) found that compression was not effective with RRF.

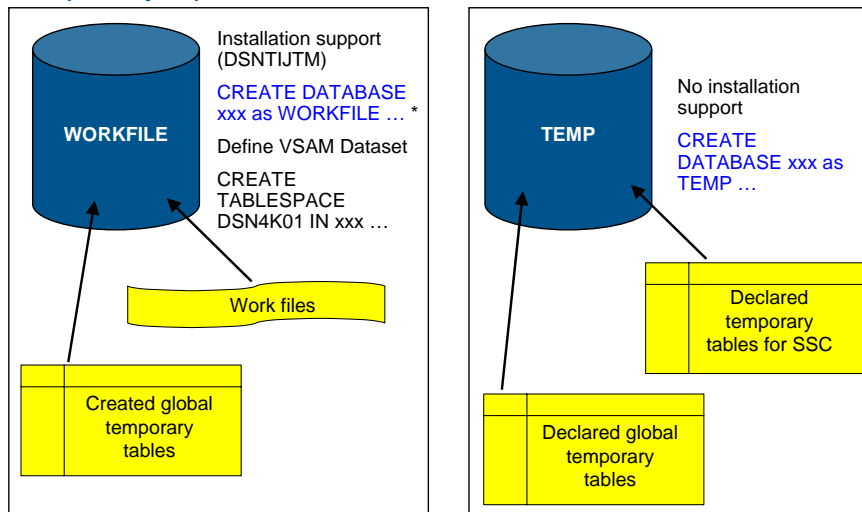
Others need to control conversion to usage of RRF to allow DSN1COPY scenarios to work. An option is provided to control conversion to RRF for data other than XML, but many benefits in DB2 9 and beyond depend upon RRF.

-ACCESS only worked for specific names, but it's being changed to allow wild cards and ranges.

OPTIOWGT was generally recommended, and becomes the default with PK75643.

If you have multiple members of a data sharing group in a single LPAR and prefer the V8 group attach logic, parameters are added to control the random group logic.

Temporary Space DB2 V8 Picture



* Only in a data sharing environment – in non-data sharing syntax is `CREATE DATABASE DSNDBO7`

75

© 2009 IBM Corporation

The work file database is used by SQL statements that require working storage, includes statements that use:

GROUP BY (without index)	IN (subselect)
ORDER BY (without index)	ANY (subselect)
DISTINCT (without index)	SOME (subselect)
UNION (except UNION ALL)	ALL (subselect)
EXISTS (subselect)	Some joins

Created Global Temp Tables

SORTs needed for processing SQL statements

Two databases for temporary files and temporary tables:

WORKFILE database: created Global Temporary Tables

Workfiles for SQL statements requiring sorts, view materialization, triggers, nested table expressions ...

TEMP database

User-defined Declared Global Temporary Tables

DB2-defined Declared Global Temporary Tables (for static scrollable cursors)

WORKFILE database

Pseudo-segmented table space organisation (can't specify SEGSIZE)

4KB and 32KB page size only

Tables can span multiple table spaces (up to 254) with maximum size of 16TB (254 * 64GB)

The number of tables used for DB2 sorts is limited only by buffer pool size.

TEMP database

Segmented table spaces (user-defined SEGSIZE or default of 4)

4KB, 8KB, 16KB and 32KB page sizes

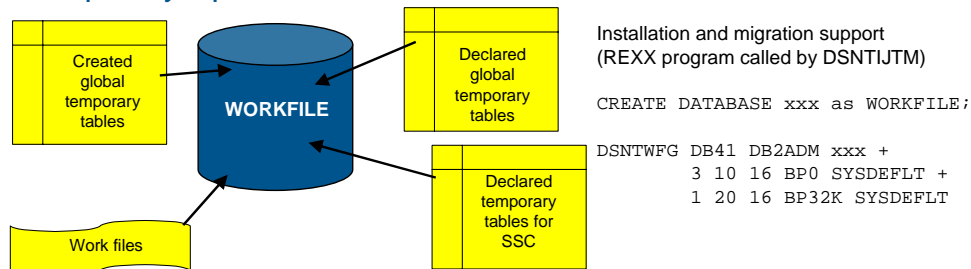
Tables cannot span multiple table spaces – maximum table size is 64GB

Maximum number of tables per DB2 subsystem is limited to < 32K.

The number of work files that can be used for DB2 sorts is limited only by the buffer pool size. The WORKFILE database does not use global DBD management to assign and release table IDs. Each work file has its own local DBD, copied from the global DBD. ('Global' here means global within a single DB2 subsystem.)

The global DBD has only the table space descriptor objects for each table space defined in the database. The work file's local DBD has the table space descriptors and the table descriptor for just that table which is associated with that work file. Since the DBD is not globally shared (within a DB2 subsystem), there is no limit to the number of tables in the WORKFILE database.

Temporary Space DB2 9 Picture



Declared Global Temporary Tables and Static Scrollable Cursors now use the WORKFILE database instead of the TEMP database

Use DB2-managed (instead of user-managed) storage in SYSDEFLT storage group

Segmented table space organisation (user-defined SEGSIZE or default of 16)

4KB and 32KB page sizes only – no 8KB or 16KB, Much more 32KB

Changes in PK70060, more in PM02528, install PM17336

DB2 9 for z/OS introduces the following enhancements to the management of temporary space:

The two temporary databases are converged, leaving a single database - the WORKFILE database – which is used as storage for all temporary files and tables, but preserving the current functionality of the temporary tables. This includes:

Discontinuing use of the TEMP database and switching the temporary tables that currently use the TEMP database to use the WORKFILE database.

Provide enhancements to the WORKFILE database to incorporate TEMP database features that are currently not in the WORKFILE database.

Provides control over temporary space utilization at the agent level, to prevent applications from monopolizing space usage at the expense of others.

Provides the ability to monitor temporary space utilization at the DB2 subsystem level, helping you to manage the space allocated to the work file database.

Provides support for in-memory work files for small sorts.

Cannot specify SEGSIZE in Compatibility Mode – preset value of 16 always used.

Tables can span multiple table spaces (up to 254) with maximum size of 16TB (254 * 64GB)

Created Global Temporary Tables can span table spaces (as before).

Declared Global Temporary Tables and Static Scrollable Cursor result tables cannot span table spaces (as before).

The number of work files that can be used for DB2 sorts is limited only by the buffer pool size. The WORKFILE database does not use global DBD management to assign and release table IDs. Each work file has its own local DBD, copied from the global DBD. ('Global' here means global within a single DB2 subsystem.) The global DBD has only the table space descriptor objects for each table space defined in the database. The work file's local DBD has the table space descriptors and the table descriptor for just that table which is associated with that work file. Since the DBD is not globally shared (within a DB2 subsystem), there is no limit to the number of tables in the WORKFILE database.

Planning For Converged TEMP Space

Migration from DB2 V8 Much more 32KB space is needed in DB2 9

- Reclaim TEMP database storage with drop TEMP database and reallocate storage
- Recommendation: If you have enough space, do not drop TEMP database until you are sure that you will not fall back to V8, to avoid having to recreate it after fallback.

New installation panel for work file database definitions (DSNTIP9)

- In migration mode, if you specify non-zero values
 - Migration job DSNTIJTM will create additional DB2-managed WORKFILE table spaces in the SYSDEFLT storage group → new REXX program DSNTWFG
 - DSNTWFG does not include the existing work file table spaces

Recommendation: Keep 'DSVCI' default ZPARM YES to allow DB2 to match VSAM CI size to table space page size

Ensure you have 32KB WORKFILE table spaces & BP for Declared Global Temporary Tables and Static Scrollable Cursors

- Use new limits, e.g. 25% of total workfile space

When migrating from DB2 V8, if you wish to reclaim the storage associated with your existing TEMP database, it is *YOUR* responsibility to drop the TEMP database and reallocate the storage for use by the WORKFILE database or something else. It is probably better not to **DROP** the TEMP database until you are sure that you will not be falling back to V8, to avoid having to recreate it after fallback.

Before DB2 9, the DB2 installation process defines user managed data sets for table spaces in the WORKFILE database. The Version 9 installation process generates the DDL to create table spaces in the WORKFILE database using DB2-managed storage in the SYSDEFLT storage group. If you want to allocate the WORKFILE table spaces on user managed data sets, then you need to provide your own AMS statements to define them, and modify the installation-generated DDL to specify the VCAT clause in place of the STOGROUP clause.

DSVCI ZPARM should be YES. This enables DB2 to create DB2-managed data sets with a VSAM control interval size that matches the page size for the table spaces.

The fact that 8KB and 16KB page sizes are not supported means that Declared Global Temporary Tables and Static Scrollable Cursors require a table space with a page size of 32KB.

The syntax for the **DECLARE GLOBAL TEMPORARY TABLE** statement or for declaring a static scrollable cursor is not affected, but some messages are. Before defining a Declared Global Temporary Table, ensure that the WORKFILE database exists with at least one 32KB page size table space. DB2 does not create an implicit table space for Declared Global Temporary Table. If the workfile database does not exist, or if a table space of the required size does not exist, then the **DECLARE GLOBAL TEMPORARY TABLE** statement will fail SQLCODE -904 RC00E7009A – watch out for changed message text.

Service Best Practice DB2 9 or 10

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>

II14401 & II14464 install service
II14474 or II14477 for 10

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 on a quarterly basis with published results. This means that when you order your current service deliverable (Show/Select) you'll receive and install one recommended 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 following PMLDs installed on our systems. For more information about our environment, see our [CST Quarterly Report](#).

We have a quarterly CST test cycle:

- 1st month: Install service from prior quarter; Run initial workloads; Run in new test scenarios
- 2nd month: Run in high stress environment; Run in recovery tests
- 3rd month: Provide recommended service package; Publish quarterly report

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.

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

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.

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



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>

Package or Access Path 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

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 of the REBIND PACKAGE command.

2. A new system parameter called PLANMGMT for specifying the default setting of PLANMGMT option of REBIND PACKAGE. Possible settings are: OFF, BASIC and EXTENDED. The default value of this parameter is OFF. To use a setting other than OFF, update your DB2 9 subsystem parameter (DSNZPxxx) modules, then either use the SET SYSPARM command or stop and start DB2 for the change to take effect.

3. A new REBIND PACKAGE option called SWITCH to allow users to revert a package to an older copy, effectively undoing a prior REBIND invocation.

SWITCH(PREVIOUS) - The PREVIOUS copy is activated

SWITCH(ORIGINAL) - The ORIGINAL copy is activated

4. A new FREE PACKAGE option called PLANMGMTSCOPE to allow users to free older copies that are no longer necessary.

PLANMGMTSCOPE(ALL) - Free the entire package including all copies. This is the default.

PLANMGMTSCOPE(INACTIVE) - Free all old copies

Package or Access Path Stability 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 - PK80375 compression
- 30% CPU overhead for REBIND
- No impact on statement runtime processing, fast switch

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 10, 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.

Package Stability - BASIC support

REBIND ... PLANMGMT(BASIC)



REBIND ... SWITCH(PREVIOUS)





Chart is to be read from bottom to top


- The diagram shows the scenario of BASIC support. It is to be read bottom up. At the left hand side, when you rebind using PLANMGMT(BASIC), we discard the previous copy, take the current copy and store it as the previous copy, and take the newly generated copy as the current copy.
- Then when you switch, you basically swap around the current copy and the previous copy.

Package Stability - EXTENDED support

REBIND ... PLANMGMT(EXTENDED)



REBIND ... SWITCH(ORIGINAL)



84

© 2009 IBM Corporation

- This diagram shows the EXTENDED support. The previous copy is thrown to the trash can. The current copy becomes the precious copy. Only the first time you execute this, when we don't have the original copy stored, do we take the current copy and store it as the original copy. And the newly generated copy becomes the current.
- For rebind, we don't touch the original but swap around current and previous. But if you do rebind SWITCH(ORIGINAL), the previous goes to the trash, the current becomes the previous, and the original becomes the current. We don't throw the original copy, it stays where it was.

There is no overhead for PLANMGMT(OFF). The CPU overhead for binds using multiple

package copies is a fully justifiable 15-20% in these measurements. REBIND with SWITCH()

is very efficient with a 2 to 3 times reduction.

Note: The PTF for APAR PK80375 provides SPT01 compression.

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



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

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

87

© 2009 IBM Corporation

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**
- **CFLEVEL 16, 15, 14, 13, 12 for best performance**



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.

Best practices for DB2 upgrade

- Get the right service level
- Be ready for performance & access path management
- Set up the work space carefully, more 32K, limit
- Use new data format
- Watch for recent and upcoming changes
- Check the web for the latest books, advice, service, and best practices information

www.ibm.com/developerworks/data/bestpractices/db2zos/



Best Practices in DB2 9 for z/OS Migration January 2011 Customers are migrating more quickly to DB2 9 and DB2 10. The DB2 9 value is helpful for most customers, and the migration experience has been positive with customer migrating for the value. DB2 9 has been generally available for three years now. Migration is a broad topic, so I'll emphasize some keys to success, point to many resources, emphasize some important points and provide the latest information. In this article, I'd like to note current service level recommendations, note some of the key changes in access path management, note changes for work space, indicate recent changes in data format, and then point to the newest function changes.

Get the right service level As you migrate to DB2 9 or 10, using the best practices and recommendations can help you have a better experience. The Consolidated Service Test or CST process is being used successfully by very large and small customers. The current service level recommendation is to get to the current quarterly CST level and some specific identified APARs as a minimum level. The January 2011 level is CST4Q10 (RSU1012). RSU1012 includes all service through the end of September 2010 not already marked RSU, PE resolution and HIPER/Security/Integrity/Pervasive PTFs and their associated requisites and supersedes through the end of November 2010. Additional service is noted in the reports. Enhanced HOLDDATA can be used to find current information.

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

Check key information APARs. Check the areas which you use most. System z zIIP processors can make parallel processing less expensive than sequential access, but can cause problems in storage management and parallel processing, so check those areas if you use parallelism.

II14401 and II14464 note the key migration and fallback APARs. II14477 and II14474 (V8) for DB2 10.

II10817 indicates the key storage APARs.

II12836 notes the APARs for parallelism.

II14219 indicates the service for zIIPs.

II14334 shows the APARs for large objects or LOBs.

II14426 shows the XML APARs.

II14441 notes recent APARs for incorrect output.

II14203 shows the key APARs related to distributed processing or DDF.

Most DB2 for z/OS customers are large in some metrics, and unique in a number of ways. After getting all of the recommended service, you need to test with your unique workloads and options to find any other needed service. Staging DB2 service levels through development, QA, and production also helps find situations, applications and techniques that are unique.

Best practices for migration notes

Get ready for access path management Recommendations for some of the subsystem parameters have changed. At this time, the recommended value for parameters is:

STATCLUS = ENHANCED ENHANCED is the default and recommended value. DB2 9 optimization is better when the matching statistics are collected, improving information for clustering and scanning. Get current RUNSTATS data with the enhanced clustering statistics, so that the DB2 9 optimizer uses DB2 9 statistics. Where you have high cardinality non-uniform distributions, use histogram statistics.

OPTIOWGT = ENABLE ENABLE is the default after APAR PK75643, PTF UK42565(0901) is installed. This change provides more appropriate weight balance for CPU and IO with faster processors, such as the zEnterprise, z10, and z9.

OPTIXOPREF = ON ON is the default after APAR PK77426 is installed. This value will favor index only access.

The DB2 package stability or access path stability function has been very helpful with customers who have access path regression and want to get back to the old access paths. Read section 10.13 of Packages Revisited for a discussion of the package stability options and techniques. If the disk space for two or three times for packages in SPT01 is a concern, then get APAR PK80375, which will allow compression of the SPT01 table space. Chapter 4 of Packages Revisited discusses conversion to packages from plans containing DBRMs. DB2 9 for z/OS Packages Revisited, SG24-7688 <http://www.redbooks.ibm.com/abstracts/SG247688.html>

The best techniques for minimizing potential access path issues are get needed reorganizations done first, followed by new statistics, and then run rebinds, followed by checking. The updated cluster ratio data, new data repeat factor and histogram statistics provide important information for optimization. Run the statistics adviser to generate the recommended statistics to be collected. Use the Optimization Service Center (OSC) to capture the SQL statements and related information. You can start using OSC on V8 to get the needed information and to help tune your queries. Data Studio enhancements provide the function in later versions.

Set up the work space with more 32KB space and understand your use of declared temporary tables. DB2 9 has substantial changes for work space. Rather than separate databases and table space for work files and temporary tables, the space has converged. Space for work files will favor 32K files when record sizes are over 100 bytes. Every customer needs to have much more space for 32K files and will move from the TEMP database to the work file database. While work files can span table spaces, declared temporary tables cannot. If you can avoid having small work files compared to the largest declared temporary tables, then most problems are avoided. Use of new MAXTEMPS parameter to limit a single user to a fraction of the total work space, such as 10% avoids problems. Some customers need more separation of declared temporary space and work files. A technique was delivered in APARs PK70060, PM02528, and PM17336 to provide separation. Some documentation and installation changes are still in process. The converged space uses DB2-managed table spaces instead of user-managed ones. Use multiple table spaces with zero secondary space for work files. DB2 gives preference to table spaces with zero secondary quantity when allocating space for work files. Multiple work file table spaces help to support efficient concurrent read and write I/O's to work files. If applications use large Declared Global Temporary Tables (DGTs), then some table spaces should be defined with non-zero secondary quantity in the WORKFILE data base. This will minimize contention for space between work files and DGTs. DB2 allocation of space for DGTs gives preference to table spaces that can grow beyond the primary space allocation (with SECQTY > 0), since DGTs are limited to one table space and cannot span multiple table spaces as the work files can. If the subsystem has only small DGTs, then having all table spaces defined with zero secondary quantity works best.

Use new data format Reordered row format occurs in NFM. A number of recent changes have adjusted the function to be more flexible. Some table spaces have reduced compression. The primary challenge is table spaces with many VARCHAR(1) columns, where the compression ratio was reduced significantly. APAR PK87348 and PK85881 provide more extensive changes, externalizing a subsystem parameter and a utility keyword to keep basic row format. PK87348 also makes changes to avoid issues with the hidden parameter. The default and general recommendation are to use reordered row format.

Watch for recent and upcoming changes

Rather than creating 60,000 databases implicitly, changes in PK62178 default to 10,000 and allow you to change the number of databases to use. If you want to pipe information to the LOAD utility, look at PK70269. PK72214 changes the install process to stop changed data capture for catalog tables as needed, asking you to restart them after switching to DB2 9 and ENFM. If you have multiple members of a data sharing group in a single LPAR and prefer the V8 group attach logic, parameters are added to control the random group logic in APARs PK79228 and PK79327. The -ACCESS command only worked for specific names, but the command is being changed to allow wild cards and ranges in APAR PK80925.

DB2 and DFSORT have been changed to allow sort CPU to be eligible for zIIP in both V8 and DB2 9 in PK85889 with z/OS V1.10 and PK85856 or V1.11. Extended addressing was added for sequential data in z/OS V1.11 with PK81151.

Note the new premigration job DSNTIIPA provided in APAR PM04968. Other service helps remove function to prepare for DB2 10. Help to remove private protocol comes in PK92339 and PK64045. The ability to convert DBRMs in plans into packages comes in PK62876 and PK79925(V8). PK85068 finds old plan table formats. This is the latest information today, but more is coming soon. See updates and the latest books in IBM Redbooks publications. DB2 9 for z/OS: Distributed Functions, SG24-6952-01 DB2 9 for z/OS: Resource Serialization and Concurrency Control, SG24-4725-01, DB2 Utilities, SG24-6289-01 Watch for more updates to DB2 books and other information on the web.

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

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

When to migrate <http://www.ibm.com/support/docview.wss?uid=swg2100695>

Next phase: Use the new function, get the value

Reduce CPU time
Save disk space


Improve availability
More flexible security

Productivity for applications & administrators

Easier regulatory compliance

- PureXML & Data Studio
- SQL & LOBs
- Portability via SQL & data definition compatibility

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. For most customers this is NFM, which is also new value mode.


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 10.

For example, native SQL procedure conversion is described in a Technote.

<http://www.ibm.com/support/docview.wss?&uid=swg21297948>



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 DSNTIJP9 or DSNTIJPA job** 
 - **On every subsystem**
 - **Current maintenance for DB2, z/OS & CF**
 - **Get the latest books: library, IBM Redbooks publications, Information Center**

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 jobs early and often. Check the CATMAINT and CATENFM 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. The books and Information Center are being updated more frequently.

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/>



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

- ✓ Administration Guide
- ✓ Data Sharing: Planning and Administration
- ✓ Performance Guide
- ✓ Utility Guide and Reference
- ✓ Application Programming Guide & Reference
- ✓ Application Programming Guide & Reference for JAVA™
- ✓ ODBC, Spatial, ...
- ✓ Redbooks
- ✓ Installation Guide
- ✓ RACF Access Control Module Guide
- ✓ Messages
- ✓ Codes



© 2009 IBM Corporation

94

You can get most of the books from the Information Center or from the DB2 Library web page. The books continue to be updated, so get the latest ones. Some of the IBM Redbooks publications will be helpful. You may need books from the z/OS Library as well.

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

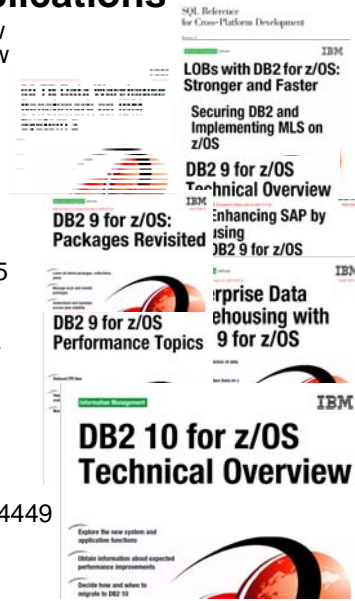
<http://www.ibm.com/support/docview.wss?rs=64&uid=swg27011656>

<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 added lately or updated (next page).

DB2 9 and 10 IBM Redbooks Publications

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



95

© 2009 IBM Corporation

DB2 library more information <http://www.ibm.com/software/data/db2/zos/library.html> Many IBM Redbooks publications, Redpapers and one cross-platform book on DB2 9 are published, in addition to the standard library, with more in the works. Check for updates. <http://www.redbooks.ibm.com/cgi-bin/searchsite.cgi?query=db2+AND+z/os>

1. DB2 9 Technical Overview, SG24-7330 <http://www.redbooks.ibm.com/abstracts/SG247330.html>
2. DB2 9 Performance Topics, SG24-7473, <http://www.redbooks.ibm.com/abstracts/SG247473.html>
3. DB2 9 Stored Procedures, SG24-7604, <http://www.redbooks.ibm.com/abstracts/SG247604.html>
4. Index Compression DB2 9, REDP4345, <http://www.redbooks.ibm.com/abstracts/redp4345.html>
5. Deploying SOA Solutions SG24-7663, <http://www.redbooks.ibm.com/abstracts/SG247259.html>
6. Cross-Platform Development Version 3, <http://www.ibm.com/developerworks/db2/library/techarticle/0206sqlref/0206sqlref.html>
ftp://ftp.software.ibm.com/ps/products/db2/info/xplatsql/pdf/en_US/cpsqlrv3.pdf
7. Enterprise Data Warehousing, SG24-7637, <http://www.redbooks.ibm.com/abstracts/sg247637.html>
8. LOBs: Stronger & Faster SG24-7270, <http://www.redbooks.ibm.com/abstracts/SG247270.html>
9. Securing DB2 & MLS z/OS, SG24-6480-01, <http://www.redbooks.ibm.com/abstracts/sg246480.html>
10. Enhancing SAP, SG24-7239, <http://www.redbooks.ibm.com/abstracts/SG247239.html>
11. Best practices SAP BI, SG24-6489-01, <http://www.redbooks.ibm.com/abstracts/sg246489.html>
12. New Tools for Query Optimization, SG24-7421, <http://www.redbooks.ibm.com/abstracts/sg247421.html>
13. Data Sharing in a Nutshell, SG24-7322, <http://www.redbooks.ibm.com/abstracts/sg247421.html>
14. DB2 9 for z/OS Data Sharing: Distributed Load Balancing and Fault Tolerant Configuration <http://www.redbooks.ibm.com/abstracts/redp4449.html>
15. Considerations on Small and Large Packages redp4424 <http://www.redbooks.ibm.com/abstracts/redp4424.html>
16. Backup and Recovery Considerations redp4452 <http://www.redbooks.ibm.com/abstracts/redp4452.html>
17. Powering SOA IBM Data Servers, SG24-7259 <http://www.redbooks.ibm.com/abstracts/SG247259.html>
18. Packages Revisited, SG24-7688 <http://www.redbooks.ibm.com/abstracts/SG247688.html>
19. 50 TB Data Warehouse Benchmark on IBM System z <http://www.redbooks.ibm.com/abstracts/sg247674.html>
20. SAP on DB2 9 for z/OS: Implementing Application Servers on Linux for System z <http://www.redbooks.ibm.com/abstracts/sg246847.html>
21. IBM Data Studio V2.1: Getting Started with Web Services on DB2 for z/OS <http://www.redbooks.ibm.com/abstracts/redp4510.html>
22. Ready to Access DB2 for z/OS Data on Solid-State Drives <http://www.redbooks.ibm.com/abstracts/redp4537.html>
23. Parallel Sysplex Operational Scenarios <http://www.redbooks.ibm.com/abstracts/sg242079.html>
24. Distributed Architecture <http://www.redbooks.ibm.com/abstracts/sg246952.html>
25. Buffer Pool Monitoring & Tuning <http://www.redbooks.ibm.com/abstracts/redp4604.html>
26. Securing and Auditing Data <http://www.redbooks.ibm.com/abstracts/sg247720.html>
27. Serialization & concurrency, SG24-4725-01 <http://www.redbooks.ibm.com/abstracts/sg244725.html>
28. Utilities SG24-6289-01 <http://www.redbooks.ibm.com/abstracts/sg246289.htm>
29. DB2 9 and Storage Management SG24-7823 <http://www.redbooks.ibm.com/abstracts/sg247823.htm>

More information and resources

Resources Now Available

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



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

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

ftp://ftp.software.ibm.com/software/data/pubs/papers/DB2_for_zOS_V9_Business_Value_White_Paper.pdf

http://www.ibm.com/software/os/systemz/newsletter/mainstreamed11_uuk.html

DB2 10 Technical Overview SG24-7892,

<http://www.redbooks.ibm.com/abstracts/sg247892.html>

Extremely pureXML in DB2 10 SG24-7915,

<http://www.redbooks.ibm.com/abstracts/sg247915.html>

The Business Value of DB2 for z/OS, SG24-6763,

<http://www.redbooks.ibm.com/abstracts/sg246763.html>

DB2 9 for z/OS Technical Overview, SG24-7330, chapter 2 on System z synergy,

<http://www.redbooks.ibm.com/abstracts/sg247330.html>

DB2 9 for z/OS Performance Topics, SG24-7473,

performance on z10,

<http://www.redbooks.ibm.com/abstracts/sg247473.html>

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

Home Solutions Services Products Support & downloads My IBM Welcome Mr.

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

DB2 for z/OS
 DB2 10 for z/OS delivers innovations in operational efficiency

• **Technical resources**
 • **Events**
 • **Support**
 • **V8, DB2 9**
 • **Developer Domain**
 • **DB2 Magazine**
 • **zIIP**
 • **Tools**

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

Add to My Interests

Announcing DB2 10 for z/OS
Savings ... right out of the box
 → Learn more

Overview
 DB2 10 for z/OS delivers innovations in operational efficiency for out-of-the-box savings, business resiliency, warehouse deployment and enhanced business analytics.

Learn more

- Features & benefits
- System requirements
- Edition comparison
- Product library
- IOD Global 2010
- IDUG EMEA 2010 (link resides outside ibm.com)

Downloads

IBM Software products
 Find products

Keyword search:
 Enter search terms
 Search

INTRODUCING DB2 10 for z/OS
 → View the demo

DB2 for z/OS is the only cost effective, simple and proven database on the market which offers 24x7 availability, reliability and security.

97

© 2009 IBM Corporation

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

<http://www.ibm.com/software/data/db2/support/db2zos/>

Click
Planning,
then the
“Product
Support
Lifecycle”
link

Planning
Support for my selected products

IBM Software support lifecycle
DB2 for z/OS 9.1.x, DB2 for z/OS 8.1.x, DB2 for z/OS 10.1.0, DB2 for OS/390 6.1.x, DB2 for OS/390 5.1.x, and DB2 for OS/390 and z/OS 7.1.x

The announcement letter dates below are US only. Information for other country announcements is available on the [IBM Offering Information page](#).

Subscribe via RSS to Product Support Lifecycle site updates.

View lifecycle dates for other products.

Product information	License type ¹ and lifecycle policy ²	Lifecycle dates ³ and announcement letters	Last updated
DB2 for z/OS Version: 9.1.x PID: 5635-DB2	License type: ICA Lifecycle policy: Standard	GA: 16-Mar-2007, 207-041	14-Oct-2010
DB2 for z/OS Version: 8.1.x PID: 5625-DB2	License type: ICA Lifecycle policy: Standard	GA: 26-Mar-2004, 204-029 EOM: 09-Sep-2009, 908-280 EOS: 30-Apr-2012, 910-169	14-Oct-2010
DB2 for z/OS Version: 10.1.0 PID: 5605-DB2	License type: ICA Lifecycle policy: Standard	GA: 22-Oct-2010, 210-380	20-Oct-2010

© 2009 IBM Corporation

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

IBM

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

99 © 2009 IBM Corporation

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/>



Info APARs II14474 8 → 10 II14477 9 → 10
II14401 & II14464 8 → 9

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

Related software

II14401: DB2 V8
MIGRATION/FALLBACK INFOAPAR
TO/FROM DB2 V9

APAR status
INTRAN

Error description

5740XYR00 R910
DB2 V8 MIGRATION/FALLBACK INFOAPAR TO/FROM DB2 V9

FOLLOWING PTFS ARE FOR DB2 V9 MIGRATION:
(Updated on Oct/4/2008)

PK11129/UK90008 (PE) (F610)
Toleration of fallback from DB2 Version 9 (R910). In a Data
Sharing environment, required for coexistence of V8R1 and
.....

100

© 2009 IBM Corporation

<http://www.ibm.com/support/docview.wss?rs=64&context=SSEPEK&q1=ii14401&uid=isg1II14401>
<http://www.ibm.com/support/docview.wss?rs=64&context=SSEPEK&q1=ii14464&uid=isg1II14464>

These informational APARs describe the needed service for successful migration.

from DB2 V8 to DB2 9 read II14401 and II14464

from DB2 V8 to DB2 10 read II14474

from DB2 9 to DB2 10 read II14477

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

← Return to IBM Support Portal

Support & downloads >

DB2 for z/OS - Technical Resources

Support & downloads
 Bookmark this page
 View my bookmarks
 Feedback

•Download 30 PDF books
•Access html books
•DB2 new December 2010
•Program Directories
•Information Center

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

Version 10 | Version 9 | Version 8 | Version 7

- Information center search
- Manuals
- Program Directories for DB2 10 for z/OS
- Program Directories for DB2 10 for z/OS Value Unit Edition
- DB2 Accessories Suite for z/OS
- Additional information resources
- Submit Feedback and suggest updates to DB2 for z/OS information

Information center search
 As part of IBM's ongoing effort to improve your information experience, we are optimizing the information center search experience.

101

© 2009 IBM Corporation

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 for z/OS Best Practices

www.ibm.com/developerworks/data/bestpractices/db2zos/

Watch recorded presentations. Read papers from your favorite DB2 experts. Get practical recommendations for areas such as:

- Virtual storage tuning
- Security
- Software maintenance strategies
- Migration
- Setting up application servers to access DB2 for z/OS
- Partitioning table spaces
- Debugging stored procedures

And more!

Send feedback and topic suggestions to db2zinfo@us.ibm.com



For more information about DB2 10

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

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

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

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

103

© 2009 IBM Corporation

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

The announcement itself is here:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

We have a white paper on DB2 10 from Dave Beulke

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

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

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 System i 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 March 2009 [working on update](#)

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

V3.1 DB2 for LUW 9.5
DB2 for i V6.1
DB2 9 for z/OS



New SQL features now included

- DECFLOAT data type Additional timestamp format
- ROW CHANGE expressions New built-in functions
- order-by-clause and fetch-first-clause in a subselect
- SELECT FROM INSERT Full outer join
- IMPLICITLY HIDDEN columns row-change-timestamp columns
- RESTRICT on DROP of functions and procedures
- CURRENT DECFLOAT ROUNDING MODE special register & SET
- ALTER FUNCTION Other portability enhancements

This is the web page for the SQL Reference for Cross-Platform Development. It has four 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 latest book is Version 3.1 for platforms DB2 9 for z/OS, DB2 for I V6.1, and DB2 9.5 for Linux, UNIX and Windows. The addition of newer versions of the DB2 products means the following great SQL features are now included:

- DECFLOAT data type Additional timestamp format
- ROW CHANGE expressions New built-in functions
- order-by-clause and fetch-first-clause in a subselect
- SELECT FROM INSERT Full outer join
- IMPLICITLY HIDDEN columns row-change-timestamp columns
- RESTRICT on DROP of functions and procedures
- CURRENT DECFLOAT ROUNDING MODE special register & SET statement
- ALTER FUNCTION 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/cpsqlrv31.pdf

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

http://www.ibm.com/developerworks/data/library/techarticle/db2common/?S_TACT=105AGX11&S_CMP=LP

IBM

ibm.com/software/data/db2imstools/

Click Support For Compatibility Matrix → → → & V8 or DB2 9 or DB2 10

Home Products Services & solutions Support & downloads My account

Software > Information Management >

IBM DB2 and IMS Tools

Library
News
How to buy
Events
Training and certification
Services
Support

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.

Get the latest news on [announcements](#), [webcasts](#), and [customer case studies](#). Link to the [library](#) to find user's guides and redbooks. A brief [overview](#) maps the functional categories of the tools to the benefits they can bring to your on demand business environment. Refer to the Highlights for the newest items of interest.

Related software

- DB2
- IMS
- DB2 QMF
- CICS Tools
- WebSphere MQ Family
- WebSphere AD Tools

Related hardware

- zSeries
- Storage Systems
- IBM eServer

• Warranty info

Select a DB2 Tool

Select an IMS Tool

Select a WebSphere Information Integration Tool

DB2 Tools for Linux, Unix, and Windows

→ Visit the DB2 Multiplatform Tools webpage

Read the DB2 Multiplatform Tools overview fact sheet.

→ NEW! DB2 Toolkit for Multiplatforms V1.1 brings

New / Enhanced DB2 z/OS Tools

Announced March 2006

→ DB2 SQL Performance Analyzer, V3.1

Announced December 2005

→ DB2 Change Accumulation Tool, V1.4

Information Management software

Portfolio of DB2 & IMS Tools

Introducing z9-109

z9-109

→ Learn more

Highlights

→ Register for EMEA IM Tech Con, The Hague May 15

→ Register for IDUG NA, Tampa, May 7

→ Register for June 8

106

© 2009 IBM Corporation

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 Information Management Tools and DB2 10 for z/OS Compatibility

News

Abstract
This table provides information regarding DB2® Tools support of DB2 10 for z/OS.

Content

DB2 Version 8 DB2 9 DB2 10

DB2 10 Support

Key:

- NS = Not Supported**
The product does not function properly on DB2 10.
- 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 10 functions.
- E = Exploit**
The product utilizes most, but not necessarily all, of the new functions and features available in DB2 10 that IBM deems the most significant for this product. It is anticipated that any gaps in utilization would be eventually closed as new releases or PTFs are issued over time.

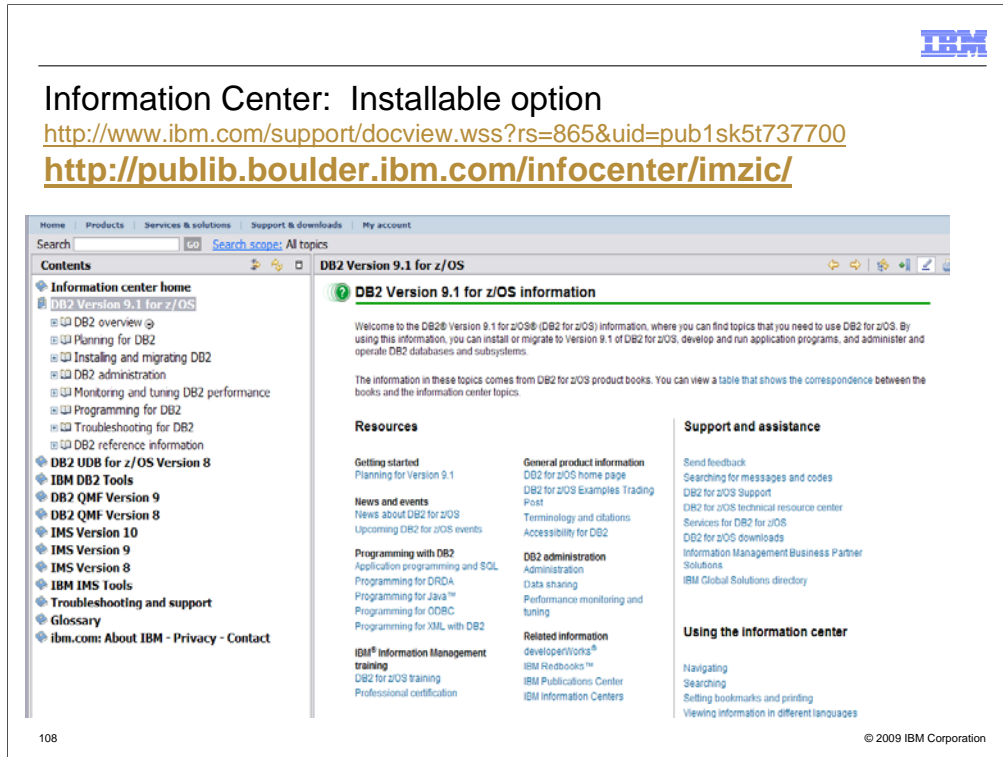
For more information on which new functions of DB2 10 are utilized by any specific product, or to find out about DB2 10 compatibility for new products that have been released after this matrix was last updated, please contact your IBM Customer Service Representative.

DB2 Tools	VRM	DB2 10 Support			Comments
		NS	T	E	
Application Recovery Tool	1.2.0		X		PM22390

107

© 2009 IBM Corporation

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.



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.

The latest change in this area is an installable Information Center, so that you can use the facility even when the internet is not accessible.
<http://www.ibm.com/support/docview.wss?rs=865&uid=pub1sk5t737700>

The online version is

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

DB2 for z/OS Information Center: Cool features

- Optimized for Google search: easily get to DB2 for z/OS information from Google
- Browser search plug-in available: search the information center from your browser's search box, anywhere on the web
- Integrated search for messages, APARs, PTFs, Redbooks
- Filter to see only the products you use
- RSS feeds for automatic notification of updates
- Three new Flash tours with tips for using the information center
- Links to related product information and community resources

Access information for all supported versions of DB2 of z/OS, most DB2 Tools, QMF, and many other IM on z/OS products

Locally installable information center now available

Wide customer adoption: 12.1M hits from 57K unique visitors in March 2009 (up 40% and 30% over March 2008)

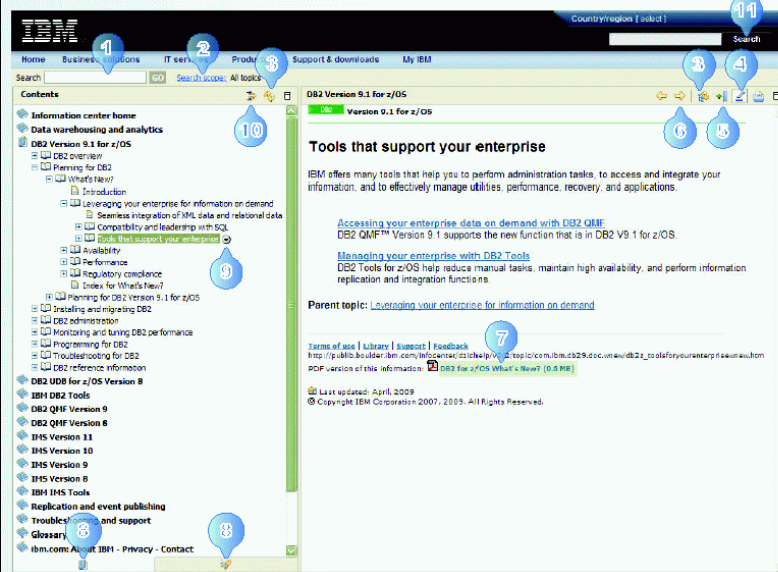
Here's a link to the browser search plugin for the information center:

http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.dzic.doc/additional_search.html

Install the browser search plugin by clicking the first link on the page.

Also included on the page is a Google search box for searching the information center using Google.

Tips for using the information center



- 1 You can perform simple and complex searches. You can search for exact words or phrases, and search with wildcards and Boolean operators. Searches are not case sensitive.
- 2 Narrow a search by clicking **Search scope**. In the Select Search Scope window, you can specify the name of your search list and the topics to search.
- 3 Synchronize the Topics view and the Contents view so that the displayed topic is highlighted in the Contents view.
- 4 Turn off and on the highlighting that displays in search results.
- 5 Bookmark a topic in your browser.
- 6 Navigate between the topics that you previously viewed.
- 7 Link to the same content in PDF.
- 8 Toggle between the Contents view and the Search Results view.
- 9 Print or search the current topic or its subtopics.
- 10 Set a filter to limit the topic titles that are displayed in the Contents view.
- 11 Search within IBM.com.

Here are many hints and tips for using the Information Center. They are available on a pdf on the web. URL is on the next page.

Resources:

DB2 for z/OS information center:

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

Launch page for the Flash tours:

<http://publib.boulder.ibm.com/infocenter/imzic/topic/com.ibm.dzic.doc/dzicviewlet.htm>

Search for APARs and PTFs from the information center:

<http://publib.boulder.ibm.com/infocenter/imzic/topic/com.ibm.support.dzic.doc/html/search.html>

Internal download for the installable information center:

<http://ibmurl.hursley.ibm.com/9UP>

Printable flyer (2-sided/color):

ftp://ftp.software.ibm.com/software/data/db2/zos/info/DZIC_2009.pdf

Browser search plugin for the information center:

http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.dzic.doc/additional_search.html

DB2 for z/OS information center:

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

Launch page for the Flash tours:

<http://publib.boulder.ibm.com/infocenter/imzic/topic/com.ibm.dzic.doc/dzicviewlet.htm>

Search for APARs and PTFs from the information center:

<http://publib.boulder.ibm.com/infocenter/imzic/topic/com.ibm.support.dzic.doc/html/search.html>

Internal download for the installable information center:

<http://ibmurl.hursley.ibm.com/9UP>

Printable flyer (2-sided/color):

ftp://ftp.software.ibm.com/software/data/db2/zos/info/DZIC_2009.pdf

Browser search plugin for the information center:

http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.dzic.doc/additional_search.html

More information: 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 fmt <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>

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>

Planning Your Migration to DB2 10 for z/OS

Questions?

Roger Miller

IBM Silicon Valley Lab
millerrl@us.ibm.com



113

© 2009 IBM Corporation

Thanks for coming.

Title – DB2 for z/OS technical evangelist, strategist, architect, designer, developer, writer, service, DB2 factotum (from the Latin for does everything or jack of all DB2 trades, master of several).

Current Projects - Roger is working to roll out DB2 9 for z/OS, to design the next improvements in DB2.

Technical accomplishments/education - Roger Miller is a DB2 for z/OS technical evangelist, architect and designer who worked on many facets of DB2, ranging from overall design issues to SQL, languages, install, security, audit, standards, performance, concurrency, and availability. He has worked for 31 years on DB2 development, product design and strategy. He often helps customers to use the product, answers many questions and presents frequently to user groups.

Fun facts - Roger likes hiking, bicycling, reading, Shakespeare, Yosemite and bears. He has learned to like working out with a personal trainer.

Upgrade to IBM DB2 10 for z/OS

CPU savings ... right out of the box

January 27, 2010

Upgrading to DB2 10

Roger Miller



DB2 V8

DB2 9

DB2 10

Is your current fish bowl getting tight? What is limiting you? Is it CPU cycles, real or virtual memory, disk storage or I/O, network access or time, object access times, locking, latching waiting, application throughput, elapsed times, transactions per unit time, catalog or directory limits, utilities or something else? Are you currently running DB2 9? V8? V7? Upgrading to a bigger fish bowl may be just what you need.

Should you upgrade to DB2 10? Yes.

To 10, or not to 10, when and how are the questions. Whether 'tis nobler in the mind to suffer the slings and arrows of outrageous limits, Or to take arms against a sea of troubles, And by upgrading, end them? To die: to sleep; No more; and by a sleep to say we end the heart-aches and the thousand natural shocks that old versions 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 versions have shuffled off this mortal coil, must give us pause: There's the respect that makes calamity of too long life for old versions. [With abject apologies to the Bard and to Hamlet act 3 scene 1]

The answer to upgrading to DB2 10 is a definite Yes. The question is not so much whether to upgrade as when and how to upgrade. If you are running DB2 9 today, then DB2 10 is in your near future, giving you more room to grow, with higher limits, lower CPU, and more for less. If you are running DB2 V8 today, then you have a choice of jumping to DB2 9 or directly to DB2 10. Both choices are fully supported with migration and fallback.

This paper discusses the key decisions, planning, options, and guidance for the upgrade. At this time, DB2 10 has recently become generally available October 22, 2010. The guidance will change over the life of the product, and so this paper has an expiration date of December 2011, although it will be updated before that.

Justify the upgrade

Two of the initial tasks for an upgrade are to make some primary choices and to justify the value for the work. For some, the strategic value is the key. Others need more detail, with descriptions and value assessment. Still others need to understand the detailed costs, the work and the return on that investment. I suggest reading the white paper by Julian Stuhler, Triton Consulting: DB2 10 for z/OS Business Value Whitepaper - A Smarter Database for a Smarter Planet.

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

As usual, DB2 10 has a price increase, but generally saves CPU. If the CPU saving is enough to reduce the capacity, then you may be able to save on the full hardware and software stack. The costs and values for individual enterprises are often very different. The pricing options and agreements differ substantially. The value delivered depends upon the constraints and limits, which differ widely.

Here is a summary of the two versions

DB2 9: Robust, Scalable, Available and Easily Manageable

DB2 9 delivers CPU reductions for utilities that are generally up to 20%. Customers report saving terabytes of disk space using index compression. More CPU time is shifted to zIIP processors, reducing costs. Security is improved with more flexible trusted contexts and roles. Resilience is improved as more changes can be made while applications keep running. One table can be replaced quickly with a clone. Indexes and columns can be renamed. Many more utilities can be online.

DB2 9 delivers seamless integration of XML and relational data with pureXML and makes big strides in SQL for productivity and portability of applications. A new storage structure is introduced for large tables. Today's complex applications include both transactions and reporting, so performing both well is required. The key improvements for reporting are optimization enhancements to improve query and reporting performance and ease of use. More queries can be expressed in SQL with new SQL enhancements. Improved statistics are provided for the optimizer, and algorithms are also improved. Reduced CPU and elapsed times can be achieved with the FETCH FIRST clause specified on a subquery. The INTERSECT and EXCEPT clauses make SQL easier to write.

DB2 10: Cut CPU and improve performance

DB2 10 for z/OS provides the best reduction in CPU for transactions and batch in over 20 years. We expect most customers to reduce transaction and batch CPU times between 5% and 10% initially, with the opportunity for more, depending upon tuning and workloads. The initial improvements in CPU and memory generally depend upon REBIND. Thread reuse with the release deallocate option is able to be used more widely, reducing CPU for high volume, short transactions and batch with frequent commits. Use of the z10 and z196 hardware large page size can be significant for some customers. The utility performance expectation is equivalence with DB2 9, but much better than V8.

Applications can have larger CPU and memory reductions if they can take advantage of additional benefits, such as constraint relief or insert performance. Scalability is the second major benefit, with the ability to run five to ten times as many threads in a single subsystem by moving 80% to 90% of the virtual storage above the bar. Online schema evolution or data definition on demand enhancements improve availability. SQL and pureXML improvements extend usability and application portability for this platform. Productivity improvements for application developers and for database administrators are very important as data grows in scale and complexity. More description is provided in the overview section at the end of this paper, and many other resources are noted below.

When and how should I upgrade to DB2 10?

The choice of migration target and timing are crucial decisions for the success of a migration project. The answers depend upon your current situation, the value delivered, and resources available. The guidance will change over the life of the version. This version is discussing the first year after general availability. First let's outline some of the dates which have been set and options.

DB2 V8 end of service is announced as April 30, 2012, 18 months from October 2010. Most customers depend upon tools or other software. Having that software work with DB2 is a prerequisite which may extend the schedules. DB2 V8 customers may need to purchase Extended Service if they don't migrate before end of service. When you upgrade DB2 V8 directly to DB2 10, single version charging may be offered for up to 18 months. Upgrade projects tend to be lengthy, so enterprises need to begin planning the needed upgrade very soon. DB2 9 does not yet have an end of service announcement, thus upgrade timing is not so constrained.

Most customers avoid running software which is out of service. Some customer policies target the second most current release. As we have a new version in service and an old version going out of service, many customers will need to move to DB2 9 or DB2 10.

While DB2 10 is expected to be better than prior versions, it will have maturity, stability, and service delivery similar to other software and versions. Normal maturity means more potential defects at first, then fewer as the software matures. Determining when the software is ready for a specific customer and when the customer is ready for the software depends upon the specific customer resources, prior experience, the ability to create and implement an appropriate upgrade plan, and the value for the improvements versus the need for stability.

Most DB2 for z/OS customers are large and have unique attributes. Testing the unique characteristics is challenging, and customers who upgrade to a new version soon after general availability need to have more comprehensive testing. Early users also need to use best practices for service.

Questions for you

The right answer is not "One size fits all." If we know the key factors for you, we can help you make a better choice. Here are some of the key objectives. Which ones are most important for you?

- Performance improves in both DB2 9 and 10, with larger CPU reductions in DB2 10 for transactions and batch, CPU reductions for utilities in DB2 9 and DB2 Sort. Insert improvements are in both DB2 9 and 10. See performance presentations and books for detail.
- Scalability is primarily constrained by memory below the bar, which is improved by about 10% to 15% in DB2 9 and improved by a factor of 5 to 10 in DB2 10.
- Availability is enhanced in both, with more online changes in both. Online utilities are a key strength DB2 9, with several more online facilities in DB2 10.

- Security is made stronger and more flexible with roles in DB2 9 and with more options, new audit capabilities, and row and column access controls in DB2 10.
- Productivity is helped in both releases, with more improvements in DB2 10. Upgrading a single version at a time is shorter and easier than skipping two versions ahead, but skipping from V8 to DB2 10 delivers faster results with less work than two separate migrations.
- Stability is better in more mature versions. If avoiding risk is the primary driver, you should avoid being out of service so migration from V8 to DB2 9 would be the solid, known, tested, safe path.
- Skills: What skill set is available within your organization? Do you have people with the right skills and time to plan and run a project? DB2 migration planning workshops can help with education. Transition classes provide more education for one or both versions. Services could be used if the skills are not presently available.
- Technology adoption model: Are you using the latest technology that is being shipped, or are the operating system and hardware back-level? Is the technology one level back, two or more? This question tells both how much work will be required and the comfort level of your organization for running the latest version.
- Platform management practices: What are your platform management practices? What type of change management practices are in place? How robust is your testing for new software? What is the inventory of software and tools? How many vendors are involved? Which ones? Almost every vendor has software ready for DB2 9. Many vendors will be ready for DB2 10 immediately, but specific tools and new function may take some time after general availability. Discuss the timeline for your tools and function you want with your vendors during your early planning.
- Numbers of servers: How many LPARs and subsystems does the organization have? An organization that has 100 subsystems has a different set of challenges than does one that has 5 subsystems. If each subsystem is unique, then the process is more complex.
- Service process: How often do you apply preventive maintenance? Best practice is using a Consolidated Service Test or RSU three or four times a year for preventive service with additional service to correct specific problems. Customers who migrate early in a version need to use the best practice.
- Capability to test: Do you have a strong testing structure? Can you run close to production volumes outside of production? Best practice is having a set of applications which cover 95% of the production. Have you identified a set of applications which reflect your workloads? Do you have separate quality assurance subsystems? Unless the development organization is engaged and involved, just staging from sandbox to development to production does not cover the unique attributes of your work that could cause problems. Early migrations require best practices for testing.
- Organizational considerations: What additional organizational factors, such as politics, policies, process and vacations must be considered?

What version are you running?

Here is a starting point for customers who are running various DB2 versions.

- DB2 9: If you are on DB2 9 today, then you are a good candidate for an early upgrade to DB2 10, especially if your custom is to move in the first year or two after general availability. Listen to reports from early customers and upgrade for the value. Early customers have presented at 2010 conferences, and will be at most 2011 conferences.
- V8 If you are on DB2 V8 today, then the next questions are on timing for you and readiness for the new version. How soon after general availability do you normally upgrade?
 - If you have the resources and can work with a new version, then you are more likely to be successful skipping to DB2 10. Some customers have usually upgraded quickly to new versions, but did not move to DB2 9. If customers usually start working with new versions in the first year, they have tests and processes for dealing with the normal maturity pattern. These customers may be able to skip DB2 9 and deliver more value in a shorter timescale.
 - Some customers who do not migrate early will be attracted to the reduced CPU time or scalability improvements and will want to skip to DB2 10. The challenge will be the need to do more robust testing and service while working with a version that is less mature. The time frame is helped if they commit early, getting additional extended service for V8 and the extended single version charging. The time frame will still be a challenge, so the team needs to be larger with more careful project management, compared to the usual single version upgrade.
 - V8 recent: Are you still in the process of moving to NFM or have you recently finished V8 upgrade? If you just finished, then you probably will wait for some time and use the skip. If you have resources for an upgrade, but DB2 10 is too new for you, then DB2 9 is probably your next move.
- V7: If you are currently on DB2 V7, then upgrade to DB2 V8. Then you can use the skip version upgrade to DB2 10 when you are ready for the next upgrade.

DB2 has several new versions and upgrade paths for you to consider. This story will be changing, but you can hear the latest at IDUG, IOD, and Share conferences and on the web. DB2 9 is ready for you now. DB2 10 became generally available in October 2010, but is delivering higher limits, lower costs, and more for less.

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

When should I upgrade to DB2 for z/OS Version 8? DB2 9? DB2 10?

DB2 V7 has been out of service for almost 3 years, since June 30, 2008. DB2 V8 has been in the field for almost 7 years, DB2 9 for almost 4 years, and DB2 10 is now generally available. End of service for DB2 V8 is April 2012. If you have not upgraded to Version 8 already, now is the time. If you are on DB2 V8 now, then this is a good time for migration. Many customers have moved to DB2 9 in the past four years. Experience from thousands of customers has mapped out the bumps and helped us to smooth the road. If you generally upgrade to a new version two years or more after general availability and don't want to go out of service, then DB2 9 is an effective upgrade.

Each customer has unique business requirements that drive the need to move to new versions. Some customers upgrade due to the delivery of needed function. Others upgrade when they find that many others have upgraded successfully. Some enterprises plan to be on the second most current release. Still others upgrade only when the prior version is going out of service. Each of these groups has new reasons for an upgrade.

Upgrade for value: Evaluate what DB2 functions you need for your business. Determine the value of those DB2 functions for your business. Weigh the costs. Ensure that you have all prerequisite hardware and software and the time to plan and execute the migration. If you want to upgrade to DB2 9 or DB2 10, then first you need to upgrade to Version 8 and to New Function Mode.

Upgrade when others have: The largest work loads and the largest customers are generally running on DB2 9 today. The key field experience measures such as the rate for PTFs in error, the number of severity 1 APARs, the problem rate and the APAR closure time are better than Version 8 and Version 7 at the same point.

Upgrade for end of service: The end of service date for Version 7 was June 2008. The end of service for V8 is April 2012. So plan your migration if your version is lower than 9. Customer budgeting, planning, ordering and migration often takes more than a year, and you need to start now, if you have not already. As you upgrade, you should be looking at the function which is removed or deprecated in later versions as well. Adjusting the applications is easier with more time. As applications are changed and tested, adjustments should be made for the functions being removed or deprecated in DB2 9 and DB2 10.

If you have been unable to upgrade because you depend upon old COBOL or PL/I compilers, that reason was avoided in DB2 V8, 9 and 10. A DSNHPC7 precompiler that is almost the same as the Version 7 precompiler is being delivered with DB2 Version 8, 9 and 10. DSNHPC7 is included in V8 with APAR PK46170, DB2 9 with APAR PK91610, and in the base for DB2 10. This technique works if you are not changing the program to use DB2 improvements. This function is deprecated, so applications should be changed to not depend upon products that have been out of service for a decade or more.

Experience from V5 to V7 skip release migration may help with some customer choices. Customers who used skip version migration generally found that the project was roughly 50% longer and 50% more work than a single version migration. Education for two versions needs twice the amount of time. Remediation for the items removed is twice as much work for a skip release. The prerequisites for skipping take more time, as customers will need to be more current. I'd estimate saving 20% to 25% of the work compared to two separate migrations. The primary negative for a skip version migration is getting the value of DB2 9 improvements later. One key value is getting to DB2 10 earlier.

Customers who are on DB2 V7 or earlier should upgrade to DB2 V8 as soon as possible to be on a solid, supported version of DB2. They can use the skip version migration process directly to DB2 10, when they are ready for the upgrade. This may be the best choice for customers who have just finished or are not yet finished migrating to DB2 V8 NFM. If you are planning on migrating to DB2 9 in 2010 or 2011, that may still be your best option.

Before you upgrade, perform the following actions:

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

To read about the features that are available with each version of DB2 for z/OS®, go to the [DB2 for z/OS homepage](#) and to the [Support page](#).

Best Practices in DB2 10 for z/OS Migration

DB2 10 is the most exciting new version in over 20 years. Migration is a broad topic, so I'll emphasize some keys to success, point to resources, emphasize some important points and provide the latest information. In this article, I'd like to discuss getting ready, rebinding, current service level recommendations, some key changes in access path management, and early customer experiences.

Get educated about DB2 10

Conferences and webcasts have provided some information about this new version, but you will need more depth to avoid learning the hard way. Start with a no charge migration planning workshop by contacting your local IBM team. Then do more reading and have key people take a transition class. If you are upgrading from DB2 V8, then the DB2 9 education is also needed. Education for all of the groups involved will help to avoid problems later.

The upgrade process has been improved in a number of ways. Customization for the stored procedures and user defined functions supplied by IBM is now performed for migration and for maintenance. WLM environments are set up. Documentation for the stored procedures and user defined functions was improved.

Standard parameters for an SAP subsystem can be set with one new option. Checklists are provided that have been very helpful to customers who use them. A new job is included to set the parameters from an existing subsystem.

How should I start getting ready for DB2 10? Start now.

Are you ready for DB2 10? Here is a checklist on getting ready to upgrade.

1. Check prerequisite software levels. If you are not yet running in NFM for DB2 9 or V8, then that is the first migration step. Ensure that you are at the required level for z/OS and other software. Use the Program Directories and announcement.
<http://www.ibm.com/support/docview.wss?rs=64&uid=swg27011656>
http://www.ibm.com/common/ssi/rep_ca/0/897/ENUS210-380/ENUS210-380.PDF
2. Attend your migration planning workshop. Get more details on DB2 10 and a project framework. The information, discussion, and project plan framework in the migration planning workshop will be very useful. Contact your local IBM representative.
3. Plan what you expect to gain and process for testing. Plan for memory and performance. Set your objectives and what tests will be made to measure them. What unique workloads do you have and what tests reflect that work?
4. Build your detailed migration plan.
5. Contact your key vendors and communicate your plan. Contact your key vendors and communicate your plan to them to understand when they will have the function you need. Start getting the needed upgrades and adjust the

schedule if needed. For many IBM tools, get the needed level and service information on the web.

<http://www.ibm.com/support/docview.wss?uid=swg21409518>

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



6. Get all of the parts out of the box. DB2 10, DB2 Utilities Suite V10, Accessories Suite V2.1 (spatial, OmniFind and international components for Unicode), Data Studio V2.2.1, client code and DB2 Connect 9.7 fixpack 3a, QMF 10, then add to the stack your performance tools, languages, Java drivers, ODBC, application tools, DBA tools, backup and recovery tools, replication, transaction management, data governance, and applications. Updated code for client and DB2 Connect V9.7 fixpack 3a are needed for a number of DB2 10 new functions: dynamic statement cache enhancements, timestamp with time zone, greater timestamp precision, binary XML format, extended indicator variables, explain mode special register, and Unicode collection and package names. Get it from <http://www.ibm.com/support/docview.wss?uid=swg24028306>
7. Read information APAR II14477 (DB2 9) or II14474 (V8). Have you looked though the information APAR II14477 (DB2 9) or II14474 (V8) recently? <http://www.ibm.com/support/docview.wss?uid=isg1II14477>
<http://www.ibm.com/support/docview.wss?uid=isg1II14474>
8. Apply required service to your current system. Do you have the required service applied to your current system? Have you checked the CST level? <http://www.ibm.com/servers/eserver/zseries/zos/servicetst/mission.html>
9. Run premigration jobs on all groups and subsystems. Have you run the premigration jobs on all groups and subsystems? The new premigration job DSNTIJPA is provided in APAR PM04968 and PM15965 on both DB2 V8 and 9. In DB2 10, the job is DSNTIJPM. http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db2z10.doc.inst/db2z_dsntijpm.htm
10. Resolve incompatible changes. Many are found in the premigration job, but some are noted in the Information Center and books. See the sections on incompatibilities in applications and SQL, utilities, commands, security, other, deprecated functions, and functions that are no longer supported. These changes are larger for V8 than for DB2 9.
DB2 V8
http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db2z10.doc.inst/db2z_premigr8checklist.htm
DB2 9
http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db2z10.doc.inst/db2z_premigr9checklist.htm

11. Get rid of private protocol. Have you gotten rid of private protocol? Help to remove private protocol comes in PK92339 and PK64045 and is described in the Distributed book, SG24-6952-01 chapter 5.2.
<http://www.redbooks.ibm.com/abstracts/sq246952.htm>
 12. Convert to packages instead of DBRMs directly in plans. Have you converted to packages instead of DBRMs directly in plans? The ability to convert DBRMs in plans into packages comes in PK62876 and PK79925 (V8) and is described in the Packages Revisited book, SG24-7688 chapter 4.
<http://www.redbooks.ibm.com/abstracts/sq247688.html>
 13. Upgrade plan table formats to Unicode V8 or DB2 9 level. Have you upgraded plan table formats to Unicode V8 or DB2 9 level? PK85068 finds old plan table formats and helps with conversion. Get ready to upgrade to DB2 10 format.
 14. Get ready to use SMS managed DB2 catalogs. Are you ready to use SMS managed DB2 catalogs? All new indexes and new table spaces in the catalog and directory created as SMS-controlled and DB2-managed with extended addressability and extended format.
http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db2z10.doc.inst/db2z_definesmsenvironmentmigr.htm
 15. Save information about your current access paths and performance for a baseline. You may want to use tools, like Data Studio, Optim Query Workload Tuner, or Optimization Service Center. If you use the older Visual Explain or Optimization Service Center, then you will need to change to newer tools. You might use access path management function. You might just save the prior information explain tables and accounting records or reports.
 16. Use DB2 Installation Guide checklists. The DB2 Installation Guide has checklists with more detail for V8 and DB2 9 migration to DB2 10. A premigration and a migration checklist is provided for each. Use the appropriate lists from this point on. For upgrade from V8 use these checklists.
http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db2z10.doc.inst/db2z_premigr8checklist.htm
http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db2z10.doc.inst/db2z_migr8checklist.htm
- For upgrade from DB2 9, use these checklists.
http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db2z10.doc.inst/db2z_premigr9checklist.htm
http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db2z10.doc.inst/db2z_migr9checklist.htm

Do I need to REBIND? When?

REBIND is not required for migration to DB2 10, but REBIND is strongly recommended. Getting the best performance improvements and eliminating regression does depend upon rebind in most situations: getting current structures, better access paths, and reusing threads. Eliminating performance regression may depend upon REBIND. Storage constraint relief depends upon REBIND. Changing to use release deallocate

requires a REBIND. The migration process has been smoother when customers REBIND, with fewer problems. All plans containing DBRMs must be rebound. All packages that were last bound on V5 or lower must be rebound. Other REBINDs can be staged over weeks of time, and REBIND is only needed once per package for the migration. Improvements in access paths can be very significant, such as stage 2 predicates that can become stage 1. REBIND in DB2 10 takes more CPU and elapsed time than in prior versions, but more concurrent REBINDs are possible in NFM. So please be kind. REBIND.

Get the right service level

As you upgrade to DB2 10, following the best practices and recommendations can help you have a better experience. The Consolidated Service Test or CST process is being used successfully by very large and small customers. The current service level recommendation is to get to the current quarterly CST level and some specific identified APARs as a minimum level. For example, the January 2011 level is CST4Q10 (RSU1012). RSU1012 includes all service through the end of September 2010 not already marked RSU, PE resolution and HIPER/Security/Integrity/Pervasive PTFs and their associated requisites and supersedes through the end of November 2010. Additional service is noted in the reports. Enhanced HOLDDATA can be used to find current information.

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

Check key information APARs. Check the areas which you use most. System z zIIP processors can make parallel processing less expensive than sequential access, but can cause problems in storage management and parallel processing, so check those areas if you use parallelism.

II14477 and II14474 (V8) note the key DB2 10 migration and fallback APARs.

II10817 indicates the key storage APARs.

II12836 notes the APARs for parallelism.

II14219 indicates the service for zIIPs.

II14334 shows the APARs for large objects or LOBs.

II14426 shows the XML APARs.

II14441 notes recent APARs for incorrect output.

II14203 shows the key APARs related to distributed processing or DDF.

Most DB2 for z/OS customers are large in some metrics, and unique in a number of ways. After getting all of the recommended service, you need to test with your unique workloads and options to find any other needed service. Staging DB2 service levels through development, QA, and production also helps find situations, applications, and techniques that have unique challenges.

Get ready for access path management

DB2 package stability or access path stability function has been very helpful with customers who have access path regression and want to get back to the old access

paths. Read section 10.13 of Packages Revisited for a discussion of the package stability options and techniques. Chapter 4 of Packages Revisited discusses conversion to packages from plans containing DBRMs. DB2 9 for z/OS Packages Revisited, SG24-7688 <http://www.redbooks.ibm.com/abstracts/SG247688.html>

The best techniques for minimizing potential access path issues are to get needed reorganizations done first, followed by new statistics, and then run rebinds, followed by checking. The DB2 9 updated cluster ratio data, new data repeat factor and histogram statistics provide important information for optimization. Run the statistics adviser to generate the recommended statistics to be collected. Use Optim Data Studio to capture the SQL statements and related information.

The default for access path management has changed in DB2 10 from none to extended. Customers who move from DB2 9 and have used access path management have some improvements. Customers who did not have access path management will note increases in BIND CPU time from this change. If you want to reduce the time for BIND, then change the subsystem parameter.

Get the latest information

This is the latest information today, but more is coming soon. See updates and the latest books in standard publications and IBM Redbooks publications. Most questions can be answered quickly and thoroughly by entering search criteria in the search window for the Information Center. You can also subscribe to RSS feeds or download the books in pdf format. Updated information and books are provided only in softcopy. The current DB2 9 books have as many as 7 revisions in just over 3 years. Not having most books in hardcopy can avoid the problem of looking in an obsolete book. Primary DB2 10 product information is viewable from the Information Center at:

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

The first IBM Redbooks publication for DB2 10 is DB2 10 for z/OS Technical Overview, SG24-7892. The second is Extremely pureXML in DB2 10 for z/OS, SG24-7915. Watch for the DB2 10 Performance Topics book, for more updates to DB2 books, and for other information on the web.

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

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

For IBM Redbooks publications, see <http://www.redbooks.ibm.com/>

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

For information about best practices for using DB2 for z/OS, check out the new DB2 for z/OS Best Practices site on DeveloperWorks:

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

Conferences have a range of sessions on DB2 10

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

Introduction to DB2 10 for z/OS

Customers need to reduce costs and adapt quickly to support business growth, without sacrificing the resiliency required for today's demanding business requirements. DB2 10 for z/OS (DB2 10) addresses those needs, building on the capabilities of DB2 9 for z/OS and the System z platform,

DB2 10 for z/OS delivers innovations in these key areas:

Reduce DB2 CPU time for out-of-the-box savings

DB2 10 delivers great value by reducing CPU usage. Most customers can achieve out-of-the-box CPU savings of 5 to 10 percent for traditional workloads and more for some workloads. Improved scalability and constraint relief can add to the savings.

Unsurpassed resiliency for business-critical information

DB2 10 innovations raise the bar on data resiliency through scalability improvements, fewer outages, and improved security. DB2 10 delivers the ability to handle up to five to ten times more active concurrent users in a single DB2 subsystem. Customers can scale-up or scale-out simply, and with less system management. Schema evolution lets you make more changes while business keeps running.

Rapid application and warehouse deployment for business growth

SQL and pureXML enhancements in DB2 10 help productivity, improve performance, and simplify application ports to DB2 for z/OS. DB2 10 adds unique capabilities to support temporal data using business and system time within the database itself, making application development and maintenance simpler and more reliable.

Now let's provide a little more detail to explain the improvements.

DB2 10 overview

Improved performance and reduced CPU for out-of-the-box savings

DB2 10 delivers by improving performance and reducing CPU usage. Most customers can achieve out-of-the-box CPU savings of 5 to 10 percent for traditional workloads and up to 20 percent for specific workloads described below. Measurements compare to previous releases of DB2 for z/OS. REBIND is needed to obtain the best performance and memory improvements. DB2 reduces CPU usage by optimizing processor times and memory access, leveraging the latest processor improvements, larger amounts of memory, and z/OS enhancements. Improved scalability and constraint relief can add to the savings. Productivity improvements for database and systems administrators can drive even more savings.

In DB2 10, performance improvements focus on reducing CPU processing time without causing significant administration or application changes. Most performance improvements are implemented by simply migrating to DB2 10 and rebinding.

You gain significant performance improvements from distributed data facility (DDF) optimization, buffer pool enhancements, parallelism enhancements, and more.

Early DB2 10 performance benchmarking and customer experience has shown a 5 to 10% CPU reduction in transactions after rebinding. Some customers may get more or some less CPU reduction depending on the workload. Customers who have scalability issues, such as virtual storage constraints or latching can see higher improvements. Opportunities for tuning can take advantage of memory improvements. More high volume, short-running distributed transactions can take advantage of CPU reductions, using release deallocate.

Concurrent sequential insert CPU time can be reduced from 5% - 40%. Queries can be improved as much as 20% without access path change, and more for better access paths. A native SQL procedure workload has shown up to 20% CPU reduction using SET statements, IF statements and SYSDUMMY1. Customers moving from DB2 9 should expect a the same CPU times for utilities, while customers moving from DB2 V8 will see CPU reductions up to 20%.

Productivity improvements

New SQL and XML capabilities improve productivity for those who develop new applications and for those who are porting applications from other platforms. Automating, reducing, or eliminating tasks and avoiding manual actions improve productivity and can help avoid problems. Resiliency improvements for virtual storage and availability increase productivity. DB2 10 improvements make the install, migration, and service processes faster and more reliable. Installation and migration information has been improved, using customer feedback.

Flexibility in migration paths

For this release, you can upgrade to DB2 10 directly from a DB2 Version 8 subsystem in new-function mode without starting the system in DB2 9. This provides customers greater flexibility to meet their business needs and to save time getting to DB2 10. Several process improvements make the upgrade simpler.

Unsurpassed resiliency for business-critical information

Business resiliency is a key component of the value proposition of DB2 for z/OS, System z hardware, the z/OS operating system, and other key System z software, like IMS and CICS. Resiliency helps to keep your business running even during unexpected circumstances. Innovations in DB2 10 drive new value in resiliency through scalability improvements and fewer outages, whether those outages are planned or unplanned. Virtual storage enhancements deliver the ability to handle five to ten times more concurrent active users in a single DB2 subsystem than in previous releases of DB2 (as many as 20,000 concurrent active threads). Improved availability is supported by

allowing more changes using schema evolution or data definition on demand. Security improvements also contribute to robust business resiliency.

Continuous availability enhancements

DB2 10 provides online schema enhancements that allow you to make changes to database objects (indexes and table spaces) while maximizing the availability of the altered objects. Through enhancements to ALTER statements, you can now change more attributes of indexes and table spaces without having to unload the data, drop and re-create the objects, regenerate all of the security authorizations, re-create the views, and reload the data. The changes are materialized when the altered objects are reorganized. DB2 10 allows fast changes of table space types, page sizes, data set sizes, and segment sizes. Conversion to universal table spaces is much simpler.

In addition, DB2 10 improves the usability and performance of online reorganization in several key ways. It supports the reorganization of disjoint partition ranges of a partitioned table space (also in DB2 9 now), and improves SWITCH phase performance and diagnostics. Also, DB2 10 removes restrictions related to online reorganization of base table spaces that use LOB columns.

Reduced catalog contention

In DB2 10, the DB2 catalog is restructured to reduce lock contention by removing all links in the catalog and directory. In addition, new functionality improves the lock avoidance techniques of DB2, and improves concurrency by holding acquired locks for less time and preventing writers from blocking the readers of data.

In DB2 10 new-function mode (NFM), you can access currently committed data to minimize transaction suspension. Now, a read transaction can access the currently committed and consistent image of rows that are incompatibly locked by write transactions without being blocked. Using this type of concurrency control can greatly reduce timeout situations between readers and writers who are accessing the same data row.

Virtual storage relief

Enhancements in DB2 10 substantially increase the capacity of a single DB2 subsystem by removing virtual storage and other constraints. This release moves most memory to 64-bit, which provides virtual storage relief and can greatly improve the vertical scalability of your DB2 subsystem while minimizing administration. In addition, a 64-bit ODBC driver is now available on DB2 9 and 10.

Security enhancements

This release of DB2 provides critical enhancements to security and auditing, strengthening DB2 security in the z/OS environment. DB2 10 provides increased

granularity for DB2 administrative authority. DB2 10 delivers a new DB2 data security that enables you to manage access to a table at the level of a row, a column, or both. In addition, you can define and create different audit policies to address the various security needs of your business.

Rapid application and warehouse deployment for business growth

SQL, pureXML®, and optimization enhancements in DB2 10 help extend usability, improve performance, and ease application portability to DB2. DB2 10 delivers significant query improvements, with better performance and CPU reductions, allowing you to manage and maintain your data in a single platform infrastructure with single audit and security processes, and, most importantly, providing a single answer based on your core operational data.

SQL improvements

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

Temporal tables and versioning

In this release of DB2 for z/OS, you have a lot of flexibility in how you can query data based on periods of time. DB2 supports two types of periods, which are the system time (SYSTEM_TIME) period and the business time (BUSINESS_TIME) period. The SYSTEM_TIME period is a system-maintained period in which DB2 maintains the beginning and ending timestamp values for a row. For the BUSINESS_TIME period, you maintain the beginning and ending values for a row. Support of business time and system time allows for significant simplification of applications, pushing the complicated handling of these concepts down to the database engine itself.

In addition, DB2 10 introduces versioning, which is the process of keeping historical versions of rows for a temporal table that is defined with a SYSTEM_TIME period, or both time periods, allowing for simple retrieval of key historical data.

pureXML improvements

DB2 10 improves DB2 family consistency and productivity for pureXML users. These improvements also deliver excellent performance improvements. DB2 10 delivers binary XML format, XML schema validation as a built-in function, XML date and time data types and functions, XML parameters in routines, and much more.

Enhanced business analytics and mathematical functions with QMF

Query Management Facility (QMF) Version 10 has new analytic and mathematical functions and OLAP support. Providing access to many more data sources via JDBC opens QMF to a wider array of information that can be combined with DB2 within the same report.

DB2 10 resource pointers

DB2 main web page <http://www.ibm.com/software/data/db2/zos/>

DB2 10 web page <http://www.ibm.com/software/data/db2/zos/db2-10/>

DB2 books, Information Center

<http://www.ibm.com/support/docview.wss?rs=64&uid=swq27011656>

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

DB2 best practices web page

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


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

DB2 presentations

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

Announcing DB2 10 for z/OS

Savings...right out of the box



Questions and Answers about DB2 10 for z/OS

January 6, 2011

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

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

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

More details are provided on the web:

DB2 main web page <http://www.ibm.com/software/data/db2/zos/>

DB2 10 web page <http://www.ibm.com/software/data/db2/zos/db2-10/>

DB2 books, Information Center

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

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

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

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

DB2 10 booklet

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

The first DB2 10 IBM Redbooks publication is the DB2 10 for z/OS Technical overview. The second book is Extremely pureXML, discussing XML in DB2 9 and 10. The third book will be Performance Topics, anticipated in the second quarter of 2011.

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

IDUG www.idug.org

IOD www.ibm.com/software/data/2010-conference/

Share www.share.org

IBM Education www.ibm.com/training/us/db2zospath

DB2 events www.ibm.com/software/data/db2/zos/events.html

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

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

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

How does DB2 10 use the zEnterprise?

Faster CPUs , more CPUs, and more memory means better DB2 performance and scalability. LSPR measurements of DB2 9 with the zEnterprise z196 show substantial reductions when compared to z10 processors. This transaction workload showed a range of 1.3 to 1.6 times CPU time reduction, with the best CPU reductions when more processors per LPAR are used. Including the larger number of faster processors (80 vs 64) and DB2 10 can mean fewer footprints. This has excellent synergy with DB2 10, which removes many single system scaling inhibitors.

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

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

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

The answer to upgrading to DB2 10 is a definite Yes. The question is not so much whether to upgrade as when and how to upgrade. If you are running DB2 9 today, then DB2 10 is in your near future, giving you more room to grow, with higher limits, lower CPU, and more for less. If you are running DB2 V8 today, then you have a choice of jumping to DB2 9 or directly to DB2 10. Both choices are fully supported with migration and fallback.

While DB2 10 is expected to be better than prior versions, it will have maturity and service delivery like other software, with more defects at first, then fewer as the software matures. Determining when the software is ready for a specific customer and when the customer is ready for the software depends upon the specific customer resources for testing, prior experience, and the value for the improvements versus the need for stability. Many customers depend upon tools or other software, and having that software work with DB2 is a prerequisite. Customers who migrate early will need more robust test and service processes.

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

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

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

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

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

When will DB2 V8 go out of support?

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

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

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

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

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

How should I start getting ready for DB2 10?

Are you ready for DB2 10? Are you running in DB2 9 NFM or DB2 V8 NFM? Have you looked though the information APARs? Have you contacted your

vendors? Do you have the required service applied to your current system? Have you run the premigration job and worked to resolve the incompatibilities? Have you gotten rid of private protocol? Have you converted to packages instead of DBRMs directly in plans? Those are first tasks, along with planning what you expect to gain and the process for testing the new version.

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

Do I need to REBIND? When?

REBIND is not required for migration to DB2 10, but REBIND is strongly recommended. Getting the best performance improvements and eliminating regression does depend upon rebind in most situations: getting current structures, better access paths, and reusing threads. Eliminating performance regression may depend upon REBIND. Storage constraint relief depends upon REBIND. Changing to use release deallocate requires a REBIND. The migration process has been smoother when customers REBIND, with fewer problems. All plans containing DBRMs must be rebound. All packages that were last bound on V5 or lower must be rebound. Other REBINDs can be staged over weeks of time, and REBIND is only needed once per package for the migration. REBIND in DB2 10

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

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

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

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

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

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

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

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

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

improvement. High-volume, short running distributed transactions can take advantage of CPU reductions, using RELEASE(DEALLOCATE). This is the scenario for a benchmark transaction that is run on DB2 9, then on DB2 10. This scenario uses some new function in DB2 9 to BIND or REBIND a package with access control management to allow three copies. These are fairly light CICS transactions that have been used for many DB2 transaction benchmarks.

In step 1, this application is moved to DB2 10 CM without a REBIND, and the result is a 3.7% reduction in CPU time.

In step 2, still in DB2 10 CM, a REBIND is performed but with exactly the same access path. With the REBIND, the CPU savings over DB2 9 was 7.4%, double that without the REBIND.

In step 3, moving to NFM, the CPU time is the same.

In step 4, these transactions are changed to use RELEASE(DEALLOCATE), saving an additional 10% of the CPU time compared to the prior RELEASE(COMMIT). This change could also be made in CM with step 2.

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

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

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

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

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

Did some customers find performance regression?

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

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

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

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

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

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

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

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

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

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

Skip release is costly for developers and complex, so it's not for every version. DB2 delivered skip for V5 to V7 and now for V8 to V10. I think that most customers are best served by staying current with DB2 versions, but some customers get behind. Skip migration is great for customers running V7 today or for those who just finished migrating to V8. Customers who plan to migrate to DB2 9 in 2010 or 2011 should generally do so. The presentation and paper on migration have much more information.

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

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

Will EXPLAIN show hash access if chosen over index access?

Explain shows hash access. See the presentation on hash access.

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

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

How is the hash access path established?

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

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

MEMBER CLUSTER is not permitted for universal table spaces (UTS) in DB2 9 and is for DB2 10. With DB2 10, you can ALTER to a UTS, then REORG to change the table space to MEMBER CLUSTER, using a pending ALTER.

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

PC=YES has been the only option since V8.

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

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

During PENDING ALTER, is the object is accessible?

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

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

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

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

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

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

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

Does DB2 10 support private protocol?

No. See the upgrade presentation and paper.

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

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

How small is a small LOB?

A small LOB is able to fit on the same page size along with its other associated data column information. This avoids creating a separate table space for the LOB data. The desired is to fit these small LOBs within their 4k, 8k, 16k or 32k page sizes so that no additional I/Os are needed to store or retrieve the LOB data. Typically these small LOBs can be anywhere from 2k to 30k so that they fit on the same page.

What are the improvements for BLOB and CLOB?

Many improvements were made in DB2 10 for LOBs. Inline LOBs, improved LOB streaming to and from DDF are key performance improvements. Indexes can be defined on the inline LOB. Better streaming and minimized LOB default sizes help improve performance when LOB materialization occurs. This is especially important for DDF applications because of the network impact of large objects. DEFINE NO for LOB and XML columns makes definition easier. REORG SHRLEVEL(CHANGE) for LOBs improves availability. Online REORG permits movement of rows between partitions when LOB columns exist, and allows DISCARD to delete associated LOB values. Non-NULL default values can be defined for inline LOBs. Now you can stream LOBs into one data set, rather than taking a data set for each LOB.

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

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

Will this change the DB2 Utilities?

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

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

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

Can solid state drives be used with hardware replication?

Yes. Here are some pointers:

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

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

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

<ftp://ftp.software.ibm.com/common/ssi/sa/wh/n/tsw03044usen/TSW03044USEN.PDF>

<https://www.ibm.com/developerworks/mydeveloperworks/blogs/InsideSystemStorage/tags/announcements?lang=en>

Can you explain the common SDSNEXIT in DB2 10?

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

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

Why one would need to select the SYSLGRNX?

SYSLGRNX has been needed for some recovery situations.

Will R10 restrict BINDING a plan to a single DBRM?

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

What happens for existing DBRMs bound into plans?

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

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

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

Packages Revisited, SG24-7688

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Does BIND time increase in DB2 10?

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

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

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

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

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

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

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

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

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

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

The DB2 10 dynamic statement cache (DSC) enhancements consolidate all types of SQL statements and their literals. So if the same statement is used

within the application with different literals DB2 will consolidate it within the DSC leaving more room for other statements and better caching overall. This behavior requires a parameter to be set, as literals can provide important information for optimization.

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

ACCUMACC, commonly called rollup, combines multiple accounting records into one. The rollup technique is being improved to work better with packages, allowing more customers to use this option. Compression uses an SMF compression option. Customers can use either one or both techniques to reduce SMF accounting volume.

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

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

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

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

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

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