


Total Solution Event for System z™ 2012

Innovate Transform Grow

IBM DB2 10 Migration Planning and Very Early Migration Experiences – Part 2

John Campbell
DB2 for z/OS Development



Notices

This information was developed for products and services offered in the U.S.A.

Note to U.S. Government Users Restricted Rights — Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to: IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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.

This information contains 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.

Copyright License:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

- This presentation contains trade-marked IBM products and technologies. Refer to the following Web site:

<http://www.ibm.com/legal/copytrade.shtml>

Objectives

- Share lessons learned, surprises, pitfalls
- Provide hints and tips
- Address some myths
- Provide additional planning information
- Provide usage guidelines and positioning on new enhancements

Agenda

- Keys to customer migration success
- Performance and Scalability
- BIND, REBIND and EXPLAIN
- Availability
- Migration Planning
- Removal of DDF Private Protocol
- Security Considerations when removing DDF Private Protocol
- Other
- Summary

Performance and Scalability ...

- High INSERT performance
 - Significant improvements for UTS
 - Now support for MEMBER CLUSTER
 - Changes to space search algorithm (like classic partitioned)
 - Goal was for UTS to be equal or better than classic partitioned (PTS)
 - Not there yet, but much closer
 - Very workload dependent
 - Some good, some worse
 - Still trade off between space vs. throughput and reduced contention
 - Work still to do on UTS PBR/PBG with RLL and sequential insert

Performance and Scalability ...



- High INSERT performance ...
 - Reduced LRSN spin for inserts to the same page
 - Works well for MRI and INSERT within loop in a data sharing environment
 - Optimization for 'pocket' sequential insert works well
 - Index manager picks the candidate RID during sequential insert (next lowest key rid)
 - Higher chance to find the space and avoiding a space search
 - Parallel index IO works very well when activated for random key inserts
 - ≥ 3 indexes
 - Prefetch and deferred write offload to zIIP to compensate

High Insert Workload Description



- 2-way data sharing
- Database schema
 - 3 tables with total of 6 indexes (4 unique, 2 non-unique indexes, 2 secondary indexes)
 - Table space types: Classic Partitioned, Classic Segmented, UTS (PBR, PBG)
- SQL
 - INSERTs contain 5, 9 and 46 columns of integer, bigint, char, varchar, decimal and timestamp data type
- Application implemented in Java
- Sequential inserts into empty tables
 - 240 concurrent threads
 - Multi-row inserts (100)
- Random inserts into populated tables
 - 200 concurrent threads
 - Single-row inserts

DB2 10 Range Defined Table Spaces

Random Inserts





Sequential Inserts


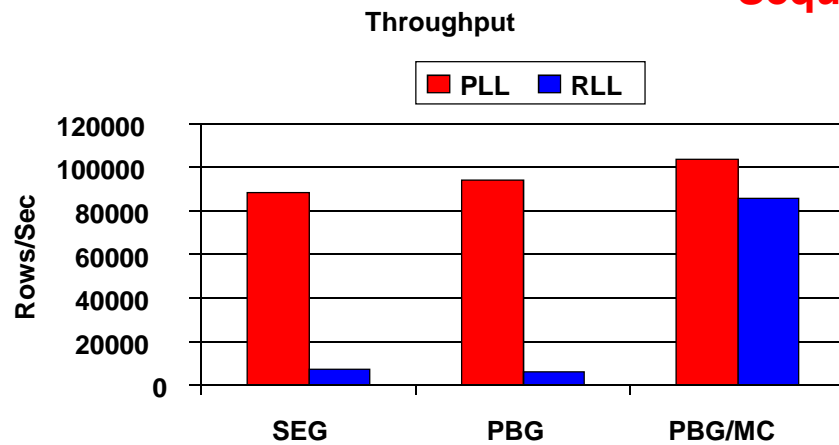


DB2 10 Non-range Defined Table Spaces

Random Inserts



Sequential Inserts



Performance and Scalability ...

- Accounting Trace Class 3 enhancement – separate counters
 - IRLM Lock/Latch waits
 - DB2 Latch waits
- Data sharing
 - Faster DB2 shut down by avoiding local buffer pool scan per GBP-dependent object
 - Avoiding scan of XXXL local Buffer pool when
 - Pageset/partition transitions into GBP-dependency
 - Pageset/partition transitions out of GBP-dependency
- Inline LOBs work very well if you hit the sweet spot
 - Potential for significant CPU and elapsed time improvement with the right inline value
 - Trade off in setting the right inline value
 - Avoiding access to auxiliary tablespace
 - Increasing base row size with fewer rows per page
 - May have to increased page size
 - Inline portion can be compressed
 - Significant space savings with small LOBs (<1000 bytes)

Performance and Scalability ...

- Compress on INSERT
 - Compression ratios almost as good compared with running REORG later
- Active log writes
 - Prior to V10, log writes are done serially when re-writing partial CIs
 - Determined that destructive writes due to IO errors no longer occur
 - Now all log write IOs are done in parallel
 - Elapsed time improvements
- Limited value of Currently Committed locking semantics
 - SQL UPDATE not supported

Performance and Scalability ...

- Hash access vs. Index only access and index lookaside
 - Competes against index only access and index lookaside
 - Advantage that index only access still provides for clustered data access
 - Can now have unique index with INCLUDE columns
 - Reduce number of indexes required for performance reasons
 - Improve insert, update and delete performance
 - Need to find the sweet spot
 - High NLEVELS in index (≥ 3)
 - Purely direct row access by primary key
 - Truly random access
 - Read intensive, not volatile
 - No range queries
 - Many rows per page etc
 - Space allocation of fixed hash space is key to control overflow
 - Too small will lead to rows in overflow
 - Too large will lead to random IO
 - REORG AUTOESTSPACE(YES) but still some rows in overflow
 - Degraded LOAD and REORG utility performance

Performance and Scalability ...

- Improved index space search when index leaf page split
 - In V8/V9, Index Manager
 - First searches the space map page covering the splitting page
 - If there is no free entry, searches the space map pages starting from the first space map page to the highest allocated page
 - If all the space map pages are full, has to extend
 - If the index is huge and all the space map pages having free entries are toward the end of the index, this process can take a very long time
 - In V10, Index Manager
 - After searching the space map page covering the splitting page, and if it is full, will start searching from the page number it last remembered having the free entry in it (page A)
 - When it reaches to the highest allocated page, it starts from the beginning and search forward till it reaches to page A
 - Then it has to extend since the entire index is full
 - The page number of the space map page having free entry is stored in an in-memory control block
 - When Index Manager finds a space map page with free entry, it is updated to be the page number of that space map page
 - Index Manager updates this value when an index page is deleted or when the index is mass deleted
 - Retrofitted back to V9 via APAR PM15474

Performance and Scalability ...

- Comparing CPU performance on V10 relative to V8 or V9
 - More difficult to do in real customer production environment
 - Uncertainty caused by application changes
 - Fluctuation in the daily application profile especially batch flow
 - Must try to normalise things out to ensure workloads are broadly comparable
 - Broadly similar in terms of SQL and getpage profile
 - Usually have to exclude the batch flow
 - Factor out extreme variation
 - Need to look at multiple data points

BIND, REBIND and EXPLAIN

- Value of REBIND under V10
 - Improved performance from new run time (avoid puffing, enable SPROC)
 - Maximize DBM1 31-bit VSCR
 - Allow RID overflow to workfile
 - Take advantage of query optimization changes (available in CM mode)
 - Reduce exposure to problems with migrated packages from earlier releases
 - INCORROUTs
 - Thread abends
- Can mitigate exposure to bad access path change introduced with REBIND which leads to degraded run time performance (regression)
 - Use access plan stability (PLANMGMT=EXTENDED|BASIC) and fallback if needed
 - PLANMGMT=EXTENDED is now the default
 - Use APREUSE and APCOMPARE

BIND, REBIND and EXPLAIN ...

- Access Path Reuse - APREUSE and APCOMPARE
 - Introduced in V10 through APARs PM33767 (function) and PM25679 (enabled)
 - Works independently of access plan stability (PLANMGMT=EXTENDED|BASIC)
 - Provides ability to generate new SQL run time but will attempt to reuse the existing access path by internally using OPTHINTS
 - Applies to both BIND REPLACE and REBIND
 - Available in CM when migrating from V9 -> V10 and then applies to subsequent BIND/REBIND in V10
 - Only works when package previously bound on V9 or higher
 - Starting with V9, EXPLAIN information is saved away inside the package in SPT01
 - Referred to as “Explain Data Block (EDB)” which is a compact representation of PLAN_TABLE information
 - Not guaranteed to succeed in all cases as OPTHINTS are not enforceable 100% of the time
 - Some incompatibilities between old and new release (restrictions)
 - Seeing 95-99% compatibility

BIND, REBIND and EXPLAIN ...

- Access Path Reuse - APREUSE and APCOMPARE ...
 - Conservative approach where minimal access path changes are required
 - Step 0 (Optional)
 - Use REBIND ... **EXPLAIN(ONLY) + APREUSE(ERROR)**
 - Perform an impact analysis before actual REBINDs
 - Step 1: REBIND PACKAGE (*)
 - Use PLANMGMT(EXTENDED) ... backup of V9 access paths, just in case
+ EXPLAIN(YES)
+ APREUSE(ERROR)
 - Step 2: For packages that failed Step 1 (i.e., leftovers)
 - 2a: Leave them as is ... they will be at the old level
 - OR
 - 2b. REBIND with PLANMGMT(EXTENDED) + APREUSE(NO)
 - > This step exposes yourself to access path changes
 - > But you have a backup

BIND, REBIND and EXPLAIN ...

- RUNSTATS/REBIND recommendations based upon on scenario
 - V8 preparation
 - If RUNSTATS will be difficult on large number of objects immediately after migration to V9/10, then REORG and/or RUNSTATS (V8) immediately prior to migration can reduce RUNSTATS need on V9/10 - as RUNSTATS INDEX under V10 can be sufficient to capture new CR/DRF
 - V9 migration
 - RUNSTATS objects as soon as possible after migration
 - Target dynamic applications first as these are exposed to new access paths immediately
 - Delay static REBINDs until associated objects have RUNSTATS run
 - V8->V10 migration
 - RUNSTATS objects as soon as possible after migration
 - Target dynamic applications first as these are exposed to new access paths immediately
 - Equal priority - target static parallelism packages to REBIND to avoid incremental bind at each execution
 - Delay non-parallelism REBINDs until associated objects have RUNSTATS run
 - V9->V10 migration
 - REBIND static parallelism packages as soon as possible to avoid incremental bind at each execution
 - Delay non-parallelism REBINDs until associated objects have RUNSTATS run
 - BIND/REBIND options APREUSE/APCOMPARE are available on V10 for packages bound on V9

BIND, REBIND and EXPLAIN ...

- RUNSTATS/REBIND recommendations based upon on scenario ...
 - V8/9 co-existence
 - Set STATCLUS=STANDARD while in co-existence with V8
 - Set ABIND=COEXIST while in co-existence with V8
 - Avoid REBIND while in co-existence with V8
 - Follow V9 migration steps after all members are V9, including resetting the following zparms
 - Set STATCLUS=ENHANCED
 - Set ABIND=YES
 - V8/10 co-existence
 - Set STATCLUS=STANDARD while in co-existence with V8
 - Set ABIND=COEXIST while in co-existence with V8
 - What to do with static parallel queries?
 - Accept incremental bind whenever executed on V10 member
 - OR, REBIND with DEGREE('1') to disable parallelism while in co-existence.
 - Follow V8-V10 migration steps after all members are V10, including resetting the following zparms
 - Set STATCLUS=ENHANCED
 - Set ABIND=YES

BIND, REBIND and EXPLAIN ...

- RUNSTATS/REBIND recommendations based upon on scenario ...
 - V9/10 co-existence
 - Set ABIND=COEXIST while in co-existence with V8
 - What to do with static parallel queries?
 - Accept incremental bind whenever executed on V10 member
 - OR, REBIND with DEGREE('1') to disable parallelism while in co-existence.
 - Follow V8-V10 migration steps after all members are V10, including resetting the following zparms
 - Set ABIND=YES

BIND, REBIND and EXPLAIN ...

- Single thread BIND/REBIND performance
 - Degraded CPU and elapsed time performance on entry to CM
 - PLANMGMT=EXTENDED is now default
 - New indexes defined for post ENFM when hash links are eliminated
 - Change in access path (index access) on entry to CM
 - No concurrency improvement until after Catalog restructure in ENFM
- Concurrent BIND/REBIND performance
 - Problems addressed
 - Performance problems related to DELETE/INSERT process
 - Space growth in SPT01 for both LOB space and base table
 - Now working well
 - Inefficient space search for out of line LOB in data sharing (APAR PM24721)
 - Inline LOB with compression for SPT01 to address SPT01 growth (APAR PM27073)
 - More efficient space reuse for base table and UTS (APAR PM27973)
 - Recommendations
 - Customers need to change existing procedures to go parallel
 - But cannot do this until post ENFM
 - Benefit from reducing application down time to implement new application releases

BIND, REBIND and EXPLAIN ...

■ EXPLAIN tables

- Format and CCSID from previous releases is deprecated in V10
 - Cannot use pre V8 format
 - SQLCODE -20008
 - V8 or V9 format
 - Warning SQLCODE +20520 regardless of CCSID EBCDIC or UNICODE
 - Must not use CCSID EBCDIC with V10 format
 - EXPLAIN fails with RC=8 DSNT408I SQLCODE = -878
 - BIND with EXPLAIN fails with RC=8 DSNX200I
- Recommendations
 - Use CCSID UNICODE in all supported releases (V8, V9, V10) due to problems with character truncation and conversion etc
 - Use the V10 extended column format with CCSID UNICODE when
 - Applications access EXPLAIN tables and can only tolerate SQLCODE 0 or +100
 - V10 column format is supported under V8 and V9 with the SPE fallback APAR PK85956 applied with the exception of
 - DSN_STATEMENT_CACHE_TABLE due to the BIGINT columns
- APAR PK85068 can help migrate V8 or V9 format to the new V10 format with CCSID UNICODE

Capturing documentation for IBM

- Methods for capturing documentation for all releases is documented here
 - <https://www.ibm.com/support/docview.wss?uid=swg21206998>
 - OSC and DB2PLI8 do not support DB2 10
- SYSPROC.ADMIN_INFO_SQL supports V8 -> V10 (**Required**)
 - Excellent developerWorks article here:
 - <http://www.ibm.com/developerworks/data/library/techarticle/dm-1012capturequery/index.html>
 - It is installed in V10 base and is subject to the installation verification process
 - *DB2HLQ.SDSNSAMP(DSNTE SR)* will create and bind it
 - calling program is DSNADMSB, and sample JCL in DSNTEJ6I
 - Ensure DB2 9 and DB2 10 have APAR PM39871 applied
- Data Studio V3.1 incorporates this procedure into a GUI (**Best Practice**)
 - <http://www.ibm.com/developerworks/downloads/im/data/>
 - No charge product, replacement for OSC and Visual Explain
 - Several versions:
 - DBA's should download the Administration Client
 - Incorporates Statistics Advisor
 - FTP doc directly to DB2 Level 2
 - Can be used to duplicate stats in TEST environment

Component / Operating system	Version	Size
Administration client	V3.1	359MB - 362MB
Red Hat Linux, SUSE Linux, Windows		

Availability

- Online Schema Evolution ('Deferred Alter')
 - Migrate from classic table space types (simple, segmented, partitioned) to UTS PBG/PBR
 - One way ticket only
 - UTS is pre-requisite for Cloned Table, Hash, Inline LOB, Currently Committed
 - Once migrated to UTS PBG/PBR can change attributes in both directions
 - DSSIZE, index page size, MEMBER CLUSTER, Hash Access, ...
 - Benefits
 - Streamed line way to move to UTS
 - Reduce administrative time and cost
 - Cuts down on errors
 - Reduce outages
 - Issue that PIT recovery to point before successful materializing REORG not possible
 - Incorrect results from REORG
 - Application change rollback

Availability ...

- Restart Light enhancement
 - LBACKOUT will now be honoured
 - LBACKOUT=YES|AUTO will cause postponed abort (PA) URs to be created
 - Restart will complete
 - DB2 will shut down
 - Retained locks will be kept on behalf of PA URs
 - Controlled via new system parameter
 - Also retrofitted back to V9 via APAR
- Online REORG with FORCE
 - Only running threads which are blocking are cancelled
 - Threads which are suspended / inactive will cause REORG to still fail
- Online REORG LOB with DISCARD
 - Cannot handle LOB columns greater than 32KB

Migration and Planning

- Migration process very similar to V8 and V9
 - Works well with few problems with migration fallback
- Migration from either DB2 for z/OS V8 NFM or DB2 9 for z/OS NFM
- These migration fallback sequences are not valid
 - V8 NFM > V10 CM8 > V8 NFM > **V9 CM**
 - V8 NFM > V9 CM > V8 NFM > **V10 CM8**
- Fallback Toleration SPE
 - APAR PK56922
- Early Code
 - For V8/V9 APAR PK87280 (superseeds APAR PK61766)
- Information APARs
 - II14474: V8 to V10
 - II14477: V9 to V10

Migration and Planning ...

V8 premigration checklist:

http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db2z10.doc.inst/src/tpc/db2z_premigr8checklist.htm

V8 migration checklist:

http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db2z10.doc.inst/src/tpc/db2z_migr8checklist.htm

V9 premigration checklist:

http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db2z10.doc.inst/src/tpc/db2z_premigr9checklist.htm

V9 migration checklist:

http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db2z10.doc.inst/src/tpc/db2z_migr9checklist.htm

Migration and Planning ...

- Use of V10 Early Code with V8
 - It will take an IPL to originally install the V10 Early Code
 - V8 Early Code does not understand –REFRESH
 - However, subsequent maintenance to the V10 Early Code can be accomplished with a -REFRESH command
- If coming from V8
 - BSDS must be reformatted for larger active / archive tracking
- IPL amounts for need to be adjusted based on number of DB2 members
 - 64-bit Private (1TB)
 - 64-bit Shared (128GB)
 - 64-bit Common (6GB)

Migration and Planning ...

- DB2 Connect
 - Minimum level
 - V9.5 FP7
 - V9.7 FP3A and for new functions
 - Start with the latest levels based on CST/RSU and stabilise

Migration and Planning ...

- DBRMs bound directly into plans no longer supported
 - If found in V10, will trigger auto bind into packages
 - For V8 and V9
 - APARs PK62876/PK79925 adds new syntax to convert from DBRMs to packages
 - REBIND PLAN option COLLID
 - Could result in access path change
 - APARs PM01821 (Version) and PM30382 (Location from * to blank) should be on
 - Best to migrate DBRMs to packages before migrating to V10
- Old plans and packages bound prior to V6 will require REBIND
- Catalog and Directory must be SMS managed (EF, EA) ahead of CM
- PDSEs required for SDSNLOAD, SDSNLOD2, ADSNLOAD
- DSNHDECP NEWFUN=V10|V9|V8

Removal of DDF Private Protocol

- Must absolutely eliminate all use of DDF Private Protocol before migrating
 - No longer supported In V10
 - Any local packages miss tagged with DDF Private Protocol will be tolerated
 - Otherwise package must exist in both local and remote sites
 - A lot of packages and plans are bound with DBPROTOCOL(PRIVATE) because this was the default (zparm DBPROTCL) when introduced in DB2 V6
 - DSNT226I is issued if DBPROTOCOL(PRIVATE) is used during REBIND

Removal of DDF Private Protocol ...

- Must absolutely eliminate all use of DDF Private Protocol before migrating
 - APAR PK92339 in V8 introduces new zparm PRIVATE_PROTOCOL=YES|NO
 - Must be in sync with zparm setting of DBPROTCL
 - DSNT225I message to indicate private protocol failure for REBIND
 - IFCID 157 can be used to identify packages which are using private protocol
 - Beware APAR PM17665 removes authorization behaviour for private protocol i.e., .
 - Execute privilege on package is required for the user running the program under DRDA
 - APAR PM37300 will provide relief (see later slides)
 - APAR PK64045 (PE, PM14816) in V8 new zparm DRDA_RESOLVE_ALIAS
 - Private protocol always performs alias resolution
 - 3-part name is sent to the server
 - If DRDA_RESOLVE_ALIAS is set to YES then also DRDA connections get this resolution e.g.,
 - For remote BIND with COPY option

Security considerations when removing DDF Private Protocol

- This section also applies to customers using DRDA exclusively
- There are fundamental differences on how authorization is performed based on the distributed protocol used
- Private Protocol (DB2 for z/OS requester)
 - Supports static SQL statements only
 - Plan owner must have authorization to execute all SQL executed on the DB2 server
 - Plan owner authenticated on DB2 requester and not on the DB2 server
- DRDA Protocol
 - Supports both static and dynamic SQL statements
 - Primary auth ID and associated secondary auth IDs must have authorization to execute package and dynamic SQL on the DB2 server
 - Primary auth ID authenticated and secondary auth IDs are associated on DB2 server
- Prior to V10, Private Protocol and DRDA Protocol can be used by same application
 - Private Protocol security semantics was used due to possible inconsistent behavior which is dependent on how programs are coded and executed

Security considerations when removing DDF Private Protocol ...

- But there is also a difference prior to V10 in the authorizations required by an incoming DRDA connection at the DB2 for z/OS server, depending on where the connection come from:
 - Dynamic SQL DRDA connection from DB2 Connect and/or DB2 client direct connection
 - Connecting userid needs authority to run the appropriate DB2 package and authority to access the DB2 table
 - Dynamic SQL DRDA connection from DB2 for z/OS requester
 - Connecting userid needs authority to access the DB2 table
 - Originating plan owner needs authority to run the appropriate DB2 package
- It is different for DB2 for z/OS requester to DB2 for z/OS server because connections were designed to use Private Protocol (PP) semantics to avoid changing authids when switching between PP to DRDA Protocol
- With the disappearance of PP in V10, DB2 have decided to bring the DRDA connection from DB2 for z/OS requester to DB2 for z/OS server in line with other DRDA requesters and to change the authorizations required
 - This was retrofitted back into V8 and V9 with APAR PM17665
 - It is very important to distinguish clearly between the behavior of DRDA before and after APAR PM17665

Security considerations when removing DDF Private Protocol ...

- APAR PK92339 introduced new zparm PRIVATE_PROTOCOL=YES|NO
 - To prevent future introduction of PP then set PRIVATE_PROTOCOL=NO
- The result of migrating to V10 or the introduction of APAR PM17665 under V8 or V9, when running with PRIVATE_PROTOCOL=NO introduces the authorization changes at the DB2 for z/OS server for DRDA connections coming from DB2 for z/OS requester
 - PP security semantics are no longer used as default for access from a DB2 for z/OS requester
 - Plan owner value is ignored and connecting userid must be granted authority to execute the package at the remote site
 - Otherwise the connection will fail with SQLCODE -551
- As a result of customer complaints, APAR PM37300 introduces PRIVATE_PROTOCOL=AUTH which allows an installation to
 - Disable PP but keep the plan owner authorization check (the "private protocol semantics")
- Migration to V10 or the application of PTF for APAR PM17665 does affect you even if you have everything already bound as DRDA

Security considerations when removing DDF Private Protocol ...

- In summary
 - Before disabling private protocol, ensure all appropriate grants are performed
 - Grant execute privilege to any user who plans to run a package or stored procedure package from a DB2 for z/OS requester, just like other DRDA clients
 - DB2 V8 and V9 can disable private protocol but still maintain private protocol authorization checks by
 - Setting system parameter PRIVATE_PROTOCOL=AUTH
 - DB2 10 does not support private protocol but can allow private protocol authorization checks for use of DRDA protocol for DB2 for z/OS requesters by
 - Setting system parameter PRIVATE_PROTOCOL=AUTH

Other

- Ability to create classic partitioned table space (PTS)
 - Classic PTS deprecated in V10
 - By default will be created as UTS PBR
 - UTS will only support table based controlled partition syntax
 - Options to be able to officially create classic PTS
 - Specify SEGSIZE 0 on CREATE TABLESPACE
 - Set new zparm DPSEGSZ=0 (default 32)
- Fast Log Apply storage
 - System parameter LOGAPSTG is eliminated
 - Total FLA storage is now 510MB
- Old COBOL and PL/1
 - V7 lookalike pre-compiler (DSNHPC7) for older COBOL and PL/I is still provided
- DDL Concurrency after Catalog restructure
 - Some help provided but concurrency issues not absolutely solved
 - Still deadlocks with parallel heavy DDL against different databases

Other ...

- SPT01 compression is back
 - Via system parameter COMPRESS_SPT01=YES (default is NO)
- Statistics Interval
 - Only the frequency of IFCIDs 105, 106, 199 are controlled via STATIME
 - IFCIDs 2, 202, 217, 225, 230 are always cut at a fixed 1 minute interval
 - Remember to normalise the data when comparing V8/V9 vs. V10

Other ...

■ RUNSTATS

– Page Sampling

- Performance improvement can be phenomenal
- Potential issues with accuracy because error rates increase as the sample size decreases
- No sampling done on indexes

– zIIP offload

- Nearly all RUNSTATS INDEX processing is offloaded, but only 'basic' RUNSTATS TABLE processing is offloaded
- Much less for advanced/complex statistics
 - Not supported:
 - > Inline stats
 - > COLGROUP
 - > DSTATS
 - > Histogram stats

Other ...

- Incompatible changes

- CHAR() scalar function

- Solution via APAR PM29124 to restore compatible behavior for CHAR

```

SELECT COL1
       ,CHAR(COL1)
       ,HEX(CHAR(COL1))
       ,SUBSTR(CHAR(COL1), 2, 4) APPL_SUBSTR
FROM DB2_TABLE;

```

COL1	CHAR_V9	HEX_CHAR_V9	APPL_SUBSTR	COL1	CHAR_V10	HEX_CHAR_V10	APPL_SUBSTR
9	0009,	40F0F0F0F96B	0009	9	9	F94040404040	
99	0099,	40F0F0F9F96B	0099	99	99	F9F940404040	9
999	0999,	40F0F9F9F96B	0999	999	999	F9F9F9404040	99
9999	9999,	40F9F9F9F96B	9999	9999	9999	F9F9F9F94040	999

Other ...

- Incompatible changes ...

- Create Trigger

- Impact: Trigger created in V10 not executable in V9
 - Workaround: drop and create trigger in V9
 - PMR open

```
DSNT408I  SQLCODE = -723, ERROR: AN ERROR OCCURRED IN A TRIGGERED SQL  
STATEMENT IN TRIGGER AWAT.OCA_TEST_TRIGGER. INFORMATION RETURNED: SQLCODE -  
904, SQLSTATE 57011, MESSAGE TOKENS 00E30305,00000801,AWAT.OCA_TEST_TRIGGER-  
.18E5950B04A23EEC, SECTION NUMBER 1
```

Other ...

- Incompatible changes ...

- SQL stored procedures

```
SQLCODE = -904, ERROR: UNSUCCESSFUL EXECUTION CAUSED BY AN  
UNAVAILABLE RESOURCE. REASON 00E7009E, TYPE OF RESOURCE 00000801,  
AND RESOURCE NAME BE2TFKT.LESE_VARIANTEN.18DEB14D119115C4
```

- Impact: Deployment has to be executed on V9 member in V10
- Workaround: run ALTER PROCEDURE REGENERATE on V9 member
- APAR PM13525

Other ...

- DB2 for z/OS Exchange on DeveloperWorks
 - Relaunch and rebranding of ' DB2 zHotline'
 - It is a DB2 Technical Newsletter for DBAs and Systems Programmers
 - Critical APAR info (more than HIPER, less than Red Alert)
 - Living undocumented incompatible change list
 - Critical issues seen in field and ways to navigate
 - Hints, tips, practical advice
 - Will be updated monthly
 - RSS Feed from DeveloperWorks into IDUG
 - Link
<https://www.ibm.com/developerworks/mydeveloperworks/groups/service/html/communityview?communityUuid=22586cb0-8817-4d2c-ae74-0ddcc2a409bc>

Summary

- Very good release in terms of the opportunities for price/performance and scalability improvement
 - Significant DBM1 31-bit VSCR after rebind
 - Use long term page fixed buffer pools
 - Exploit 1MB real storage page frames on z10 and z196
 - Reduced latch contention, log manager improvements, etc
 - Opportunity for further price performance improvements
 - More use of persistent threads
 - CICS, IMS/TM, High Performance DBATs
 - More use of RELEASE(DEALLOCATE) with persistent threads
 - More use of RELEASE(DEALLOCATE) is a trade off
 - Increased storage consumption
 - > Need to plan on additional real memory
 - Reduced concurrency
 - > BIND/REBIND and DDL
 - Increase MAXKEEPD to reduce short prepares for dynamic SQL
 - Opportunity for scale up and LPAR/DB2 consolidation

Summary ...

- Carefully plan, provision and monitor real storage consumption
- Any customer migrating from either V8 or V9 to V10 should make a solid plan, take extra care to mitigate against the risk and set themselves up for success
 - Regular full ‘major’ maintenance drops
 - Exploitation of CST/RSU recommended maintenance
 - Augment by regular use of Enhanced HOLDDATA
 - Perform application regression and stress testing to keep ‘fires’ away from production
 - Plan should allow some contingency for some ‘bumps in the road’

More information on zEnterprise

- IBM zEnterprise / System z Redbooks Portal: <http://www.redbooks.ibm.com/portals/systemz>
- IBM zEnterprise Announcement Landing Page: ibm.com/systems/zenterprise196
- IBM zEnterprise HW Landing Page: ibm.com/systems/zenterprise196
- IBM zEnterprise Events Landing Page: ibm.com/systems/breakthrough
- IBM Software: ibm.com/software/os/systemz/announcements
- IBM System Storage: ibm.com/systems/storage/product/z.html
- IBM Global Financing: ibm.com/financing/us/lifecycle/acquire/zenterprise/
- Global Technology Services:
<http://www.ibm.com/services/us/index.wss/offerfamily/gts/a1027714>

Bedankt!

