

# DB2 10 for z/OS Planning and Very Early Experiences

Kevin Harrison  
Sr. IT Architect / Sr Software Engineer  
DB2 for z/OS Development  
[kharrison@us.ibm.com](mailto:kharrison@us.ibm.com)



## 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

- **Introduction**
- **Highlights**
- **Performance and Scalability**
- **Availability**
- **Other**
- **Migration and Planning**
- **Migration Fundamentals**
- **Security considerations when removing DDF Private Protocol**
- **Items planned for post-GA delivery**
- **Summary**

## Highlights

### ■ Good Results

- DBM1 virtual storage constraint relief
- Insert performance
- Hash Access good when hitting the smaller than expected sweet spot
- Complex queries
- Inline LOBs (SLOBs)
- Latch contention reduction
- Quality of problems and issues found
- Improved reliability and confidence as program progressed

## Highlights ...

- **Mixed Results**

- OLTP performance, mostly good, some bad
- Single thread BIND/REBIND performance
- DDL concurrency
- Access path lockdown

## Highlights ...

- **Mainly positive customer experience and feedback about the program**
- **Majority of customers planning to start migration to V10 in 2011**
- **Incremental improvement over V8 and V9 programs**
- **No single voice / messages across the customers**
- **Significant variation in terms of customer commitment and achievement**
  - Subset of customers did a very good job on regression and new function testing
    - Good give back
  - Other customers
    - Limited qualification about what they were going to do
    - Limited qualification about what they did and what they achieved
- **Need to appreciate that it is hard for customers to sustain effort over 6 period based on business and technical priorities**
  - People / hardware resources and time are constrained
- **No customers in true business production by end of program**
- **Need to appreciate difference between QPP/Beta vs. ESP**

## Highlights ...

- **Many opportunities for price/performance (cost) improvements**
  - Major theme of this release
  - Most welcome to our customers
- **Customers intimidated by the marketing noise about improved performance**
  - Expectation of their CIO
  - For some of their workloads not seeing improvements in CPU and elapsed time
  - Conversely see big improvements for certain workloads
  - Small workloads can skew expectations on savings
  - Some measurements and quotes are insanely positive
    - Should be ignored
  - How to extrapolate and estimate for production mixed workload?
    - Estimation with accuracy and high confidence not practical
    - Benchmarking effort would be required

## Performance and Scalability

- Plan on **additional 10-30% real memory** (estimate)
- Many traditional OLTP workloads saw **5-10% CPU reduction in CM mode after** (some more, some less)
- Prerequisites
  - REBIND packages to generate new SQL run time
  - Use PGFIX=YES on buffer pools to exploit 1MB real storage frames available on z10 and z196 (100% backed)
- **But there were some exceptions < 5% CPU savings for OLTP with very light transaction, skinny packages with few simple SQL**
  - Package allocation/deallocation cost overrides benefit from SQL optimizations
  - APAR PM31614 may solve this by improving package allocation performance
  - Use of persistent threads with RELEASE(DEALLOCATE) will compensate



## Performance and Scalability

- Query Performance enhancements
  - No REBIND required for
    - Index list prefetch
    - INSERT index I/O parallelism
    - Workfile spanned records
    - SQLPL performance
    - High performance DBATs
    - Inline LOBs

## Performance and Scalability

- Query Performance enhancements ...
  - REBIND required for
    - Use of RELEASE(DEALLOCATE)
    - Early evaluation of residual predicates
    - IN-list improvements (new access method)
    - SQL pagination (new access method)
    - Query parallelism improvements
    - Index include columns
    - More aggressive view/table expression merge
    - Predicate evaluation enhancements
    - RID list overflow improvements
  - Execute RUNSTATS before REBIND
    - When coming from V8, to collect improved index statistics including CLUSTERRATIOF
    - When coming from V9, if do not already include the KEYCARD option of RUNSTATS

## Performance and Scalability ...

### ■ Customers measurements

- Not always consistent and repeatable
- Wide variation on measurement noise especially elapsed time
- In most cases not running in a dedicated environment or scale/size of production
- Many cases running subset of production workload
- Sometimes use of synthetic workload to study specific enhancements
- Do not trust some of the very big numbers on CPU and especially elapsed time savings
- Recommendation: customers should not spend the savings until they see them in production

## Performance and Scalability ...

Workload	Customer Results
CICS online transactions	Approx. 7% CPU reduction in DB2 10 CM after REBIND, additional reduction when 1MB page frames are used for selective buffer pools
CICS online transactions	Approx 10% CPU reduction from DB2 9
CICS online transactions	Approx 5% CPU reduction from DB2 V8
CICS online transactions	10+% CPU increase -> investigating
Distributed Concurrent Insert	50% DB2 elapsed time reduction, 15% chargeable CPU reduction after enabling high performance DBAT
Data sharing heavy concurrent insert	38% CPU reduction
Queries	Average CPU reduction 28% from V8 to DB2 10 NFM
Batch	Overall 20-25% CPU reduction after rebind packages

## Performance and Scalability ...

- **Use of 1MB real storage page frames on z10 and z196**
  - Potential for reduced for CPU through less TLB misses
  - Buffer pools must be defined as PGFIX=YES
  - Buffer pool page fix introduced in V8 to reduce CPU
  - Many customers reluctant to use PGFIX=YES because of potential for real storage
    - Running too close to the edge of the amount of real storage provisioned
    - Understand the value but only applies for 1-2 hours per day
    - But page fix is a long term decision
    - In most cases requires DB2 recycle to change attribute
    - 75% cost reduction on real storage on z196 (USD1.5K vs. USD6K)

## Performance and Scalability ...

- **Use of 1MB real storage page frames on z10 and z196 ...**
  - Must partition real storage between 4K frames and 1MB frames
    - Specified by LFAREA xx% in IESYSnn parmlib member and only changeable by IPL
    - 1MB frames are non-pageable
    - If 1MB page frames are overcommitted, will use 4K page frames
    - Recommendation: to add 20% in size to allow for growth and tuning
  - Be careful
    - Make sure critical z/OS maintenance applied before using 1MB pages
  - Benefit based on customer experience 0 to 6% reduced CPU
  - Requirement for new parameter to separate use of PGFIX=YES from use of 1MB page size
    - Plan to address in the next release of DB2

## Performance and Scalability ...

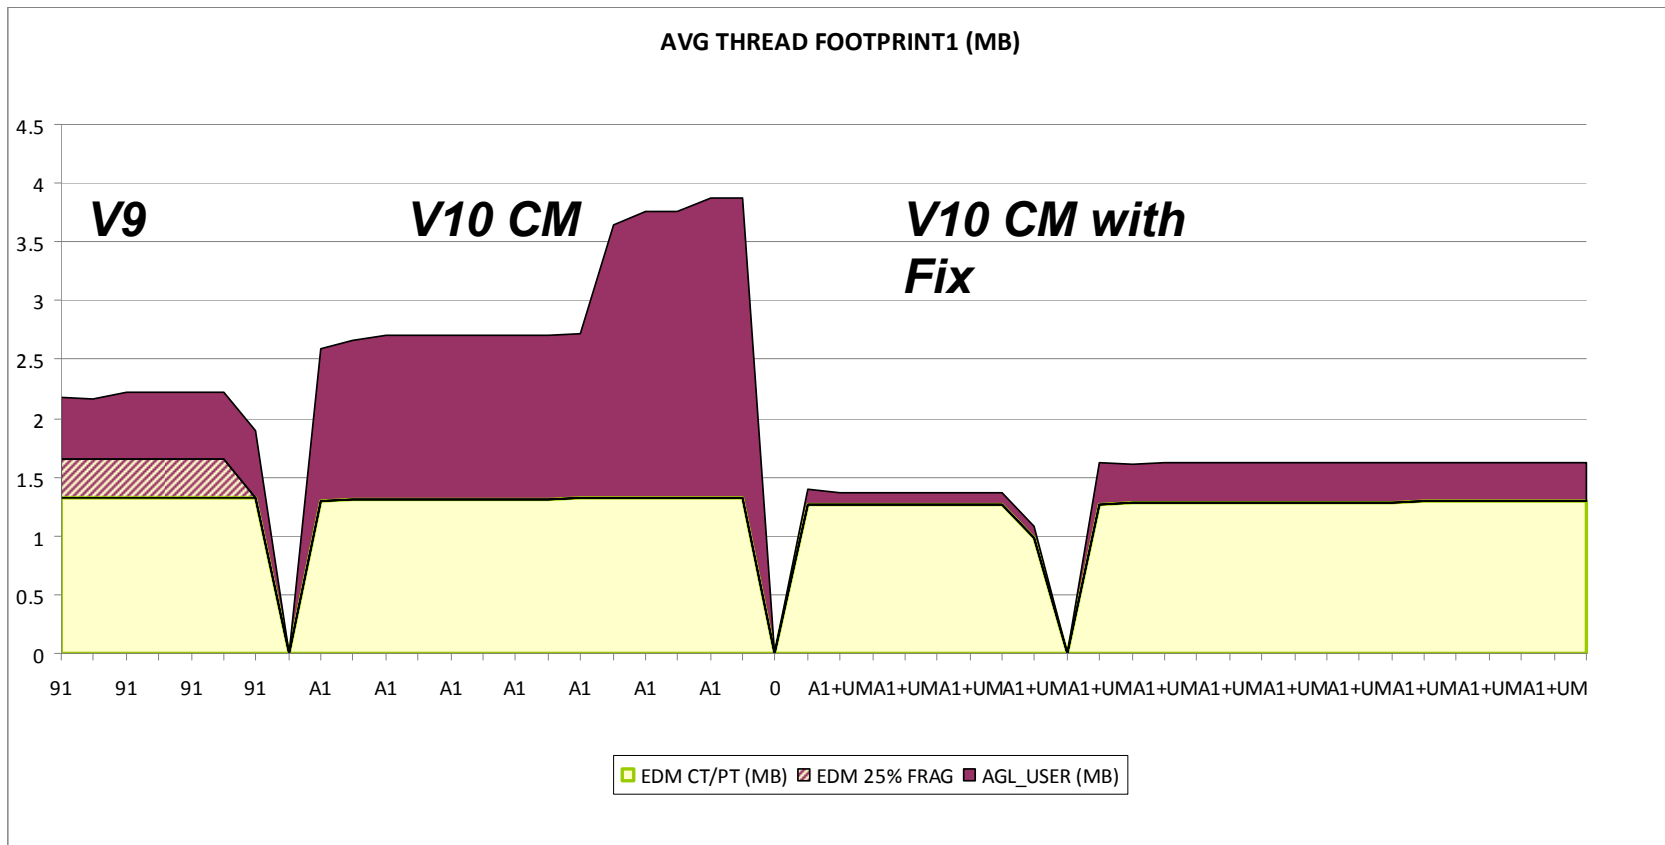
- **DBM1 Virtual Storage Constraint Relief with 64-bit SQL run time**
  - Available in CM
  - Requirement to REBIND static SQL packages to accrue maximum benefit
  - Very good results achieved
  - Have high degree of confidence that problem addressed
    - Real world proposition: 500 -> 2500-3000 threads plus
  - Limiting factors now on vertical scalability (# number of threads, thread storage footprint)
    - Amount of real storage provisioned
    - ESQA/ECSA (31-bit) storage
    - Active log write





## Performance and Scalability ...

- DBM1 31-bit Thread Storage V9 vs. V10 – as at GA after Fix



## Performance and Scalability ...

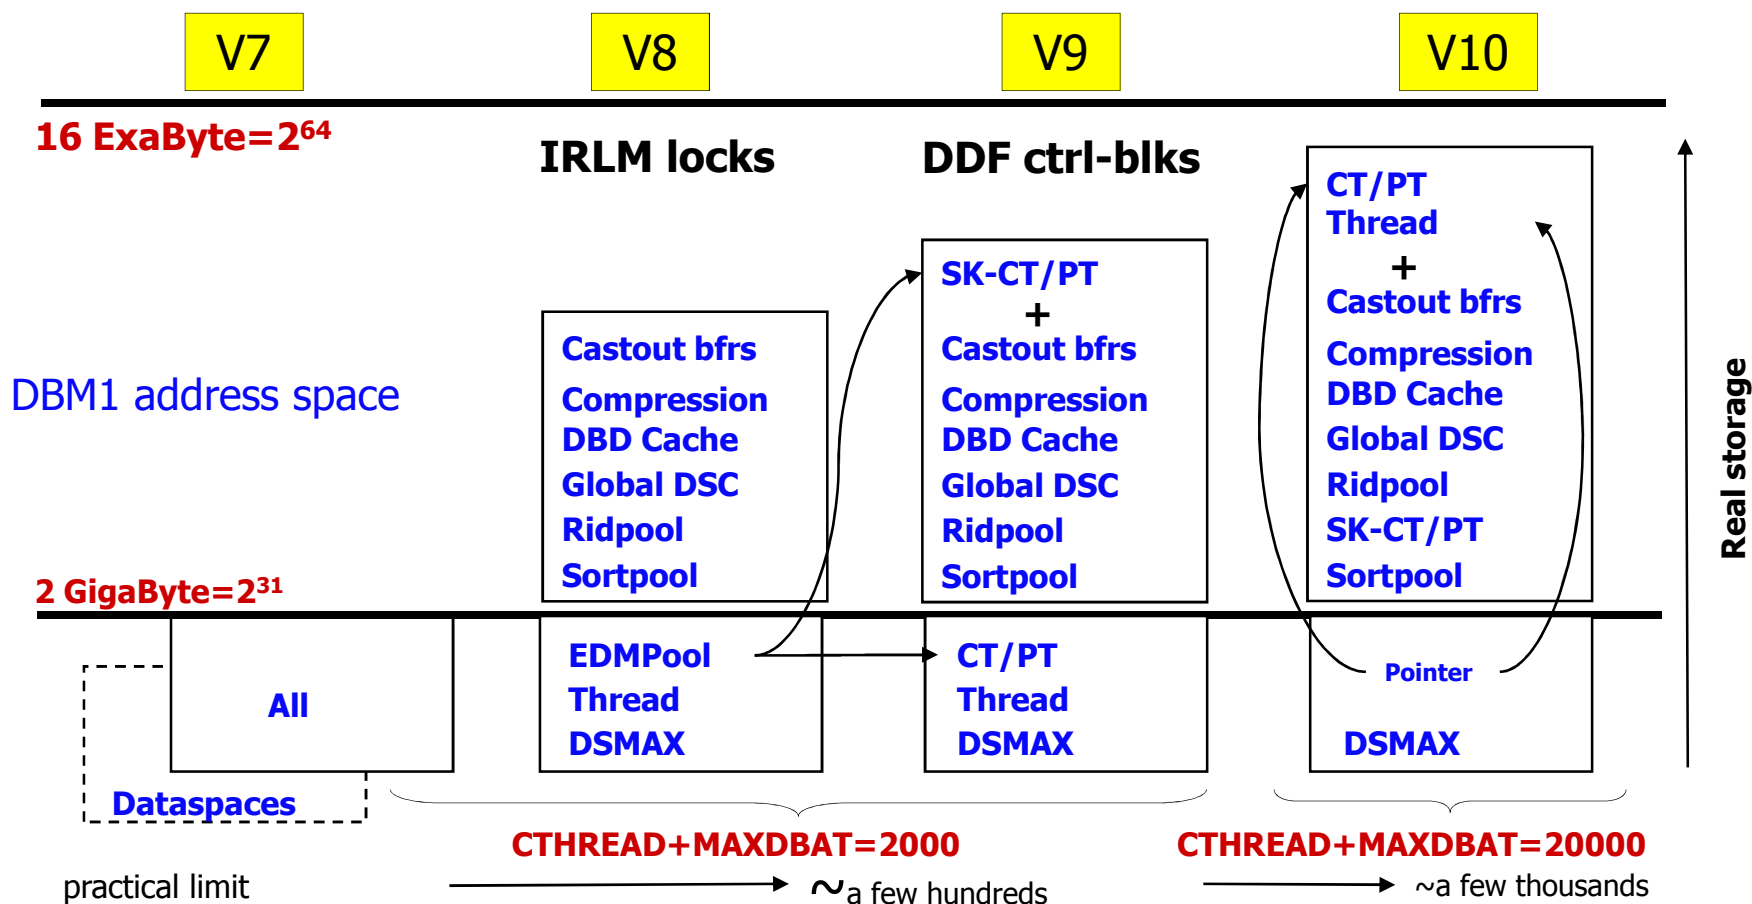
- **DBM1 Virtual Storage Constraint Relief with 64-bit SQL run time**
  - Major customer opportunities here for 31-bit VSCR and improved price/performance
  - Potential to reduce legacy OLTP transaction CPU cost through use of
    - More CICS protected ENTRY (persistent) threads
    - More use of RELEASE(DEALLOCATE) with persistent threads
    - Must provision **additional real storage** to back the requirement
  - Potential to reduce CPU for DRDA transactions by using High Performance DBAT
    - Must be using CMTSTAT=INACTIVE so that threads can be pooled and reused
    - Packages must be bound with RELEASE(DEALLOCATE) to get reuse for same connection
    - MODIFY DDF PKGREL(BNDOPT) must also be in effect
    - Do not to overuse RELEASE(DEALLOCATE) on packages
      - Will drive up the MAXDBAT requirement
      - Will need **additional real storage** to support increased number of threads
  - More persistent threads with RELEASE(DEALLOCATE) is also trade off with BIND/REBIND and DDL concurrency
  - CICS-DB2 accounting for cost of thread create and terminate, or avoidance thereof
    - CICS uses the L8 TCB to access DB2 irrespective of whether the application is thread safe or not
    - Thread create and terminate cost will clock against the L8 TCB and will be in the CICS SMF Type 110 record
    - Note: prior to OTE did not capture the thread create in the SMF Type 110

## Performance and Scalability ...

- **DBM1 Virtual Storage Constraint Relief with 64-bit SQL run time**
  - Potential to reduce the number of DB2 subsystems
    - Collapse multiple DB2 members running on the same LPAR
    - Reduce the total number of DB2 members
    - May be able to reduce the number of LPARs
    - Consider the increase of logging rate per DB2 member
    - Consider the increase in SMF data volume per LPAR
      - Can enable DB2 compression of SMF data to reduce SMF data volume
        - Experience is that Accounting records compress 70-80%
        - Tiny CPU overhead at ~1%
      - Re-consider use of accounting roll up for DDF and RRSAF workload (default)
        - Compromises performance PD/PSI as lose information on outlying transactions
        - Significant enhancements to package level accounting so it is now useful
    - Consider the increased DUMPSRV and MAXSPACE requirement
  - Re-emphasize the continued value of data sharing to differentiate the platform
    - Support avoidance of planned outages
    - Avoid humongous single points of failure
    - Minimum of 4-way for true continuous availability

# Performance and Scalability ...

- DBM1 31-bit Virtual Storage Constraint Relief with 64-bit SQL run time



## Performance and Scalability ...

### ▪ Real storage

- Need to carefully plan, provision and monitor real storage consumption
- Currently cannot monitor 64-bit shared and common via IFCID 225 when running multiple DB2 subsystems on the same LPAR
  - How much is in REAL?
  - How much is in AUX?
- DB2 APAR PM24723 will provide the capability that is needed
  - MVS APAR OA35885 will provide new callable service to RSM to provide REAL and AUX for addressing range for shared objects
  - APAR PM24723 will probably close mid June
- Will need to re-evaluate the hidden zparm SPRMRSMX ('real storage kill switch')
  - Need to factor in 64-bit shared and common use to establish new footprint
- 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)

## Performance and Scalability ...

- Real storage ...
  - DB2 APAR PM24723 is very important and will probably close in June
    - Monitoring issue is addressed and new extensions to IFCID 225 provided
      - Pre-req is new MVS APAR OA35885 which provides a new callable service to RSM to provide REAL and AUX used for addressing range for shared objects
    - SPRMRSMX hidden zparm now becomes an opaque parameter REALSTORAGE\_MAX
    - Will also introduce DISCARD mode to contract storage usage to protect against excessive paging and use of AUX
      - New zparm REALSTORAGE\_MANAGEMENT controls when DB2 frees storage frames back to z/OS
        - > ON -> Discard unused frames all the time - discard stack, thread storage, keep footprint small
        - > OFF -> Do not discard unused frames unless things are getting out of hand
        - > AUTO (default) -> Detect whether paging is imminent and try to reduce the frame counts to avoid system paging
      - With AUTO, DB2 monitors paging rates, switches between ON/OFF and decides when to discard frames based on
        - > 80% of SPRMRSMX reached
        - > 50% of AUX (ENF55 condition) used
        - > Hitting AVQLOW (available real storage frame)
      - New messages (DSNV516I, 517I) for when paging rate thresholds cause DB2 to free real frames
    - Strong recommendation to apply PTF for APAR PM24723 before going into business production and to run with REALSTORAGE\_MANAGEMENT=AUTO

## 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
  - 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
    - $\geq 3$  indexes
    - Prefetch offload to zIIP to compensate

## Performance and Scalability ...

- **Accounting Trace Class 3 enhancement – separate counters**
  - IRLM Lock/Latch waits
  - DB2 Latch waits
- **Data sharing**
  - Faster DB2 shut down
    - Avoid local buffer pool scan per GBP-dependent object
  - XXXL Buffer pool
    - Avoids local buffer pool scan when pageset/partition transitions into GBP-dependency
    - Avoids local buffer pool scan when pageset/partition transitions out of GBP-dependency
- **Inline LOBs**
  - Potential for significant CPU and elapsed time improvement with the right inline value
- **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



## Performance and Scalability ...

### ▪ Hash access vs. Index only access

- Competes against index only access
  - 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 ( $\geq 3$ )
  - 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

- High Performance DBATs (Hi-Perf DBATs) – new type of distributed thread
  - Must be using CMTSTAT=INACTIVE so that threads can be pooled and reused
  - Packages must be bound with RELEASE(DEALLOCATE) to get reuse for same connection and -MODIFY DDF PKGREL(BNDOPT) must also be in effect
  - When a DBAT can be pooled after end of client's UOW
    - DBAT and client connection will remain active together
      - Still cut an accounting record and end the enclave
    - After the Hi-Perf DBAT has been reused 200 times
      - DBAT will be purged and client connection will then go inactive
    - All the interactions with the client will still be the same in that if the client is part of a sysplex workload balancing setup, it will still receive indications that the connection can be multiplexed amongst many client connections
    - IDTHTOIN will not apply if the Hi-Perf DBAT is waiting for the next client UOW
    - If Hi-Perf DBAT has not received new work for POOLINAC time
      - DBAT will be purged and the connection will go inactive
    - If # of Hi-Perf DBATs exceed 50% of MAXDBAT threshold
      - DBATs will be pooled at commit and package resources copied/allocated as RELEASE(COMMIT)
    - Hi-Perf DBATs can be purged to allow DDL, BIND, and utilities to break in
      - Via -MODIFY DDF PKGREL(COMMIT)

## Thread Management

### ■ CICS (V8/V9)

- Selective use of protected entry thread while keeping RELEASE(COMMIT)
  - Aim at high volume transactions
    - Consider collapsing several transactions into a single plan to drive multi-programming
  - Avoid the repetitive cost of thread creates and terminates
  - Reduce CPU resource consumption as reported in the CICS SMF 110 records
- If sufficient DBM1 31-virtual storage headroom, could consider selective use of RELEASE(DEALLOCATE) for high used packages to complement thread reuse
  - Trade-off with concurrency for BIND, REBIND, DDL
  - Should not be used in conjunction with LOCK TABLE and lock escalation

### ■ CICS (V10 CM)

- Assuming adequate additional REAL memory
  - More extensive use of protected entry threads with RELEASE(DEALLOCATE)
  - See reduction in the CICS SMF 110 records and the DB2 Class 2 CPU

## Thread Management ...

### ■ Batch (V8/V9)

- Selective use of BIND option RELEASE(DEALLOCATE) for long-running batch programs that take frequent intermediate commits
  - Sustain storage allocation, index look-aside, dynamic prefetch, fast column processing
  - Need to monitor closely the impact on DBM1 31-virtual storage usage
  - Trade-off with concurrency for BIND, REBIND, DDL
  - Should not be used in conjunction with LOCK TABLE and lock escalation

### ■ Batch (V10 CM)

- V10 will allow more use of BIND option RELEASE(DEALLOCATE) for batch
  - Assuming adequate additional REAL memory

## Availability

- **Online Schema Evolution ('Deferred Alter')**
  - Note: UTS is pre-requisite for Hash, Inline LOB, Currently Committed
  - Migrate from classic table space types (simple, segmented, partitioned) to UTS PBG/PBR
    - One way ticket only
  - Once migrated to UTS PBG/PBR can change attributes such as
    - 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

- **Online REORG with FORCE**
  - Customers perceive limited value
  - Only running threads which are blocking are cancelled
  - Threads which are suspended / inactive will cause REORG to still fail
- **Online REORG LOB with DISCARD**
  - Customers perceive limited value
  - Cannot handle LOB columns greater than 32KB

## 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 create classic PTS
    - Specify SEGSIZE 0 on CREATE TABLESPACE
    - Set new zparm DPSEGSZ=0 (default 32)
- **Old COBOL and PL/I**
  - V7 lookalike pre-compiler (DSNHPC7) for older COBOL and PL/I is still provided
- **DDL Concurrency after Catalog restructure**
  - Concurrency issues not solved
  - Still deadlocks with parallel heavy DDL against different databases

## Other ...

### ▪ **Single thread BIND/REBIND**

- Degraded CPU and elapsed time performance on entry 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 in data sharing mode**

- 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 line LOB in data sharing (APAR PM24721 closed)
  - Inline LOB with compression for SPT01 to address SPT01 growth (APAR PM27073 target 1Q/2011)
  - More efficient space reuse for base table and UTS (APAR PM27973 target 1Q/2011)
- 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



## Other ...

### ▪ Incompatible changes

#### – CHAR() scalar function

- Solution via APAR PM29124 to restore compatible behavior for CHAR SQL scalar function CHAR() for COL1 DEC(4,0)

```
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
9	0009,	40F0F0F0F96B	0009
99	0099,	40F0F0F9F96B	0099
999	0999,	40F0F9F9F96B	0999
9999	9999,	40F9F9F9F96B	9999

COL1	CHAR_V10	HEX_CHAR_V10	APPL_SUBSTR
9	9	F94040404040	
99	99	F9F940404040	9
999	999	F9F9F9404040	99
9999	9999	F9F9F9F94040	999

## 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

#### – Create Trigger

```
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
```

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

## 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**
- **Cannot migrate**
  - 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 ...

- **If coming from V8**
  - BSDS must be reformatted for larger active / archive tracking
- **DB2 Connect**
  - Minimum level V9.1 FP1
  - V9.7 FP3A required for new functions
- **Must absolutely eliminate all use of DDF Private Protocol before migrating**
  - Local packages miss tagged with DDF Private Protocol will be tolerated
- **DBRMs bound directly into plans no longer supported**
  - Will trigger auto bind into packages
- **Old plans and packages bound prior to V6 will require REBIND**
- **Catalog and Directory must be SMS managed**
- **PDSEs required for SDSNLOAD, SDSNLOD2, ADSNLOAD**

## Migration and Planning ...

- **DSNHDECP NEWFUN=V10|V9/V8**
- **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 use not use CCSID EBCDIC with V10 format
      - EXPLAIN fails with RC=8 DSNT408I SQLCODE = -878
      - BIND with EXPLAIN fails with RC=8 DSNX200I
  - Recommended to use the V10 extended format with CCSID UNICODE
    - If application access EXPLAIN tables and only tolerate SQLCODE 0 or +100
  - APAR PK85068 can help migrate V8 or V9 format to the new V10 format with CCSID UNICODE

## Migration Fundamentals: DB2 10 for z/OS – Pre-requisites

- **The following transaction management products work with DB2 10:**
  - Information Management System (IMS)
    - IMS V11 (5635-A02)
    - IMS V10 (5635-A01)
- **Customer Information Control System (CICS)**
  - CICS Transaction Server for z/OS V4.1 (5655-S97), or later
  - CICS Transaction Server for z/OS V3.1 (5655-M15)
    - For Version 3.1 and Version 3.2, you also need APAR PM01880 to return the correct version and release number for DB2 10.
- **Migration Process:**
  - From Version 8 or 9 NFM
  - Data sharing coexistence in CM8 or CM9
  - DSNTIJPA in V8 or V9 with PM04968

## DB2 10 Migration Restrictions:

- **If you start a Version 8 subsystem migration to Version 10 and then fall back to Version 8, you cannot subsequently migrate the subsystem to Version 9.1.**
- **If you start a Version 9 subsystem migration you must complete the migration to New function Mode(NFM).**
- **Existing DB2 systems must be in DB2 8 New Function Mode (NFM) or DB2 9 New Function Mode**
- **A data sharing group that is migrated from Version 8 to Version 10 cannot have any Version 9.1 members.**
- **If you start a Version 8 subsystem migration to Version 10, you can fall back to only Version 8.**
- **A data sharing group that is migrated from Version 9.1 to Version 10 cannot have any Version 8 members.**
- **If you start a Version 9.1 subsystem migration to Version 10, you can fall back to only Version 9.1.**

Source: DB2 10 for z/OS - *Installation and Migration Guide* / GC19-2974-00

## Migration and Planning ...

- **Early customer adopters of V10 migrating from either V8 or V9 should make plans and take extra care to mitigate against the risk of instability**
  - Perform application regression and stress testing to keep ‘fires’ away from production
  - Need to be more aggressive on planned continual application of preventative service
    - Will have to stay a lot more current than 2 full ‘major’ drops a year
    - Regular full ‘major’ maintenance drops including HIPERs/PEs essential required for the first year or so
      - May be 4 ‘major’ drops in the first year
    - Can move to 2 ‘major’ and 2 ‘minor’ maintenance drops as the release passes the early adopter curve
    - Exploit CST/RSU recommended maintenance as opposed to the PUT route
      - Recommended maintenance after successful testing for a least one month
      - Testing performed over and above that performed by DB2 Development
      - CST testing still does not replace customer regression/stress testing
  - Must be prepared to tolerate hit some ‘bumps in the road’
  - Customer who are not prepared to take mitigating actions and have no tolerance for ‘bumps in the road’ should not be early adopters
    - V8 customers should migrate to V9 quickly as it is relatively stable



## Skip Release Considerations – 21 pages in v10, 12 in v9 for Appl.

### Infrastructure

- *Maintenance currency – IBM and vendors*
- *Hardware dependency and REAL storage*
- DDF 64-bit support enabled / 128GB shared private

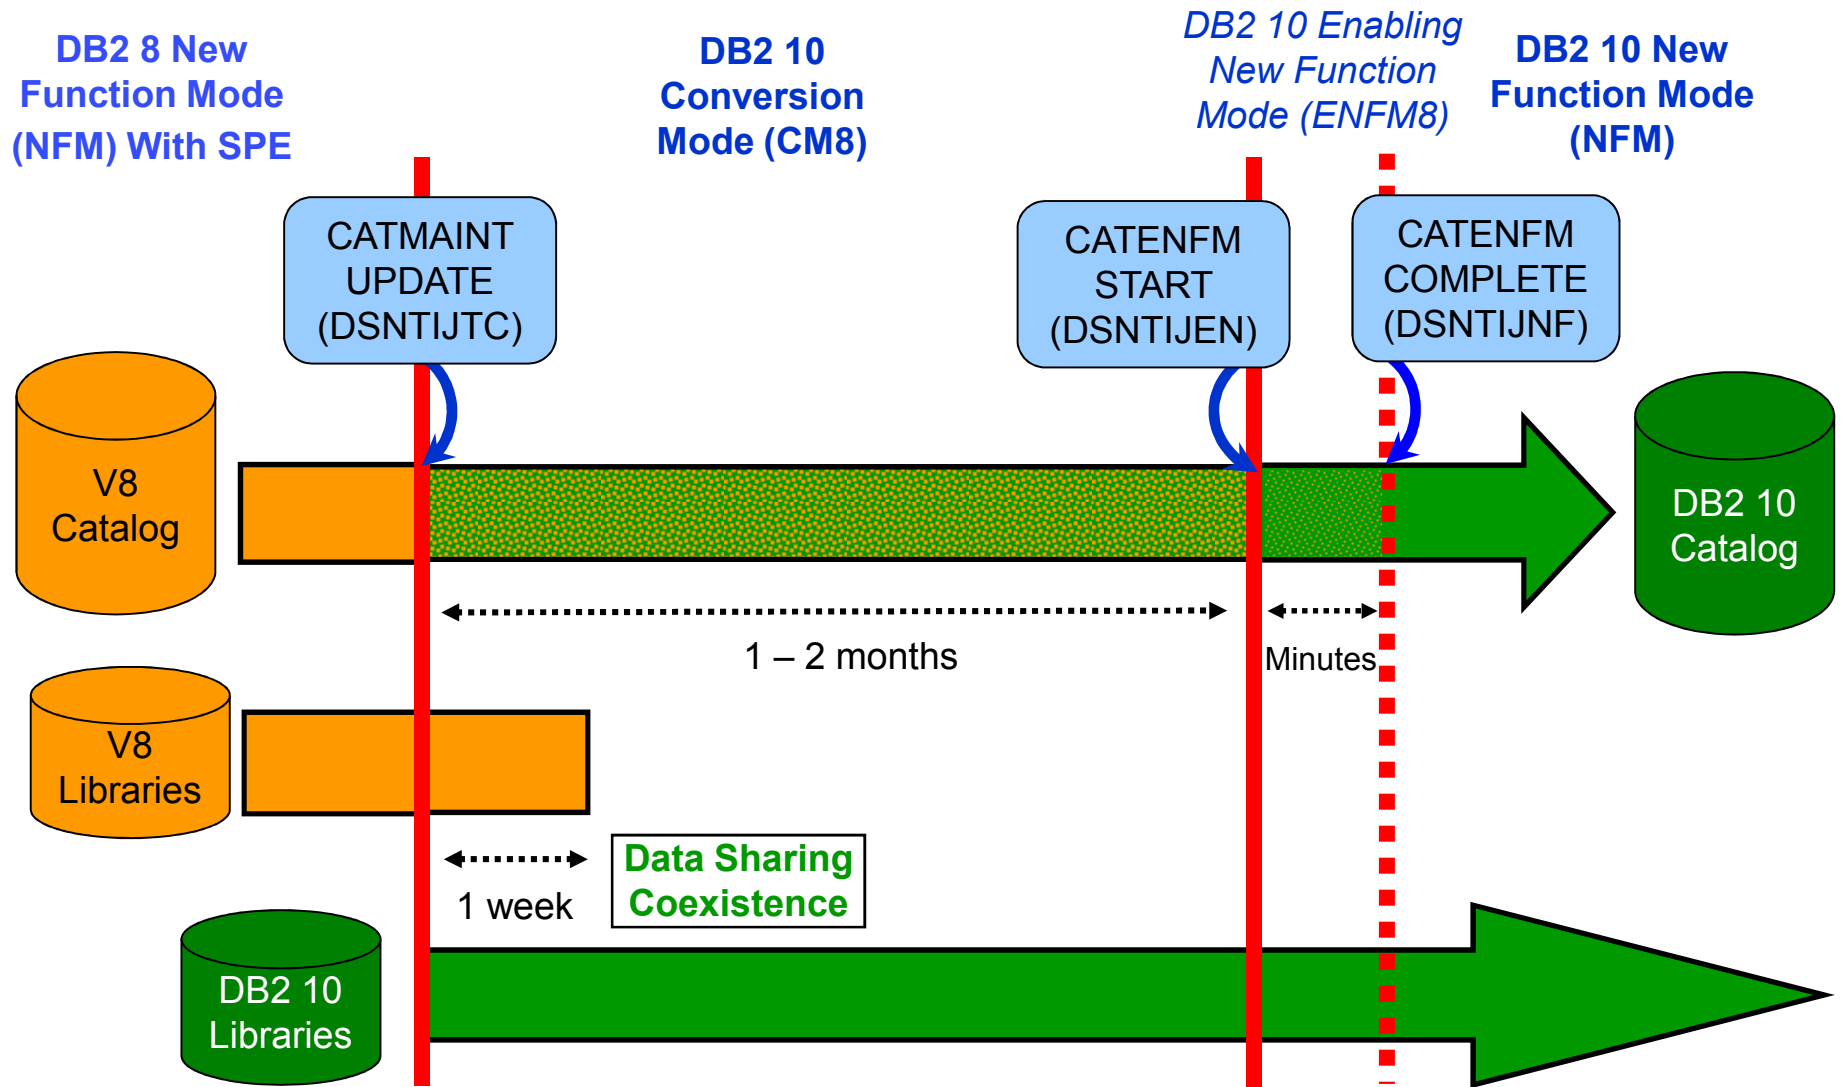
### DB2 9 changes included

- Subsystem parameters (MGEXTSZ, OPTIOWGT, MAXTEMPS, IGNORE\_FREESPACE, WFDBSEP)
- TCP/IP ports INADDR\_ANY suggested
- Parallel REORG jobs / AND Real Time Stats is utilized across the board
- SORTNUM elimination
- Reordered Row considerations.....(compression, HONOR\_KEEPDICTIONARY, SPRMRRF)
- DB2 Connect 9.7 fixpak 3a ++
- REXX support required
- DSNJCNVB run on v8 -> expand BSDS
- PDSEs for SDSNLOAD, SDSNLOD2, ADSNLOAD, ADSNLOD2
- DB2-managed storage for Catalog and directory, and PBG -> SMS-managed

### Application considerations

- Private protocol: DSNTDP2DP
- Workfile DB revamp
- Implicit DB and TS – DPSEGSZ (32)
- TEXT extender, XML extender, AV extender, DSNWZP
- Management Client's features, OSC
- Incomplete objects defined
- CHAR() with Substring
- Need Order By, not just Group By

# Skip Migration Overview V8 → 10



## 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.

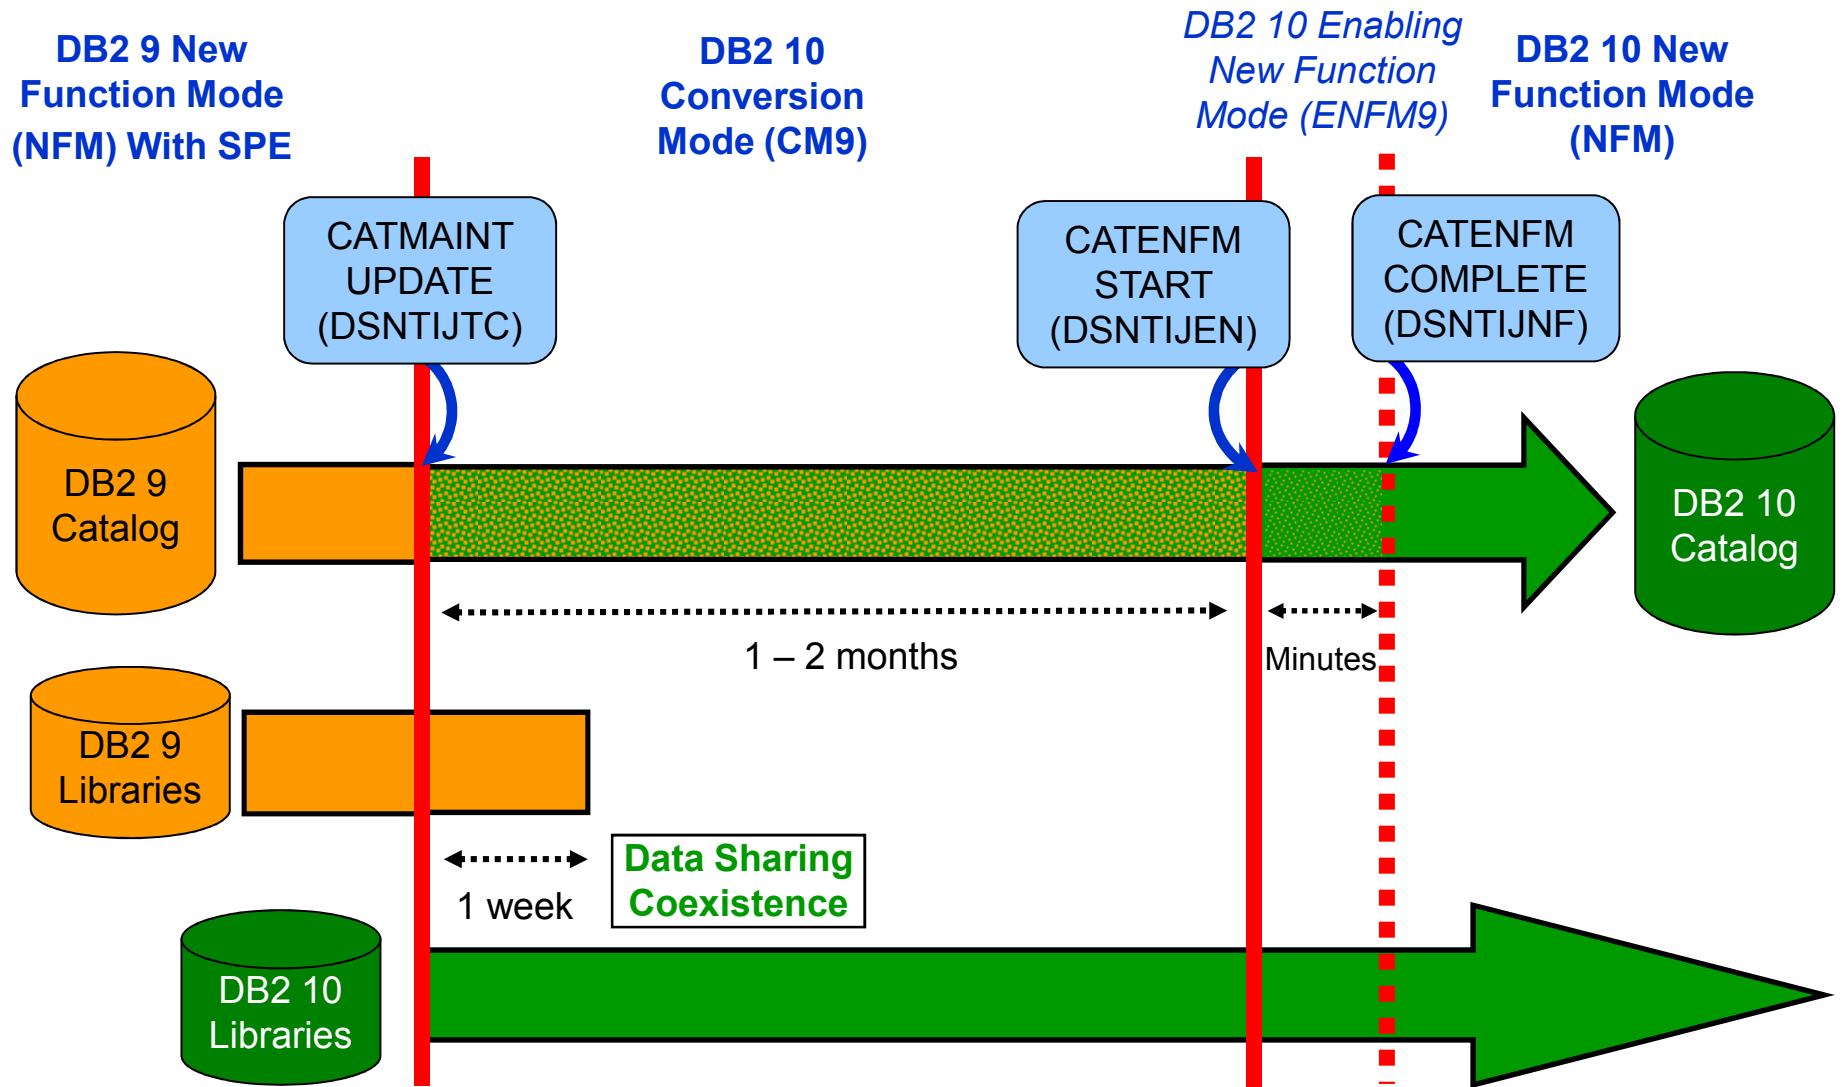
**ENFM8 Enabling New Function Mode** - This mode is entered when CATENFM START is executed (the first step of job DSNTIJEN). DB2 remains in this mode until all the enabling functions are completed. Data sharing systems can only have DB2 10 members in this mode.

**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.

# Normal Migration Overview DB2 9 → DB2 10



## 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.

## Migrating to DB2 10

- **The DB2 10 Catalog & Directory must be managed by DB2 & SMS**
  - Job DSNTIJSS provided to set up the SMS environment
  - SMS environment must be set up before migration
  - Not required to convert catalog / directory before migrating
    - Will be converted during the next REORG
  - Data Class attributes of
    - Extended Format
    - Extended Addressability
- **Reestablish V8 or V9 IVP to test DB2 10 before NFM**
- **Assess ISV Requirements / Readiness**
  - Tools and applications
  - Some vendors may add instructions for migration and / or require maintenance
- **Assess the training requirements for your organization**

## Maintenance and Operations Recommendations

- **Read the Installation and Migration Guide, 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.**
- **Start with latest RSU + Identified Hipers**
- **Leverage CST/RSU process**
  - Apply 2 to 3 preventative service drops annually (Early adoption of V10 may require applying maintenance more frequently)
  - Exploit Enhanced HOLDDATA to be vigilant on HIPERs and PEs - review weekly
- **Use the DB2 9 'Package Stability' function for static SQL**
  - Offers access path preserving option. Recovers to prior access path if regression is encountered
- **Minimize potential query performance issues**
  - Use Data Studio's Query Tuner to capture SQL statements and use RUNSTATS Stats Advisor to generate the recommendation for statistics collection
  - Run RUNSTATS to ensure critical stats are collected as recommended by the advisor before rebind of package
- **Ensure a PMR is opened prior to migration start**
  - Should a problem arise during the migration, "call" the Support Center and escalate the PMR
  - If the problem is serious, ask to be transferred to a Duty Manager

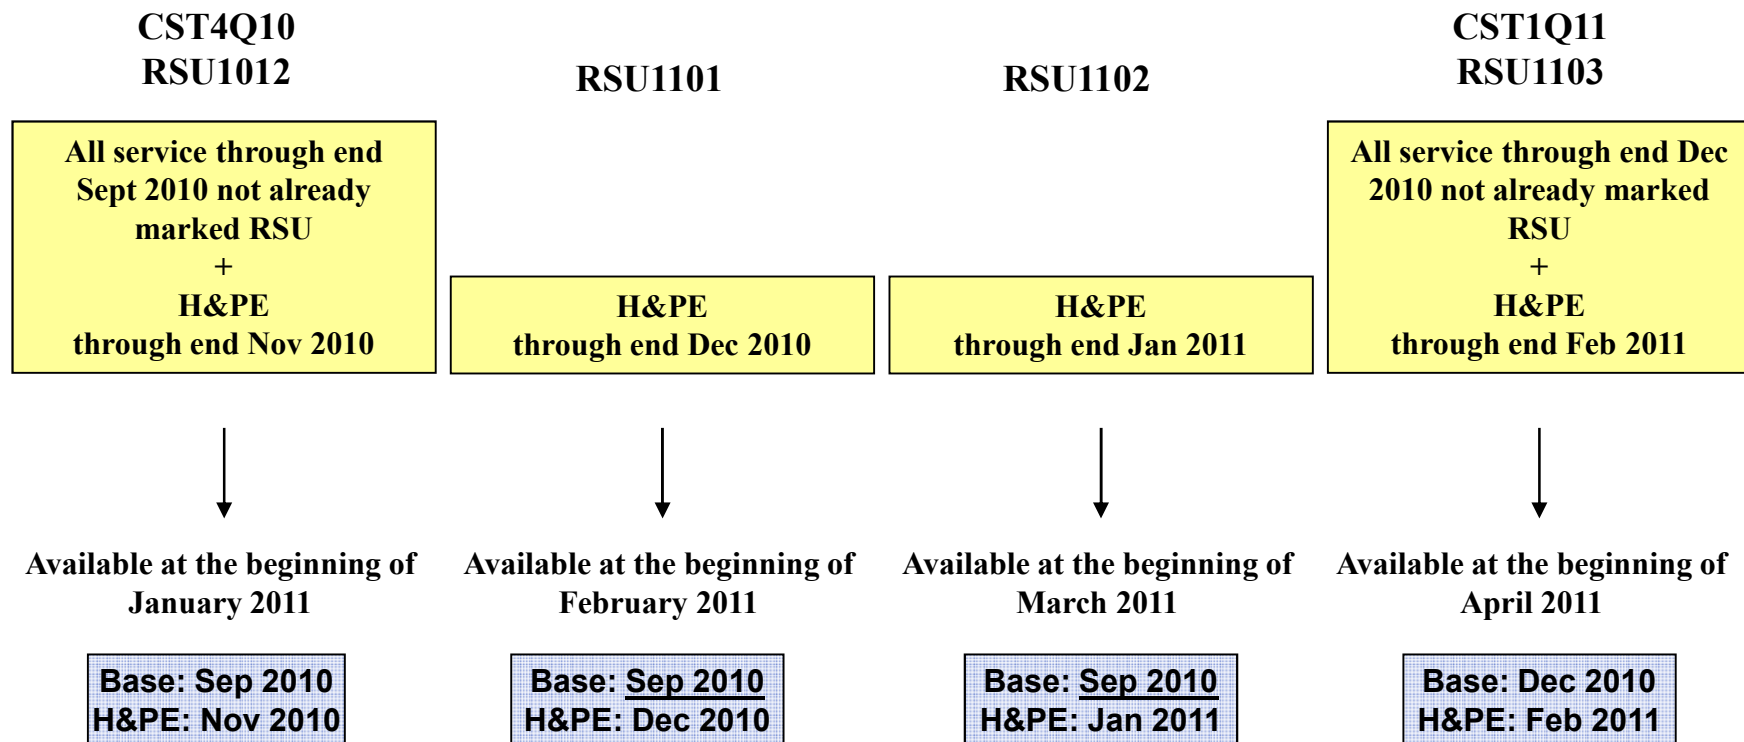
## Recommended Early Adopter Skip Migration Best Practices

1. Attend both DB2 9 and 10 for z/OS Migration Planning Workshops(MPWs)
  - Although only one migration is being done, two versions of DB2 z/OS are being adopted
  - Learn and understand all of the new basics features and incompatibilities (example: WORKFILE, Private Protocol, DBRMs bound into Plans... )
2. Establish performance baselines for important workloads and applications
  - Explain / PLAN\_TABLE output - IBM Tools, like Path Checker, can assist in identifying access path changes and in using optimization hints if needed.
  - Save SMF reports and key job runtime statistics
3. Customers must be prepared to aggressively apply DB2 maintenance, as they need be prepared to uncover new issues early in the product life cycle.
  - Even though IBM has done extensive testing, each of our customers seems to use our software in ways that other customers don't, whether it is features of the code they exploit or the infrastructure that they have in place to run their business. See migration information APARs II14474 (V8) or II14477 (V9).
4. Test Plans - Customers must have rigid testing plans in place, so that they can find issues in their test environment BEFORE they get to production.
  - Test in-house and ISV programs and applications
5. Understand that customers will be moving to an early version of DB2 10, and they must be prepared to experience bumps along the way.



## Migration and Planning ...

- CST and RSU example**



*H&PE = HIPER/Security/Integrity/Pervasive PTFs + PE resolution (and associated requisites and supersedes)*

## Maintenance Rollout thru environments...

### Test Environment



Install new Qtrly + .....



Begin Qtrly RSU + monthly RSU + H&PE

Install QA Qtrly +  
....

- Apply RSU updates throughout test cycle
- At the end of the quarter the new quarterly RSU should be equivalent to the preceding quarterly plus maintenance.
- Re-iterate the cycle thru the required time period.
- PROD receives a quarterly RSU plus the maintenance to a specified point
  - Receives a tested quarterly plus maintenance
  - Target is to be within 3 months or better..
    - Which is already 3 months old

## Where do the DBRMs go? (SG24-7688)

Table 4-2 Multiple PLAN binds

PLAN	COLLID	PKLIST	
Individually listed	Explicitly named	Not specified	Packages are bound into the named COLLID.
	*	Not specified	Packages are bound into the DSN_DEFAULT_COLLID_ <i>planname</i> collection
	Explicitly named	Specified	Packages are bound into the named COLLID. Created packages are at the beginning of the PKLIST, specified PKLISTs follow.
	*	Specified	Packages are bound into the DSN_DEFAULT_COLLID_ <i>planname</i> collection. Created packages are at the beginning of the PKLIST, specified PKLISTs follow.
*	Explicitly named	Not specified	Packages are bound into the named COLLID
	*	Not specified	Packages are bound into the DSN_DEFAULT_COLLID_ <i>planname</i> collection.
	Explicitly named	Specified	Packages are bound into the named COLLID. Created packages are at the beginning of the PKLIST, specified PKLISTs follow.
	*	Specified	Packages are bound into the DSN_DEFAULT_COLLID_ <i>planname</i> collection. Created packages are at the beginning of the PKLIST, specified PKLISTs follow.

## DBRMs...

- **PK62876**
- **Plan= Employee**
  - DBRM=Raise
  - REBIND Plan Employee COLLID(\*)
    - Employee (plan)
    - DSN\_DEFAULT\_COLLID\_Employee (collection)
    - Raise (package)
  - New package goes to top of PKLIST (DBRM entry deleted)
- **Caution: ways to lose a DBRM**
  - DBRM with same name in multiple plans and you override default COLLID
  - DBRM has same name as package that already exists in that plan

## Security considerations when removing DDF Private Protocol


- 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
- Until DB2 10 - 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
- DRDA Migration (PK64045)

## DB2 10 Catalog...

- **V10 catalog restructure provides greater concurrency for catalog operations**
  - Concurrent binds
- **V8 migrations will include new Catalog table spaces for**
  - Real-Time Statistics
  - New page size for SYSOBJ
  - XML
  - Trusted Context
  - Extended Index definitions
- **V8 and V9 migrations see many table space changes for the catalog restructure**
- **Several LOB columns**
  - Using Inline LOBs
  - They should get their own bufferpool and set DWQT = 0

## DB2 10 Catalog...

- **SYSDBASE, SYSPLAN, SYSDBAUT, SYSVIEW, SYSGROUP and DBD01 had links**
  - These table spaces used page level locking because of the links.
- **SPT01, SYSOBJ, and SYSPKAGE are also processed in ENFM.**
- Combine SYSUTIL and SYSUTILX into a single table SYSUTILX
- SYSUTIL, SYSUTILX, SCTR, SYSLGRNX = UNICODE
- **All of these table spaces will be removed and the tables within each will be moved to new PBG table spaces**
  - Row level locking
  - New row format
  - Partition-by-growth
  - One table per table space
  - Referential Integrity in place of links
  - DSSIZE 64 G
  - MAXPARTS 1



## DB2 10 Catalog...

### ▪ Summary

- 55 spaces are UTS Partition By Growth
- All are Unicode, except
  - SYSCOPY
  - SYSEBCDC
  - SYSTSASC
- All are row level locking, except the following (page level):
  - SYSEBCDC
    - Contains SYSDUMMY1 & SYSDUMMYE
  - SYSSEQ
    - Has MAXROWS set to 1
- 68 are RRF
- 18 LOB Columns
  - 6 use Inline LOBs

TABlename	COLNAME	INLINE_LENGTH
SYSPACKSTMT	STATEMENT	15360
SYSPACKSTMT	STMTBLOB	7168
SYSCONTROLS	RULETEXT	16000
SYSCONTROLS	DESCRIPTOR	12000
SYSQUERY	STMTTEXT	2048
SYSVIEWS	PARSETREE	27670



## Recent zParm additions

- **WFDBSEP in macro DSN6SPRM**
  - YES – workfiles go to spaces with '0' SECQTY
    - DGTTs go to spaces with non - '0' SECQTY
    - Or the process ends with negative SQLCODE
  - NO – if no preferred table spaces are available, they use whatever is available
- **MAXTEMPS – default 0MB**
  - Max amount of storage for a single agent
- **REORG\_IGNORE\_FREESPACE in macro DSN6SPRM**
  - Tells the utility to ignore FREESPACE quantity on PBG table in order to ensure rows fit back into original partition
- **COMPRESS\_SPT01 in macro DSN6SPRM**
  - PK80375
- **Opaque DSNZPARM: SPRMRRF=ENABLE|DISABLE – PK87348**
- **PRIVATE\_PROTOCOL(NO) – PK92339**
  - Will fail any BIND or REBIND of a package or plan with message DSNT225I
  - Will fail any remote BIND/REBIND package request with SQLCODE -30104

## Items Planned for post-GA Delivery via APAR

- **APREUSE and COMPARE - PM25679**
  - <https://www-304.ibm.com/support/entdocview.wss?uid=swg1PM25679>
- **Delete data sharing member - PM31009**
  - <https://www-304.ibm.com/support/entdocview.wss?uid=swg1PM31009>
- **Compress SPT01 with inline - PM27811, PM27073**
  - Compression, BIND performance
- **Enhancements for new DBA authorities (PM28296)**
  - Prevent privileged users from stopping audit traces
  - No implicit system privileges for DBADM
- **Online REORG concurrency for materializing deferred ALTERs (PM25648)**
- **Temporal enhancements**
  - TIMESTAMP WITH TIMEZONE support (PM31314)
  - Enhancement for data replication (PM31315)
  - ALTER ADD COLUMN, propagate to history table (PM31313)

## Items Planned for post-GA Delivery via APAR ...

- **New system profile filters based on “client info” fields (PM28500)**
  - Three new columns for userid, appname, and workstation
  - Wildcard support: if column is ‘\*’ then all threads pass that qualification
- **Zparm to force deletion of CF structures on group restart (PM28925)**
- **Relief for incompatible change in CHAR of decimal data (PM29124)**
- **Real storage monitoring enhancements (PM24723)**
- **Hash LOAD performance (PM31214)**
- **DSSIZE > 64GB – APAR coming**
- **REORG REBALANCE SHRLEVEL CHANGE – APAR coming**

## Misc. APARS

- **PM10726** – add job to SDSNSAMP to allow custom DSNTIDxx member to be built from current DSNTIJUZ
- **PK62178** – allows ALTER of SYSIBM.DSNSEQ\_IMPLICITDB
- **PM11941 - SYSPROC.ADMIN\_INFO\_SQL**
  - Collects same information as OSC service SQL and SQL Environment Capture in Data Studio
- **PM12256** – zIIP offload improvement up to 60%, and less overhead
- **II14464** – info APAR on v8- 9 migration issues
- **SQL migration performance best practices**
  - [http://www-01.ibm.com/support/docview.wss?rs=64&context=SSEPEK&q1=Db2+9+migration&uid=swg27015988&loc=en\\_US&cs=utf-8&lang=en](http://www-01.ibm.com/support/docview.wss?rs=64&context=SSEPEK&q1=Db2+9+migration&uid=swg27015988&loc=en_US&cs=utf-8&lang=en)
- **II14477** – info APAR DB2 9 -> DB2 10
  - <http://www-01.ibm.com/support/docview.wss?uid=isg1II14477>
- **II14474** – info APAR DB2 v8 -> DB2 10
  - <http://www-01.ibm.com/support/docview.wss?uid=isg1II14474>

## More Misc. APARs v10

- **z/OS APAR to reduce ESQA shared private storage above the bar OA33126**
- **Reduce LC25 contention for package allocation in DB2 10 – PM31614**
- **Efficient space reuse for UTS - PM27973**
- **Online REORG materializing deferred alters - PM25648**
- **IFCID 225 measures real storage - PM24723**
  - <https://www-304.ibm.com/support/entdocview.wss?uid=swg1PM24723>
  - Also needs z/OS APAR OA35885
- **IBM tools compatibility chart**
  - <http://www-01.ibm.com/support/docview.wss?uid=swg21409518>
- **II14620 OEM tools compatibility**
  - <https://www-304.ibm.com/support/docview.wss?mynp=OCSSEPEK&mync=R&uid=isg1II14620&myns=swgimgmt>

# 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](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\\_migr8cm8checklist.htm](http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db2z10.doc.inst/src/tpc/db2z_migr8cm8checklist.htm)

## V9 premigration checklist:

[http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db2z10.doc.inst/src/tpc/db2z\\_premigr9checklist.htm](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\\_migr9cm9checklist.htm](http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.db2z10.doc.inst/src/tpc/db2z_migr9cm9checklist.htm)

## Summary

- **Very good release in terms of the opportunities for price/performance and scalability improvements**
  - 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
  - Opportunity for scale up and LPAR/DB2 consolidation

## Summary ...

- **Carefully plan, provision and monitor real storage consumption**
- **Early customer adopters of V10 migrating from either V8 or V9 should make plans and take extra care to mitigate against the risk of instability**
  - Regular full ‘major’ maintenance drops
  - Exploitation of CST/RSU recommended maintenance
  - Perform application regression and stress testing to keep ‘fires’ away from production
  - Must be prepared to tolerate some ‘bumps in the road’



## IDUG DB2 10 Migration Experiences Forum

- **A new IDUG forum dedicated specifically to DB2 10 for z/OS migration experiences**
  - Regular updates and surveys on DB2 10 migration topics
  - Share and discuss your upgrade plans and experiences with your peers and IBM and other industry experts
  - Links to valuable DB2 10 technical information from IBM, IDUG and other sources
- **Forum can be found at <http://www.linkedin.com/groups?gid=3797589>**

## Questions

