DB2 12 — The ultimate enterprise database for business-critical transactions and analytics

DB2 12 for z/OS: Migration Considerations

Mark Rader
WSC: DB2 for z/OS

Disclaimer/Trademarks

DB2 12 — The ultimate enterprise database for business-critical transactions and analytics

DB2 12

Acknowledgements

- Special thanks to John Campbell for the original material
- Thanks to all who participated in the DB2 12 for z/OS Early Support Program (ESP)

DB2 12 — The ultimate enterprise database for business-critical transactions and analytics

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

- Help customers migrate as fast as possible, but safely

Agenda

- DB2 11 prerequisites for migration to DB2 12

- DB2 12 Migration – Quick Hits

- Maintenance recommendations for early adopters of DB2 12

- DB2 12 Risk Mitigation

- Understand Continuous Delivery starting with DB2 12

- Understanding new function levels

- DB2 12 Greatest Hits

    – Fast Un-clustered INSERT

    – Index Fast Traversal

    – …

DB2 12 — The ultimate enterprise database for business-critical transactions and analytics

DB2 11 prerequisites for migration to DB2 12

- Ensure catalog consistency
  - REPAIR DBD TEST/DIAGNOSE + CHECK DATA/LOB/INDEX + DSNTESQ + ….
- Run pre-migration check queries
  - DSNTIJPM (V12) or DSNTIJPC (APAR PI58254 for V11)
- Apply fallback SPE PTF to all data sharing members
  - APAR PI33871 / II14794
- Make sure DB2 11 PTF level is reasonably current and all maintenance is applied related to DB2 12 migration
  - Use SMP/E Fix categories
    - IBM.Migrate-Fallback.DB2.V12 and
    - IBM.Coexistence.DB2.SYSPLEXDataSharing

DB2 12 — The ultimate enterprise database for business-critical transactions and analytics

DB2 11 prerequisites for migration to DB2 12 …

- Convert BSDS to 10 byte log RBA before beginning migration to DB2 12
  - For data sharing, convert single member at a time
  - Things to consider before converting the BSDS (DSNJCNVT)
    - Stop the DB2 subsystem that owns the subject bootstrap data set
    - Any utility (e.g., RECOVER, REORG) that reads from peer BSDS must be terminated in data sharing
      - Special considerations for Data Replication
        » Stop any data replication process to ensure BSDS is successfully renamed and replaced
        » Best practice is to stop data replication process first, then stop the DB2 subsystem
    - RACF user ID running DSNJCNVT must have read/write access on the new BSDSs, and read access on the old BSDSs
  - After converting the BSDS, will see increased logging volume (3 <-> 40%)
    - May need to increase size/number of active log pairs to maintain recommended 6 hours of recovery log data across active log configuration

DB2 12 — The ultimate enterprise database for business-critical transactions and analytics

DB2 12

7

DB2 11 prerequisites for migration to DB2 12 …

- Avoid autobind on pre-V10 <u>plans</u> and packages under V12
  - Explicitly rebind under V11 NFM before beginning migration
  - Use plan management for packages to keep a backup copy
  - Resolve any authorization issues
- Remember to set ZPARM ABIND=COEXIST if planning to use mixed release coexistence (V11, V12)
- FREE inactive package copies (access plan management)
- Upgrade EXPLAIN tables to V12 format (should be at least V11 version)
  - Can be done in V11 NFM with fallback SPE applied
  - Use of sample batch job DSNTIJXA with REXX DSNTXTA can help
- Apply PTFs for APARs PI69589 (V11) & PI69584 (V12)
  - Reduce catalog contention during "online" migration to V12
- Plan for activation of DB2 12 ERLY code
  - Activation via IPL or Command -REFRESH DB2,EARLY

DB2 12 — The ultimate enterprise database for business-critical transactions and analytics

DB2 12

IBM

8

DB2 12 Migration – Quick Hits

- Minimum OS level lifted from z/OS V1R13 to V2R1

- Minimum hardware level lifted from z10 to z196/z114

- Replication

  - DB2 12 (with DB2 APAR PI70998) and DB2 11 for z/OS require the Q Capture and Capture programs from IBM InfoSphere Data Replication for DB2 for z/OS Version 10.2.1

  - Q Apply and Apply programs at architecture level 1001 will work with DB2 12 and 11 for z/OS

  - APAR PI70998 for DB2

  - APAR PI66768 for IIDR 10.2.1 Q and SQL

  - APAR PI61562 for CDC

DB2 12 — The ultimate enterprise database for business-critical transactions and analytics

DB2 12 Migration – Quick Hits …

- DB2 Connect
  - Any level of DB2 Connect drivers should work with V12, both before and after new function is activated with no behavior change
  - But … Data server clients and drivers must be at the following levels to exploit DB2 for z/OS function-level application compatibility of V12R1M501 or greater:
    - IBM® Data Server Driver for JDBC and SQLJ: Versions 3.72 and 4.22, or later
    - Other IBM data server clients and drivers: DB2 for Linux, UNIX, and Windows Version 11.1 Modification 1 Fix Pack, or later
  - New clientApplCompat setting allows you to control the capability of the client when updated drivers ship changes to enable new server capability
    - You might want specific control of driver capability when:
      - DB2 client driver introduces new behavior currently not controlled by DB2 application compatibility
      - Change needs to be controlled at the application level to ensure compatibility with new behavior
  - clientApplCompat V12R1M500 is required to access V12 Server capability shipped after GA at function levels beyond DB2 V12R1M500

DB2 12 — The ultimate enterprise database for business-critical transactions and analytics

DB2 12

DB2 12 Migration – Quick Hits …

- Changes to Utilities Suite installation
  - Requires registration in SYS1.PARMLIB(IFAPRDxx)
    - CBPDO is being sunset, and SystemPac is the strategic direction
    - Any separately orderable product using only F or J FMIDs has to be changed to use an E or H base FMID
  - Documented in DB2 Utilities Suite program directory
    PRODUCT OWNER('IBM CORP') NAME('DB2 UTIL SUITE') ID('577-AF4')
    VERSION(12) RELEASE(1) MOD() FEATURENAME('V12R1') STATE(ENABLED)
  - Failure to register Utilities Suite results in utility errors
    DSNU3333I 012 14:35:50.01 DSNUGPRS - THE DB2 UTILITIES SUITE FOR Z/OS HAS
    NOT BEEN ENABLED
    DSNU3330I 012 14:35:50.09 DSNUGPTS - THE xxxxxxxx UTILITY HAS RESTRICTED FUNCTION
    IT IS PART OF THE DB2 UTILITIES SUITE FOR Z/OS WHICH HAS NOT BEEN ENABLED

- REORG MAPPING TABLE format change for longer length XRID columns (7 byte RIDs)
  - No toleration logic in V11 NFM
    - V11 NFM REORG running with the V12 mapping table format will fail
  - REORG under V12R1M100 tolerates V11 format mapping table format
  - REORG under V12R1M100 and V12R1M5xx supports the V12 mapping table format

DB2 12 Migration - Quick Hits …

- RACF changes
    - DB2 now utilizes TCP/IP DROP API through the EZBNMIFR callable service (NMI)
    - DB2 requires that RACF security profiles be defined to permit DB2 to successfully utilize this API
    - RACF PERMIT ACCESS(CONTROL) on MVS.VARY.TCPIP.DROP(OPERCMDS) for userid for xxxxDIST started task
- HVSHARE should be 510 TB (default)
    - DB2 12 requires 1 TB of 64-bit shared private storage in z/OS (same as DB2 11)
    - Virtual, not real
    - Monitor with IFCIDs 217 and 225
- Plan for real memory increase
    - Trend continues … using larger size REAL memory to deliver performance improvements
    - Expect ~ 10% increase
    - Expect up to 30% increase if taking advantage of new in-memory function
        - Largest percentage from use of Fast Traverse Block area – 20% increase on allocated VPSIZE
- Consider current zIIP utilization
    - Trend to extend zIIP offload continues
        - REORG and LOAD RELOAD phase
        - SQL query parallelism

DB2 12 — The ultimate enterprise database for business-critical transactions and analytics

DB2 12

DB2 12 Migration - Quick Hits …

- Deprecation of Basic Row Format (BRF)
  - Pagesets in BRF will continue to be supported for the time being …
  - zparm SPRMRRF will be removed in V12
  - Any REORG or LOAD REPLACE will convert BRF to RRF
  - New objects created will always be RRF
  - ROWFORMAT option in REORG will be removed from the documentation
    - Still supported from a utility syntax perspective
- Invalidation of prepared SQL statements in dynamic statement cache
  - Prior to V12, RUNSTATS would always invalidate prepared statements dependent on the object that the utility was run against
  - In V12, RUNSTATS by default will <u>not</u> invalidate the prepared statements (incompatible change)
    - Use new INVALIDATECACHE YES option to force the invalidation of prepared statements
  - Invalidation of prepared statements will still occur when
    - RUNSTATS ... INVALIDATECACHE YES
    - RUNSTATS after SQL DDL (CREATE/DROP INDEX) and statistics profile updated
    - RUNSTATS … UPDATE(NONE) REPORT(NO)
    - For other utilities, if the object was in an invalid state before the utility began e.g., rebuild pending or reorg pending

DB2 12 — The ultimate enterprise database for business-critical transactions and analytics

Maintenance recommendations for early adopters of DB2 12

- Apply regular drops of preventative service plus corrective fixes as needed
- Stabilize maintenance level before production cutover and try to keep it
- During first year, look at applying drops monthly staying 2 months back from latest level
- Pull Enhanced HOLDDATA weekly to review HIPERs and PEs, and apply where applicable
- Must be continued going forward after production cutover
- Run with stable level for a month before production cutover
- Might have to take APAR fixes

DB2 12 Risk Mitigation

- Regression testing is critical piece to keep "fires away from production"
  - Test all critical and custom processes, and scale them up
  - Run performance measurements and establish DB2 11 baseline for comparison
  - Go / No Go decision for V12 migration of production system should be based on positive results from proper testing
    - Be prepared to postpone migration as opposed to forcing in V12
  - Practice migration fallback from V12 to V11 and back to V12
  - Design fallback strategy and practice it in pre-production environments
- Minimize change and use of new function in and around when DB2 12 is first introduced into production
- For production systems, stay on V12R1M100 for at least a month until it runs smoothly
  - Leaves back door to go back to V11 NFM open for emergencies
- Disable certain DB2 12 functions – point-in-time statement until sufficient corrective maintenance is available and applied to fix high impact defects and provide additional serviceability
  - Fast Index Traversal
  - Insert Algorithm 2 (aka "Fast Insert" or "Smart Insert") for fast un-clustered insert
  - Active Log Dataset Size > 4G
  - UTS PBR RPN

Understand Continuous Delivery starting with DB2 12

- With Continuous Delivery, there is a single delivery mechanism for defect fixes and enhancements
  - PTFs (and collections of PTFs like PUTLEVEL and RSU) → same as today
- With Continuous Delivery, there are four DB2 levels
  - Maintenance level (ML) – lifted by applying maintenance
    - Also known as code level
    - Contains defect and new enhancement fixes
  - Catalog level (CL) - vehicle to enable new FL - cumulative (skip level possible)
    - DB2 Catalog changes that are needed for some FLs
  - Function level (FL) – needs to be activated - cumulative (skip level possible)
    - Introduces new DB2 features and functionality
  - APPLCOMPAT level (APPLV) – set by application - provides an "island of stability" for a given application
    - Determines SQL level of applications – can increase FL (and fallback)
    - Activates new SQL syntax
    - Freezes SQL syntax even if FL is later moved back to earlier level
    - Relies on BIND/REBIND since APPLV level in package rules
  - Minimum starting point for Continuous Delivery is DB2 12 GA level: V12R1M500

DB2 12 — The ultimate enterprise database for business-critical transactions and analytics

DB2 12

Understanding new function levels

- CM / ENFM / NFM no longer used
- Function Level V12R1M100
  - Analogous to CM
  - DB2 12 engine and catalog / directory
  - DSNTIJTC (CATMAINT) to get there
  - Fallback to V11 NFM possible
- Function Level V12R1M5xx analogous to NFM
  - New functionality available
  - Command –ACTIVATE FUNCTION LEVEL(V12R1M5xx) to get there
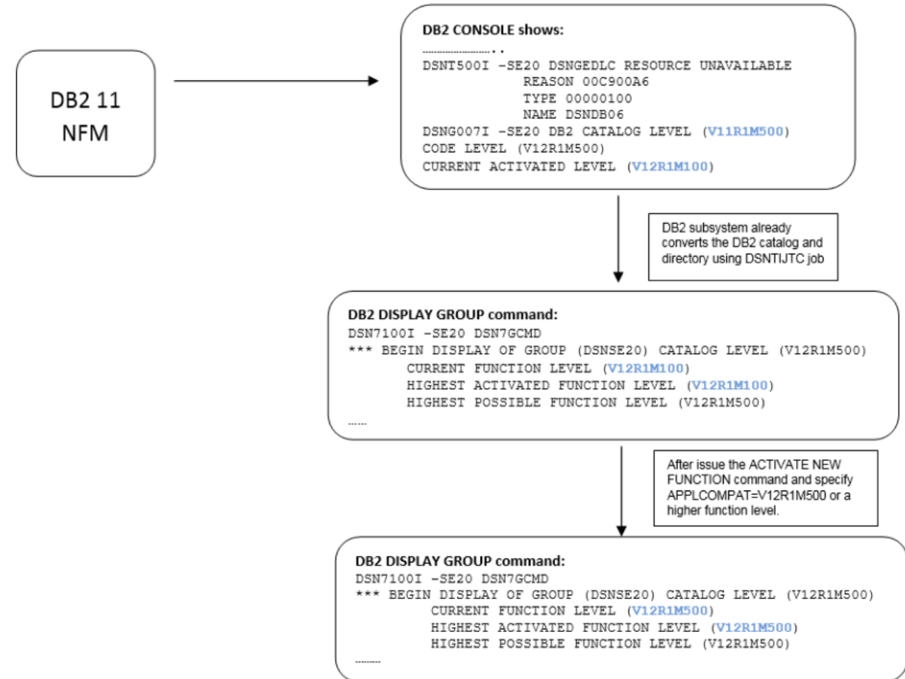 - Fallback to V11 NFM no longer possible (PIT recovery would be required)
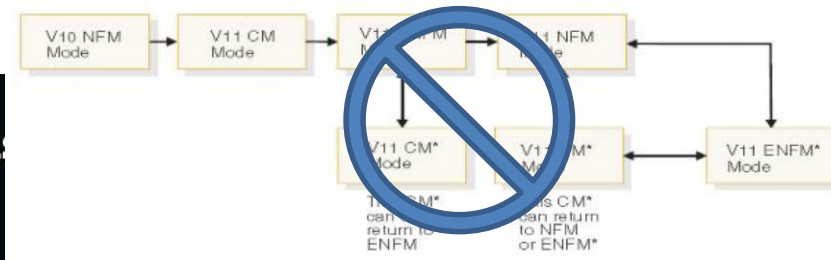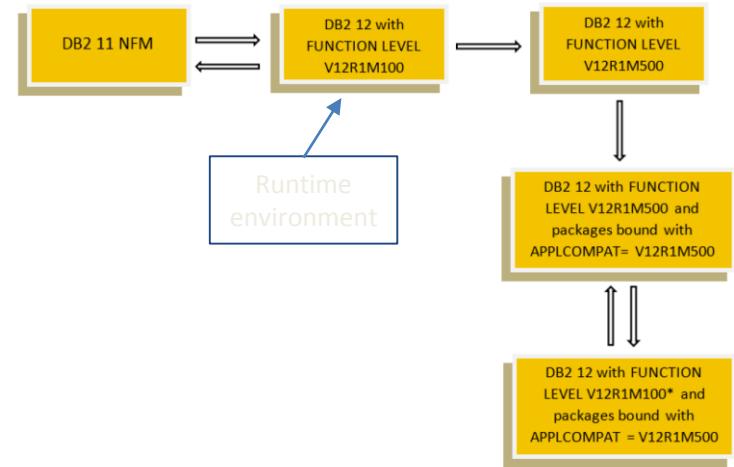
Understanding new function levels …

1. Set zparm APPLCOMPAT to V12R1M100
   - Activated Level = V12R1M100
   - Keeps the application running at the current level
2. Apply new maintenance (PTFs) to DB2 libraries
   - Maintenance Level (Code Level) = V12R1M500
   - New function exists but not active
3. CATMAINT to update Catalog Level
   - Catalog Level = V12R1M500
   - DB2 12 Catalog and Directory updated
   - Fallback to DB2 11 possible
4. -ACTIVATE FUNCTION LEVEL command
   - Function Level = V12R1M500
   - New functions available
   - No fallback possible to DB2 11
5. Bind packages to set APPLCOMPAT for DB2 Connect
   - Activated Function Level = V12R1M500
   - Set zparm APPLCOMPAT to V12R1M500
   - Applications can use new SQL syntax

DB2 11
NFM

DB2 CONSOLE shows:
```
-----------......
DSNT500I -SE20 DSNGEDLC RESOURCE UNAVAILABLE
                REASON 00C900A6
                TYPE 00000100
                NAME DSNDB06
DSNG007I -SE20 DB2 CATALOG LEVEL (V11R1M500)
CODE LEVEL (V12R1M500)
CURRENT ACTIVATED LEVEL (V12R1M100)
```

DB2 subsystem already converts the DB2 catalog and directory using DSNTIJTC job

DB2 DISPLAY GROUP command:
```
DSN7100I -SE20 DSN7GCMD
*** BEGIN DISPLAY OF GROUP (DSNSE20) CATALOG LEVEL (V12R1M500)
      CURRENT FUNCTION LEVEL (V12R1M100)
      HIGHEST ACTIVATED FUNCTION LEVEL (V12R1M100)
      HIGHEST POSSIBLE FUNCTION LEVEL (V12R1M500)
......
```

After issue the ACTIVATE NEW FUNCTION command and specify APPLCOMPAT=V12R1M500 or a higher function level.

DB2 DISPLAY GROUP command:
```
DSN7100I -SE20 DSN7GCMD
*** BEGIN DISPLAY OF GROUP (DSNSE20) CATALOG LEVEL (V12R1M500)
      CURRENT FUNCTION LEVEL (V12R1M500)
      HIGHEST ACTIVATED FUNCTION LEVEL (V12R1M500)
      HIGHEST POSSIBLE FUNCTION LEVEL (V12R1M500)
--------
```

Example of how to get to a new function level

Apply recommended RSU or PUT Level to all DB2 members (V12R1M504)

FL 503 requires new Catalog Level. Run CATMAINT to update Catalog Level.

BIND with higher APPLCOMPAT is only necessary if new SQL features are required

state:	state:	state:	state:
ML 500	ML 504	ML 504	ML 504
CL 500	CL 500	CL 503	CL 503
FL 500	FL 500	FL 500	FL 503

DB2 group migrated to V12R1M500

FL 503 **needed** as customer would like to use new DB2 functions

Run DB2 command -ACTIVATE FUNCTION LEVEL(V12R1M503) Functions are now available

DB2 12 — The ultimate enterprise database for business-critical transactions and analytics

DB2 12

DB2 12 Greatest Hits

- Fast Index Traversal (FTB)
- Dynamic Plan Stability
- Granular global commit LSN and read LSN
- Enhanced SQL MERGE
- SQL pagination syntax LIMIT / OFFSET
- Online ALTER to increase DSSIZE
- Lifting partition size limit (1 TB)
- Insert Partition
- LOB compression
- DRDA Fast Load
- Asynch CF lock Duplexing
- REORG (and LOAD) use of statistics profile

DB2 12 — The ultimate enterprise database for business-critical transactions and analytics

DB2 12

IBM

Fast Un-clustered INSERT

- Insert workloads are amongst the most prevalent and performance critical
- Performance bottleneck will vary across different insert workloads
  - Index maintenance?
  - Log write I/O?
  - Data space search (page contention, false leads)
  - Format write during dataset extend
  - PPRC disk mirroring
  - Network latency
  - etc
- Common that Index insert time may dominate and mask insert speed bottleneck on table space

Fast Un-clustered INSERT …

- Often referred to as "Insert Algorithm 2" or "Smart Insert" or even "Fast Insert"
- May potentially deliver significant improvement for un-clustered inserts (e.g., journal table pattern) where both
  - Heavy concurrent insert activity (many concurrent threads)
  - Space search and false leads on data is the constraint on overall insert throughput
- Applies to any UTS table space defined with MEMBER CLUSTER
  - Applies to tables defined as APPEND YES or APPEND NO
- Implements advanced new insert algorithm to streamline space search and space utilisation
  - Eliminates page contention and false leads
  - Default is to use the new fast algorithm for qualifying table spaces
    - INSERT ALGORITHM zparm can change the default
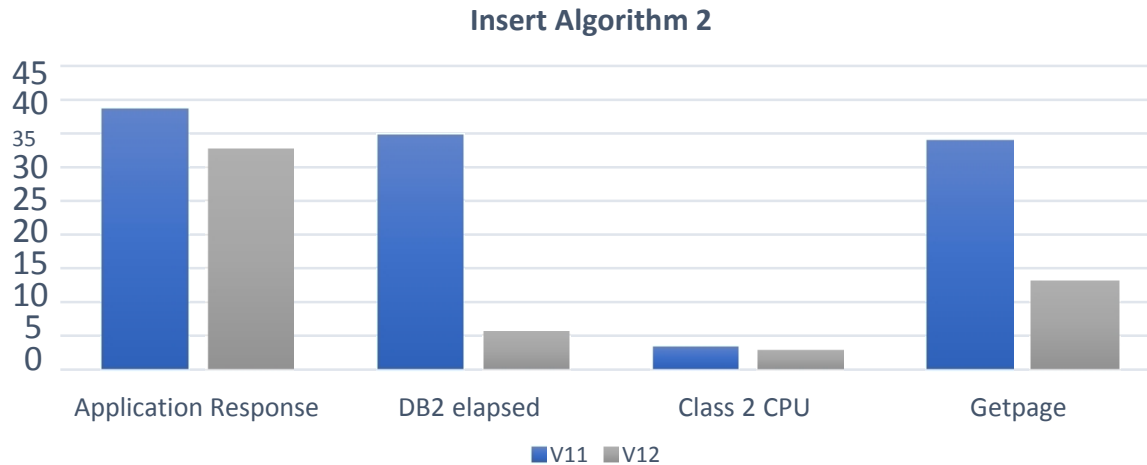   - INSERT ALGORITHM table space attribute can override zparm

Fast Un-clustered INSERT …

- Your mileage will vary
  - Some insert workloads will see no improvement and is to be expected
  - Some specific insert workloads may see significant improvement
- Will shift the bottleneck to the next constraining factor
- LOAD SHRLEVEL CHANGE can also benefit from Fast Un-clustered INSERT
- Fast Un-clustered INSERT will not be used when lock escalation occurs or use of SQL LOCK TABLE
- Available after new function activation (V12R1M5xx)

Fast Un-clustered INSERT – Shifting The Bottleneck …

Insert Algorithm 2

Bar chart comparing V11 and V12 across Application Response, DB2 elapsed, Class 2 CPU, and Getpage.

Fast Un-clustered INSERT - DB2 11 PMR Recreate …
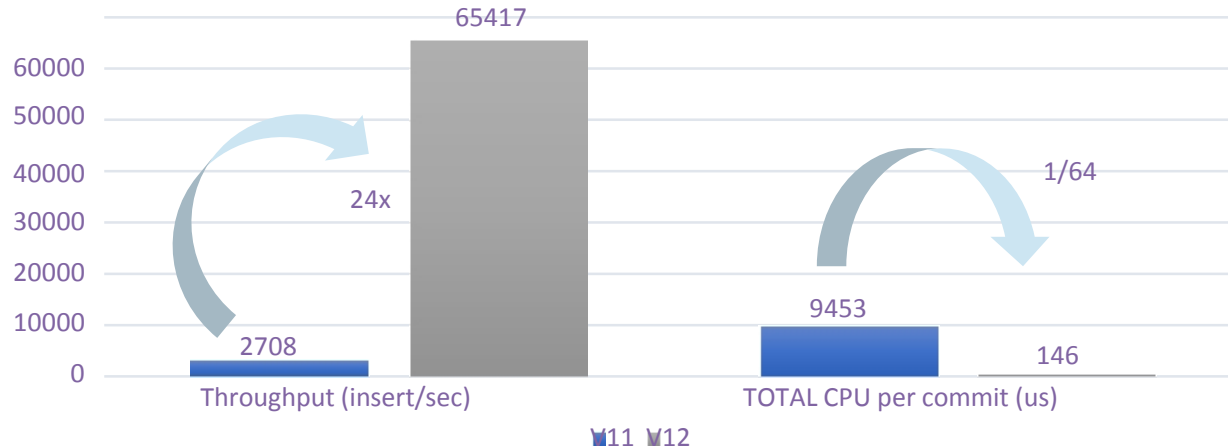
UTS PBG with MEMBER CLUSTER, RLL, with 400 bytes per row, one index,
800 concurrent threads, 10 insert per commit

DB2 12 — The ultimate enterprise database for business-critical transactions and analytics

Fast Index Traversal

- One of the most important performance features in V12
- Used for fast index lookup by avoiding expensive index B-tree traversal
- Access must be random (index traversal) pattern to benefit
- SELECT, INSERT, DELETE, UPDATE, … can all benefit
- Separate Fast Traversal Block (FTB) memory area allocated outside of bufferpool
  - Uses a concatenated structure, containing only non-leaf pages, uses relative structure
- Does not use bufferpool
  - Non-leaf pages (except root page) are not fixed in the bufferpool
  - Pages are eligible for stealing and can be LRUed out of the bufferpool when the index non-leaf pages are stored in FTB memory
- Improved performance
  - Fast traverse block is L2 cache aware B-Tree like structure
  - Each page is equal to one cache line in size (256 bytes)
- ESP customer example with
  - 9.1% CPU reduction with 3 level index, 22.9% CPU reduction with 4 level index
- Your mileage in terms of CPU reduction will vary

Fast Index Traversal …

- zparm INDEX_MEMORY_CONTROL = <u>AUTO</u>, DISABLE, x (MB)
  - AUTO = 20% allocated bufferpool size (min 10 MB)
- Each DB2 member will determine independently the good candidate indexes (daemon)
  - Index must be unique
  - INCLUDE COLUMNS supported
  - Index entry length (key + additional columns) has maximum size of 64 bytes
  - Re-evaluates every 2 minutes and adjusts priority queue
    - Index traversal (+)
    - Index only access (++)
    - Index leaf page splits (----)
    - Index lookaside (-)
  - Internal threshold then applied
- Control by SYSIBM.SYSINDEXCONTROL
  - Indicate preference for specific indexes
  - Disable for specific indexes

DB2 12 — The ultimate enterprise database for business-critical transactions and analytics

DB2 12

Fast Index Traversal …

- Monitor
    - -DIS STATS(IMU) LIMIT(*) command
        - DSNI070I -DETA FTB 0 MB FTBN 0,1 C 1
- Trace
    - -START TRACE (PERFM) DEST(SMF) IFCID(477)
    - -START TRACE (STAT) DEST(SMF) CLASS(8) IFCID(389)
- Free FTB for an index
    - Pageset close
    - SQL mass delete
    - ALTER INDEX, RECOVER INDEX, REBUILD INDEX
    - Trick: ALTER INDEX from COPY YES to COPY NO (and the other way around)

DB2 12 — The ultimate enterprise database for business-critical transactions and analytics

DB2 12

Fast Index Traversal …

- Migration
  - Available in V12 upon migration
    - Available before new function activation (V12R1M100)
    - Mixed release coexistence or V12 before new function activation (V12R1M100)
      - FTB only used while index object is not GBP-dependent
      - If index object becomes GBP-dependent, the FTB will be deleted/bypassed
  - After new function activation (V12R1M5xx)
    - FTB can now also be used when index object is GBP-dependent

DB2 12 — The ultimate enterprise database for business-critical transactions and analytics

Data dependent vs. numeric based pagination syntax

- Works very well as advertised
- Data dependent pagination syntax e.g., SELECT … FROM … WHERE (LASTNAME, FIRSTNAME) >= (:lname, :fname)
  - Given correct index design, can go <u>directly</u> to the needed rows
  - Can exploit range-list index scan (ACCESSTYPE='NR')
- **Numeric based pagination syntax e.g.,** SELECT … FROM … OFFSET 10 ROWS FETCH FIRST 10 ROWS ONLY
  - Will have to <u>skip</u> through the unneeded rows
  - If rows are deleted/inserted from other applications in between
    - May see the same rows twice or not see the rows at all
- Many static scroll cursors can be replaced by SQL pagination
  - Result set is no longer materialized
  - Read-only applications will not create long running unit of recoveries
 - Performance can be improved

More use of list prefetch

- Expected to see an increase in list prefetch (and potentially hybrid join), but not necessarily changes in the access plan where DB2 would previously chose a sort avoidance plan

- Enhancement to the Optimizer cost model to more closely reflect the true cost (and benefit) of list prefetch

- Trying to be careful not to select list prefetch (with sort) as an access path when there was an alternative access path that could use an index to avoid a sort i.e., for pagination type SQL

Example: increase on log record size after converting BSDS in V11 and entry to V12

- About 50 byte increase after converting BSDS under V11 NFM
- Further increase in log record size in V12: about 20 bytes for table space and about 28 bytes for index
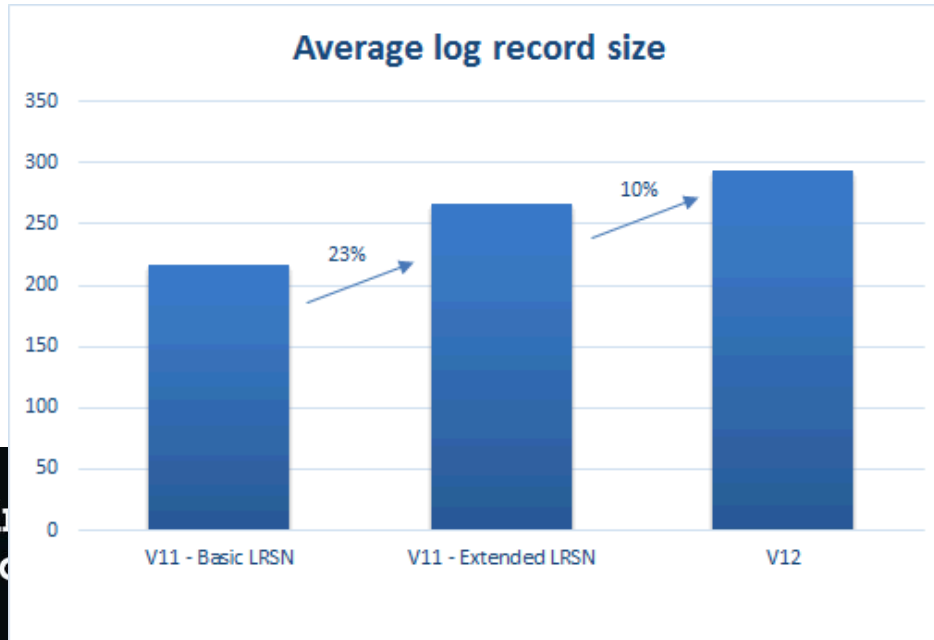
Dynamic Plan Stability

- Welcome new feature that will bring some relief in the area of dynamic SQL
- In V11 a miss in dynamic statement cache requires a new full prepare
  - DB2 subsystem recycle
  - Release migration
  - RUNSTATS
  - …
- In V12 can stabilize a query statement from the dynamic statement cache
  - No new full prepare needed
  - Statement is loaded into the dynamic statement cache from the Catalog
  - Statement is invalidated by SQL DDL like a static SQL package
- Can stabilize
  - Specific dynamic query statement
  - Dynamic query statements with more than a certain number of executions

Dynamic Plan Stability …

- No REBIND capability to "repair" after invalidations
  - Need to wait for new stabilization
- Restrictions
  - No support for concentrated statements
  - Query statements against temporal and transparent archive
- FREE stabilized dynamic query STBLGRP(x)
  - Will also invalidate the statements in the dynamic statement cache
  - May result in waves of full prepares
- Stabilized dynamic query statements do consume more CPU than the equivalent static query statement

More granular global commit LSN and global read LSN

- DB2 does not actually track "more current" value for all individual objects
- Each member maintains two global lists of the 500 objects which have the <u>worst</u> (oldest) CLSN and read-LSN values
- Global lists built by a system task that wakes every 2 seconds (subject to change)
- Rebuilds it's own list
- Merges it with every other member's list to create the global list
- When it comes time to pick up an object's CLSN or read-LSN value
  - Check the appropriate global list for the object
    - If it is on there, then we know what it's LSN is
    - If not, then use as an "alternate" LSN for the newest object (as object's LSN cannot be worse then this value)
  - Either way DB2 will compare the LSN picked with the old global value (from SCA), and use that if it is better
- Very nice enhancement that has great potential to improve lock avoidance and/or space reuse on LOB insert when the inevitable long running reader-UR is in play

LOB compression

- Requires zEDC hardware feature
  - Will decompress existing compressed LOB if zEDC not available
  - Will not compress a LOB if zEDC not available
- Inline LOB is completely separate from LOB compression
  - LOB compression only applies to the the out-of-line portion
  - Split and compressed independently
- Aimed at textual
  - Not video and audio as these are already heavily compressed outside of DB2 e.g., MP3 or MP4

SQLCODE -109 Issue

- Problem:
  - Non-documented and illegal use of SELECT ... INTO ... UNION ALL syntax
    - Customer complaints, can produce wrong results, defect
  - Loophole closed in V12
- Solution:
  - APAR PI67611 produced for V11
  - New zparm: DISALLOW_SEL_INTO_UNION
    - NO (V11 default)
      - Allows usage of this illegal SQL syntax when such usage is encountered during execution of a BIND or REBIND command
      - DB2 will write an incompatibility trace record to IFCID 376
      - Use these trace records to identify and correct applications that are using the illegal SQL syntax
    - YES (V12 default)
      - Disallow usage of this illegal SQL syntax
      - Statements that include syntax will fail with SQLCODE -109
      - Running IFCID 376 under V11 will help identify problem applications

Enhanced SQL MERGE

- V12 delivers ANSI compliant MERGE capability
- SQL MERGE is now very powerful
  - Source can now include TABLE, VIEW and full Select
  - Additional predicates on MATCHED/NOT MATCHED
  - Can do DELETE
  - Can do multiple UPDATE, INSERT and DELETE phrases
    - But not on same row
    - Can accept SIGNAL and IGNORE
- Benefits
  - Development productivity
  - Improved performance
  - Application porting to DB2

Enhanced SQL MERGE …

- But SQL MERGE is now so powerful …
  - Input can be a SELECT (JOIN) returning many rows (millions, billions)
  - # UPDATEs, INSERTs and DELETEs could explode
  - Considerations
    - No intermediate commit points
    - Long rollback time
    - Lock escalation and impact on concurrency
    - No SQL pagination support

DRDA Fast Load

- It is super fast
- Some complication to format the input records correctly
- Problem area is missing restart after failure
  - Must terminate Utility
  - RECOVER and REBUILD objects
  - Restart the DRDA Fast Load

UTS Relative Page Number (RPN)

- Motivation
  - Tremendous improvement in terms of availability and usability
    - DSSIZE can vary for different partitions
    - DSSIZE can now be increased for an individual partition with zero application impact
      - No REORG required to increase DSSIZE
  - Scalability
    - Maximum partition size increases to 1 TB
    - Maximum table size increases to 4 PB
    - Maximum number of rows in a table increases from 1.1 Tn to 280 Tn

UTS Relative Page Number (RPN) …

- Migration possible from either classic partitioned and UTS Partition By Range (PBR) table spaces
  - Steps for conversion
    1. ALTER TABLESPACE … SEGSIZE n
       - If starting from classic partitioned
    2. ALTER TABLESPACE … PAGENUM RELATIVE
       - Table space put into AREOR state
    3. REORG TABLESPACE …
  - Base and XML table spaces can be migrated separately
    - Can "coexist" running with mixed RELATIVE/ABOLUTE attributes
  - One-way ticket – no fallback to absolute page numbering (PAGENUM ABSOLUTE)
  - Extended Addressability (EA) must be used for UTS PBR RPN datasets
  - DASD space for large datasets can lead to problems (e.g. running out of volumes)
    - Datasets can only be spread across 59 volumes
    - For example, a 1TB dataset will require 3390 Model 27 or above

DB2 12 — The ultimate enterprise database for business-critical transactions and analytics

UTS Relative Page Number (RPN) …

- Migration issues
  - Part-level inline copies are required for REORG to RPN
    - Relief for tape unit constraint planned for 2Q2017
  - Pre-V6 range partitioned tablespaces with limit key values truncated at 40 bytes cannot be converted over
    - Solution planned for 1H2017

Questions

- ? ? ?

Summary

- DB2 12 migration is different
  - One phase migration (no ENFM)
  - Should reduce overall time to take advantage of new features
- Real memory exploitation required for most performance improvements
- zIIP usage likely to increase
- DB2 12 has rich set of functions that were popular with ESP customers
- DB2 12 early adopters should be prepared to:
  - Test thoroughly
  - Keep current with DB2 maintenance

DB2 12 — The ultimate enterprise database for business-critical transactions and analytics

DB2 12 — The ultimate enterprise database for business-critical transactions and analytics