

DB2 for z/OS: Availability Enhancements

Haakon Roberts
IBM

Session Code: B03

Tuesday May 3, 2011, 3:15 pm | Platform: DB2 for z/OS





Disclaimer/Trademarks

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

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

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

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

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

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

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

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

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

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



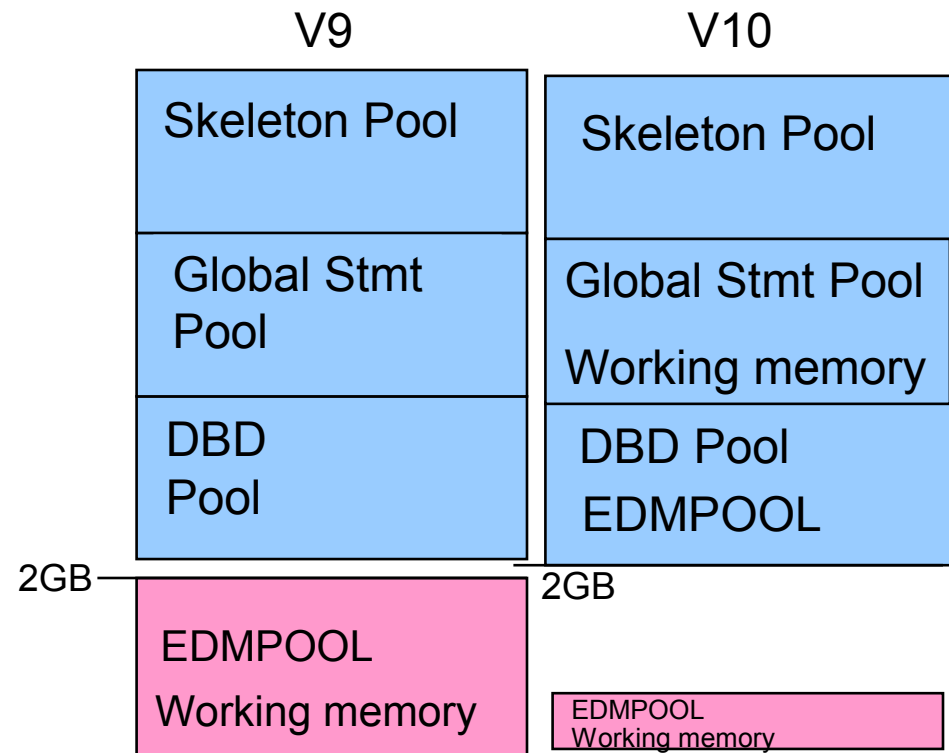
Agenda

- VSCR
- Other Scalability & Contention Relief
- Logging & Restart
- Online Schema
- Utilities
- Summary



DB2 10: 64 bit Evolution Virtual Storage Relief

- DB2 9 helped (~ 10% – 15%)
- DB2 10: 5 to 10 times more threads, up to 20,000
 - Move 80% - 90% above bar
 - More concurrent work
 - Reduce need to monitor
 - Able to consolidate LPARs
 - Reduced cost
 - Easier to manage
 - Easier to grow



Scalability: Virtual storage constraint is still an important issue for many DB2 customers in 9



DB2 10: Other Scalability & Contention Improvements

- Restructure catalog & directory
 - Improved concurrency for BIND, PREPARE and DDL
 - Increased use of LOBs
 - Eliminate 64Gb limit on SPT01 and other pagesets
 - Eliminate hash chains & contention on DBD01 hash anchors
 - No need for DSN1CHKR any more
 - Enable row level locking on key catalog tables
 - Enable REORG SHRLEVEL CHANGE on entire catalog & directory
 - Eliminate UTSERIAL lock for improved utility concurrency
 - Easier management: DB2 managed & SMS controlled



Other System Scaling & Contention Improvements

- Access currently committed data for application lock avoidance
 - New USECURRENTLYCOMMITTED BIND/PREPARE option
 - Return committed data without waiting for update locks
 - Available for UTS only
 - Supports retrieval of uncommitted deletes (pseudo-deletes)
 - Skips uncommitted inserts
 - No support for uncommitted updates
 - New instrumentation in statistics IFCID 2
 - QISTRCCI – uncommitted inserted rows skipped
 - QISTRCCD – uncommitted deleted rows read



Logging & restart

- Ability to checkpoint based on both time and number of log records
 - Meaning of CHKFREQ is unchanged
 - Minimum # of log records raised from 200 to 1000
 - New ZPARMs to control new behavior
 - CHKLOGR – number of log records between checkpoints
 - 1000 – 99,999,999
 - CHKMINS – number of minutes between checkpoints
 - 1-1439
 - CHKTYPE SINGLE|BOTH – govern old/new
 - Set by dynamic ZPARM or –SET LOG command
 - -SET LOG change does not persist across restart
 - -DIS LOG command indicates settings and if mode is SINGLE or BOTH



Logging & restart

- Dynamic add of active logs
 - New –SET LOG NEWLOG option
 - New active log must be IDCAMS defined & preformatted by DSNJLOGF
 - Only a single log dataset at a time
 - Issue command twice for dual logging
 - Limit is still 93 active log pairs
 - No dynamic delete of active logs
- Pre-emptable backout
 - Pre-V10, abort/backout schedules non-preemptable SRB
 - On single CPU system may give impression of DB2 hang
 - V10: Create enclave at restart for preemptable SRB backout processing



Logging & restart

- New `DEL_CFSTRUCTS_ON_RESTART` ZPARM for auto-delete of CF structures on restart if no active connections exist
 - Useful for DR environments
 - Delivery via APAR in 2011 (V9 & above)
- New `LBACKOUT` option to allow restart light with postponed abort
 - Retained locks will be kept for PA URs
 - Delivery via APAR 1Q 2011 (V9 & above)

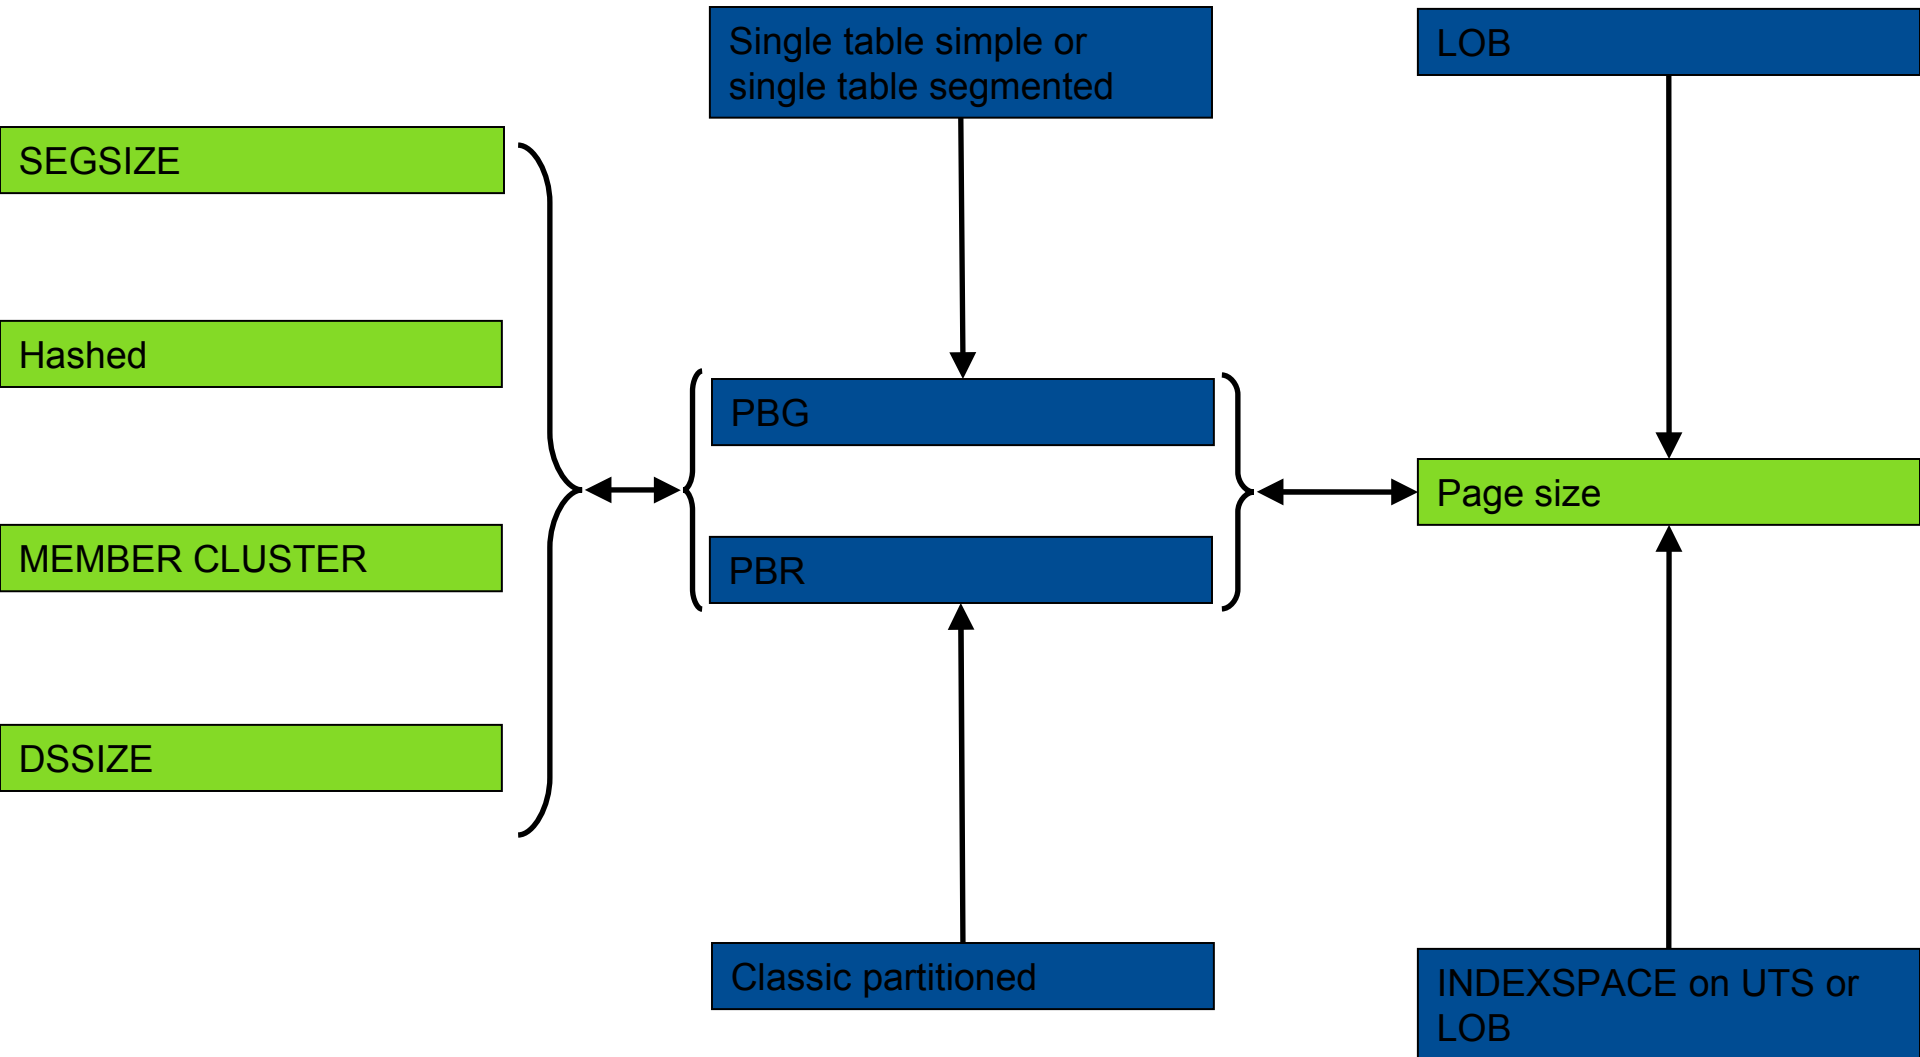


Online Schema

- V9:
 - Change of table - or index space attributes can require an outage
 - Unload data
 - Drop table space
 - Recreate table space, tables, indexes, views
 - Re-establish authorization & RI
 - Reload data
 - Undo of changes requires same process
- V10:
 - Execute ALTER statement
 - Changes are cached and materialized on next table/index space-level REORG
 - Pending changes can be dropped
 - Some restrictions exist
 - Mixing immediate & deferred DDL
 - PIT recovery to prior to a materializing REORG



Online Schema – DB2 10





DB2 10 REORG: improve availability & remove restrictions

- Reduced need for REORG INDEX
 - List prefetch of index leaf pages based on non-leaf information for range scans
- Improved performance for part-level REORG with NPIs & REORG INDEX
 - Index list prefetch results in up to 60% elapsed time reduction
- Reduced need for REORG with compress on insert
- Reduced need for REORG if clustering not important
 - New REORGCLUSTERSENS RTS column
 - If no clustering-sensitive queries then avoid REORG to restore clustering
 - DSNACCOX enhanced




DB2 10 REORG: improve availability & remove restrictions

- REORG SHRLEVEL CHANGE for all cat/dir pagesets
- REORG SHRLEVEL REFERENCE|CHANGE to remove REORP
- REORG SHRLEVEL CHANGE for LOBs
 - Independent of whether LOBs are LOG NO or LOG YES
 - No mapping table required
 - Base table space must be LOGGED
 - REORG SHRLEVEL NONE deprecated
 - Will run, but no REORG will be performed in NFM
- REORG FORCE option to cancel blocking threads
 - Same process as –CANCEL THREAD so requires thread to be active in DB2 for it to be cancelled
 - Threads cancelled on final drain attempt
- Reduce application outage on REORG with inline stats
 - Update catalog statistics after dedrain
- REORG of multiple part ranges
 - Reduce # of REORGs, process NPIs once, maximise exploitation of parallelism
 - Retrofitted to V9 in PK87762
 - LISTDEF support is not retrofitted



DB2 10 REORG: improve availability & remove restrictions


- New AUX keyword on REORG of partitioned base for improved LOB handling
 - Allows rows to flow between partitions even though LOB columns exist
 - Allows REORG REBALANCE with LOB columns
 - Allows ALTER of LIMITKEY with LOB columns
 - Permits move of rows between parts on PBG REORG
 - Permits deletion of corresponding LOBs on REORG DISCARD
 - Default is AUX NO unless LOB objects required to complete REORG
 - No XML column support for classic partitioned or PBR
 - No mapping table change





DB2 10 LOAD & UNLOAD

- Remove MAX_UTIL_PARTS zparm
 - Restriction removed for REORG in V9
- Improved performance for LOAD REPLACE with LOB data
 - Up to 50% elapsed time reduction
- Spanned record support for LOB/XML data
 - LOBs & XML documents inlined in SYSREC with base data
 - Option in addition to FRVs
 - Performance & portability





LOAD & UNLOAD performance

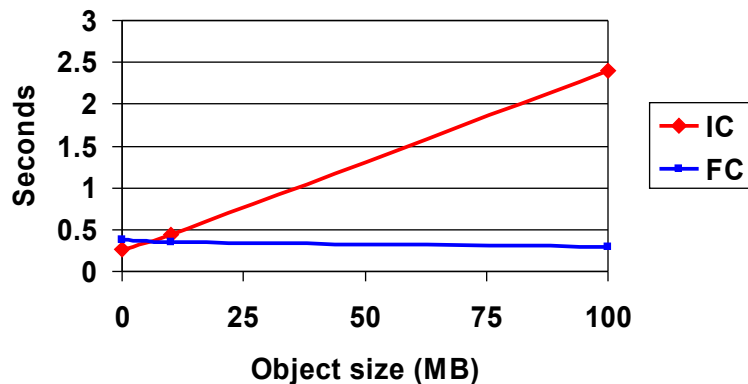
- PM08585 (V9 -)
 - REPAIR SET ... RBDPEND|PSRBDPEND
 - May be useful for improving heavy insert performance by skipping updates to non-unique indexes
- PM27962 (V9 -)
 - Allow LOAD REPLACE or RESUME of a partition even though NPI is in PSRBD
 - Allow LOAD RESUME of a partition even though NPI logical partition is in RBDP
 - Can be used with REPAIR, which can set PSRBD or RBDP
- PM19584 (V9 -)
 - LOAD PRESORTED option for improved performance when data is presorted in clustering order
 - LOAD/UNLOAD FORMAT INTERNAL
 - 85% CPU & elapsed time reduction on UNLOAD
 - 77% elapsed time, 56% CPU reduction on LOAD




DB2 10 Backup & Recovery

- Dataset-level Flashcopy support
 - COPY, RECOVER, REORG, LOAD, REBUILD INDEX, REORG INDEX
 - New zparms & utility parms to govern
 - Significant CPU & elapsed time reduction
 - Eliminate impact on buffer pools
 - Create consistent imagecopies from SHRLEVEL CHANGE
 - Provide recoverability after REBUILD INDEX
 - Consistent image copies when LOAD data is discarded

CPU time per object (z10)



Elapsed time per object (z10)





DB2 10 Backup & Recovery

- New BACKOUT YES option for point in time recovery
 - Requires COPY YES for index support
- VERIFYSET option to fail PIT recovery if entire set not included
 - Base, LOB, XML, history – not RI
- ENFORCE option to avoid CHKP/ACHKP when subset of set recovered
 - Improved performance due to avoidance of set checking (RI, aux)
- PM24237 (V9 -)
 - MODIFY to validate existence of SLBs when deciding whether to set COPY-pending



DB2 10 Data Verification

- CHECK utilities will no longer set CHKP/ACHKP
- CHECK SHRLEVEL CHANGE default changed to fail if Flashcopy not available
 - Utilities must meet availability requirement
- CHECK DATA enhanced for XML support
 - Document validation
 - Schema validation
- Automated exception table processing for XML documents
- New CHECK_FASTREPLICATION zparm
 - Allow CHECK SHRLEVEL CHANGE to use FASTREPLICATION(REQUIRED)
 - PM19034 (V9)
 - Fail CHECK utility rather than incur application outage
 - Recommendation: Set REQUIRED rather than default PREFERRED in V9



Summary

- Consistent & continuing focus
- Remove scalability inhibitors
- Reduce or eliminate contention
- Eliminate planned & unplanned outages

Haakon Roberts

IBM

haakon@us.ibm.com

B03

DB2 for z/OS: Availability Enhancements

