




# IBM zEnterprise Technology Summit

Presenter – Title

Date





## Agenda

- **Recap of DB2 10 System Availability & Scalability Features**
- **Recent Updates**
- **Recap of DB2 10 Utilities Suite Availability & Scalability Features**
- **Recent Updates**
- **A Peek into the Future**

An abstract graphic in the top left corner composed of several 3D cubes and rectangular blocks in shades of red, pink, orange, and black, arranged in a cluster.

# System Availability & Scalability Features

# DB2 10 Virtual Storage Constraint Relief

## ▪ DBM1 below 2GB

- 75-90% less usage in DB2 10 compared to DB2 9
- Some of working storage (stack, xproc storage) stays below 2GB

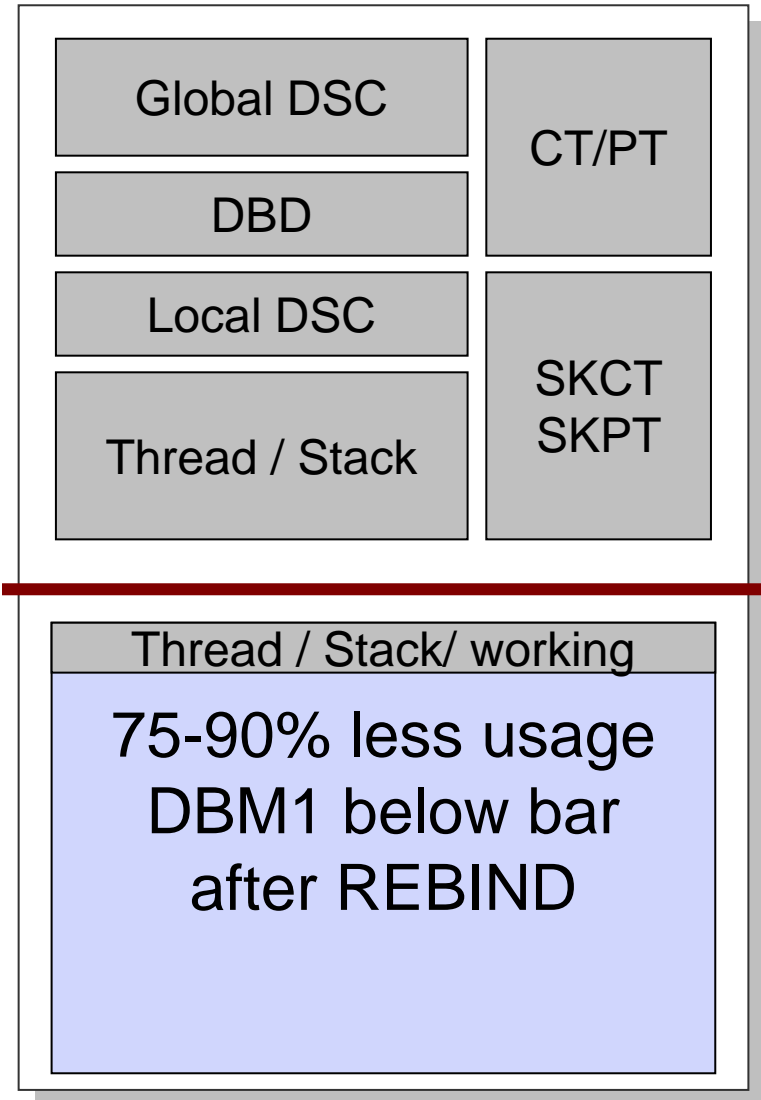
## ▪ Larger number of threads

- Possible data sharing member consolidation

## ▪ Improve CPU with storage

- More release deallocate
- Larger MAXKEEPD values for KEEP DYNAMIC=YES

## DB2 10






## Other Availability & Scalability Improvements

- **Reduced latch contention**
  - Improved efficiency for latch suspend/resume
- **New access currently committed option for readers to avoid waiting for inserters or deleters**
- **Use 64-bit common storage to avoid ECSA constraints**
- **DB2 10 NFM catalog restructure improves BIND / Prepare concurrency**
- **Remove SPT01 64Gb constraint**
- **Improved accounting rollup, compress SMF option**
  - Reduced SMF data volume
- **Lower overhead for very large buffer pools**

## Major changes in DB2 10 catalog & directory

- **Improve availability and productivity**
- **Increase maximum size substantially**
- **Reduce contention: BIND, Prepare, Utilities**
  - DDL concurrency also improved from removal of DBD01 hash anchor locks
- **Catalog changes: Remove links, hashes**
  - Many more table spaces, partition by growth
  - Row level locking, reordered row format
  - CLOB and BLOB columns for long strings
    - Inline for performance
  - Online reorganization and check
  - More automatic: DB2-managed SMS-controlled





## Logging & Restart in DB2 10

- **Ability to checkpoint based on both time and number of log records**
- **Dynamic add of active logs**
- **Pre-emptable backout**
- **Auto rebuild CF lock structure on long IRLM waits during restart**
  - Can avoid group-wide shutdowns



## Online Schema in DB2 9

- **Change of table - or index space attributes can require an outage**
- **Change of table space attributes**
  - Unload data
  - Drop table space
  - Recreate table space, tables, indexes, views
  - Re-establish authorization & RI
  - Reload data
- **Change of index space attributes**
  - Alter index
    - Index placed in RBDP
  - Rebuild index
- **Undo of DDL changes**
  - Same as above

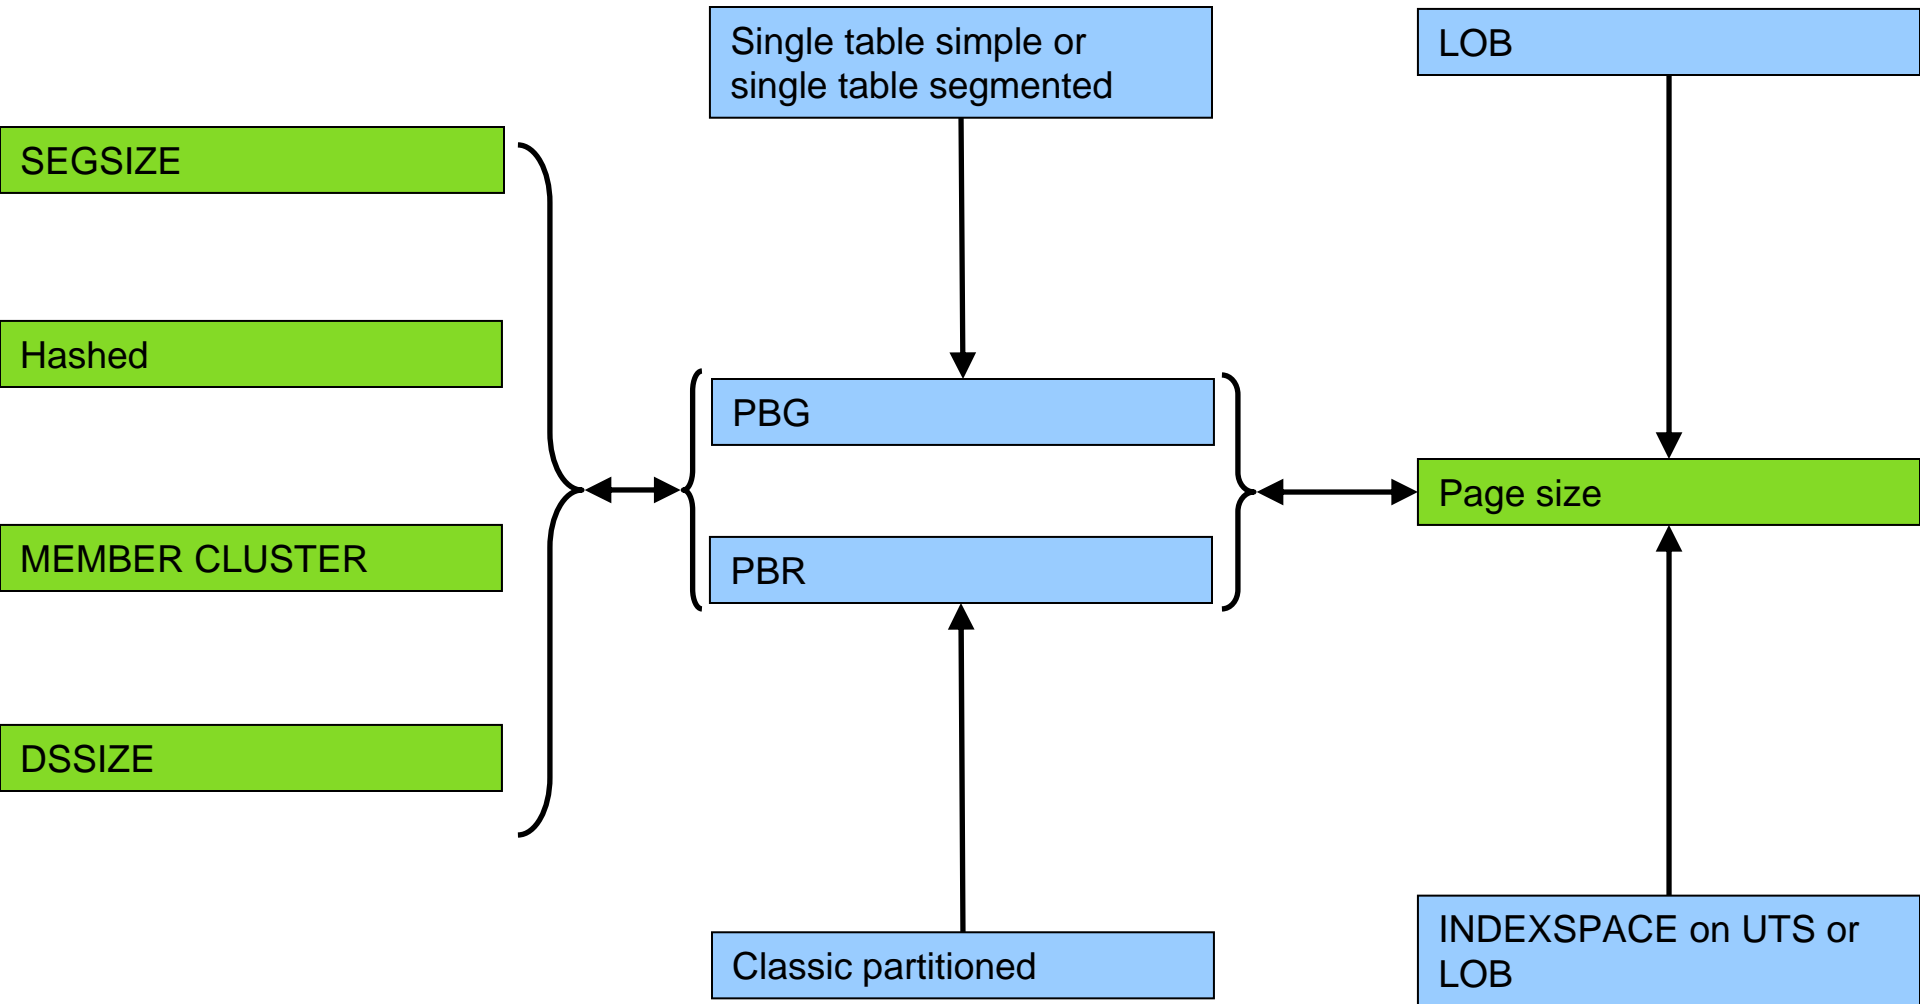




## Online Schema in DB2 10

- **Execute ALTER statement**
- **Changes are cached & materialized by next REORG**
  - SHRLEVEL REFERENCE|CHANGE
- **Undo of DDL changes if not materialized**
  - ALTER TABLESPACE... DROP PENDING CHANGES
  - All pending changes are removed
  - AREOR is not reset
- **Undo of DDL changes if materialized**
  - Perform compensating ALTER & schedule REORG
    - Assumes no dependencies on prior ALTER have evolved

# Online Schema – DB2 10





## Recent System APARs

- **New DEL\_CFSTRUCTS\_ON\_RESTART ZPARM for auto-delete of CF structures on restart if no active connections exist**
  - Useful for DR environments
  - APAR PM28925 (V10)
  - APAR PM31807 (IRLM support)
- **New LBACKOUT zparm options LIGHT and LIGHTAUTO to allow restart light with postponed abort**
  - Retained locks will be kept for PA URs
  - APAR PM34474 (V9)
  - APAR PM37675 (V10)
  - Previously, LBACKOUT would be ignored and all URs would be backed out if LIGHT specified on restart
- **New DDLTOX zparm to provide a timeout multiplier for DDL & DCL**
  - PM32921 (V9) & PM37660 (V10)



## Recent System APARs

- **DDF Availability Enhancements**

- Change DDF location alias names online
  - New MODIFY DDF ALIAS command
  - New “dynamic alias” concept in addition to existing 8 static aliases
- Online DDF CDB changes
  - LOCATIONS, IPNAMES, IPLIST
- PM26480 (V10)

- **Ability to SELECT from SYSLGRNX**

- PM35190 & PM42331 (V10)
- ISO(UR) enforced



## Recent System APARs

- **Increase FLA storage from 100Mb to 510Mb**
  - PM31641 (V10)
  - 5X increase in recommended limit on concurrent RECOVER jobs per DB2
  
- **ALTER MAXROWS to set AREO\* rather than AREOR**
  - PM43597 (V10)
  - Allows MAXROWS to take effect at partition level

An abstract graphic in the top left corner composed of several 3D cubes and rectangular blocks in shades of red, pink, orange, and black, arranged in a cluster.

# Utilities Availability & Scalability Features



## DB2 10: REORG

### ▪ REORG avoidance

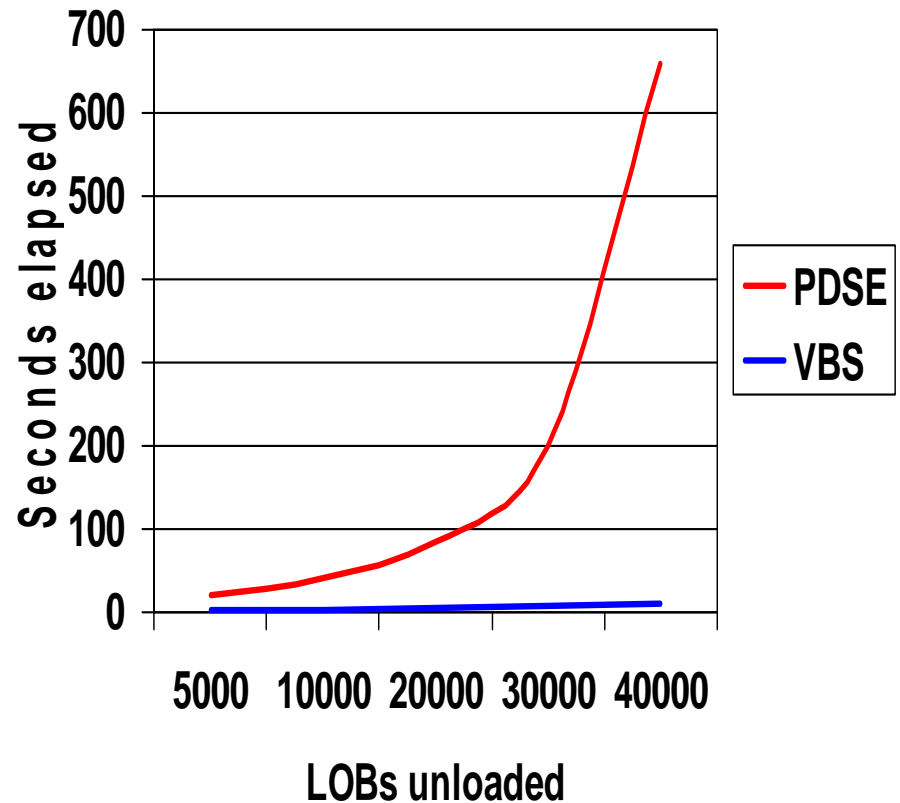
- Reduced need for REORG INDEX with index leaf page list prefetch
- Reduced need for REORG with compress on insert
- Easily identify whether REORG to regain clustering is important

### ▪ Improved availability & removed restrictions

- Reduced application outage for REORG with inline stats
- REORG support for multiple part ranges
- REORG SHRLEVEL CHANGE for all cat/dir page sets
- REORG SHRLEVEL REFERENCE|CHANGE to remove REORP
- REORG SHRLEVEL CHANGE for LOBs
- REORG FORCE option to cancel blocking threads
- REORG REBALANCE with LOB columns
- ALTER LIMITKEY permitted with LOB columns

## DB2 10: LOAD & UNLOAD

- **Improved performance for LOAD REPLACE with LOB data**
  - Up to 50% elapsed time reduction
- **Spanned record support for LOB/XML data**
  - Performance, availability & scalability



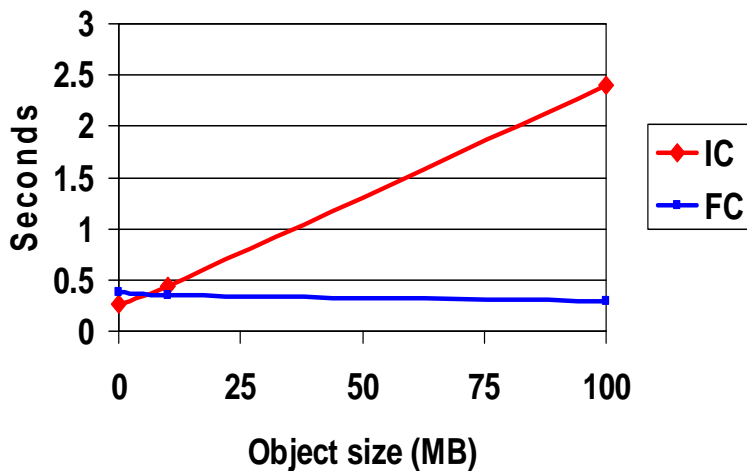


# DB2 10: Backup & Recovery


## ▪ Dataset-level Flashcopy support

- Remove recoverability constraint for indexes after REBUILD
- Create transaction-consistent image copies with no application impact
- Create partition-level inline image copies from REORG

CPU time per object (z10)




Elapsed time per object (z10)



## DB2 10: Backup & Recovery

- **VERIFYSET** option to fail PIT recovery if entire set not included
- **ENFORCE NO** option to avoid **CHKP/ACHKP** on PIT recovery of subset of set
  - Improved performance due to avoidance of set checking (RI, aux)
- **Fast recovery to point in time through new BACKOUT option**
  - Include indexes in RECOVER list to avoid the need to rebuild them
  - PIT recovery always with consistency since V9





## DB2 10: Other Utilities

### ▪ CHECK DATA

- CHECK utilities will no longer set CHKP/ACHKP
- CHECK DATA enhanced for XML support
  - Document validation
  - Schema validation
- Automated exception table processing for XML documents

### ▪ Spatial index support for REORG, REBUILD INDEX & CHECK INDEX

- Avoid DROP/CREATE of spatial indexes

### ▪ Removed UTSERIAL lock for greater utility concurrency



## Recent Utility APARs

- **LOB pageset support for RECOVER BACKOUT YES**
  - PM45650 (V10)
  - Lost LOBs will be marked invalid if necessary
  - Tablespace will be placed in AUXW
- **Improved MODIFY RECOVERY support for SLBs**
  - PM24237 (V9)
  - MODIFY will check for existence of SLB before setting copy-pending
    - REPORT RECOVERY support for SLBs delivered in base V10
- **Improved online REORG concurrency for materializing deferred ALTERs**
  - PM25648 (V10)

## Recent Utility APARs

- **REORG PARALLEL parameter & REORG\_LIST\_PROCESSING zparm to control whether LISTDEF partitions are processed in single REORG**
  - PK87762, PM13259, PM25525, PM37293 (V9)
- **Online REORG materialization of inline LOBs**
  - PM29037 (V10)
- **REORG to correct BRF/RRF mismatches after DSN1COPY**
  - PM40646 (V9)
- **Faster RECOVER BACKOUT processing**
  - PM70181 (V10)
- **REORG outage reduction**
  - Reduce final incremental processing in last log iteration
  - PM46632 (V9)
  - One case: 7 minutes to zero



## Recent Utility APARs

- **LOAD/UNLOAD FORMAT INTERNAL**

- PM19584
- 85% CPU & elapsed time reduction on UNLOAD
- 77% elapsed time, 56% CPU reduction on LOAD
- Supported by High Performance Unload

- **LOAD PRESORTED option if data already in clustering order**

- PM19584
- Up to 25% CPU reduction, 33% ET reduction depending on no. of indexes
- Complements Utility Enhancement Tool PRESORT option

- **Fast LOAD through index avoidance**

- PM27962 (V9)
- New INDEXDEFER option to skip index key insert

# Recent Utility APARs

- **Retrofit utility spatial index support to V9**
  - PM35200 (V9)
- **Prevent read-only outage from CHECK SHRLEVEL CHANGE utility when Flashcopy support not available**
  - PM19034 (V9)
  - New CHECK\_FASTREPLICATION zparm
  - Recommend set to REQUIRED
- **Improve DSN1COPY for partition-level replace with table versioning**
  - PM27940 (V9)
- **Allow utility restart when template uses &UNIQ or &UQ**
  - PM19063 (V9)
- **Extend DFSORT variable length record sort limit from 2bn to 4bn**
  - PM43006 (V9)



## Recent Utility APARs

- **Remove scalability constraint for multi-billion row online REORG**
  - PM58177 (V9)
  - Allow creation of mapping tables in PBGs
- **Remove initial drain of PBG tablespace for part-level PBG REORG**
  - PM61752 (V9 only)
- **Forward-fit REORG\_IGNORE\_FREESPACE to DB2 10 for part-level PBG REORG**
  - PM53254 (V10)
- **Partition-level REORG performance enhancement with NPSIs**
  - PM55051 (V9)





# A Peek into the Future

## DB2 11 – Some Early Thoughts on Availability

- Unmatched reliability, scalability, and availability
- Expanded logging capacity & scalability
- Address 6 byte RBA/LRSN constraint
- More deferred ALTER online schema changes
- BACKUP SYSTEM & RESTORE SYSTEM improvements
- Availability improvements in REORG
- Faster recovery from inline image copies
- Faster recovery of catalog & directory
- Improved recoverability for long-running REORGs
- Remove some recoverability restrictions associated with online schema
- Improved data integrity checking on pageset open



## Summary

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