IBM System z Technology Summit

How to Streamline Your DB2 for z/OS Utility Processing

Feb 24, 2011

Jim Brogan

Certified Consulting IT Specialist

jambrog@us.ibm.com

860-275-5647

Trends in Database/Utility Management

- Data growth puts pressure on IT infrastructure, SLAs, staff, and performance
- According to IDC, the amount of data is exploding. **Structured data is growing 32% per year, unstructured data is growing 63% and replicated data is growing 49%.** Companies are compelled to take the right steps to protect their valuable data and maintain high database availability
  - Average data growth per year is approximately 30%
  - Large critical application data growth rate is > 50%
- In the last 10 years the number of objects needing performance management has increased:
  - The number of objects that need management has increased 3X
  - the number of objects per DBA has increased 4X
- Focus on reducing CPU and elapsed time
- Running multiple databases on a server has become the norm
- 90% of customers have more than one DBMS ➔ Resource/skill issues, consistent administration efforts, increased cost in administration, greater need to automate routine operations

IBM Investment Areas for Managing Utilities

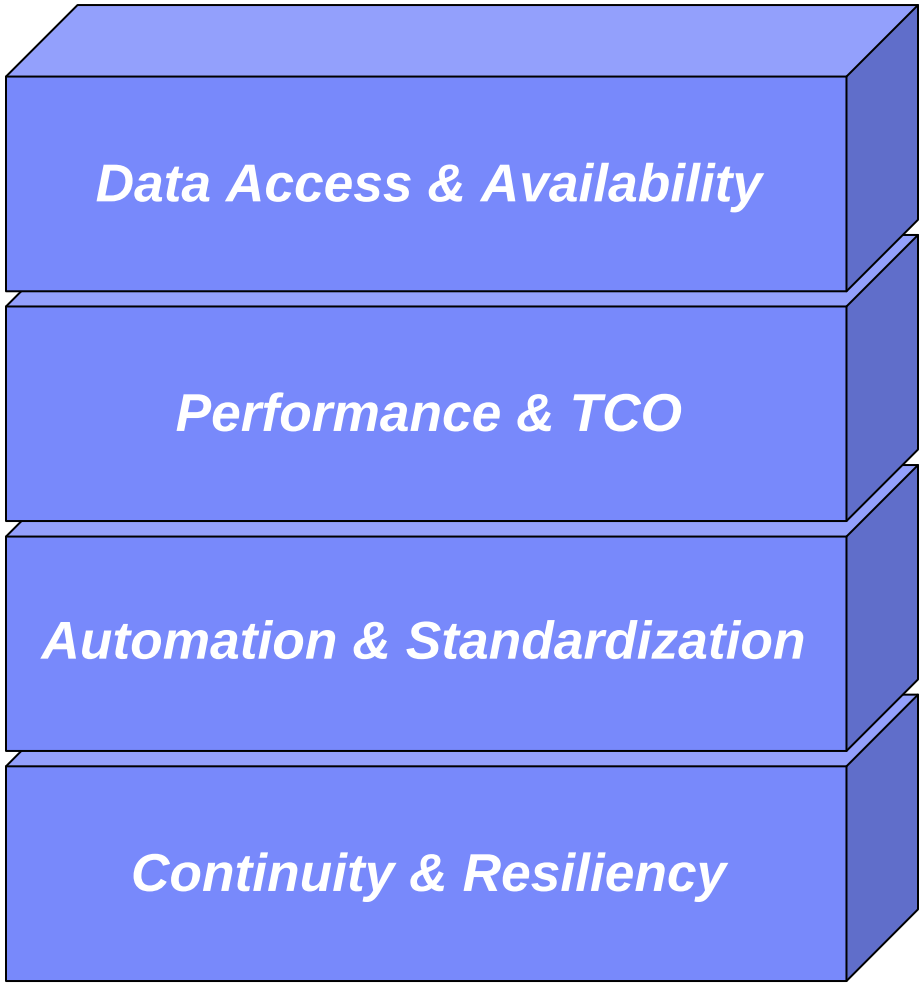

- **Data Access & Availability**
  - Fast retrieval of information
  - Reducing the amount of down time or minimizing batch window for maintenance
- **Performance & TCO**
  - Meeting or exceeding SLA's and/or chargeback
  - Reducing CPU and ET to achieve lowest TCO
- **Automation & Standardization**
  - Reducing repeated tasks, manual effort and error
  - Ensuring consistency at company level
- **Continuity & Resiliency**
  - Ensuring data integrity
  - Ensuring Day-1 support of new versions of DB2 for z/OS

DB2 10 for z/OS Utilities

DB2 10 for z/OS Utility Support

- **Topics for discussion**
- **Utility support of core DB2 10 for z/OS function:**
  - Catalog restructuring
  - Online schema
  - Inline LOBs
  - Hash access
- **REORG - - Lots of new enhancements!**
- **ALSO significant enhancements**
  - LOAD/UNLOAD
  - COPY
  - RECOVER
  - RUNSTATS
  - CHECK

Catalog restructuring

- **Reduce contention for DDL/BIND/PREPARE**
- **Eliminate 64Gb limit on SPT01**
- **Eliminate hash chains & contention on DBD01 hash anchors**
- **Performed by REORG during ENFM**
- **Enables REORG SHRLEVEL CHANGE for complete catalog & directory**
- **Enables row level locking for further concurrency**

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
    - SHRLEVEL REFERENCE or CHANGE
    - LOAD REPLACE will not materialize
  - Pending changes can be dropped
  - Some restrictions exist
    - Mixing DDL
    - PIT recovery

Online Schema – V10

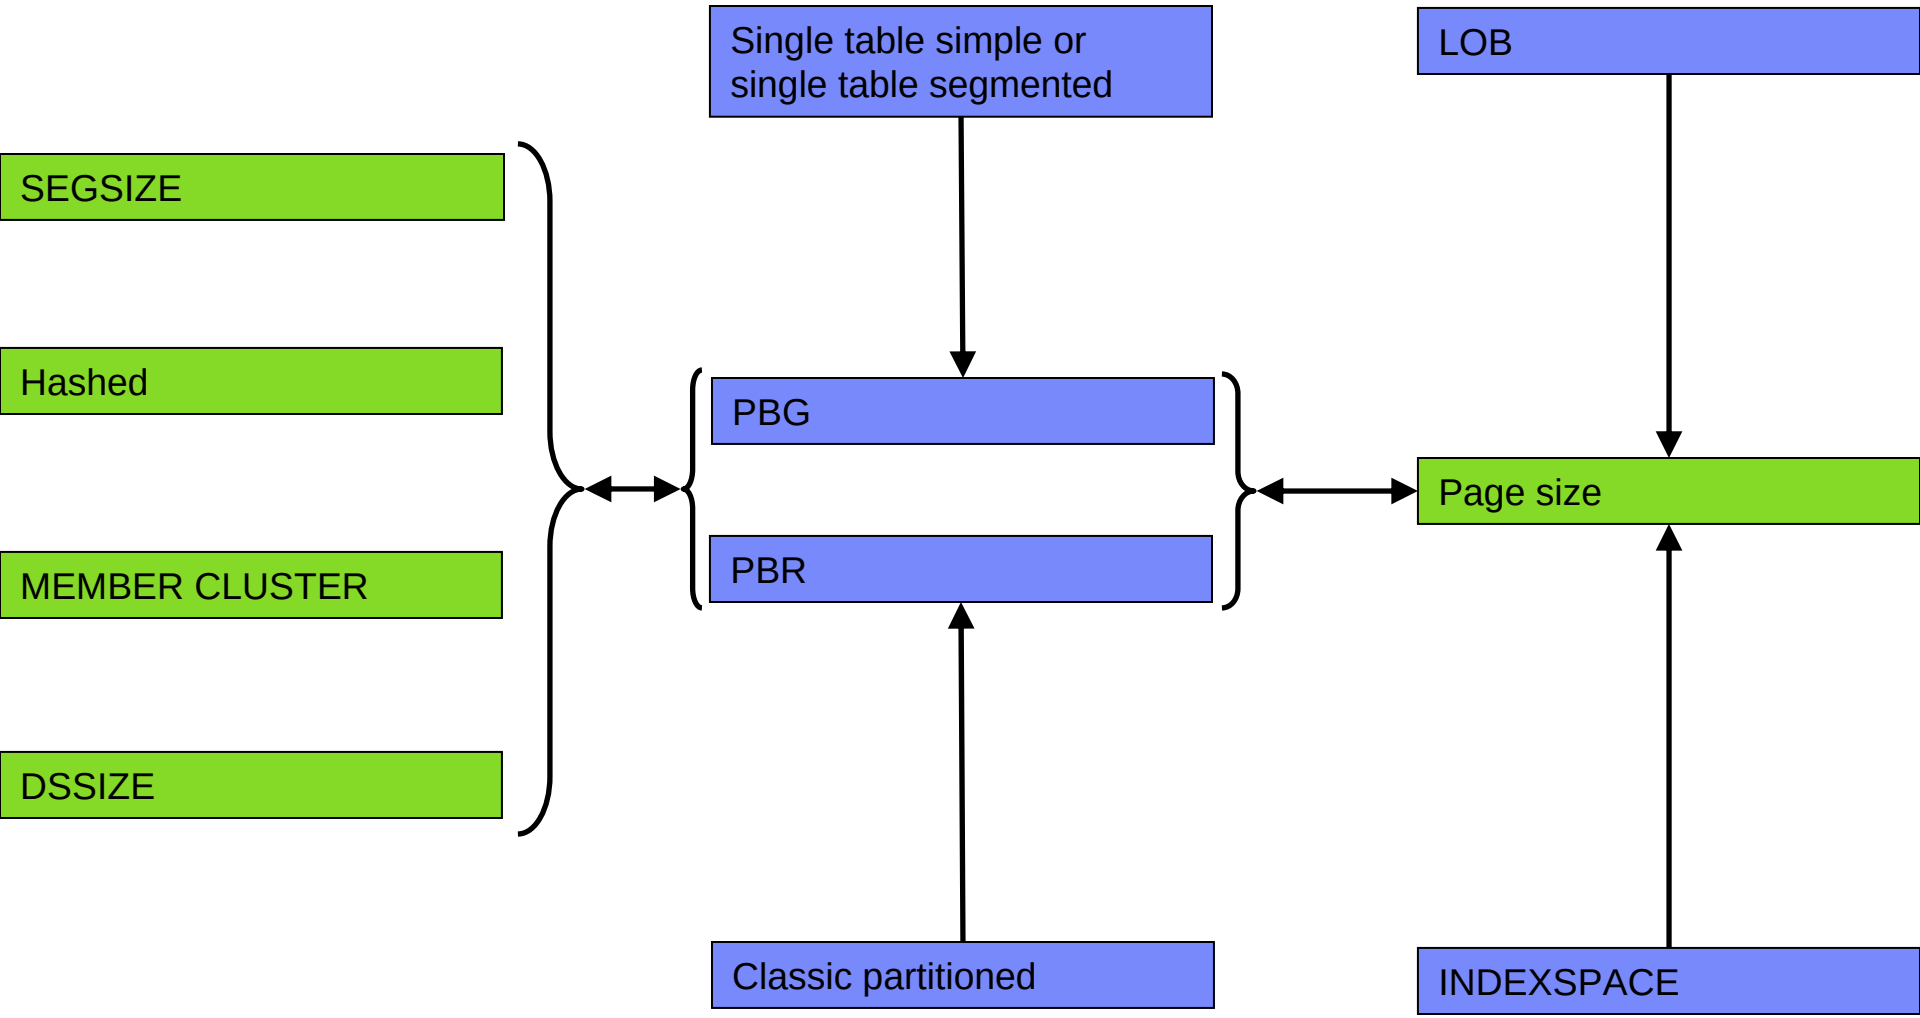

Inline LOBs

- **Inline whole or portion of LOB with base row**
- **Improved performance for small LOBs**
- **Enables compression for inlined LOB data**
- **Immediate ALTER**
  - ALTER to increase size of inline portion sets AREO*
  - ALTER to decrease size of inline portion sets REORP
- **Requires REORG to inline existing LOB data**
  - SHRLEVEL REFERENCE supported today
  - Expect delivery of SHRLEVEL CHANGE via maintenance stream early 2011

REORG – improved availability & removed 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**
- **New REORGCLUSTERSENS RTS column**
  - If no clustering-sensitive queries then avoid REORG to restore clustering
  - DSNACCOX enhanced
- **REORG SHRLEVEL CHANGE for all cat/dir page sets**
- **REORG SHRLEVEL REFERENCE|CHANGE to remove REORP**

REORG – improved availability & removed restrictions

- **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 FORCE option to cancel blocking threads**
  - FORCE ALL or just READERS
  - Same process as –CANCEL THREAD so requires thread to be active in DB2 for it to be cancelled
  - Threads cancelled on final drain
- **Reduced application outage for REORG with inline stats**
  - Update catalog after dedrain
- **REORG support for multiple part ranges**
  - REORG support retrofitted to V9 in PK87762
    - LISTDEF support is not retrofitted

REORG – improved availability & removed restrictions

- **New AUX keyword on REORG of partitioned base for improved LOB handling**
  - Permit rows to flow between partitions
  - 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

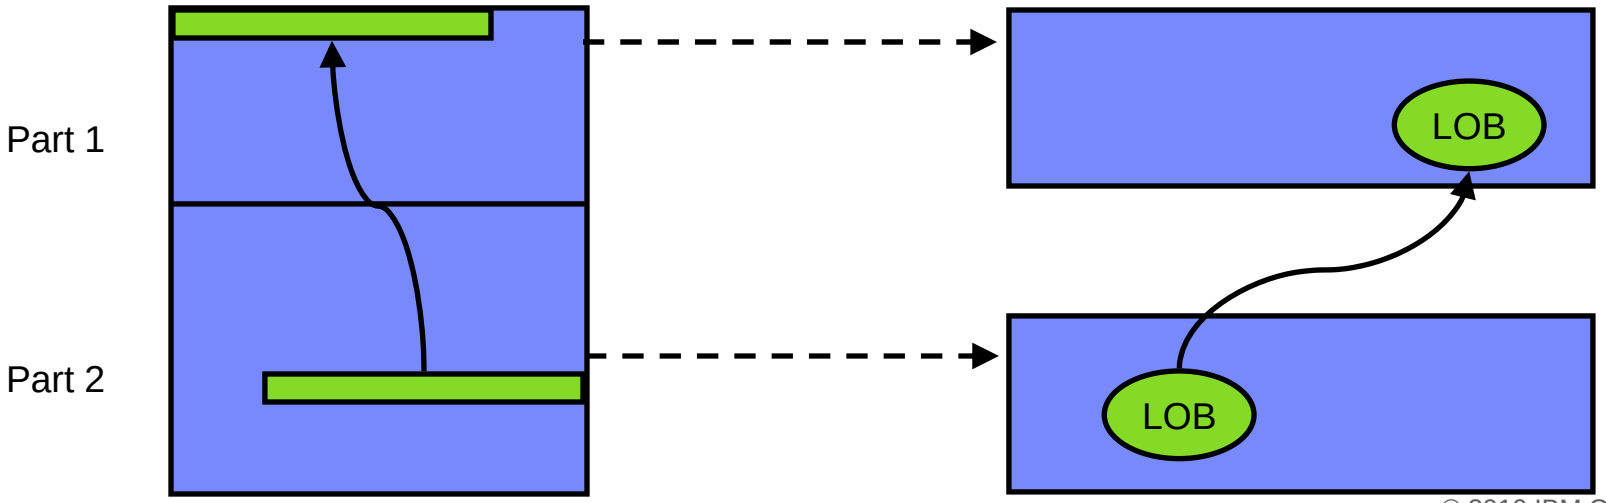

LOAD & UNLOAD

- 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

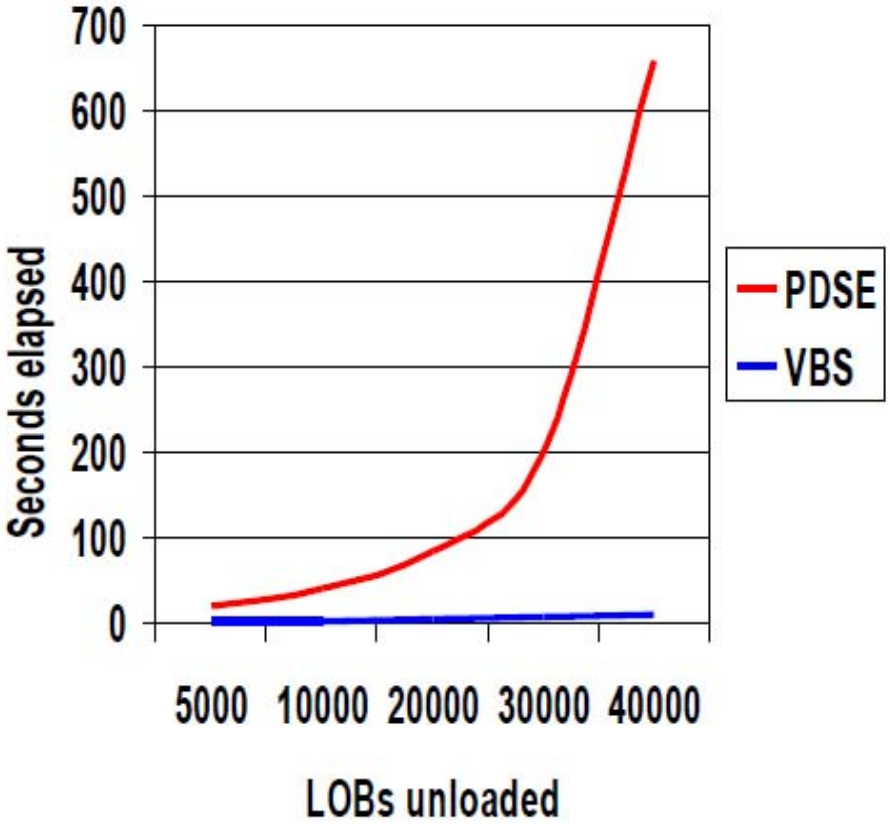

FlashCopy

- Dataset-level Flashcopy support
  - COPY, RECOVER, REORG, LOAD, REBUILD INDEX, REORG INDEX
  - New zparms & utility parms to govern
  - Significant CPU & elapsed time reduction for large page sets
  - Create transaction-consistent image copies from SHRLEVEL CHANGE

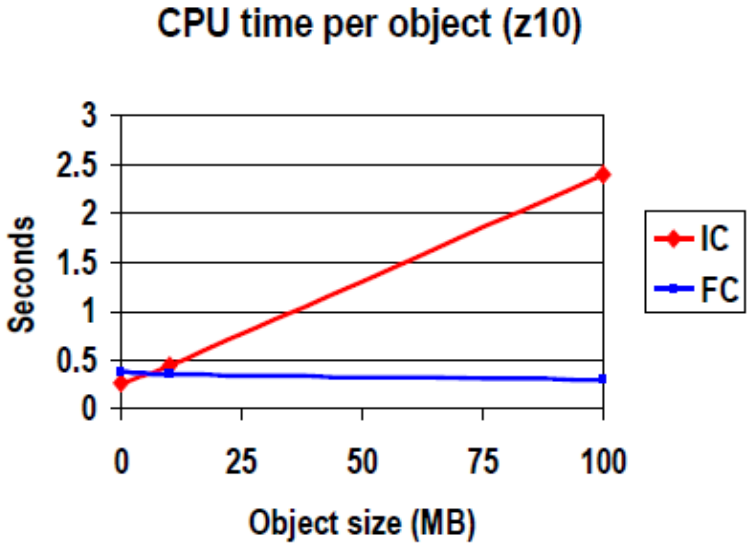

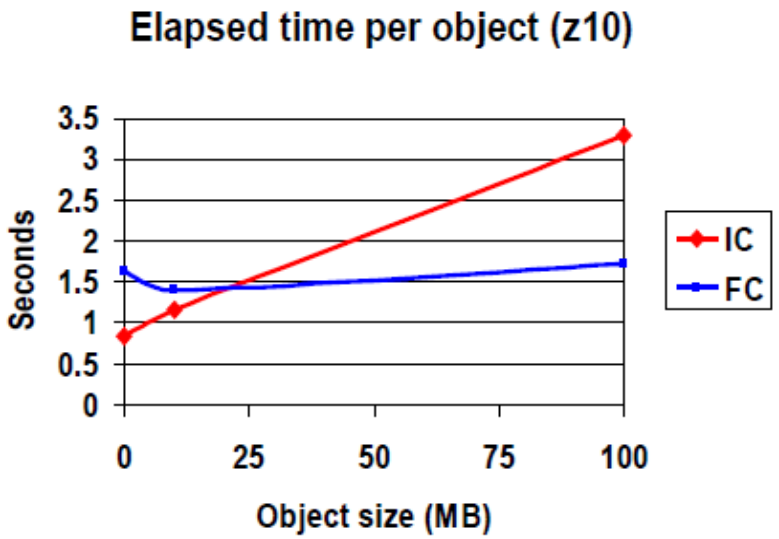

COPY

- Improved dataset management & performance
  - CHANGELIMIT will not allocate copy dataset unless copy taken
    - &ICTYPE now reflects actual image copy type
  - CHANGELIMIT to use RTS to determine full vs. incremental copy
  - Incremental copy will not allocate copy dataset unless pages changed
  - Delivery in maintenance stream post-GA

Recovery

- New BACKOUT YES option for point in time recovery
  - True rollback, not run of generated SQL undo statements
  - Requires COPY YES for indexes
- 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)

Stats

- RUNSTATS PROFILE support for simplification
- RUNSTATS on views
- Autonomic features through new stored procedures & catalog tables
- All catalog statistics columns made updatable
- RUNSTATS SHRLEVEL REFERENCE updates RTS
  - TOTALROWS & TOTALENTRIES columns
- zIIP-enablement for RUNSTATS
- Auto sampling rates & page sampling instead of row sampling
  - Significant CPU & ET savings
  - TABLESAMPLE SYSTEM AUTO

CHECK

- CHECK utilities will no longer set CHKP/ACHKP
- CHECK SHRLEVEL CHANGE default changed to fail if Flashcopy not available
  - Zparm to govern
- CHECK DATA enhanced for XML support
  - Document validation
  - Schema validation
- Automated exception table processing for XML documents

Other

- Removed UTSERIAL lock for greater utility concurrency
- LISTDEF & TEMPLATE enhancements
  - LISTDEF support for CHECK DATA
  - LISTDEF support for multiple part ranges on REORG
  - LISTDEF support for DEFINED YES|NO|ALL
    - Improved utility performance since unnecessary to build & then discard structures for undefined objects
    - Default changed to DEFINED YES
    - Empty lists result in RC4 post-GA
- REPORT RECOVERY support for SLBs
- DSNACCOX enhancements
  - Support hashed pagesets
  - New RTS columns for SSD, other

Summary

- Eliminate outages
- Improve performance
- Reduce resource consumption
- Reduce complexity & improve automation
- Day 1 utility support for core DB2 10 function
- Continued delivery of performance improvements & features of real value

Very recent enhancements

- PK88970/PK88972/PK88974 (V8 -)
  - Improved LOAD/UNLOAD processing with NUMRECS parameter
  - Replaces SORTKEYS at table space level with NUMRECS at table level
  - Simpler, eliminates risk of LOAD failure for load of multiple tables with skewed data distribution
- PM12286 (V9 -)
  - Avoid FRVs for LOAD/UNLOAD of zero length LOBs
  - Performance
- PM08223 (V8 -)
  - Switch SYSUTILX to MAXROWS 1 to avoid contention on SYSUTILX pages
    - Contention on pages in SYSUTILX caused by release of UTSERIAL lock prior to commit
    - Requires del/redef/reinit of SYSUTILX to gain full benefit
- PK96023 (V9 -)
  - Use LBI on UNLOAD to USS pipe
    - 60% CPU reduction, 50% ET reduction

Very recent enhancements

- PM26762 (V9 -)
  - Improved FlashCopy support for PPRC DASD mirroring & BACKUP SYSTEM contention
  - New zparms
    - FLASHCOPY_PPRC
      - Blank – no preserve mirror command passed to DFSMSdss (default for V9)
      - NONE – DFSMSdss will not preserve mirror
      - PREFERRED – allows duplex pending state if necessary
      - REQUIRED – prevents duplex pending (default for V10)
    - REC_FASTREPLICATION
      - V10 only
      - Prevent contention between FlashCopy recovery and BACKUP SYSTEM
      - NONE – FlashCopy will not be used for recover restore from FlashCopy
      - PREFERRED – use FlashCopy for recover restore if available (default)
      - REQUIRED – requires use of FlashCopy for restore from FlashCopy. RECOVER will fail if FlashCopy cannot be used
- PM08585 (V9 -)
  - REPAIR SET ... RBDPEND|PSRBDPEND
    - May be useful for improving heavy insert performance by skipping updates to non-unique indexes

Very recent enhancements

- PM27940 (V9 -)
  - Improve REPAIR VERSIONS to harvest version 0 information from any partition and preserve it for use by any partition
    - Useful when DSN1COPYing and not all parts have been updated since a table was first altered at the source
- 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
- PM23786 (V9 -)
  - 40% elapsed time improvement for COPY of small datasets to tape due to improved tape mark handling
- 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

Very recent enhancements

- PM27099 (V10)
  - LISTDEF empty lists change from RC8 to RC4
  - Particularly important with the LISTDEF change in 10 to default to DEFINED YES
- PM25652 (V10)
  - DSNACCOX enhanced to recommend REORG of hashed pagesets based on overflow index ratio vs. total rows.
- PM25648 (V10)
  - Improved REORG concurrency on online schema materialization
- PM30991 (V10)
  - Prohibit RECOVER BACKOUT on table space or index after mass delete on segmented or UTS
- PM24237 (V9 -)
  - MODIFY to validate existence of SLBs when deciding whether to set COPY-pending

Very recent enhancements

- PM25525 (V9 -)
  - New PARALLEL keyword on REORG with LISTDEF PARTLEVEL
  - Default is to process parts in parallel
- PM31243 (V10)
  - REORG FORCE improvement to better mimic –CANCEL THREAD
- PM29037 (V10)
  - REORG SHRLEVEL CHANGE to materialize changes in inline LOB length

Questions?