

IBM System z Technology Summit



How to Streamline Your DB2 for z/OS Utility Processing

Haakon Roberts

DB2 Development

March 2011



Disclaimer

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY.

WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED “AS IS”, WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

IN ADDITION, THIS INFORMATION IS BASED ON IBM’S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE.

IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION.

NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, OR SHALL HAVE THE EFFECT OF:

- **CREATING ANY WARRANTY OR REPRESENTATION FROM IBM (OR ITS AFFILIATES OR ITS OR THEIR SUPPLIERS AND/OR LICENSORS); OR**
- **ALTERING THE TERMS AND CONDITIONS OF THE APPLICABLE LICENSE AGREEMENT GOVERNING THE USE OF IBM SOFTWARE.**

Agenda

- **Trends in Database/Utility Management**
- **IBM's Investment in Utility Management**
- **Utilities Discussed in Detail**
- **Summary**

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

Performance & TCO

Automation & Standardization

Continuity & Resiliency

▪ 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

The case for tools & intelligent utility invocation

- **Focus on utility avoidance in addition to resource, CPU & elapsed time reduction**
- **New parameters introduced via APARs & releases**
- **Exploitation of new features critical for improved availability & efficiency**

- **Intelligent defaults, but the determination of what to run, when and with what parameters lies with you – or a tool**

- **Automation Tool, Recovery Expert & Utility Enhancement Tool provide automation, simplification & control**

Monitoring Utility Syntax

- **DB2 Utilities Enhancement Tool provides new Utility Syntax Monitor**
 - Can establish and enforce company-wide utility syntax practices
 - IT staff can control who executes which IBM DB2 utilities with what parameters on which objects
 - Can fail utility if rules are violated
 - Supports DB2 V8, DB2 9 and DB2 10 Utility Syntax
 - ADD parameters that are not present in the utility syntax
 - REMOVE parameters that are present and should not be
 - SUBSTITUTE given parameters with different parameters
 - FAIL the utility based on object name, or user ID
- **Each action is logged or JOURNALED in UET's tables for future reference**
 - Audit who is doing what
 - See what syntax was originally specified
 - See what the original syntax was changed to
- **Delivered via PTF**
 - UK60173 for all versions of DB2 for z/OS

SORTNUM Elimination

- **CHECK INDEX, REBUILD INDEX, REORG, RUNSTATS**
- **PK45916 (V8) & PK41899 (V9)**
- **Better performance, more robust, simpler**
- **SORTNUM no longer required**
 - Difficult to estimate: failure if too low, excessive sort work allocation if too high
- **New zparms UTSORTAL & IGNSORTN (online changeable)**
 - UTSORTAL YES|NO
 - Use RTS data to estimate number of rows to sort
 - DB2 will dynamically allocate sort work datasets
 - If SORTWK DD cards not hard coded
 - IGNSORTN YES|NO
 - Override utility job setting of SORTNUM
- **Recommendation**
 - Turn on UTSORTAL, test it, then consider turning on IGNSORTN

DSNU3340I 168 08:13:52.66 DSNUGLSR - UTILITY PERFORMS DYNAMIC ALLOCATION OF SORT DISK SPACE

Other recent enhancements

- **Improved LOAD/UNLOAD processing with NUMRECS parameter**
 - PK88970/PK88972/PK88974 (V8 & V9)
 - 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
- **Support REORG of multiple part ranges**
 - PK87762 & PM13259 (V9)
 - E.g. REORG PART 1,45:71,500:503,4010
 - More efficient, improved availability, exploit parallelism

Other recent enhancements

- **LOAD and UNLOAD to/from virtual file**
 - USS named pipe support with templates
 - PK70269 (V8 & V9)
 - PK96023 (V8 & V9)
 - LBI on UNLOAD – 60% CPU reduction, 50% ET reduction
- **LOAD COPYDICTIONARY**
 - PK63324/PK63325 (V9)
 - REORG avoidance – prime empty partitions with compression dictionary
- **Allow CHECK SHRLEVEL CHANGE to use FASTREPLICATION(REQUIRED)**
 - PM19034 (V9)
 - Fail CHECK utility rather than incur application outage
- **Faster image copy to tape**
 - PM23786 (V9)
 - 40% elapsed time improvement on copy of small datasets to tape due to improved tape mark handling

Performance – utility CPU consumption

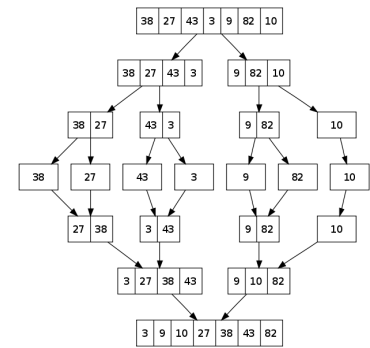
- **Focus on real CPU reduction & zIIP exploitation**
- **DB2 utilities have been zIIP-enabled since 2006**
- **Real CPU cost reduction in V9**
 - 10-20% for COPY & RECOVER
 - 5-30% for LOAD, REORG, REBUILD INDEX
 - 20-60% for CHECK INDEX
 - 35% for LOAD partition
 - 30-40% for RUNSTATS INDEX
 - 40-50% for REORG INDEX
 - 70% for LOAD REPLACE partition with dummy input
- **Flashcopy exploitation in DB2 10 dramatically reduces CPU consumption for COPY & reduces CPU for RECOVER & inline copies**
- **More zIIP offload in DB2 10 with RUNSTATS**

Performance – zIIP exploitation for sort processing

- **In spite of CPU reduction in V9, there is continued focus on CPU consumption for utilities**
- **Sort can consume ~60% of total utility CPU time**
- **DB2 in concert with DFSORT provides zIIP offload of DB2 utility memory-object fixed-length record sort processing**
- **Requirements:**
 - DB2 APAR PK85889 (V8 or V9)
 - DFSORT APAR PK85856
 - z/OS 1.10
- **PTFs can be applied independently of each other**
- **Exploitation is automatic**

DB2 Sort for z/OS v1.1

- **Announced Aug 10th, GA Sep 24th**
- **Provides high speed utility sort processing for DB2 for z/OS data**
- **Provides CPU & elapsed time reduction**
 - Up to 30% reduction in elapsed time
 - Up to 50% reduction in CPU consumption
- **zIIP-enabled for further CPU cost reduction**
- **Improved resilience, resource management & data availability**



**Customer results may vary. Results based on analysis done at SVL I*

New solutions for DB2 9

- **LOAD/UNLOAD FORMAT INTERNAL**

- PM19584
- Unload and load data in true internal format
- 85% CPU & elapsed time reduction on UNLOAD
- 77% elapsed time, 56% CPU reduction on LOAD
- Supported by High Performance Unload

- **LOAD PRESORTED**

- PM19584 – delivery in V9 & V10 post-GA
- Avoid sort overhead
- Up to 25% CPU reduction, 33% ET reduction depending on no of indexes
- Works well with Utility Enhancement Tool PRESORT option

New solutions for DB2 9

- **REPAIR SET ... RBDPEND|PSRBDPEND**
 - PM08585 (V9)
 - May be useful for improving heavy insert performance by skipping updates to non-unique indexes

- **Allow LOAD REPLACE or RESUME of a partition even though NPI is in PSRBD**
 - PM27962 (V9)
 - Also 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

DB2 10 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
 - Automation Tool enhanced
- **REORG SHRLEVEL CHANGE for all cat/dir page sets**
- **REORG SHRLEVEL REFERENCE|CHANGE to remove REORP**


DB2 10 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

DB2 10 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



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




DB2 10: COPY


▪ 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 pagesets
- Create transaction-consistent image copies from SHRLEVEL CHANGE

CPU time per object (z10)



Elapsed time per object (z10)



DB2 10: RECOVER

- **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)

DB2 10: Stats

- **RUNSTATS PROFILE support for simplification**
- **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

DB2 10: 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**

DB2 10: 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
- **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**
- **Synergy with IM Tools**