

DB2 11 for z/OS Utilities Update

Andy Lai

DB2 Utilities Development

atlai@us.ibm.com



#SHAREorg



SHARE is an independent volunteer-run information technology association
that provides education, professional networking and industry influence.

Copyright (c) 2014 by SHARE Inc.  Except where otherwise noted, this work is licensed under
<http://creativecommons.org/licenses/by-nc-sa/3.0/>

Disclaimer

© Copyright IBM Corporation 2014. All rights reserved.

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.

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

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

Agenda

- Strategy
- Overview
- REORG
- Statistics
- Backup & Recovery
- UNLOAD & LOAD
- Compression Dictionaries
- General Enhancements
- Deprecation
- Recent News
- Questions & Feedback

Strategy

- Support core function
- Reduce CPU, ET & resource consumption
- Maximize availability
- Remove constraints & limitations
- Simplify data management

Overview

- **Better availability**
 - Outage reduction for online REORG through improved drain processing & SWITCH phase elapsed time reduction
 - Better control of SWITCH phase timing
 - Online rebalance of data across partitions
 - Faster LOAD processing
- **Better performance & reduced resource consumption**
 - Reduce/eliminate need for RUNSTATS through inline stats
 - Faster LOAD processing
 - Faster recovery from part-level inline image copies
 - Faster REORG processing
 - More utility parallelism & greater parallelism control
 - More zIIP exploitation, more MRU
 - Optimizer input to RUNSTATS

Overview

- Simplified data management & improved usability
 - Automated REORG mapping table management
 - Intelligent default settings for REORG
 - Better PBG management through deletion of unused partitions
 - Improved data set management with part-level inline image copies
 - Improved XML handling with Crossloader support
 - System cloning improvements
 - Lift point in time recovery restrictions
 - More information & greater transparency through –DISPLAY UTILITY improvements
 - Enhanced statistics profile support

Key performance numbers

- CPU improvement for utilities with EXCLUDE NULL KEYS
 - LOAD 12%
 - REORG 25%
 - REORG INDEX 88%
 - REBUILD INDEX 72%
 - CHECK INDEX 79%
 - RUNSTATS 90%
- Increased parallelism elapsed time improvement
 - REORG 21%
- Inline stats vs. separate RUNSTATS
 - 40% elapsed time reduction for inline histogram stats
 - 28% elapsed time & 19% CPU reduction with inline distribution stats

Key performance numbers


- Up to 81% zIIP-eligible CPU with RUNSTATS COLGROUP
- Up to 40% zIIP-eligible CPU in REORG & LOAD with inline distribution stats
- REORG SWITCH phase reduced by 90%
- Up to 71% elapsed time reduction for REORG of subset of partitions
 - SORTNPSI option retrofitted to V9 & V10
- RECOVER from part-level image copies reduced CPU by up to 50%, elapsed by up to 40%
- LOAD from single input dataset elapsed time reduced by up to 70%
- Crossloader support for FETCH CONTINUE for LOB & XML data
 - 28% CPU reduction

Improve performance of part-level REORG with NPSIs

- New option to defer shadow index build until all keys passed through sort
- New parm & zparm to govern
 - AUTO/YES/NO options
- Retrofit to DB2 9 & 10 in PM55051
- Result:
 - Customer test of REORG of 40% of partitions showed 55% ET reduction & 22% CPU increase
 - DB2 Sort gives additional ET reduction & cuts CPU to less than original starting point

SWITCH phase impact relief – reduced application impact

- Easier drain acquisition
- Prevent new claims on all target partitions whilst waiting for drains
 - Faster drain acquisition for part-level REORG
- New DRAIN_ALLPARTS option to momentarily drain all data parts
 - Eliminates claim-drain “deadlocks” for part-level REORG with NPSIs
- Restructure SWITCH phase processing for outage reduction
 - SWITCH phase ET reduction of 91% measured when reorging 20 parts



- ```

.-SWITCHTIME--NONE-----
>+-----+>X
| .-NEWMAXRO--NONE----. |
'-SWITCHTIME--+timestamp-----+-----+-----+
 '-labeled-duration-expression-' '-NEWMAXRO--integer-'

```

# Physically delete empty PBG partitions

- Ability for REORG to physically delete empty PBG partitions
- New zparm REORG\_DROP\_PBG\_PARTS
  - DISABLE – keep V10 behavior (default)
  - ENABLE – Delete empty PBG partitions on table space-level REORG
- Considerations:
  - Cannot be specified on REORG statement
  - If PBG created using Numparts or ALTER ADD partition used, REORG may prune to a lesser number of partitions
  - No PIT recovery to prior to a pruning REORG
    - No facility to resurrect deleted partitions

# Automated mapping table handling

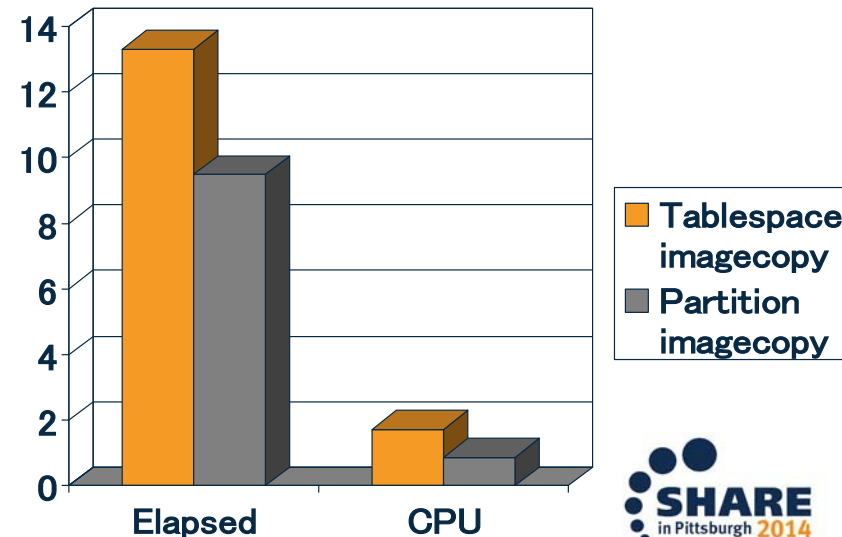
- Support mapping tables in PBG
  - Increases mapping index limit from 64Gb to 16Tb
  - Retrofitted to V9 in PM58177
- Mapping table DDL must change in 11 due to RBA/LRSN change
- Requirements to automate mapping tables
- So... New automated mapping tables in REORG
  - Automatically create new format mapping table if required
    - If mapping table specified & correct format then honor specification
    - Else if specified but incorrect format then create new in same db as original
    - Else if not specified and zparm DB specified then create in zparm DB
    - Else create in implicitly created DB
    - DROP at end of REORG or end of last REORG if multiple REORGs in job step
  - NFM requires new format mapping table, CM supports both old and new
  - No additional auth requirements necessary for creation of mapping tables

# REORG without sorting data

- Increasingly REORGs are performed for reasons other than to regain clustering of data, yet no ability to avoid cost of recluster
- Prior to DB2 11, REORG SHRLEVEL CHANGE did not support SORTDATA NO
- DB2 11 supports SORTDATA NO with SHRLEVEL CHANGE
- New RECLUSTER YES/NO option on SORTDATA NO
  - RECLUSTER NO – Do not unload data through clustering index **and do not sort data records in clustering order**

# Partition-level inline image copy

- Faster partition-level recovery from inline image copy
- Create partition-level inline image copies if using TEMPLATE with &PA or &PART
  - No new option or keyword on REORG
  - PM93611:
    - Support substring notation with &PA as long as substring ensures uniqueness
    - Support writing to tape as long as STACK YES not specified
- RECOVER of single partition of a 20 partition table space
  - ET reduced by 28%
  - CPU reduced by 49%



# Improved REORG LISTDEF processing

- PARALLEL YES/NO option introduced in APAR in V9
  - NO – Prevent REORG from processing multiple partitions in single REORG when input is partlevel LISTDEF
  - Zparm REORG\_LIST\_PROCESSING at zparm level
- Need compromise option for customers who want to take advantage of REORG parallelism but cannot afford to shadow many partitions at a time
- New option LISTPARTS n to limit # of partitions to be processed in a single REORG if input is a part-level LISTDEF
- In DB2 11, PARALLEL YES/NO is superseded by LISTPARTS
  - PARALLEL YES/NO IS deprecated but still supported in 11



# REBALANCE enhancements

- Improved availability & failure prevention
- Support REORG SHRLEVEL CHANGE REBALANCE
  - Complements online ALTER of partition limit keys
- Improve resiliency with enhanced distribution algorithm & improved handling of empty partitions
- Build compression dictionary for all partitions
  - Previously, partitions that were empty at the start of REORG would not have a dictionary built, requiring a subsequent REORG to gain compression
- New SORTCLUSTER option to sort data in clustering as well as partitioning order to avoid AREO\*
  - Occurred when partitioning key not a superset of clustering key
  - SORTCLUSTER YES – sort in partitioning and clustering order, avoid AREO\*
  - No specification – keep existing behavior, set AREO\*
  - SORTCLUSTER NO – keep existing behavior, but do not set AREO\*

# REORG of LOB data

- Support REORG of LOB data even though aux index is unavailable
  - Problem in V10 if LOB tablespace is REORP and index is RBDP
    - LOBs can't be reorged and index can't be rebuilt
- REORG SHRLEVEL NONE for LOBs changed to RC8 from 11 CM onwards
  - Not supported in 10 NFM, but returns RC0 with MSGDSNU126I

# Improved REORG serviceability

- Need ability to use online REORG even when SYSLGRNX cannot be relied upon
- Support LOGRANGES NO option for REORG SHRLEVEL CHANGE

# REORG change of defaults to match best practices

- Change default options:
  - DRAIN WRITERS to DRAIN ALL
  - DISCARD to DISCARD NOPAD YES
  - UNLOAD EXTERNAL to UNLOAD EXTERNAL NOPAD YES

# Statistics Enhancements

- More zIIP offload for RUNSTATS distribution statistics
  - Up to 80% zIIP-eligible
- zIIP offload for inline statistics
  - Additional 30% offload to zIIP
- Enhance inline statistics for RUNSTATS avoidance
  - Inline statistics collection on NPSIs during REORG with SORTNPSI
  - Inline histogram statistics
  - Inline DSTATS
- New RUNSTATS RESET option to reset existing statistics
- Improved PROFILE usability for LISTDEF processing
  - Gather default statistics if no profile exists for table

# Statistics Enhancements

- Optimizer determination of missing statistics
  - Optimizer identifies missing statistics & writes information to new catalog table DSN\_STAT\_FEEDBACK
  - OQWT modifies statistics profile
  - Automation Tool detects profile change & builds RUNSTATS job
  - DSNACCOX similarly enhanced to recommend RUNSTATS
- -ACCESS DATABASE ... MODE(STATS) option to externalize RTS statistics
- RTS overhead reduction

# Backup & Recovery Enhancements



- Faster catalog/directory recovery
  - Enhanced SYSLGRNX recording
- New VCAT name translation for RESTORE SYSTEM for system cloning
  - Support logapply when RESTORE SYSTEM used for cloning purposes
- Improved recoverability with COPY-REORG concurrency
  - Permit COPY to run concurrent with long-running REORGs
- Avoid allocating empty image copy datasets for incremental or CHANGELIMIT copies



# Backup & Recovery Enhancements



- Lifted many restrictions on point-in-time recovery prior to materializing REORG
  - PIT recovery restrictions lifted for
    - LOB table spaces
    - XML table spaces
    - PBR table spaces
    - Including when immediate alters have occurred since materializing REORG
  - PIT recovery restrictions still in place
    - Table space conversion
    - PBG table spaces
    - PBG partition pruning
    - Online DROP COLUMN
- Faster index recovery with FLA support for index log records
  - APAR PI07694 for V9 and above



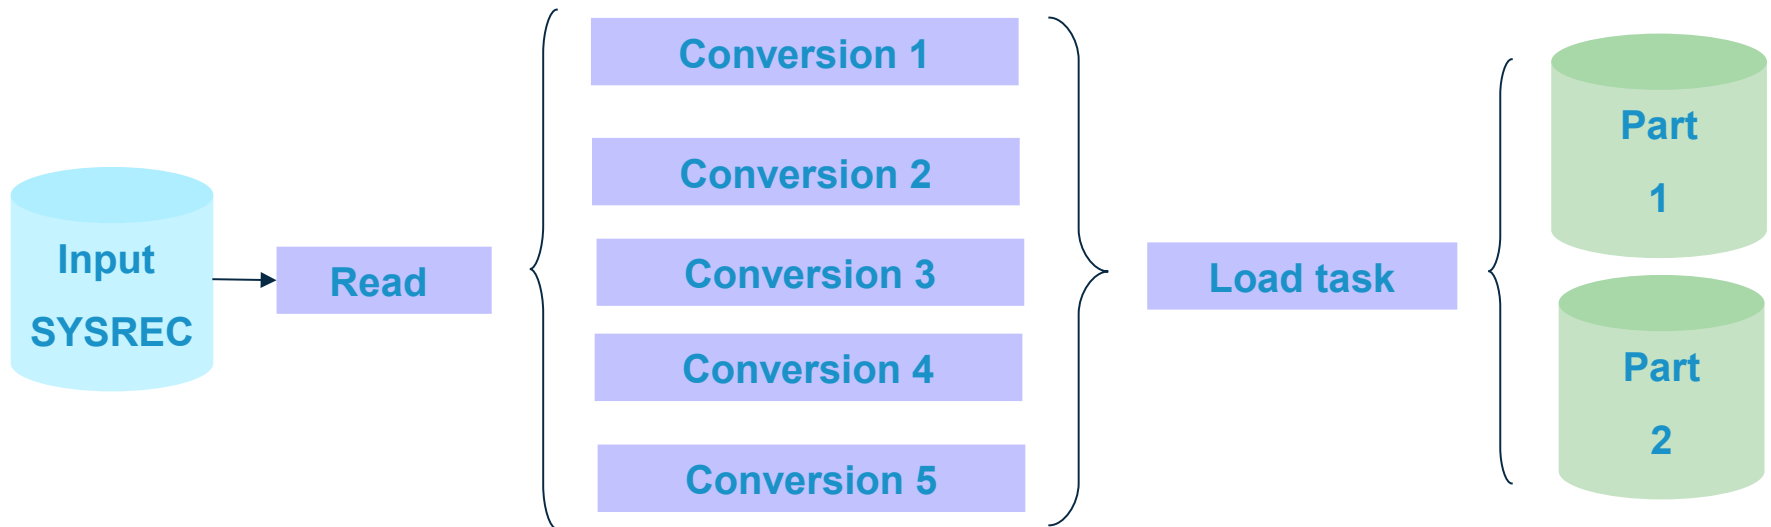


# LOAD & UNLOAD Enhancements

- Crossloader support for XML data
- Exploit FETCH CONTINUE for processing large LOBs & XML data in Crossloader
  - Reduce vstor requirement
  - Avoid DSNU1178i errors
  - 28% CPU reduction
    - Load of 1Mb LOBs
- zIIP offload for LOAD REPLACE PART clearing of NPSIs
  - 100% offload to zIIP for LOAD REPLACE with dummy input

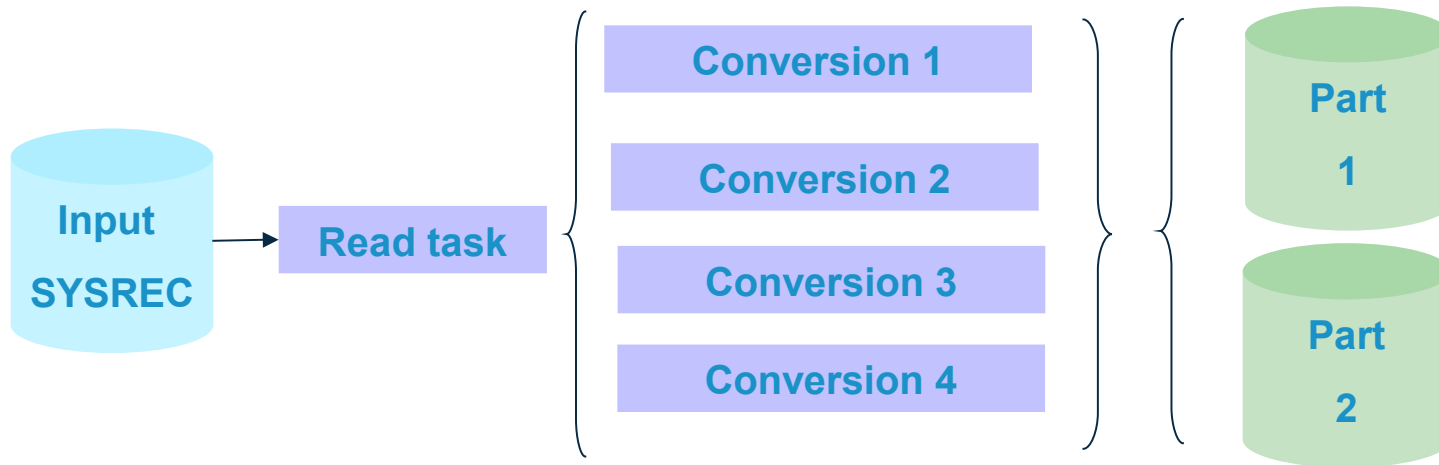
# LOAD & UNLOAD Enhancements

- **LOAD SHRLEVEL NONE PARALLEL** with single input data set
  - Parallel data conversion
  - Not supported for PBGs
  - 50% ET reduction possible on single SYSREC load



# LOAD & UNLOAD Enhancements

- **LOAD SHRLEVEL CHANGE PARALLEL**
  - Supports non-partitioned as well as partitioned
  - Single input dataset
  - Not supported for PBGs
  - >80% ET reduction



# Compression Dictionaries

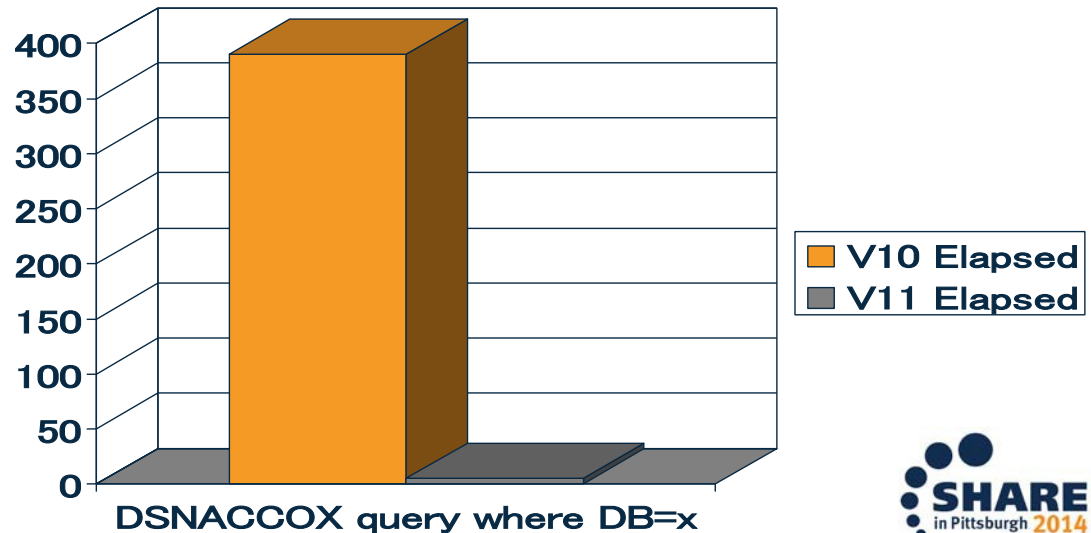
- Avoid decompression failures for IFI 306 readers when new compression dictionary built by REORG/LOAD
- Old compression dictionary stored on log
- New SYSCOPY record written pointing to old compression dictionary for CDC tables
- IFI 306 read automatically retrieves old compression dictionary if necessary
- Avoid need for replication target refresh when dictionary changes

# General Enhancements

- Greater parallelism for faster utilities
  - 11% elapsed time reduction measured for REORG, LOAD, REBUILD INDEX
- PARALLEL option for parallelism control for LOAD, REORG, REBUILD INDEX, UNLOAD, CHECK INDEX
- -DISPLAY UTILITY enhancements
  - Remove serialization between -DIS UTIL and -TERM UTIL
  - Jobname, start timestamp
  - Late addition: SWITCHTIME and NEWMAXRO
- Utility impact reduction on bufferpools
  - Extend MRU for UNLOAD, REORG TABLESPACE, RUNSTATS TABLESPACE, RUNSTATS INDEX, REBUILD INDEX, CHECK INDEX, CHECK DATA
- Improved dataset cleanup in utility stored procedures
  - Previously, datasets remained allocated on utility failure, preventing cleanup

# General Enhancements

- Improved TEMPLATE support for large / EF datasets and local time values
  - DSNTYPE LARGE, EXTREQ, EXTPREF
  - New EATTR option on TEMPLATE to request extended attributes
  - New TIME LOCAL|UTC option
- Enforce NUMTCB=1 for stored procedures
- DSNACCOX performance



# Deprecation

- REORG SHRLEVEL NONE for LOBs changed to RC8 from 11 CM onwards
  - Not supported in 10 NFM, but returns RC0 with MSGDSNU126I
- Still supported in 11, but no longer documented:
  - REORG
    - PARALLEL YES|NO
      - Superseded by LISTPARTS
    - INDREFLIMIT
    - OFFPOSLIMIT
    - LEAFDISTLIMIT
    - UNLOAD ONLY
    - UNLOAD PAUSE
    - UNLOAD EXTERNAL
  - COPY
    - CHANGELIMIT

# Recent News

- Part-level REORG NPSI insert performance improvement
  - PM87403 (V9)
  - 100m row table, 6 indexes
  - LOAD RESUME – 66% CPU reduction, 30% ET reduction
  - REORG PART 9 – 45% CPU reduction, 26% ET reduction
- Fast log apply for faster index recovery
  - PI07694 (V9)
- Retrofit REORG SWITCH phase performance to V10
  - PI09303 (V10)
  - Cut SWITCH phase duration by 90%
- Retrofit REORG REBALANCE SHRLEVEL CHANGE to V10
  - PI11839 (V10)
- LOAD REPLACE SHRLEVEL REFERENCE, LOAD RESUME SHRLEVEL REFERENCE, LOAD prevalidation
  - Delivered in Utilities Enhancement Tool / Utilities Solution Pack
  - PI04864



# Summary

- Unparalleled investment in utilities
- IBM Utilities Suite is essential for exploitation of major DB2 enhancements
- Support of core DB2 function from day 1 of GA
- Expect continued delivery of enhancements on release boundary, and in maintenance stream when prudent
- Continued focus on:
  - Elimination of application impact from utilities
  - Elapsed time & CPU consumption reduction
  - Resource consumption reduction
  - Reduction in complexity & automation improvements
  - Solutions through DB2, Utilities & Tools

Thank  
You