



| IBM Software Group

DB2 for z/OS Utilities Update

Haakon Roberts
DB2 Development
IBM Silicon Valley Lab

Aug 2008



Agenda

- Recent maintenance stream improvements
- DB2 9 for z/OS Utilities enhancements
 - ▶ Support for non-utility enhancements
 - ▶ Availability enhancements
 - ▶ Performance enhancements
 - ▶ Backup and recovery enhancements
 - ▶ Other enhancements



IBM Software Group

Recent maintenance stream improvements



SORTNUM elimination in V8/V9

- PK45916 (V8) & PK41899 (V9)
- SORTNUM no longer required
 - ▶ Correct value hard to determine, resulting in utility failure if too low or excessive sort work allocation if too high
- New zparms UTSORTAL & IGNSORTN (online changeable)
 - ▶ UTSORTAL
 - Yes
 - Use RTS data to estimate number of rows to sort
 - No (or if RTS data not available)
 - Use old method
 - DSNU3343I csect-name REAL-TIME STATISTICS INFORMATION MISSING FOR obj-type obj-qual. obj-name partition-information
 - ▶ IGNSORTN
 - Yes
 - Override utility job setting of SORTNUM
 - No
 - Utility job setting takes precedence
 - Utilities:
 - ▶ CHECK INDEX, REBUILD INDEX, REORG TABLESPACE, RUNSTATS with COLGROUP

Other recent maintenance changes

- PK56334 (V8 & V9)
 - ▶ Allow REORG to reset LPL & WEPR
- PK59399 (V9)
 - ▶ Allow DB2 to use DFSORT parameters from PARMLIB
 - ▶ Requires z/OS r10
 - ▶ Install options can be changed without reassembling ICEMAC module
- PK60612 (V9)
 - ▶ Allow unload from a non-segmented image copy when the table is segmented
 - ▶ Useful for retention/archiving since unable to create simple table spaces in V9
- PK60956 (V8 & V9)
 - ▶ SORTBLD performance improvement for indexes with small SECQTY
 - ▶ SORTBLD elapsed up to 20x improvement!!!
 - ▶ Recommendation: Set appropriate PRIQTY/SECQTY to avoid extend processing
- PK61759 (V8 & V9)
 - ▶ LOAD & REORG performance improvement
 - 10% CPU & ET improvement in reload phase
 - ▶ Sort processing performance improvement
 - 10% CPU reduction in sort phase

Other recent maintenance changes

- PK51853 (V8 & V9)
 - ▶ Allow REORG or LOAD of >254 compressed parts
 - ▶ New zparm governs LOAD & REORG in V8, only LOAD in V9
- PK31632 & PK56003 (V8 & V9)
 - ▶ REORG REBALANCE to use row balancing instead of page balancing algorithm
- PK63324 & PK63325 (V9)
 - ▶ LOAD COPYDICTIONARY
 - Allow priming of a partition with a compression dictionary
- PK65909 (V8 & V9)
 - ▶ SORTNUM elimination improvement to avoid multi-volume sort work datasets
 - ▶ DFSORT can only use the first volume
 - ▶ Still open at this time
- PK47083 (V8 & V9)
 - ▶ Invalidate cached dynamic statements on completion of LOAD



| IBM Software Group

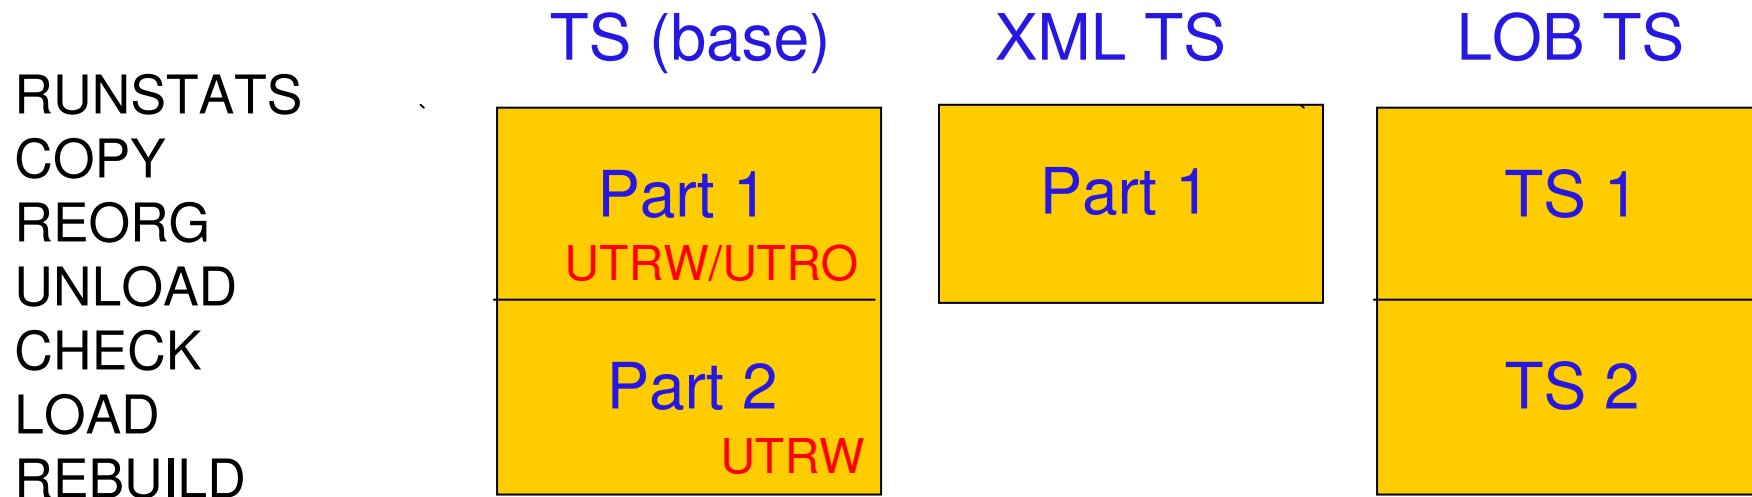
Support for non-utility enhancements



DB2 9 Utilities

- Support for all new functions in DB2 Version 9 for z/OS product
 - ▶ Universal Table Spaces (UTS)
 - Partition By Growth (PBG)
 - Partition By Range (PBR)
 - ▶ XML table spaces (PBG or PBR)
 - ▶ Not logged tables/table spaces
 - ▶ Clone tables
 - ▶ Index on expression
 - ▶ New data types (BIGINT, VARBINARY, DECFLOAT XML)

Utilities and Partition By Growth Table Spaces concurrency



Utilities and PBG objects

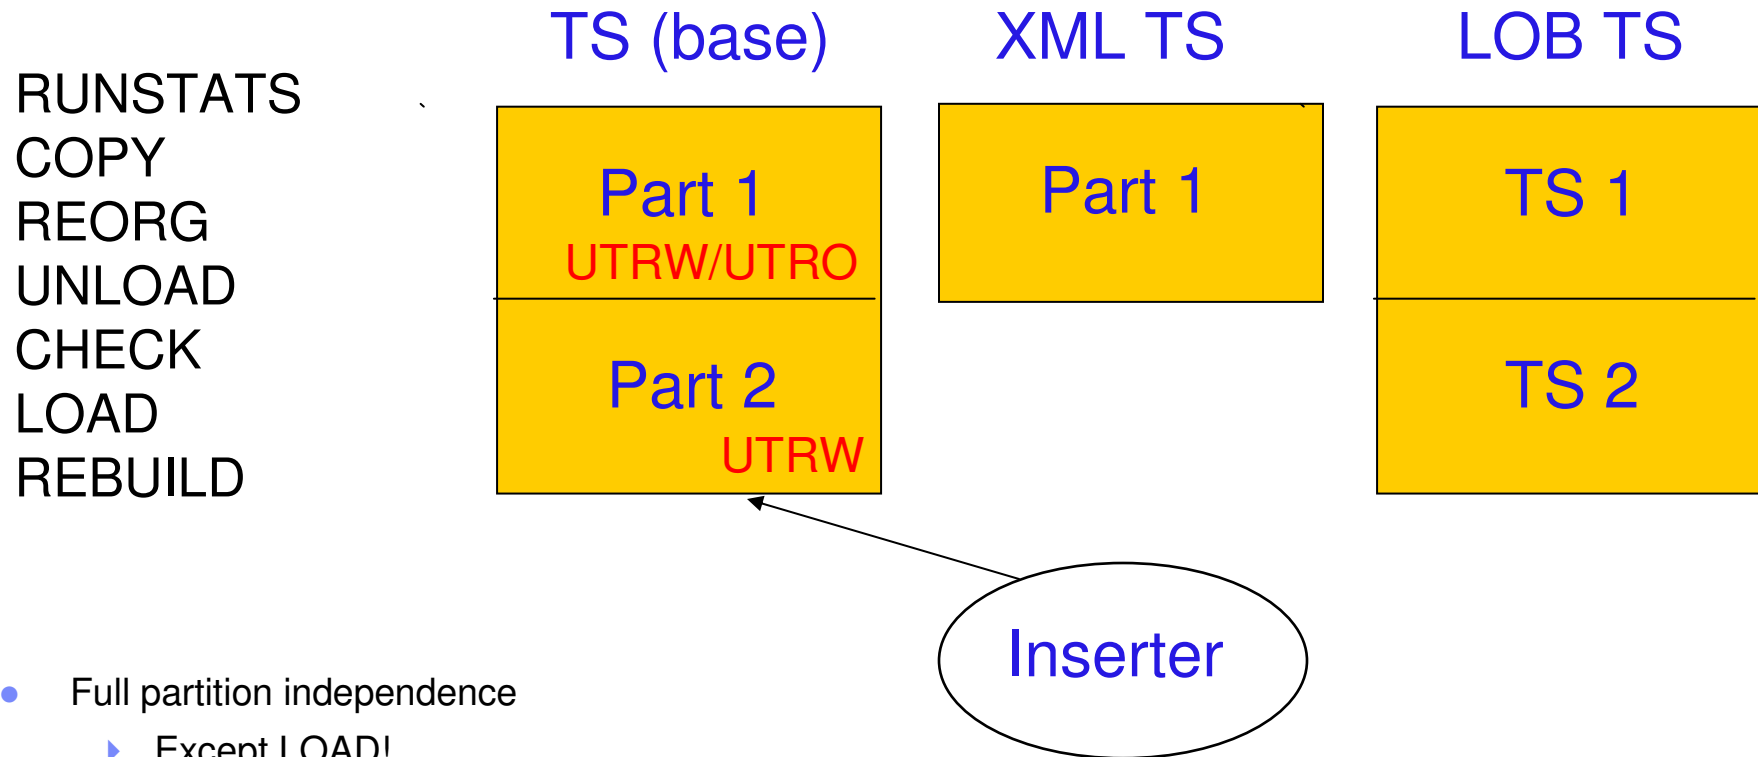
- Table space or part level supported

RUNSTATS TABLESPACE TS SHRLEVEL CHANGE

COPY TABLESPACE TS DSNUM 1 SHRLEVEL REFERENCE



Utilities and Partition By Growth Table Spaces concurrency




- Full partition independence
 - ▶ Except LOAD!
- Utility & SQL concurrency
 - ▶ New part may be added by DML even though cannot claim last part
 - ▶ But: Insert fails if unable to claim earlier part during space search



Utilities and PBGs

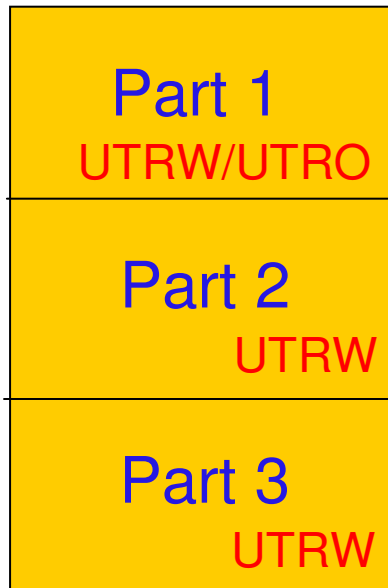
PBG



- REORG
 - ▶ May create new parts
 - ▶ If hit MAXPARTITIONS then REORG fails
 - ▶ Dictionary copied from previous part
 - ▶ REORG of single part
 - No new part creation
 - Rows must fit back into part
 - ▶ REORG of part range
 - Data can flow from one part to another within range
 - If LOB column exists then rows will not move between parts
 - ▶ SHRLEVEL CHANGE
 - New part added to both base and shadow
 - ▶ No parallelism
 - ▶ No shrinking
 - ▶ No REBALANCE

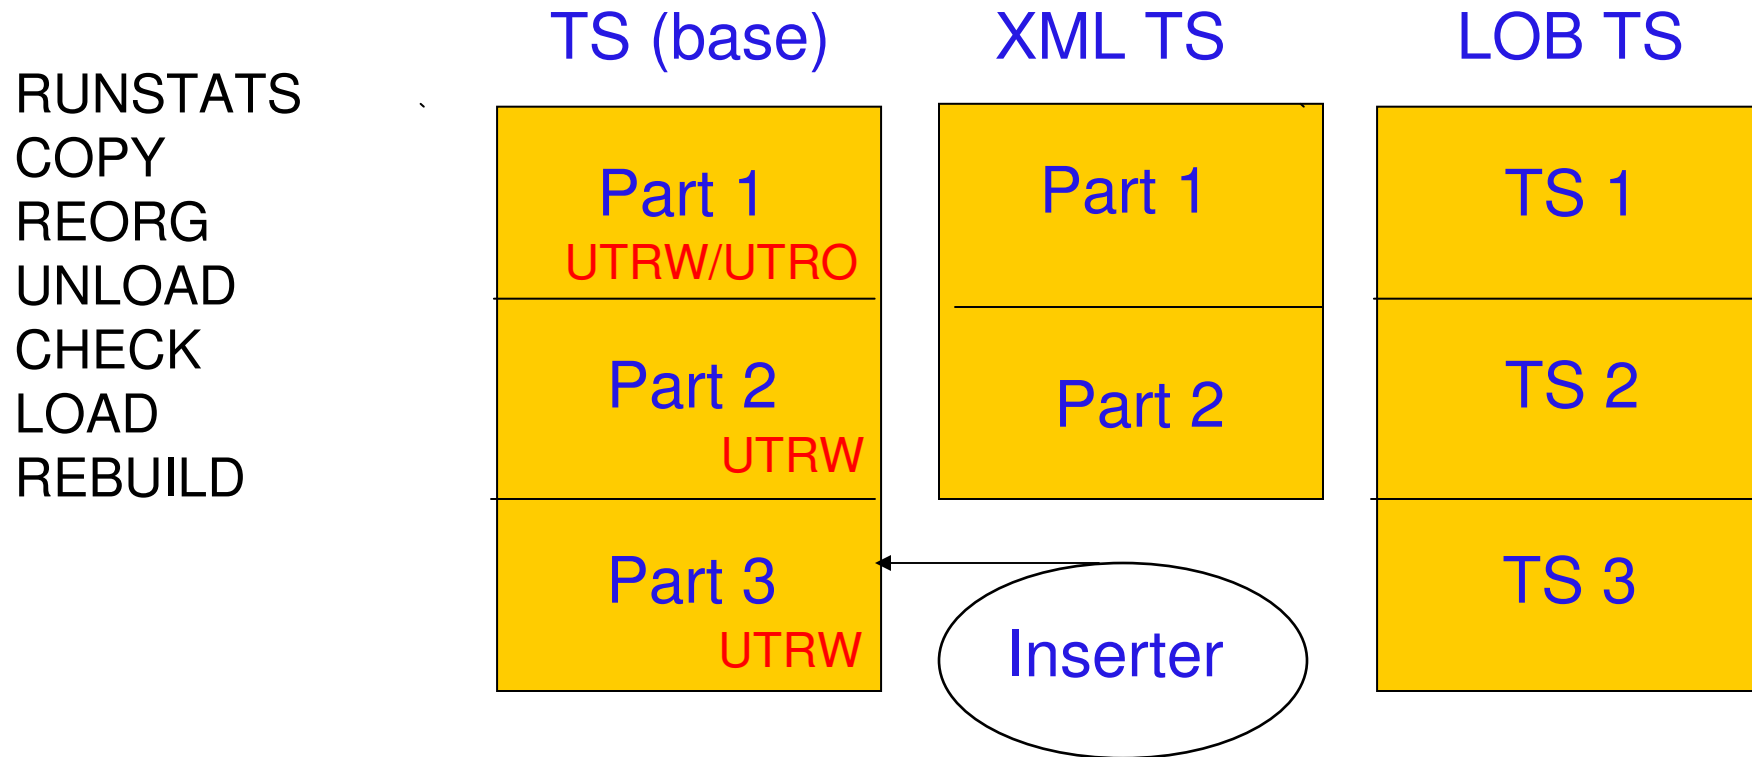
Utilities and PBGs

PBG



- LOAD
 - ▶ No LOAD at part level
 - ▶ No parallelism
 - ▶ Can grow parts
 - ▶ Dictionary copied from previous part to new part
 - ▶ Excess parts remain empty
- COPY
 - ▶ Part or tablespace level
 - ▶ Will copy new parts if SHRLEVEL CHANGE at tablespace level
- REBUILD INDEX SHRLEVEL CHANGE
 - ▶ Index entries for new part included in log apply phase

Utilities and Partition By Growth Table Spaces concurrency



PBG base and XML partitions grow independently
 PBG base and LOB table spaces grow together





| IBM Software Group

Online Utility Enhancements



More online utilities

Rebuild Index SHRLEVEL CHANGE

- ▶ Great for building new non-unique indexes or when index is in RBDP
- ▶ Use REORG INDEX SHRLEVEL CHANGE to move indexes to different volumes, not REBUILD INDEX SHRLEVEL CHANGE if availability is important (i.e. no shadows... the index is built in place!)
- ▶ Table space must be in LOGGED state (new log phase)

More online utilities (cont...)

REORG enhancements

- Partition parallelism (UNLOAD/RELOAD) in a single utility statement (performance)
- Parallel log apply (performance)
- Elimination of BUILD2 phase (availability)
 - NPIs also shadowed
 - NPIs implicitly reorged
 - REORG SHRLEVEL REFERENCE PART now has a log phase
 - Improved availability
 - ◇ But concurrent REORGs at part level for same tablespace not permitted

More online utilities (cont...)

- CHECK DATA SHRLEVEL CHANGE
- CHECK LOB SHRLEVEL CHANGE
- REPAIR LOCATE ... SHRLEVEL CHANGE
- REORG LOB now supports SHRLEVEL REFERENCE
- Clones effectively provide LOAD REPLACE SHRLEVEL CHANGE
 - ▶ UTS only
- UNLOAD with ISO(CS) supports skipping rows that are locked for transaction updates



IBM Software Group

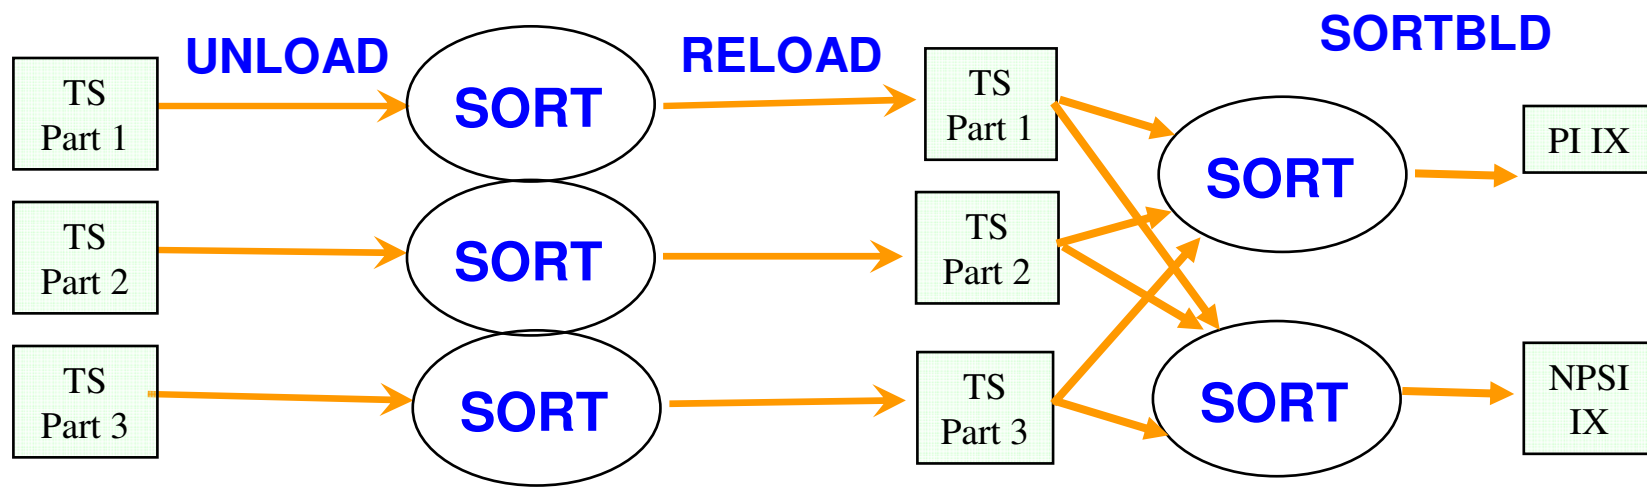
Performance Enhancements



REORG Partition Parallelism

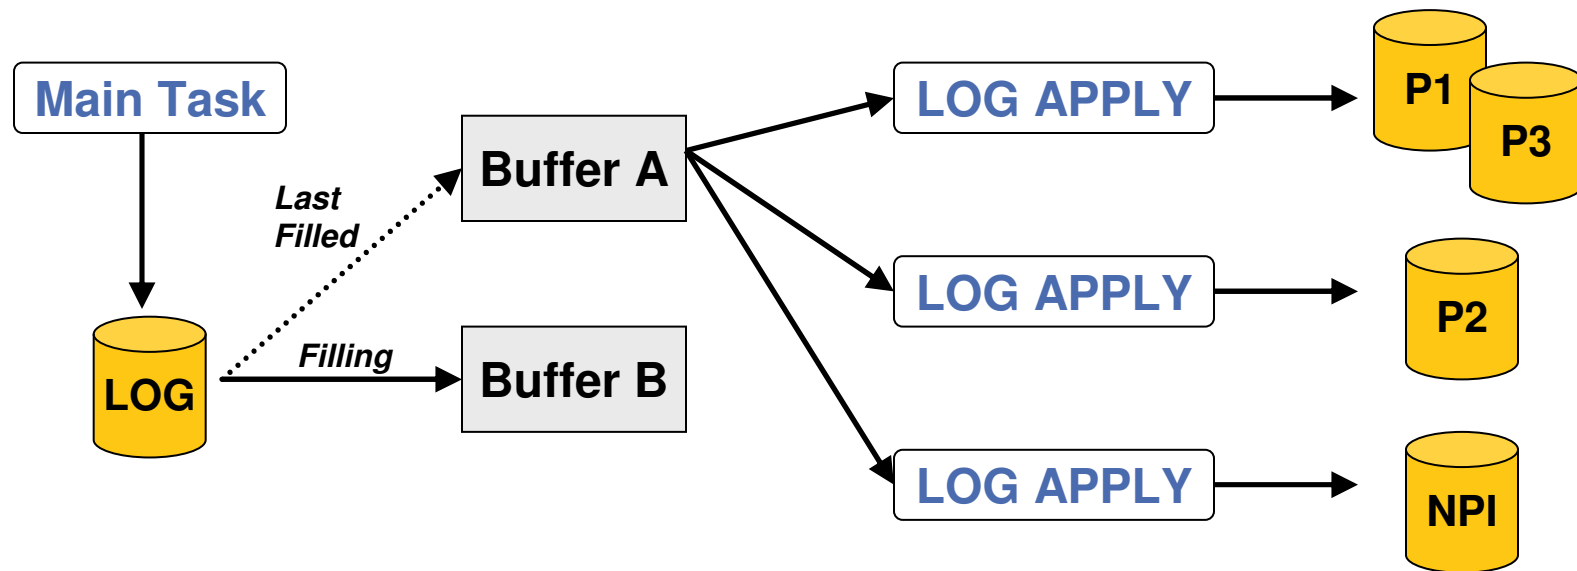
REORG TABLESPACE TS

Three partitions, one PI, one NPSI



- Multiple concurrent jobs no longer needed
- No longer allowed with PART specification and NPSIs

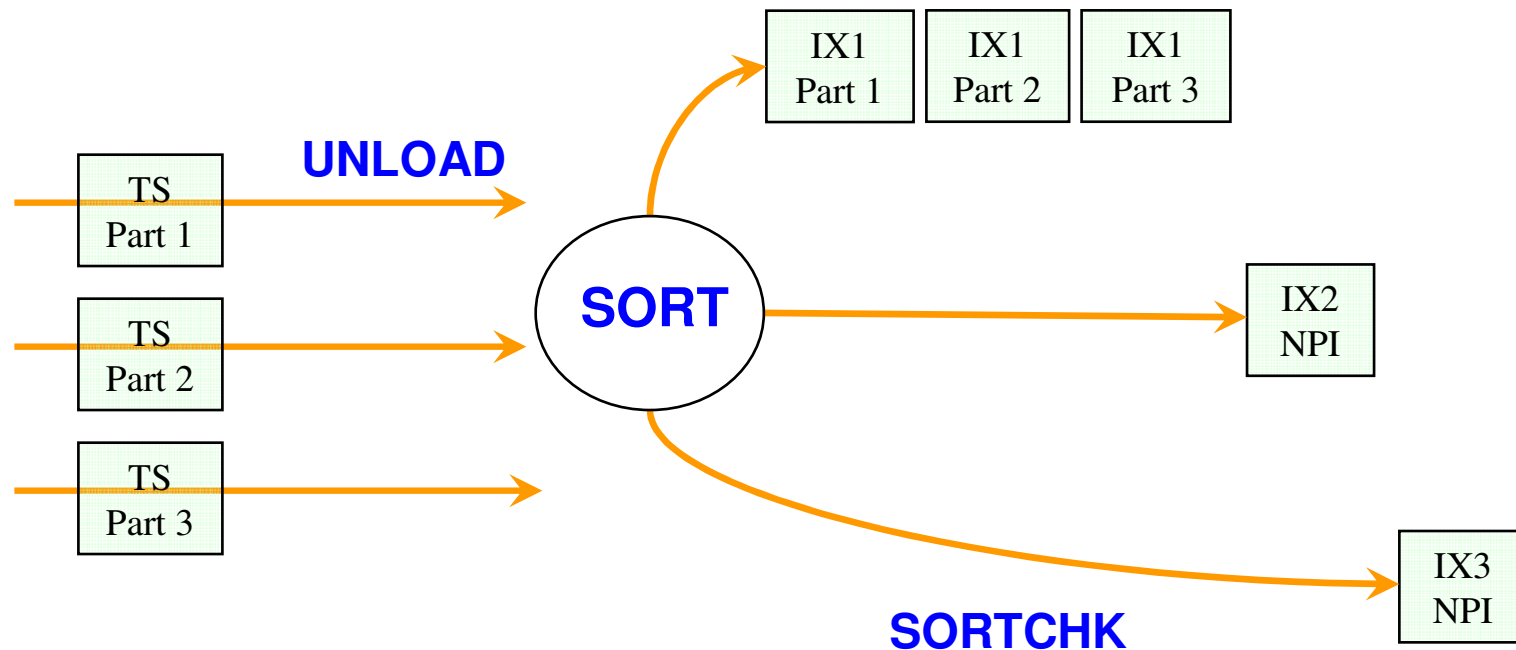
REORG Log Apply Parallelism



- Double Buffering
- Parallel log apply for SHRLEVEL CHANGE and SHRLEVEL REFERENCE PART if multiple NPIs

Check Index Shrlevel Reference Parallelism

CHECK INDEX(ALL) TABLESPACE TS SHRLEVEL REFERENCE
three partitions, three indexes



Improved elapsed time performance (like REBUILD INDEX)

Performance: Utility CPU time reduction

* primarily from index processing

- **10% to 20% in Copy, Recover table space**
- **5% to 30% in Load*, Reorg*, Rebuild Index***
- **20% to 60% in Check Index***
- **35% in Load Partition***
- **30% to 40% in Runstats Index***
- **40% to 50% in Reorg Index***
- **70% in Load Replace Partition with dummy input***

Performance: Runstats Histogram Statistics

- Summarizes data distribution on an interval scale
 - ▶ Extends non-uniform distribution stats
- DB2 uses equal-depth histograms
 - ▶ Each quantile has about the same number of rows
 - ▶ Example - 1, 3, 3, 4, 4, 6, 7, 8, 9, 10, 12, 15 (sequenced), cut into 3 quantiles

Seq No	Low Value	High Value	Cardinality	Frequency
1	1	4	3	5/12
2	6	9	4	4/12
3	10	15	3	3/12

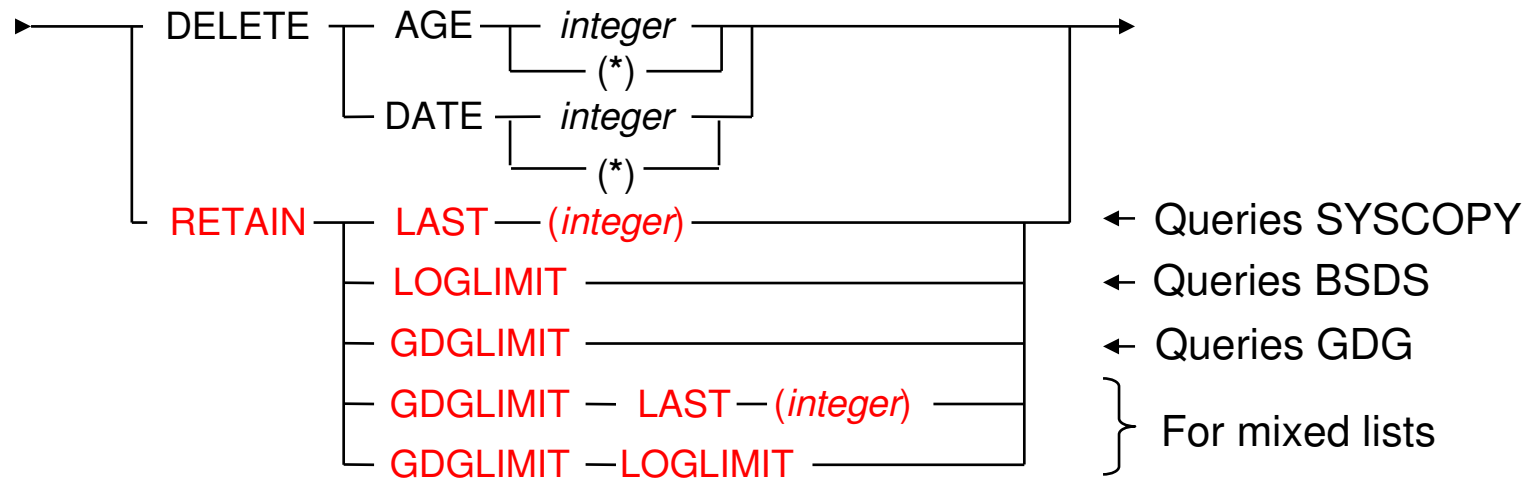


IBM Software Group

Backup and Recovery Enhancements



Backup and Recovery Enhancements



- Modify Recovery simplification and safety



Copy Utility Changes

- Always perform CHECKPAGE on the COPY utility
- The COPY utility includes SCOPE PENDING support to improve usability
- COPY utility bufferpool usage uses MRU management of those pages read by the COPY utility
- Template switching (e.g., copy to tape if large; to disk if small)

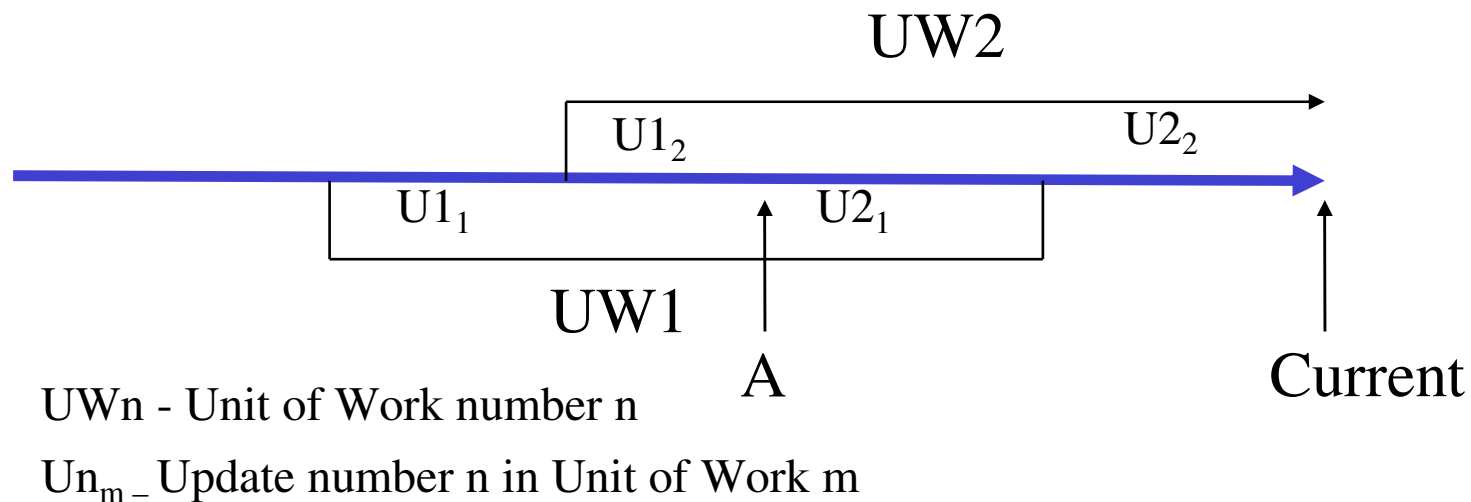
```
TEMPLATE LRG DSN &DB..&TS..D&DA..T&TI. UNIT=TAPE
TEMPLATE SML DSN &DB..&TS..D&DA..T&TI. UNIT=SYSALLDA LIMIT(20 CYL, LRG)
COPY TABLESPACE SMALL.TS COPYDDN(SML)
COPY TABLESPACE LARGE.TS COPYDDN(SML)
```

RECOVER utility changes

- Object level recovery from a system level backup
 - ▶ New zparms
 - SYSTEM-LEVEL BACKUPS
 - YES – Permit RECOVER to use system level backups as a base
 - RESTORE/RECOVER
 - YES – RESTORE & RECOVER to use tape backup, not disk
 - RECOVER statement override: FROMDUMP
 - DUMP CLASS NAME
 - Name of DFSMSHsm dump class if using backup on tape
 - RECOVER statement override: DUMPCLASS
 - ▶ No REPORT RECOVERY support at this time

Recover Utility Changes

- Recovery to a point in time with consistency (NFM mode)
 - ▶ Uncommitted changes are backed out
 - ▶ Significantly reduces (eliminates?) the need to run QUIESCE
 - ▶ Does not apply to RECOVER TOCOPY, TOLASTCOPY and TOLASTFULLCOPY using SHRLEVEL CHANGE copy (consistency is not ensured – use RBA/LRSN after COPY point)
 - ▶ Include all relevant objects in same RECOVER to ensure data consistency from the application point of view



More Backup and Recover Enhancements...

- RECOVER RESTOREBEFORE x'rba/lrsn' directs RECOVER to use a recovery base prior to the rba/lrsn
- Progression of RECOVER log apply shown via -DIS UTIL
 - ▶ To help see progress and estimate elapsed time:

DSNU116I csect-name RECOVER LOGAPPLY PHASE DETAILS:

STARTING TIME = timestamp

START RBA = ss START LRSN = rr

END RBA = ee END LRSN = nn

LAST COMMITTED RBA = cc LAST COMMITTED LRSN = ll

ELAPSED TIME = hh:mm:ss

More Backup and Recover Enhancements

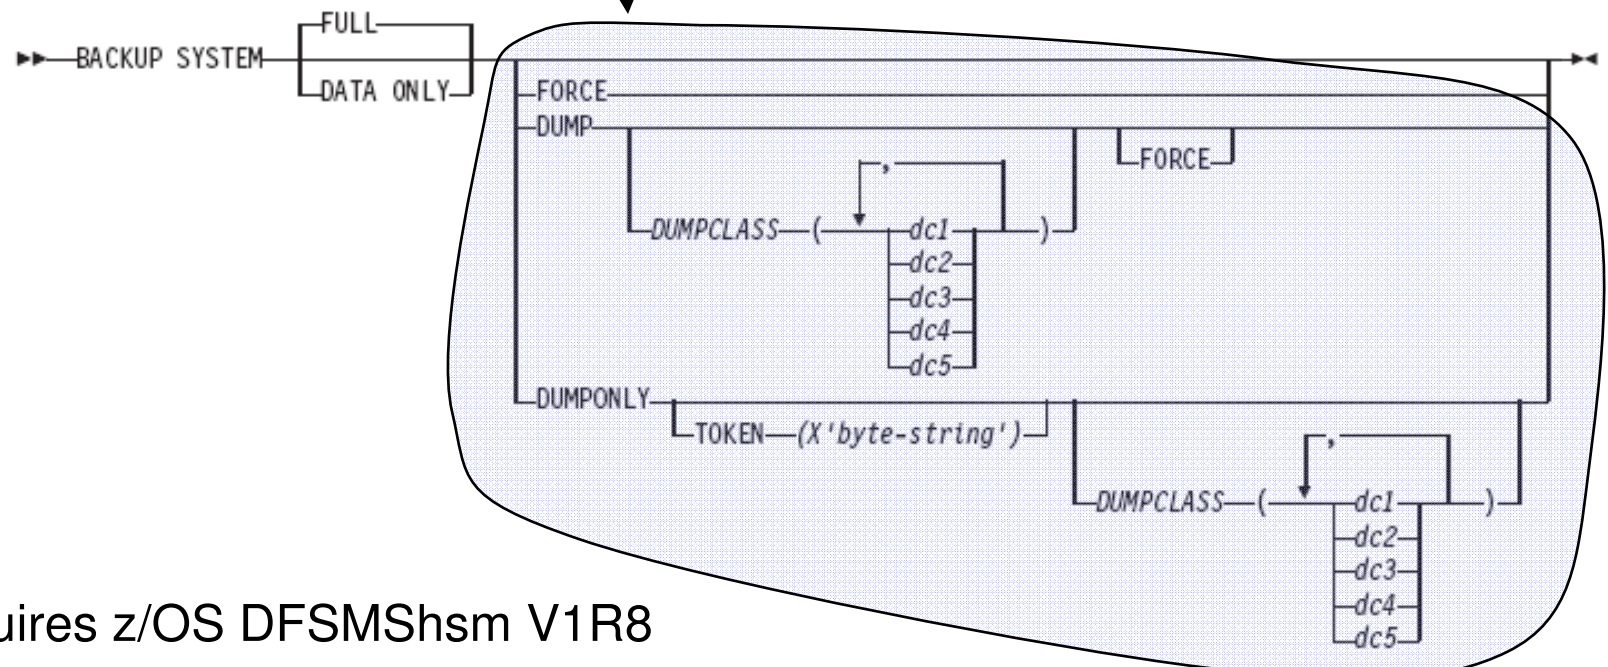
- Support large block interface (allows blocksize > 32,760 bytes) for tapes
 - ▶ Up to 40% reduction elapsed time for COPY and RECOVER RESTORE phase
- Support large format datasets (allows > 65,535 tracks per DASD volume)
 - ▶ Handy for COPYing large tablespaces with TEMPLATES
 - ▶ Requires NFM for creating large format datasets

BACKUP/RESTORE SYSTEM

- What's new in DB2 9 for BACKUP and RESTORE SYSTEM?
 - ▶ The ability to recover at the object level using system-level backups
 - RECOVER to point in time or to current
 - COPY YES indexes can be included in this recovery
 - Will restore from the previous system level backup or image copy
 - Requires z/OS V1R8
 - ▶ Tape Support
 - ▶ Support for Incremental FlashCopy
 - PK41001 & PK42014
 - ▶ RTS COPY columns also updated for BACKUP SYSTEM

DB2 9 Utilities

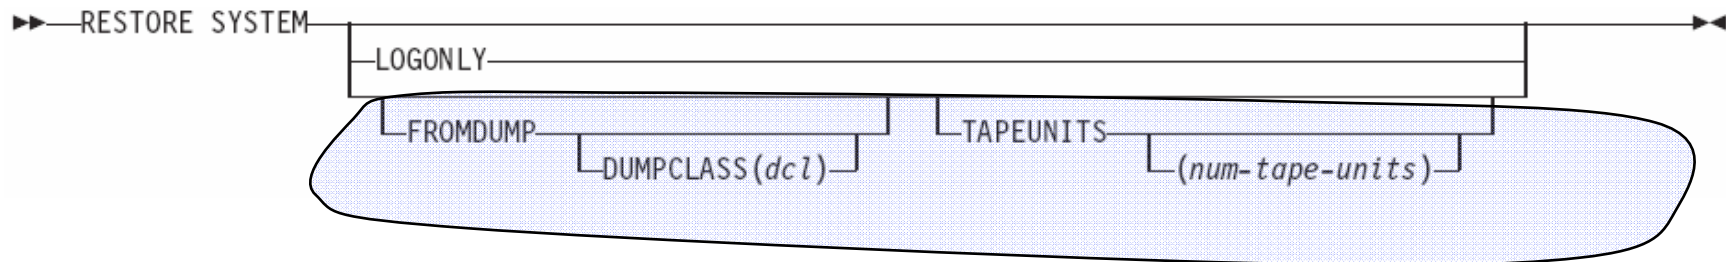
- Tape control



- ▶ Requires z/OS DFSMSHsm V1R8
- ▶ FORCE indicates that oldest copy can be overwritten even if dump to tape is in progress
- ▶ Use LIST COPYPOOL with DUMPVOLS option to verify dump status

DB2 9 Utilities

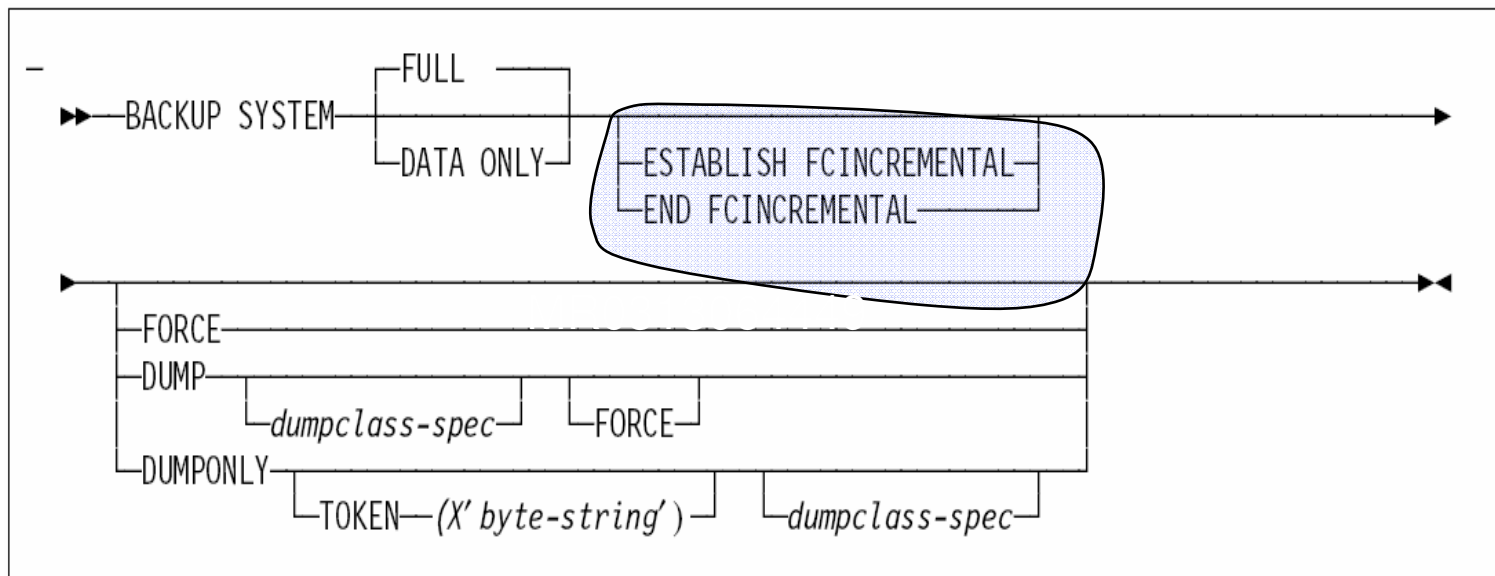
- Tape control



- ▶ Restore from dumps will use parallelism; limited by number of distinct tape volumes that the dump resides on; capped by TAPEUNITS
- ▶ ZPARMs to override tape options (DSNTIP6):
 - ▶ FROMDUMP ↔ RESTORE/RECOVER FROM DUMP
 - ▶ DUMPCLASS ↔ UTILS DUMP CLASS NAME
 - ▶ TAPEUNITS ↔ RESTORE TAPEUNITS

DB2 9 Utilities


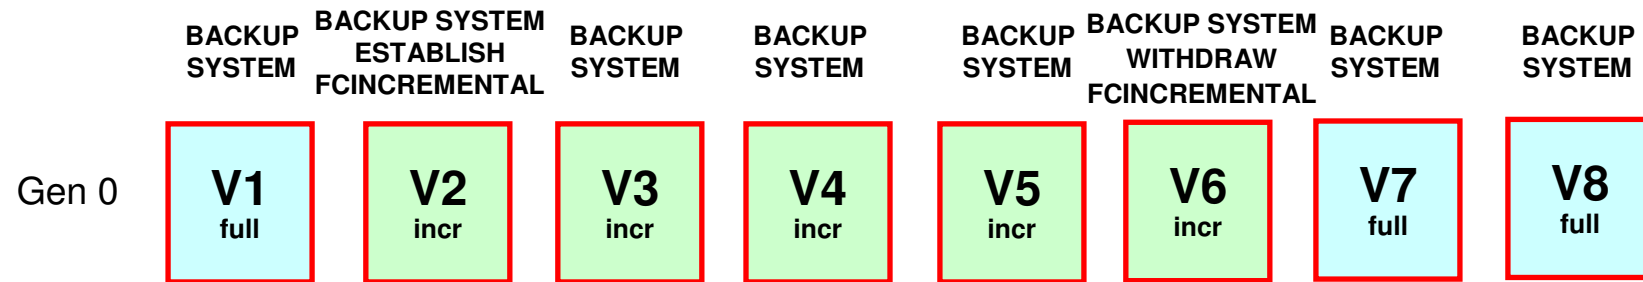
- Support for Incremental FlashCopy (Post GA via APAR PK41001)



- Incremental FlashCopy can also be activated using HSM commands; **BACKUP SYSTEM** will then use it


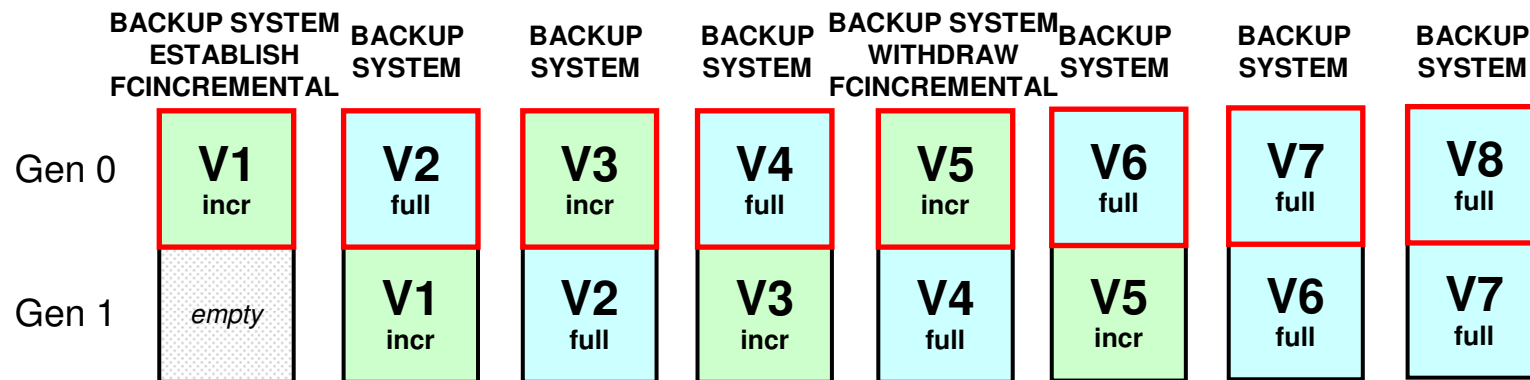
DB2 9 Utilities

- Incremental Flashcopy with a single version:



DB2 9 Utilities

- Incremental Flashcopy with two FC versions:



Miscellaneous Enhancements

- All utility messages have timestamp and julian date (day)

DSNU000I 158 22:52:30.07 DSNUGUTC...

- Limit of 254 parts per REORG on a compressed table space lifted; storage consumption reduced
- New DSNACCOX stored procedure to gather stats from the catalog (including new Real Time Stats) & recommend REORG, RUNSTATS, and COPY actions for objects ([see APAR PK44133](#))

Future enhancements

- Simplification, automation, availability, performance, flexibility
- Web-based administrative console
- Utilities companion product
- High priority requirements
 - ▶ Autonomic stats
 - ▶ Autonomic compression
 - ▶ Recover to a different table space in a different DB2 with consistency
 - ▶ Reduce RUNSTATS resource consumption
 - ▶ Data set-level FlashCopy support
 - ▶ Online consistent image copies
 - ▶ Policy-based recovery
 - ▶ UTSERIAL elimination
 - ▶ REORG enhancements (LOBs, etc.)
 - ▶ ...

References

- DB2 UDB for z/OS home page
<http://www.ibm.com/software/data/db2/zos/index.html>
- DB2 9 for z/OS Performance Topics, SG24-7473-00
- DB2 UDB for z/OS and OS/390 Version 7: Using the Utilities Suite, SG24-6289
- DB2 9 for z/OS Technical Overview, SG24-7330-0
- Recommendations for Tuning Large DFSORT Tasks
<http://www.ibm.com/servers/storage/support/software/sort/mvs/tuning/index.html>
- DFSMSrmm SMS ACS Support
<http://www.redbooks.ibm.com/abstracts/TIPS0530.html?Open>

DB2 UDB for z/OS information resources

- Information center
<http://publib.boulder.ibm.com/infocenter/dzichelp/index.jsp>
- Information roadmap
<http://ibm.com/software/db2zos/roadmap.html>
- DB2 UDB for z/OS library page
<http://ibm.com/software/db2zos/library.html>
- Examples trading post
<http://ibm.com/software/db2zos/exHome.html>
- DB2 for z/OS support
<http://ibm.com/software/db2zos/support.html>
- Official Introduction to DB2 for z/OS
<http://ibm.com/software/data/education/bookstore>