



IBM Information
ON Demand 2009

DB2 9 for z/OS Utilities Enhancements

Florence Dubois
DB2 for z/OS Development

Disclaimer

© Copyright IBM Corporation 2009. All rights reserved.

U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

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, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS AND/OR SOFTWARE.

IBM, the IBM logo, ibm.com, and DB2 are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml

Agenda

- **But first...**
- **DB2 9 Utilities Enhancements**
 - Availability enhancements
 - Performance enhancements
 - Backup and recovery enhancements
 - Other enhancements



IBM Information
ON Demand 2009

But first...

Dynamically Allocated Sort Work Data Sets

- **DB2/DFSORT determine data set size**
 - No hard coded DDs with size specification that need to be maintained
 - Single JCL (template) can be used for most utility jobs
 - Degree of parallelism determined according to available resources
- **BUT...**
 - Need to specify SORTNUM, but one size does NOT fit all
 - Different objects being processed by same job template
 - Different sorts within same utility, e.g. REORG with data and index sorts
 - DASD situation varies, SORTNUM 4 might work today, but tomorrow even SORTNUM 8 might fail
 - DB2's estimate sometimes not good enough

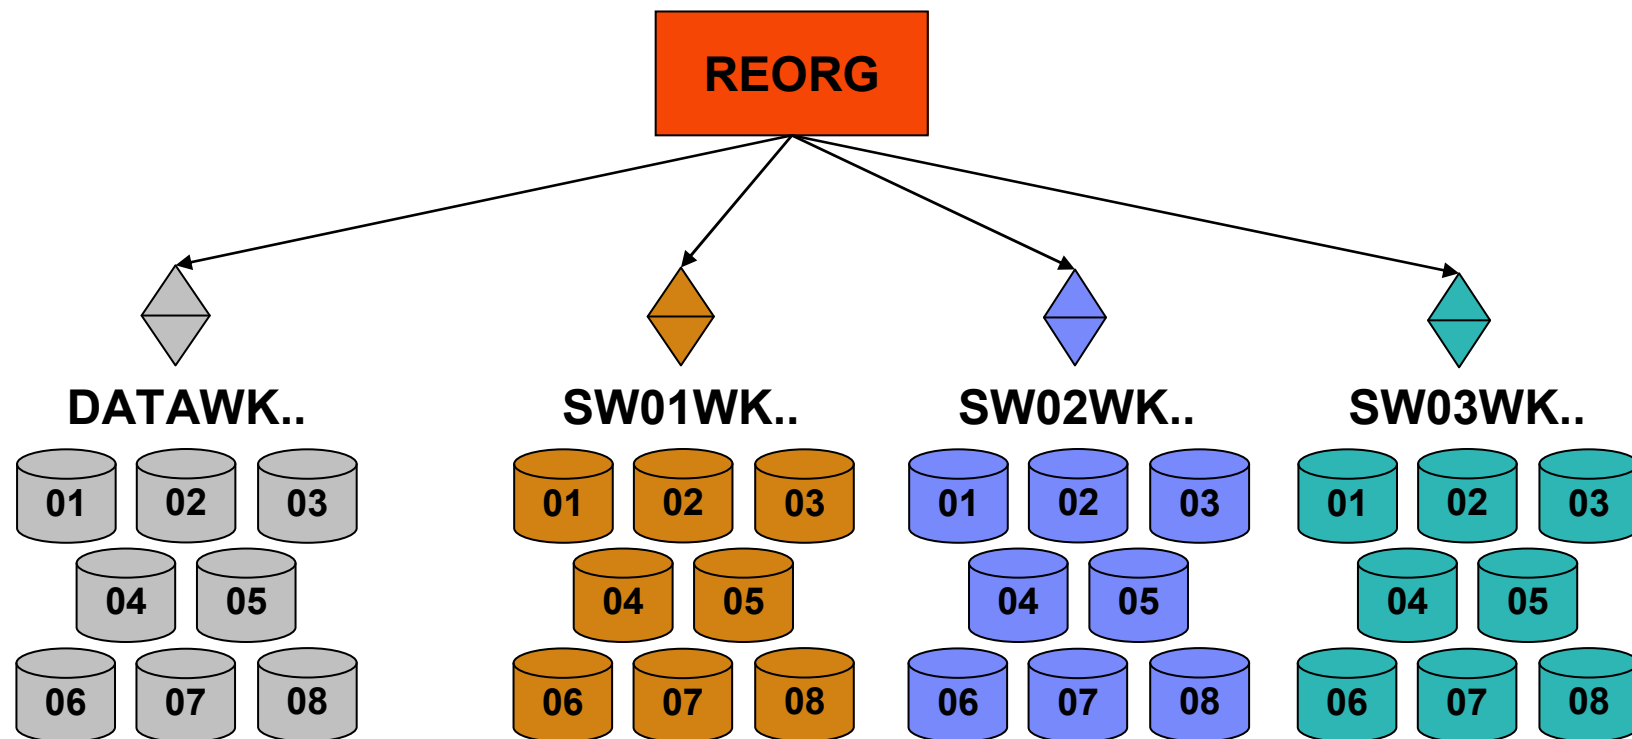
DB2 Allocated Sort Work Data Sets

- **New functionality to enable DB2 to dynamically allocate sort work data sets in utilities**
 - PK45916 (V8) & PK41899 (V9)
- **Enabled with UTSORTAL=YES**
 - No more need to specify SORTNUM
 - Existing SORTNUM specification can be honored (IGNSORTN=NO – default) or ignored (IGNSORTN=YES)
 - New behavior ignored if hard coded DD cards are found
- **Used for all sorts in utilities: LOAD, REORG, CHECK INDEX, REBUILD INDEX, CHECK DATA, RUNSTATS**
- **Message “DSNU3340I - UTILITY PERFORMS DYNAMIC ALLOCATION OF SORT DISK SPACE” indicates use**
- **Data sets for largest sorts are allocated first**
- **Attempts to allocate data sets as large as possible (starting with 2 data sets per sort task, more data sets allocated if necessary)**



DB2 Allocated Sort Work Data Sets

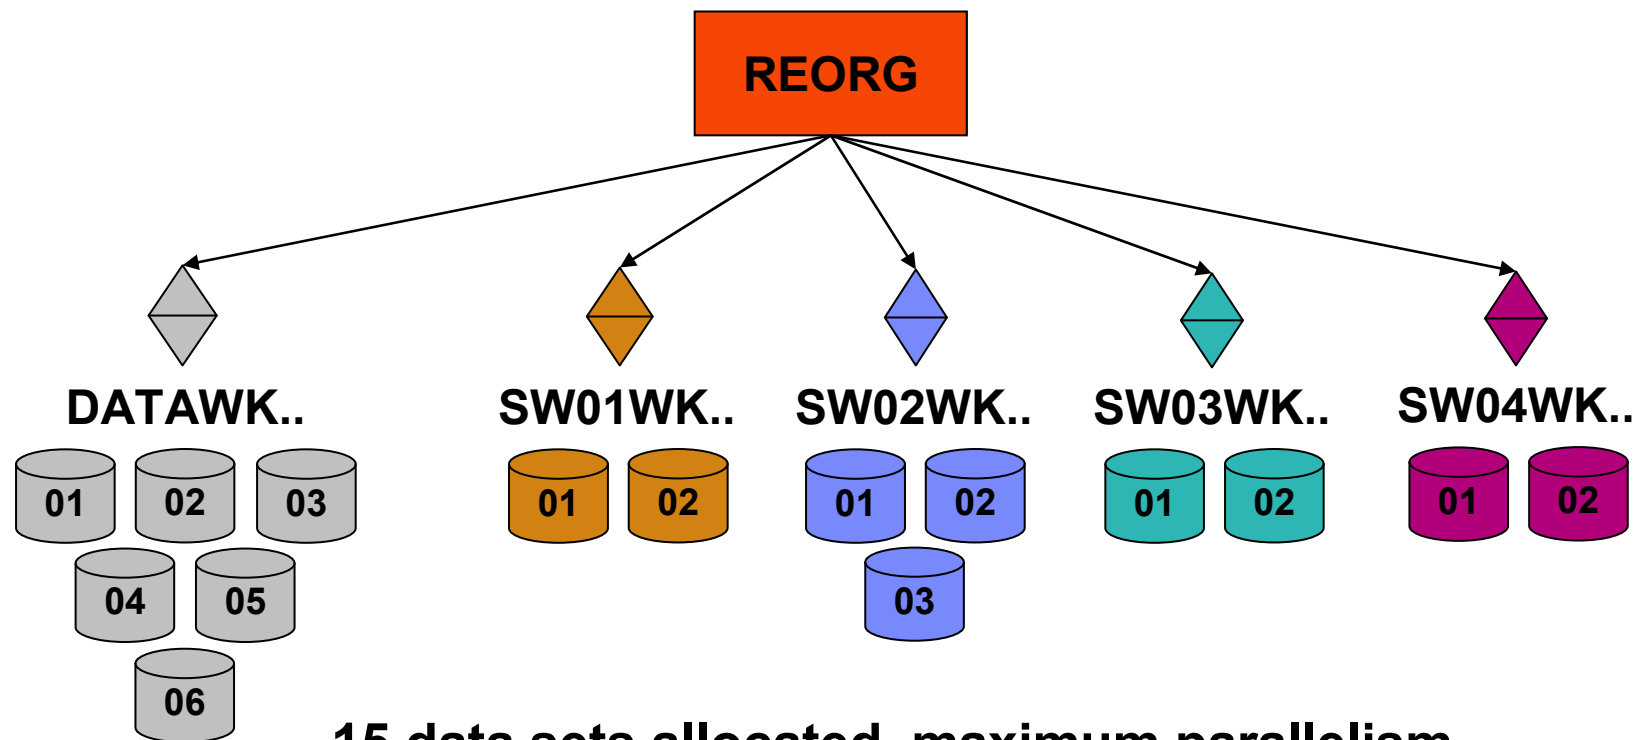
- **Sample: REORG TABLESPACE ... SORTNUM 8 with 4 indexes, index build parallelism constrained to 3 sort tasks**



32 data sets allocated, parallelism constrained

DB2 Allocated Sort Work Data Sets

- **Sample: REORG TABLESPACE with 4 indexes and UTSORTAL=YES, index build parallelism with 4 sort tasks**



DB2 Allocated Sort Work Data Sets

- **Uses Real-Time statistics for size estimates**
 - If RTS data not available, DB2 reverts to the old method
 - DSNU3343I csect-name REAL-TIME STATISTICS INFORMATION MISSING FOR obj-type obj-qual. obj-name partition-information
 - Start using RTS on V8 if not already done (always active in DB2 9)
 - RTS can benefit you in many ways
 - Required values in RTS are initialized by REORG TABLESPACE and REBUILD INDEX
 - Don't replace DB2 objects outside DB2's control to maintain RTS accuracy (e.g. DSN1COPY) or let DB2 "know":
 - Set TOTALROWS to NULL in SYSIBM.(SYS)TABLESPACESTATS or TOTALENTRIES to NULL in SYSIBM.(SYS)INDEXSPACESTATS to invalidate existing statistics if replacing with significantly different data

DB2 Allocated Sort Work Data Sets

- **Recommended maintenance**

- PK64624 - LOAD with multiple INTO TABLE
- PK64915 - Improve estimates for REBUILD and CHECK INDEX with segmented table spaces with missing RTS
- PK66597 - LOAD ABEND0C4 RC00000011 when SYSTEMPL DD specified but not used
- PK70001 - ICE046A SORT CAPACITY EXCEEDED when REORG is restarted in UNLOAD phase, improved fall back estimates for multi table table spaces
- PK71733 - DB2 overallocates sort work data sets when data sets are written in many fragments
- PK75022 - REORG TABLESPACE ABEND0C4 when no index defined on partitioned table space with unload/reload parallelism
- PK77774 - Reduce secondary quantity size for DB2 allocated sort work data sets
- DFSORT APAR PK63409 - ICE046A SORT CAPACITY EXCEEDED when estimate is slightly below actual value
 - Fix for DFSORT in z/OS 1.10 is provided in PK66899

Other Recent Maintenance Changes

- **PK61759 (V8 & V9)**
 - LOAD & REORG performance improvement
 - 10% CPU & ET improvement in reload phase
 - Sort processing performance improvement
 - 10% CPU reduction in sort phase
- **PK60956 (V8 & V9)**
 - SORTBLD performance improvement for indexes with small SECQTY
 - SORTBLD elapsed up to 20x improvement
- **PK47093 (V8 & V9)**
 - LOAD and REORG invalidation of dynamic statement cache



IBM Information
ON Demand 2009

DB2 9 for z/OS Utilities – Availability Enhancements

Improved concurrency

REBUILD INDEX Enhancements

- **Rebuild Index SHRLEVEL CHANGE**
 - Great for building new non-unique indexes or when index is in RBDP
 - Not suited for unique indexes as uniqueness can not be checked during rebuild process, so no INSERT/ UPDATE/ DELETE allowed that affects unique index
 - Index is built in place with no shadow
 - To move indexes to different volumes with availability, you still need to use REORG INDEX SHRLEVEL CHANGE
 - Table space must be in LOGGED state (new log phase)
 - No parallel jobs on different indexes of the same table space -> use single job with multiple indexes specified

REORG Enhancements (Availability)

- **Elimination of BUILD2 phase**
 - NPIs also shadowed
 - NPIs implicitly reorged
 - REORG SHRLEVEL REFERENCE PART now has a log phase
 - Improved availability, but...
 - Concurrent REORGs of parts in the same table space are not permitted if an NPI exists on that table space
 - REORG of a single part with an NPI can take longer
 - Recommendation:
 - REORG multiple (adjacent) parts in the same REORG statement
 - Expect to provide ability to REORG disparate parts via APAR in V9
- **Display claimers on drain failure – PK41182 (V8 & V9)**

REORG Enhancements (Availability) ...

• **Before** DSNUGUTC - REORG TABLESPACE DB1.TS1 COPYDDN(SYSCOPY) SHRLEVEL CHANGE STATISTICS
 DSNUGBAC - RESOURCE UNAVAILABLE
 REASON 00C200EA
 TYPE 00000200
 NAME DB1.TS1

• **After** DSNUGUTC - REORG TABLESPACE DB1.TS1 COPYDDN(SYSCOPY) SHRLEVEL NONE STATISTICS
) *****
) * DISPLAY DATABASE SUMMARY
 * GLOBAL CLAIMERS
) *****
 ...

NAME	TYPE	PART	STATUS	CONNID	CORRID	CLAIMINFO
TS1	TS		RW,UTRO	BATCH	CAPP	(WR,C)
-			AGENT TOKEN 10			
-			MEMBER NAME V91A			
TS1	TS		RW,UTRO	BATCH	CAPP	(RR,C)
-			AGENT TOKEN 10			
-			MEMBER NAME V91A			

***** DISPLAY OF DATABASE DB1 ENDED *****
) DSNTDDIS 'DISPLAY DATABASE' NORMAL COMPLETION
 DSNUGBAC - RESOURCE UNAVAILABLE
 REASON 00C200EA
 TYPE 00000200
 NAME DB1.TS1

} Culprits

CHECK INDEX Enhancements (Availability)

- **CHECK INDEX SHRLEVEL CHANGE (V8 & V9)**
 - Drain writers for target data and indexes
 - Usual drain parameters supported
 - Copy datasets to shadows with DFSMSdss ADRDSSU with FCNOCOPY
 - Uses dataset-level FlashCopy if available, else traditional media copy
 - If using DASD mirroring, consider ZPARM UTIL_TEMP_STORCLAS (PK41711)
 - After logical complete for datasets
 - Dedrain target data and indexes - Allows write access to the objects
 - Run CHECK on shadow data and indexes
 - At utilterm delete shadow datasets when DB2 managed
 - CHKP/ACHKP/AUXW no longer set if errors detected
 - Not reset either – use REPAIR
 - Look for messages and generated REPAIR statements

REORG LOB Enhancements

- **REORG LOB now supports SHRLEVEL REFERENCE**
 - Reclamation of unused space
 - Full read access to LOB data except during SWITCH phase
 - Inline imagecopy required to maintain recoverability
 - No restart capability
 - Shadow pageset discarded in event of failure
 - SHRLEVEL NONE still supported
 - Remains default, but will be deprecated in future
 - Use only if disk space is constrained

More Online Utilities

- **CHECK DATA SHRLEVEL CHANGE**
- **CHECK LOB SHRLEVEL CHANGE**
- **REPAIR LOCATE ... SHRLEVEL CHANGE**
- **Clones effectively provide LOAD REPLACE SHRLEVEL CHANGE**
 - UTS only
- **UNLOAD SKIP LOCKED DATA**
 - ISO(CS) option that skips locked rows/pages
 - Gives consistent data, but not necessarily all rows



IBM Information
ON Demand 2009

DB2 9 for z/OS Utilities – Performance Enhancements

Reducing elapsed and CPU times

Performance: Utility CPU Time Reduction

** primarily from index processing*

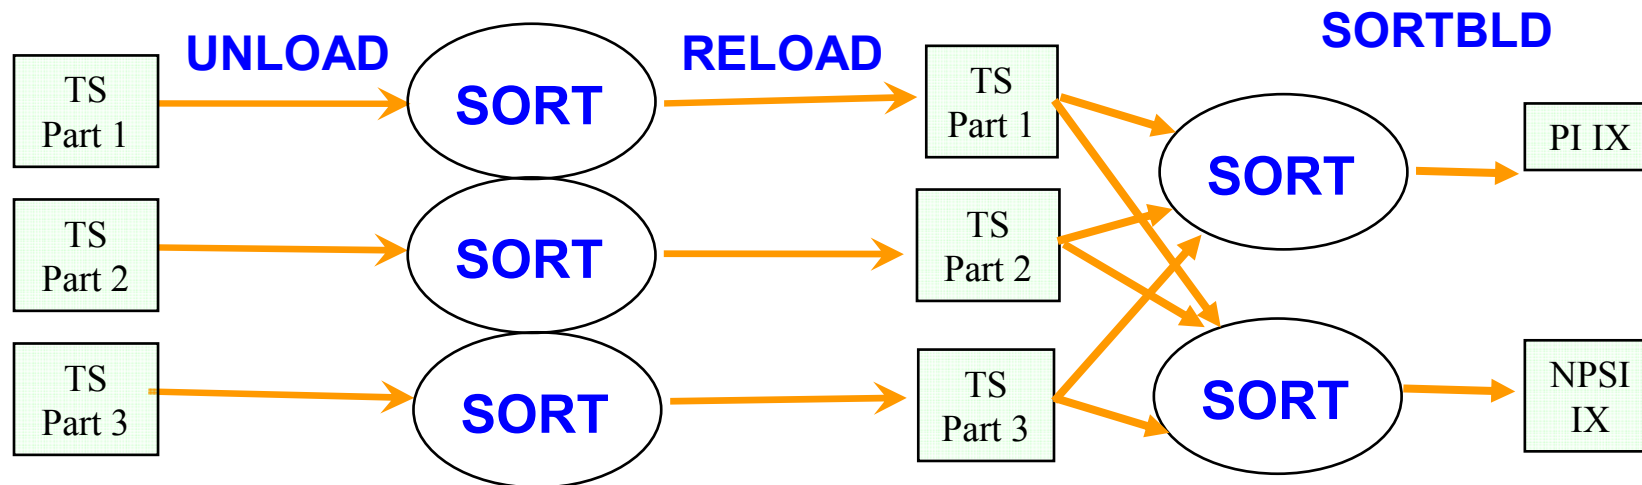
- **10% to 20% in Copy (even with forced CHECKPAGE YES)**
- **5% to 30% in Load*, Reorg*, Rebuild Index***
 - Except REORG TABLESPACE SHRLEVEL CHANGE PART with NPSIs
- **20% to 60% in Check Index***
- **35% in Load Partition***
- **30% to 40% in Runstats Index***
- **40% to 50% in Reorg Index***
- **70% in Load Replace Partition with dummy input***

REORG Enhancements (Performance)

- **REORG Partition Parallelism**

REORG TABLESPACE TS

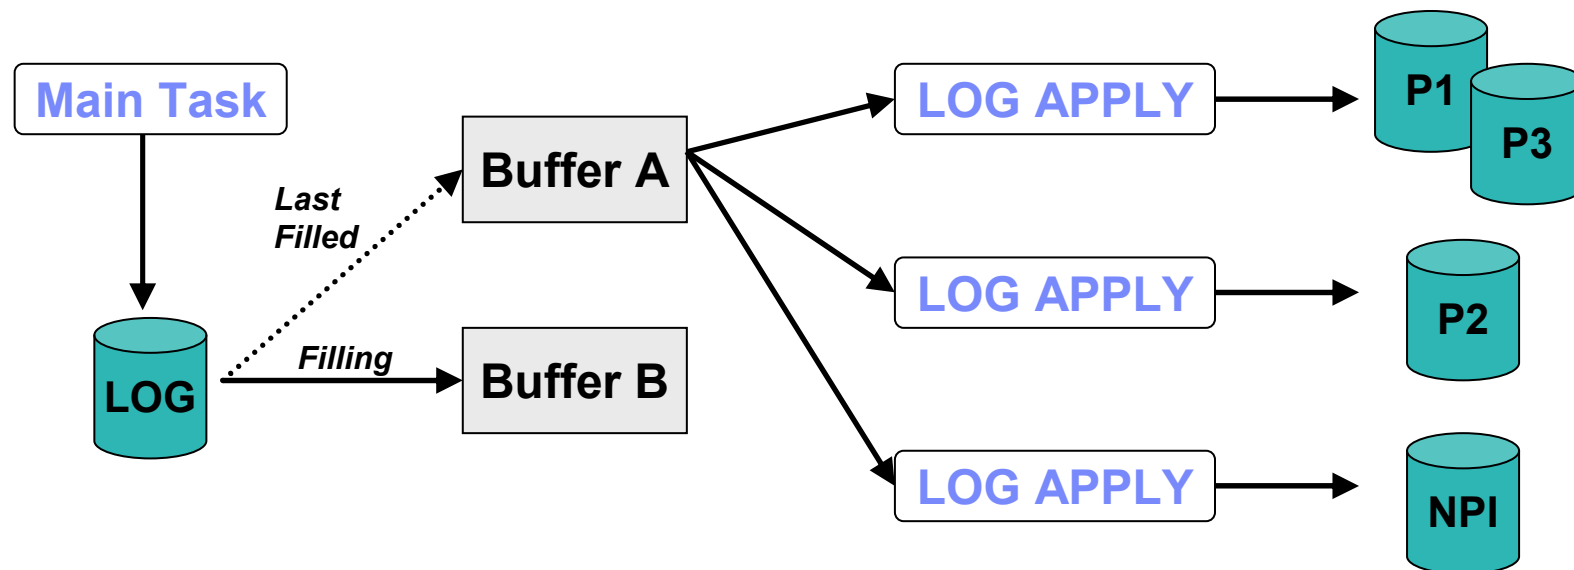
Three partitions, one PI, one NPSI



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

REORG Enhancements (Performance) ...

- REORG Log Apply Parallelism

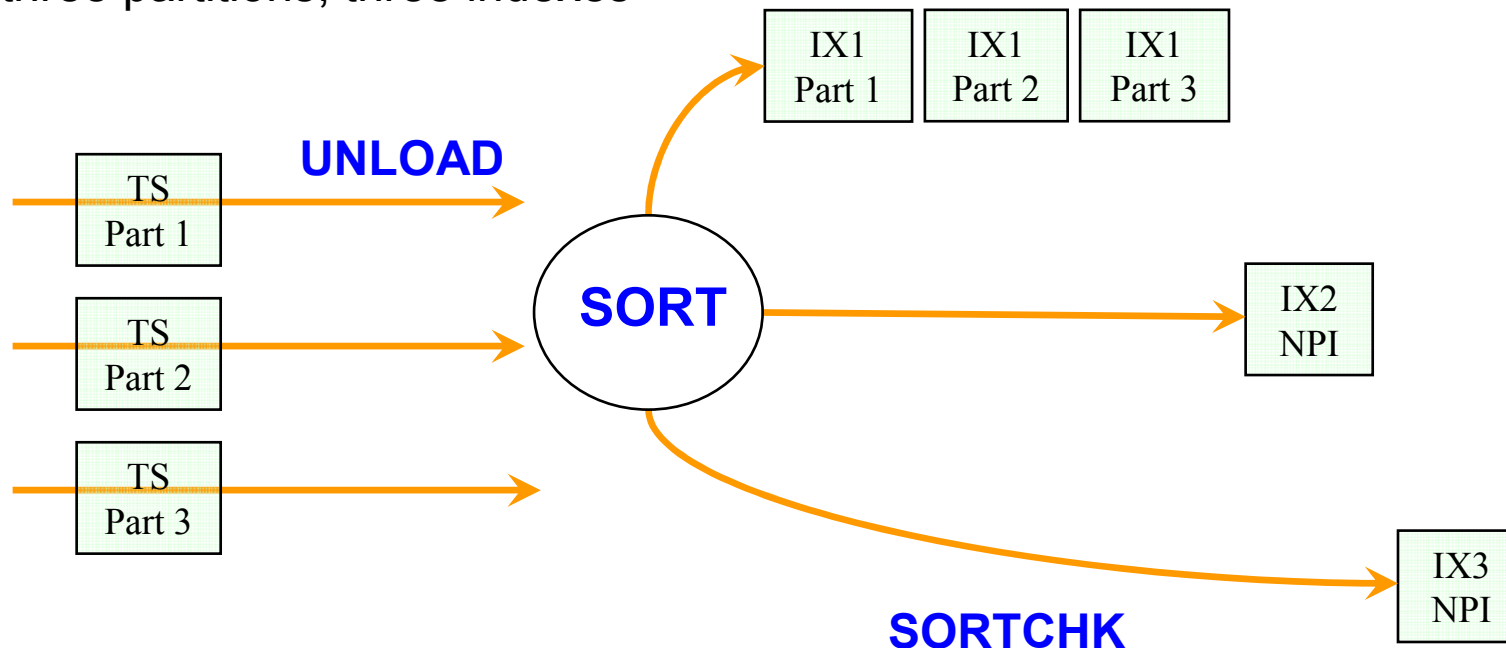


- Double Buffering
- Parallel log apply for SHRLEVEL CHANGE and SHRLEVEL REFERENCE PART if multiple NPIs

CHECK INDEX Enhancements (Performance)

- **CHECK INDEX SHRLEVEL REFERENCE Parallelism**

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



Improved elapsed time performance (like REBUILD INDEX)

What's New in DB2 for z/OS Utilities and zIIP?

- **Additional zIIP engine offload for the DB2 Utilities**
 - For DB2 Version 8 and DB2 9
- **What is eligible?**
 - Almost all DB2 utilities sorting of fixed-length records in the memory object sort path
- **System requirements**
 - System z9 or z10 with specialty engines and z/OS V1.10 with PTF UK48846 from APAR PK85856 and PTFs for either DB2 V8 or DB2 9
 - For DB2 V8
 - PTF UK48911 from APAR PK85889
 - For DB2 9
 - PTF UK48912 from APAR PK85889

Customer Value

- **Almost all utility sort processing involves fixed-length records, so this new support will apply to most utilities**
 - Exception is REORG data sorts that handle variable length records and is not eligible for the offload
- **Most utility processing involves sorting index keys**
 - Can represent a significant amount of data
 - Can consume up to 60% of utility CPU time in sort processing
- **Initial testing determined approximately 50%* of sort CPU time will be offloaded with this enhancement depending on**
 - if the sort processing selects the memory object sort path
 - the size of data being sorted
 - available system resources
- **Who will benefit?**
 - Customers who are constrained on CPU during utility processing or who would like to reduce internal CPU charge back costs for the utilities

**Actual results may vary*



IBM Information
ON Demand 2009

Backup and Recovery Enhancements

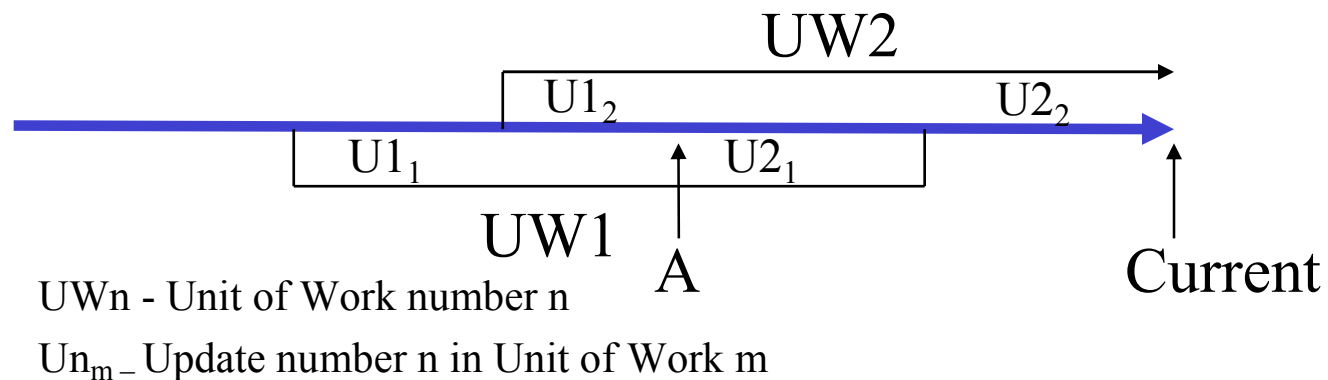
COPY Enhancements

- **CHECKPAGE** option always active
 - “Copy Pending” avoided if broken page encountered
 - Look for RC=8!
- **COPY SCOPE PENDING** to copy only objects in “Copy Pending”
- **Buffer Pool** usage uses MRU management of pages read by COPY
- **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(LRG)
```

RECOVER Enhancements

- **Recovery to any point in time with consistency (NFM mode)**
 - Uncommitted changes are backed out
 - Does not apply to RECOVER TOCOPY, TOLASTCOPY and TOLASTFULLCOPY if COPY is SHRLEVEL CHANGE
 - 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



RECOVER Enhancements ...

- **Recovery to any point in time with consistency (NFM mode) ...**
 - Two new phases in RECOVER
 - LOGCSR and LOGUNDO (will each be executed by member)
 - RECOVER can be restarted in any LOG phase
 - LOGCSR will restart at beginning, others at last commit point
 - Fast Log Apply not used for LOGUNDO phase
 - Progress of LOGUNDO phase can be monitored by regular status message DSNU1555I showing current and target log RBA
 - e.g. to detect backout of long running URs
 - This function can greatly reduce your need to run QUIESCE
 - If you want to capture a log point non-disruptively
 - -DISPLAY LOG, or
 - QUIESCE against SYSEBCDC (or some other dummy tablespace), or
 - SELECT CURRENT_TIMESTAMP from SYSDUMMY1 and convert it to an LRSN

RECOVER Enhancements ...

- **RECOVER RESTOREBEFORE x'RBA/LRSN'**
 - Directs RECOVER to use a recovery base prior to the rba/lrsn
 - Avoids need to rename undesired image copies
- **Progression of RECOVER log apply shown via -DIS UTIL**
 - Shows progress and helps estimate recovery completion time

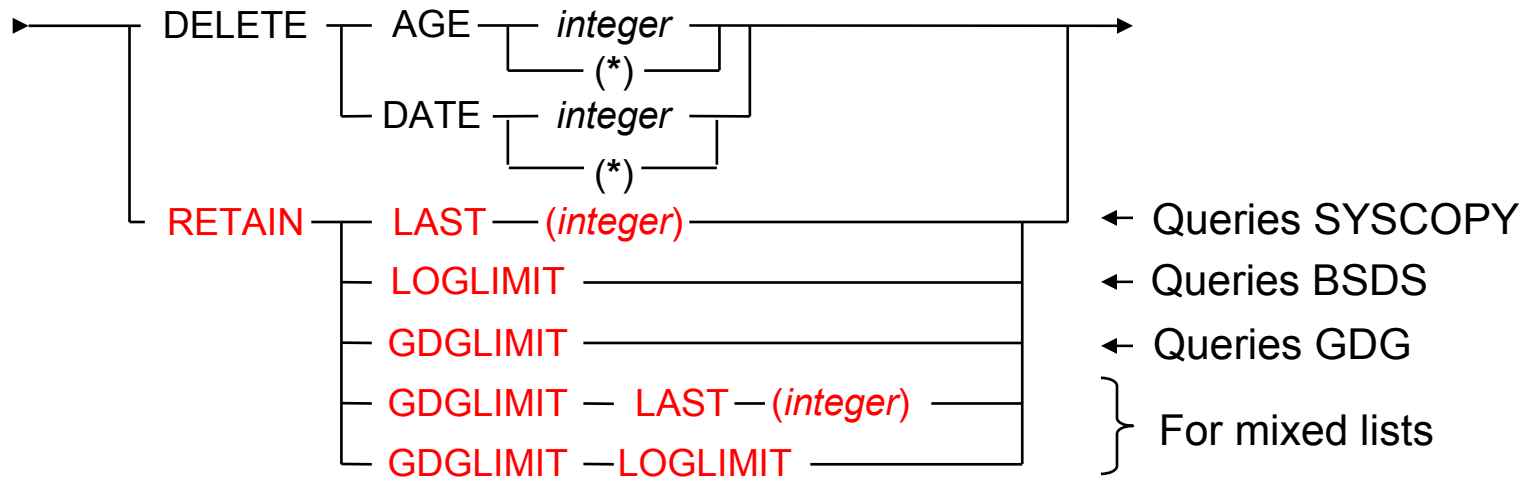
```
DSNU116I csect-name RECOVER LOGAPPLY PHASE DETAILS:  
  STARTING TIME = timestamp  
  START RBA = ss START LRSN = rr  
  END RBA = ee END LRSN = nn  
  LAST COMMITTED RBA = cc LAST COMMITTED LRSN = ll  
  ELAPSED TIME = hh:mm:ss
```

COPY and RESTORE Enhancements

- **Support large block interface for tapes**
 - Allows blocksize > 32,760 bytes
 - 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 those large tablespaces with TEMPLATES
 - Requires NFM for creating large format datasets

MODIFY Enhancements

- **Simplification and safety**




- SYSLGRNX records deleted even if no SYSCOPY records deleted
- Deletion works on dates not on timestamps, so more entries than requested might be kept

BACKUP/RESTORE SYSTEM Enhancements

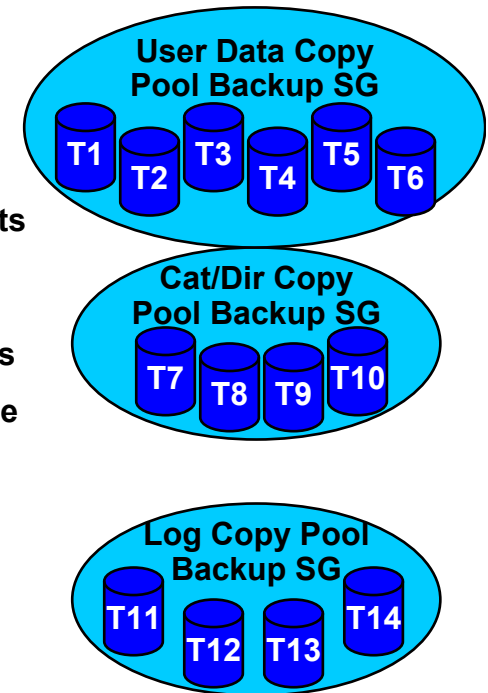
- **But first, a quick primer... DFSMS terms explained**

POOL Storage Groups



- Similar in concept to DB2 STOGROUPs
- A collection of (source) volumes
- Stores user data, catalog/directory, log data sets
- This example uses source volumes: S1 – S7

COPY POOL BACKUP Storage Group(s)

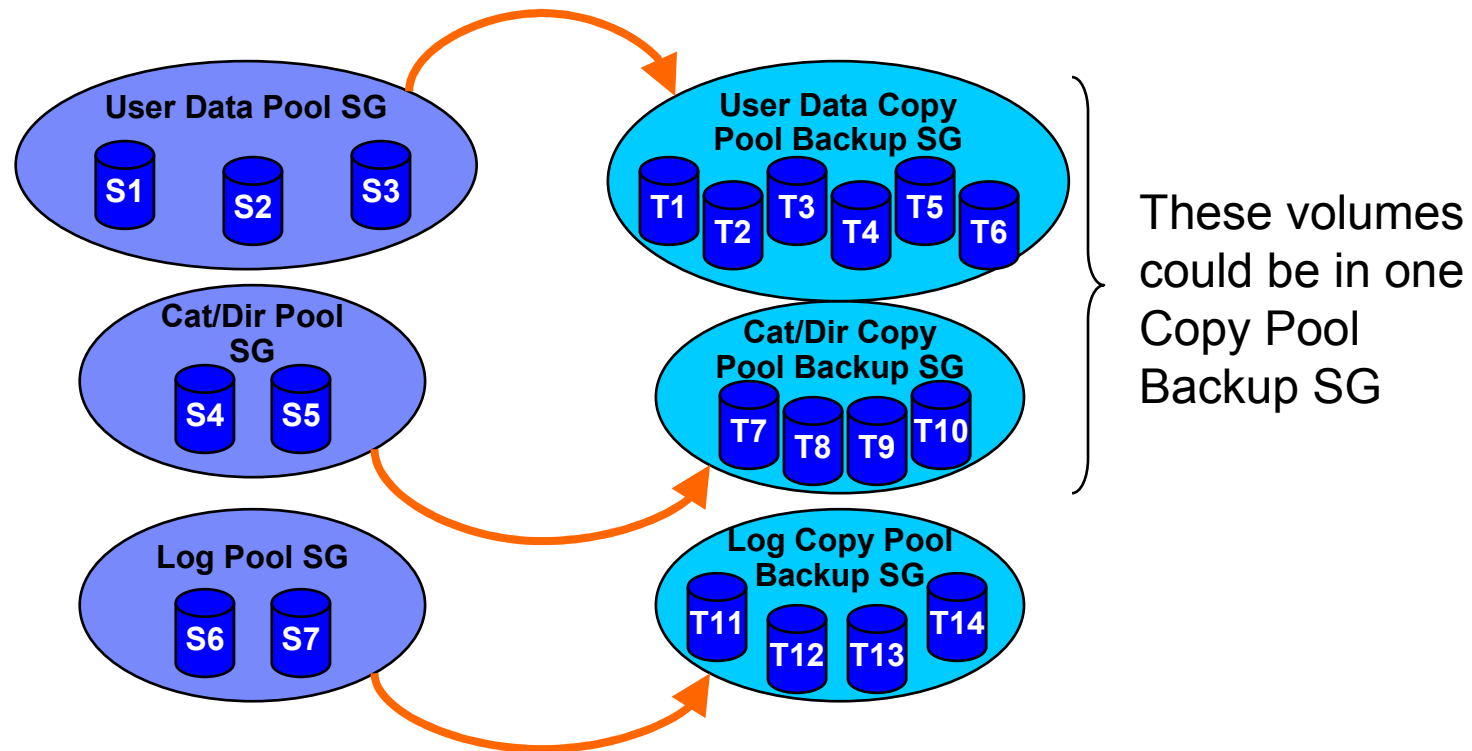


- A collection of volumes dedicated as FlashCopy targets for a source storage group
- Stores backups of user data, catalog/directory, log data sets
- Should contain a fixed multiple of volumes of the source storage group
- This example uses target volumes: T1 – T14

BACKUP/RESTORE SYSTEM Enhancements ...

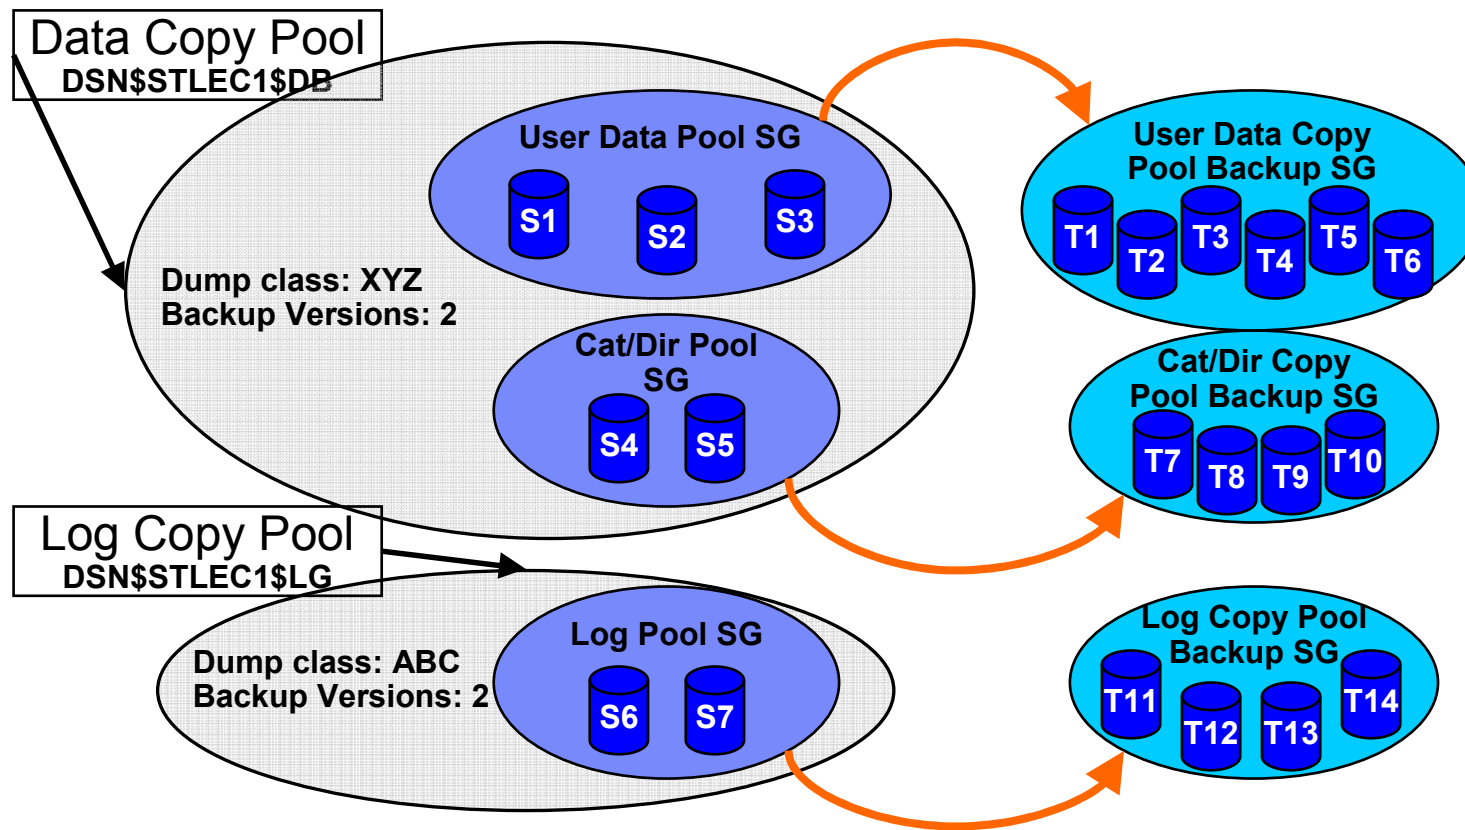
- But first, a quick primer... Defining the FlashCopy target SG

The Copy Pool Backup SG is an attribute of the Pool SG



BACKUP/RESTORE SYSTEM Enhancements ...

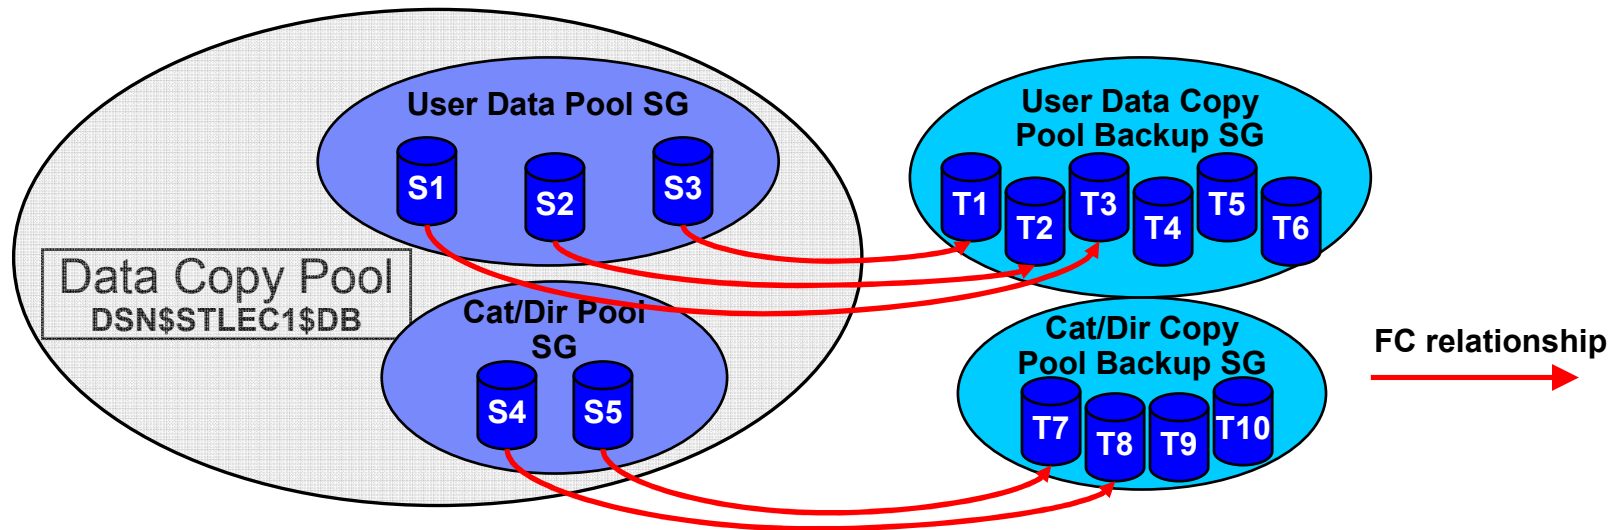
- But first, a quick primer... Copy Pools are a collection of source SGs



- A Copy Pool contains up to 256 (source) storage groups
- BACKUP SYSTEM/ RESTORE SYSTEM operate on Copy Pool Level

BACKUP/RESTORE SYSTEM Enhancements ...

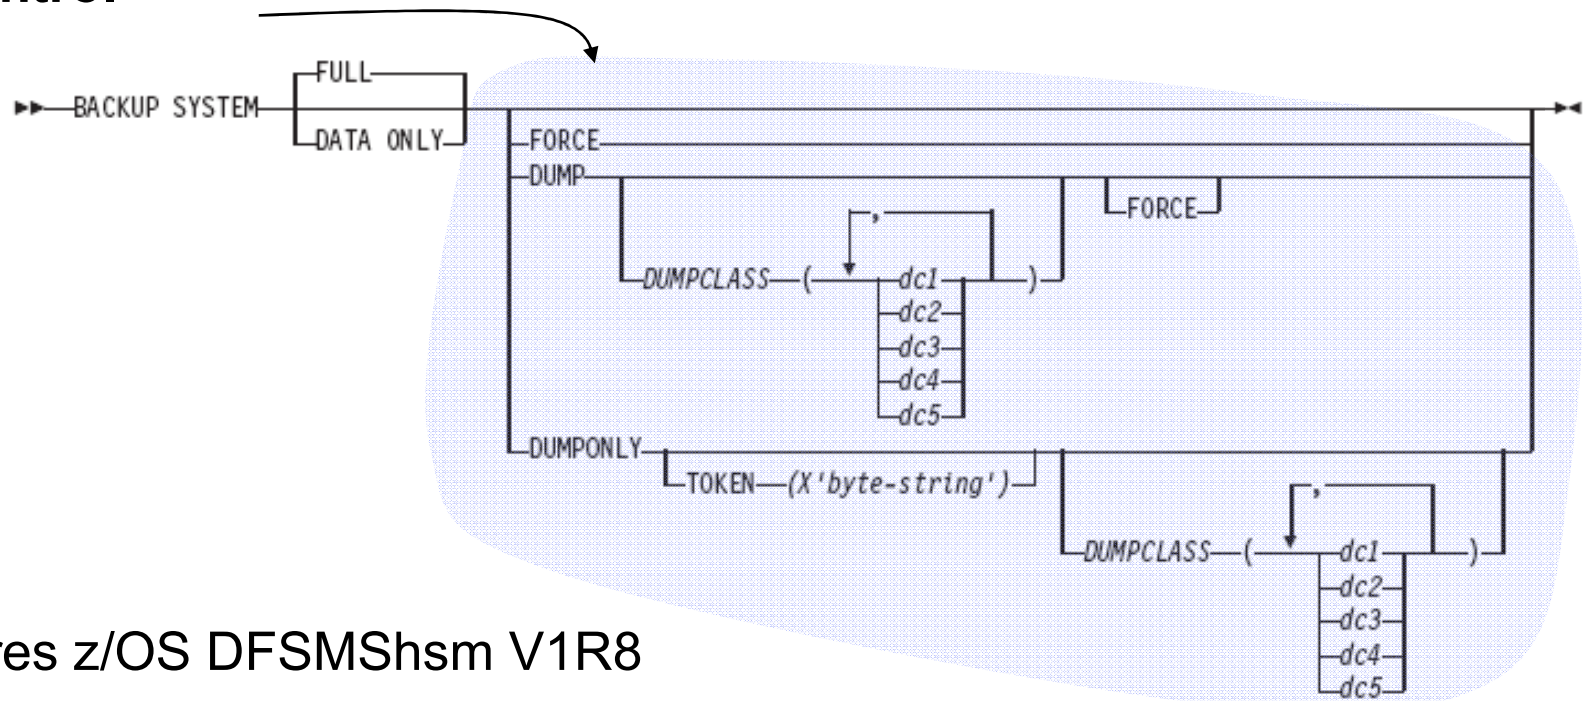
- **But first, a quick primer... FlashCopy (via the BACKUP SYSTEM DATA ONLY utility statement)**
 - A relationship is established between source volume and target volume
 - The copy is considered successful when the relationship is established.



- A background copy is then started which will result in a target volume that looks like the source volume when the relationship was established.
- The relationship goes away when the background copy completes.
- Restrictions and limitations when used in conjunction with PPRC and XRC

BACKUP/RESTORE SYSTEM Enhancements ...


- **Tape control**



- Requires z/OS DFSMSshm 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

BACKUP/RESTORE SYSTEM Enhancements ...

- **Tape control ...**



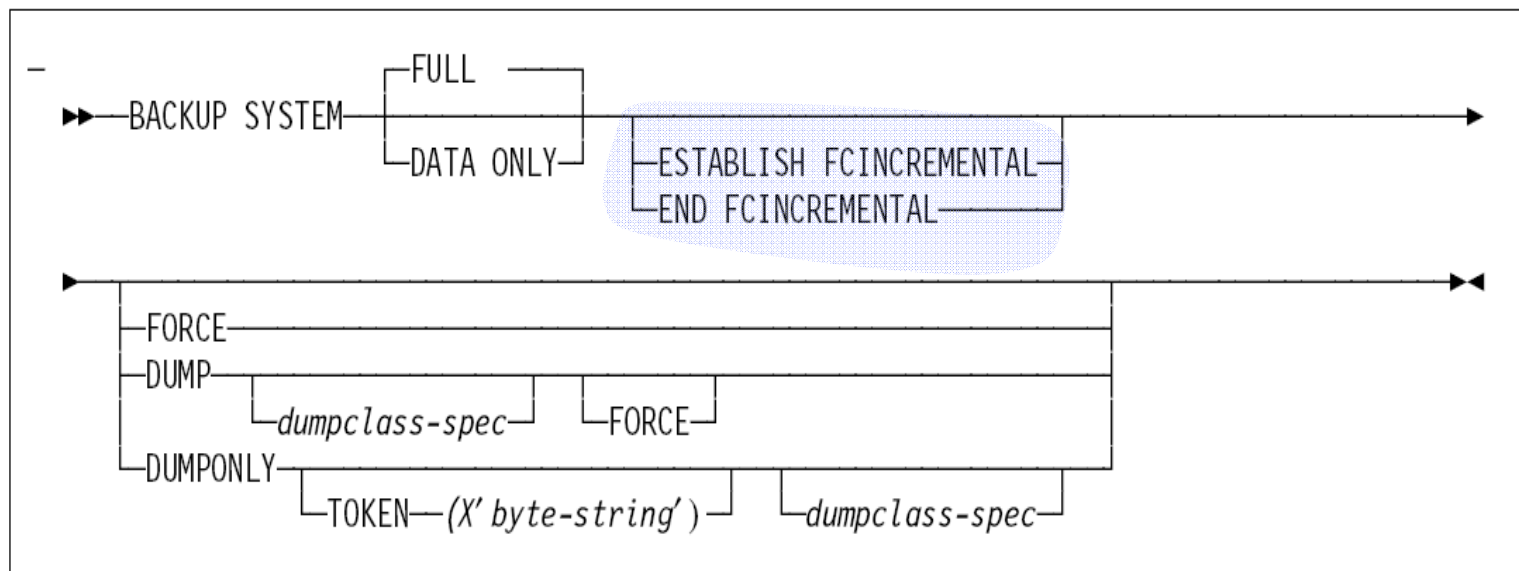
- Restore from dumps will use parallelism; limited by number of distinct tape volumes that the dump resides on; capped by TAPEUNITS
- ZPARMs to override tape options (DSNTIP6)
 - FROMDUMP ↔ RESTORE/RECOVER FROM DUMP
 - DUMPCLASS ↔ UTILS DUMP CLASS NAME
 - TAPEUNITS ↔ RESTORE TAPEUNITS

BACKUP/RESTORE SYSTEM Enhancements ...

- **A last primer... this time on Incremental FlashCopy**
 - A persistent relationship is established between two DASD devices
 - Incremental FlashCopy \neq DB2 Incremental IC
 - Initial incremental FC physically copies all tracks on the source volume
 - Subsequent incremental FC physically copy only the tracks that have changed on the source volume since the last incremental FC was taken (overriding the previous version)
 - The resulting backup is always a “full” backup
 - Requires the same number of target volumes as FlashCopy
 - Advantages
 - Minimizes I/O impact of the background physical copy
 - Considerably reduces elapsed time of the physical copy
 - A DASD volume can have only one incremental relationship
 - If a Copy Pool has more than 1 version then the remaining versions will be full backups

BACKUP/RESTORE SYSTEM Enhancements ...


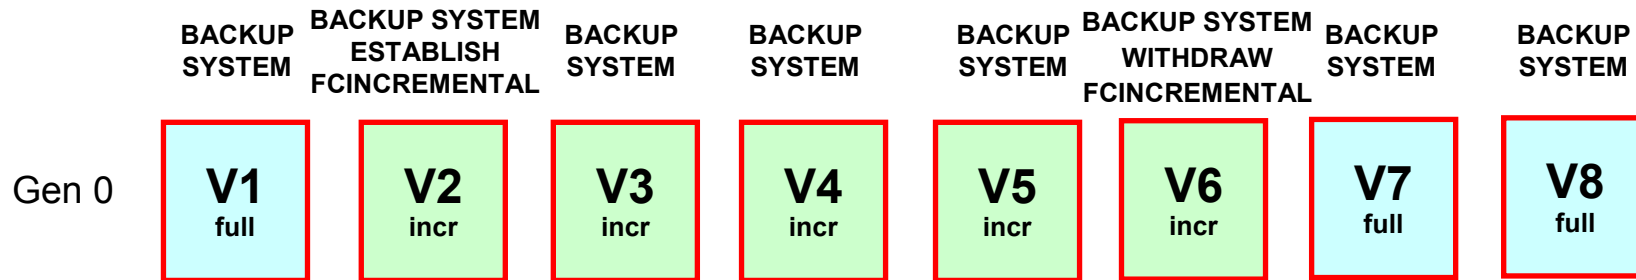
- Incremental FlashCopy support (Post GA via APAR PK41001)



- Requires z/OS DFSMSHsm V1R8
- Incremental FlashCopy can also be activated using HSM commands; `BACKUP SYSTEM` will then use it

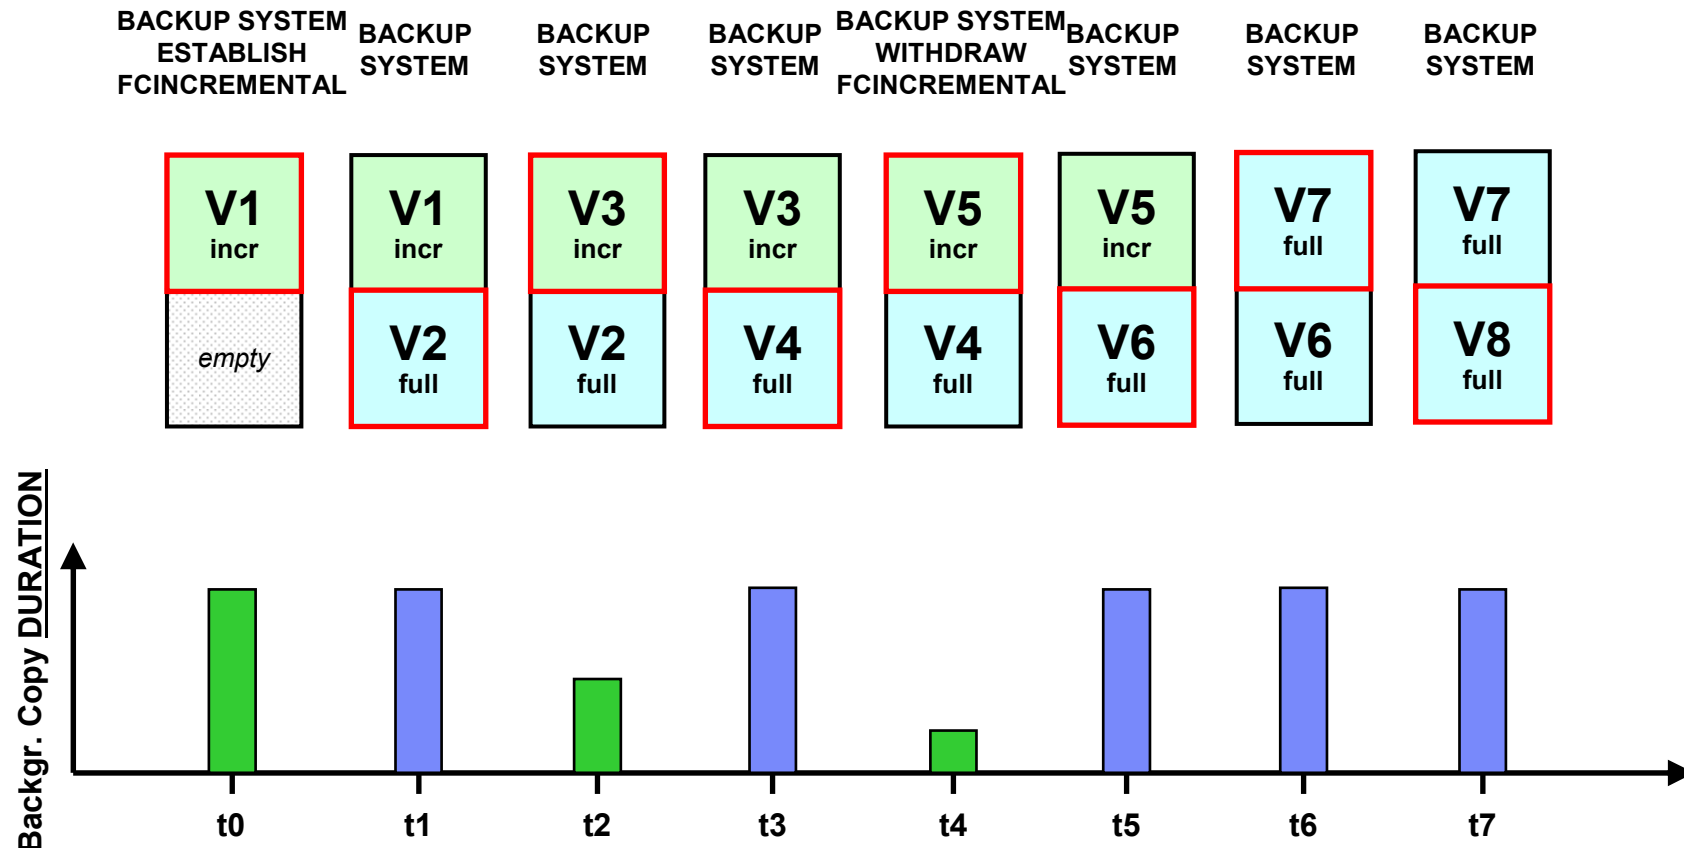
BACKUP/RESTORE SYSTEM Enhancements ...

- Incremental Flashcopy with a single version:



BACKUP/RESTORE SYSTEM Enhancements ...

- Incremental Flashcopy with two FC versions:



RECOVER Enhancements ...

- **Ability to recover at the object level using system-level backups**
 - Requires z/OS DFSMSHsm V1R8
 - 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
 - New ZPARM SYSTEM_LEVEL_BACKUPS=YES
 - Some restrictions
 - “Traditional” copy (instead of FlashCopy) will be used if
 - The physical background copy is not complete
 - The object to be recovered is located on an XRC primary volume
 - The object is located on a PPRC primary volume without the Preserve Mirror function
 - If underlying data sets are deleted or moved to a different set of volumes, then object-level recovery from volume-based backup is not possible
 - Recovery Expert for DB2 for z/OS can help overcome this restriction
- **RTS COPY... columns also updated for BACKUP SYSTEM**
- **No REPORT RECOVERY support at this time**

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 (PK44133)**
 - Uses additional real-time statistics
 - Improved recommendations for REORG, RUNSTATS, and COPY

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>