



# ***DB2 for z/OS Backup and Recovery Update - V9 and V10***

**James Teng, Ph.D.**

**Distinguished Engineer**

**IBM Silicon Valley Laboratory**

**August 9, 2011**

# Disclaimer



© Copyright IBM Corporation [current year]. 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, DB2 and DB2 for z/OS 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](http://www.ibm.com/legal/copytrade.shtml)

Other company, product, or service names may be trademarks or service marks of others.



# Agenda




- *DB2 Managed FlashCopy Solution in V8*
  - f *DB2 System Level Backup using Copy Pools*
  - f *Backup System and Restore System Utilities*
- *Backup and Recovery Enhancements in DB2 9*
  - f *Automatically manage Copy Pool backups to tapes*
  - f *Allow table space recovery using System Level Backups*
  - f *Incremental FlashCopy*
- *Backup and Recovery Enhancements in DB2 10*
- *Copy Pool Enhancements in z/OS 1.12*

# DB2 Managed FlashCopy Solution



- *Provide an easier and less disruptive way for fast volume-level backup and recovery*
  - f *Use FlashCopy to backup DB2 data and logs*
  - f *No longer need to suspend logs*
  - f *Backups are managed by DB2 and DFSSMShsm to support system level PIT recovery, Disaster Recovery and Cloning*



# DB2 Managed FlashCopy Solution

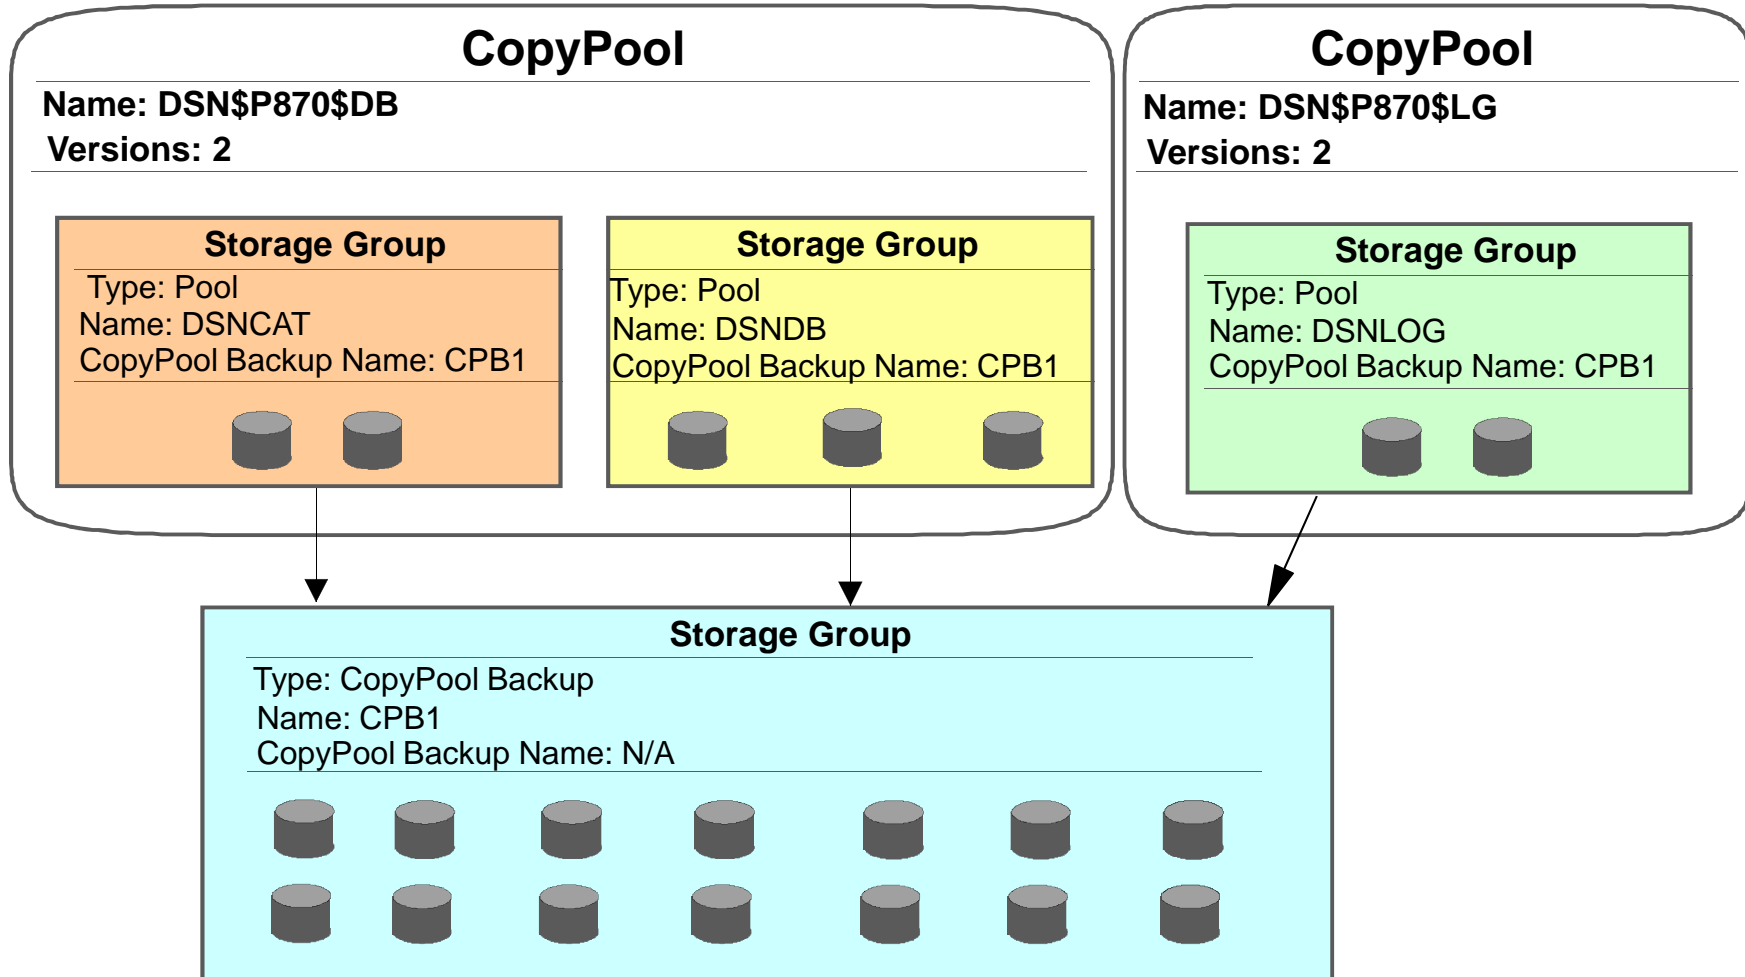
- *new utilities in DB2 for z/OS V8:*
  - f *BACKUP SYSTEM*
  - f *RESTORE SYSTEM*
  
- *Takes system-level copies of data and logs*
  - f *Exploits SMS Copy Pool*
  - f *DB2 data and logs must be SMS-managed*
  - f *Write log activity is NOT suspended*
  - f *Suspends data set creation, deletion, rename, and extend operations*

# SMS COPYPOOL



- *SMS construct*
- *Set of SMS storage groups - maximum 256*
- *Has a VERSIONS attribute - maximum 85*
- *Each DB2 system has two SMS COPYPOOLS*
  - f *DATA COPYPOOL (DSN\$location\_name\$DB)*
  - f *LOG COPYPOOL (DSN\$location\_name\$LG)*
  
- *Copy Pool Backup*
  - f *New storage group type*
  - f *Used to hold volume copies of DASD defined in the COPYPOOL*

# SMS COPYPOOL - Example



# Backup and Recovery Enhancements in DB2 9




# DB2 9 Enhancements

- **BACKUP SYSTEM** and **RESTORE SYSTEM** utilities will manage Copy Pool backups to/from tapes
  - f Up to five tape copies can be created
- Support Incremental FlashCopy
- **RECOVER** utility can use Copy Pool backups as the source for recovery of DB2 tables and indexes
  - f Backups can be on DASDs or Tapes
- Prereq. DFSMSHsm and z/OS 1.8

# CopyPool Backups on Tapes



- *Customer benefits – Reduce disk space for maintaining multiple Copy Pool backups*



- *Integrated tape management between DB2 and DFSMSHsm*
- *Retaining Copy Pool backups for long term use*
- *Providing a means of recovery from media failure*
- *Remote site recovery*


# ***RESTORE SYSTEM – Tape Support***

- *Restore the DB copy pool from tapes in parallel*
- *If the data copy pool backup resides on DASD and on tape, then the DASD version is chosen as the recovery base*
- *Install ZPARM options*
  - f *FROMDUMP - user can specify that they don't want to use the DASD version*
  - f *DUMPCCLASS (dc) - user can specify a specify HSM dump class to restore from*

## **RECOVER utility – use Copy Pool backup**

- *RECOVER utility enhancements permit using a backup taken at the system-level as the recovery base for a subset of objects in the system*
- *Need to set ZPARM System\_Level\_Backups = YES*

### **Connecting system-level backups with object level recoveries**



## **RECOVER utility – use Copy Pool backup ...**



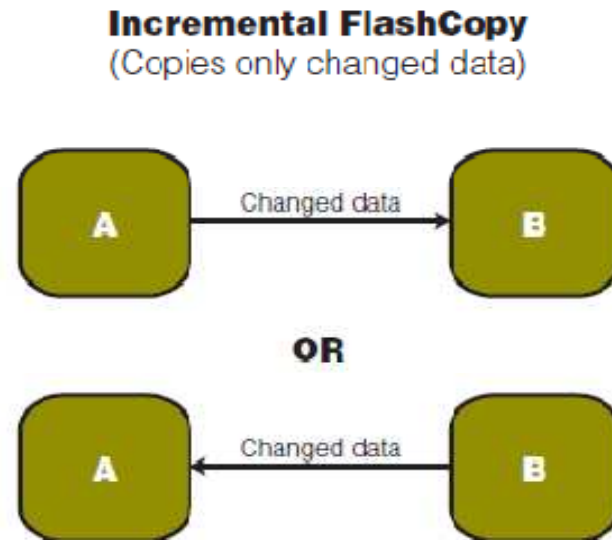

- *Most recent recovery base (prior to the recovery point) is chosen:*
  - f *could be image copy, concurrent copy, log yes event, or Copy Pool backup*
- *Takes sub-second to restore a data set if the backup is on DASD (independent of its size)*
- *Use normal I/O (i.e. not FlashCopy)*
  - f *If FlashCopy background copy is not complete*
  - f *If the production volume is the source of PPRC/XRC*
    - *Use FlashCopy with Remote Pair Flash Copy on DS8K 4.2.*
  - f *Make sure SETSYS FASTREPLICATION(PREFERRED) is used, which is the default*
- *When restoring a list of objects, the restore process is done in parallel*

## ***RECOVER utility – use Copy Pool backup ...***

- *Data set must be cataloged and allocated on the same volumes that it resided on at the time of the backup*
  - f *Support for data sets that have extended to new volumes*
  - f *DB2 Recovery Expert V2 can alleviate this problem*
  - f *DFSMS support is in z/OS 1.11 (require DB2 apar PK92725)*
    - *SMS option to capture ICFCTLG for Copy Pool*
    - *Allow recovery for moved/deleted data sets*
    - *Still need to have sufficient free space on the original DASD volume*
  
- *If the restore of datasets from DASD fails, then the recovery of the object will not proceed*
  - f *Use the RECOVER RESTOREBEFORE option to direct the utility to use a recovery base prior to the system-level backup*
  
- *If FROMDUMP is specified:*
  - f *Data sets are restored from tapes*

# Incremental Flash Copy

- *Introduced by DFSMSHsm in z/OS 1.8*
  - f *Initial incremental FlashCopy creates full base backup*
  - f *Subsequent incr. FlashCopies copy changed tracks to backup volumes only (overriding initial backup)*
- *Minimizes I/O impact (vs DB2 incremental IC)*
- *Considerably reduces elapsed time of physical copy*



# ***RECOVER to PIT with consistency***



*DB2 RECOVER utility will:*

- f Automatically detect the uncommitted transactions running at the recover PIT*
  - f Roll back their changes on the recovered objects.*
  - f Thus ensuring data consistency after PIT recoveries.*
  - f Recovered objects left in a transaction consistent state.*
- 
- *Avoid the need to regularly run the QUIESCE utility*
    - f Reduces disruption to DB2 users and applications*



# Backup and Recovery Enhancements in DB2 10

# ***DB2 10: Copy/Recover Enhancements***


- Improved COPY CHANGELIMIT performance
  - Use RTS instead of scanning space map pages
- Dataset level FlashCopy option
  - Inline COPY is optional for REORG (post-GA enhancement)
- FlashCopy backups with consistency and no application outage
- FlashCopy backups as input to:
  - RECOVER (fast restore phase)
  - COPYTOCOPY
    - Create sequential copies from FlashCopy
    - Can be used as input for UNLOAD
  - DSN1COPY/DSN1PRNT
- RECOVER “back to” log point
- REPORT RECOVERY support for system level backups
- MODIFY RECOVERY improved performance

# V10: 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)




# V10: COPY




- Post-GA Flashcopy enhancements – PM34776
  - Support for REORG AUX YES to FlashCopy aux objects
  - Remove requirement for sequential inline copy for REORG
  - Add a new zparm parameter FLASHCOPY\_PPRC which controls whether FlashCopy can be taken against PPRC primary volumes
    - REQUIRED - Require DS8K Remote Pair FlashCopy function to take FlashCopy
    - PREFERRED - Allow to take FlashCopy even if DS8K RPFC is not available (volume pairs will be in GDPS duplex pending state when the background copy is in progress)
    - NONE - FlashCopy does not preserve mirror

CPU time per object (z10)



Elapsed time per object (z10)



# V10: COPY



- Improved dataset management & performance
  - CHANGELIMIT will not allocate copy dataset unless copy taken
    - &ICTYPE now matches actual image copy type
  - Use RTS for CHANGELIMIT performance
    - No longer need to scan space map pages
  - Incremental copy will not allocate copy dataset unless pages changed
  - Insert dummy SYSCOPY record for incremental copy even though no pages changed
  - Delivery post-GA

# V10: RECOVER



- New BACKOUT YES option for point in time recovery
  - True rollback, not run of generated SQL undo statements
  - Requires COPY YES for indexes
  - Not for media recovery
- Allow fast data restore from data set level FlashCopy
- New ZPARAM REC\_FASTREPLICATION parameter
  - PREFERRED (default) – use FlashCopy if available
  - REQUIRED – use FlashCopy for recovery
  - NONE – use standard I/O for recovery
- 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)

# PPRC/XRC Restrictions

- *RESTORE SYSTEM & RECOVER from SLB fails if target is a primary in a PPRC relationship or part of an XRC relationship*
  - f *DFSMSHsm APAR OA23849 permits FlashCopy to a PPRC primary*
  - f *But volumes will be in “duplex pending” state until background copy completes*
    - *GDPS / Hyperswap failover fails for volumes in Duplex Pending*
  - f *Resolve this issue by shipping function rather than data through PPRC (DS8K FlashCopy Preserve Mirror solution)*
  - f *No XRC solution in the short term*

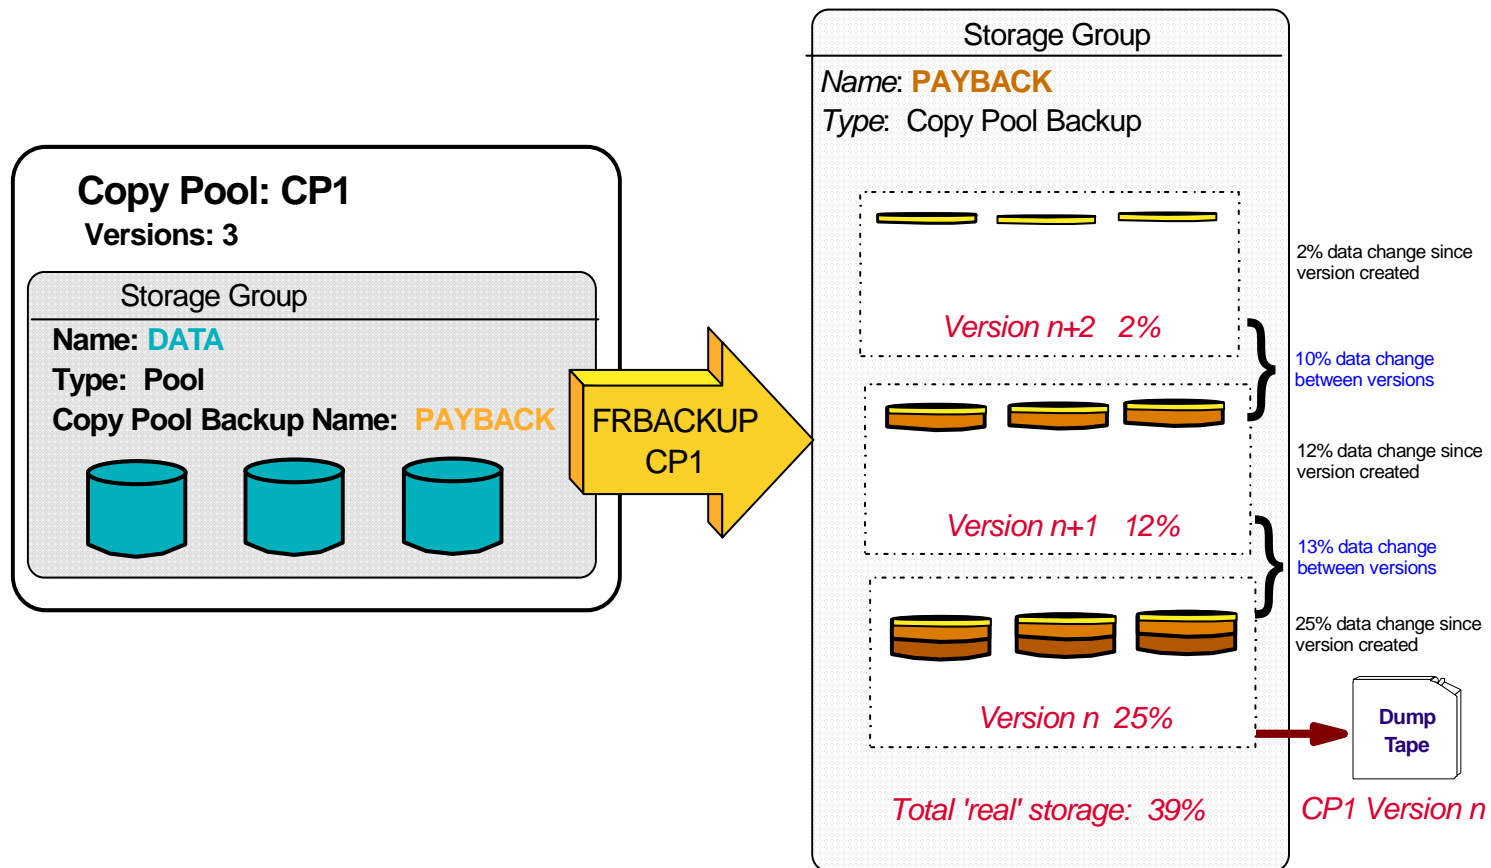
# HSM Enhancements in z/OS 1.12



- *QUERU COPYPOOL* command to display background copy percent-complete information
- *Allow RESTORE to use System-level backups without waiting for FlashCopy background copy to complete*
  - f *Use DS8K Fast Reverse Restore (i.e. FRR) feature*
  - f *New SMS Copy Pool setting to enable FRR*
  - f *Source volumes can't be in any other FlashCopy relationships*
- *LIST COPYPOOL* command to display whether FRR was enabled and the recovery status of using FRR



# Multiple Space Efficient Versions



# HSM Enhancements in z/OS 1.12



- *Allow Backup System to use Space Efficient FlashCopy to keep backups on tapes*
  - *CopyPool with Version = 0*
- *Space Efficient Volumes*
  - *Not consume physical space until writes take place*
  - *Requires FlashCopy SE licensed feature*
  - *HSM selects SE target volumes when Copy Pool is defined with NOCOPY (VERSIONS = 0)*
- *If FRR is enabled, recovery of the entire Copy Pool can be performed from the disk backup (in addition to a dump tape version)*

# Summary

- *A fast and non-disruptive backup solution using
  - f *FlashCopy and DB2 Backup System Utility*
  - f *Support Incremental FlashCopy**
- *Copy Pool backups can be used as the source for DB2 table/index recovery*
- *Automatically manage Copy Pool backups on tapes*
- *Restore System Utility can recover DB2 system from DASD or Tapes*
- *Recover Utility can recover table space and index to PIT with transaction level consistency*
- *Remote Pair FlashCopy to alleviate PPRC restriction*
- *Tight collaboration between DB2 and Storage teams*

*Thank You !*