



IBM Software Group

DB2 for z/OS Backup and Recovery Strategy

Dr. Jim Teng
IBM Distinguished Engineer



DB2 for z/OS Technical Conference
October 5-6, 2009
Taipei, Taiwan


 ON DEMAND BUSINESS™

©2009 IBM Corporation



Disclaimers & Trademarks*

Information in this presentation about IBM's future plans reflect current thinking and is subject to change at IBM's business discretion. You should not rely on such information to make business plans. Any discussion of OEM products is based upon information which has been publicly available and is subject to change. The opinions expressed are those of the presenter at the time, not necessarily the current opinion and certainly not that of the company.



The following terms are trademarks or registered trademarks of the IBM Corporation in the United States and/or other countries: AIX, AS/400, DATABASE 2, DB2*, Enterprise Storage Server, ESCON*, IBM, iSeries, Lotus, NOTES, OS/400, pSeries, RISC, WebSphere, xSeries, z/Architecture, z/OS, zSeries, System p, System I, System z

The following terms are trademarks or registered trademarks of the Microsoft Corporation in the United States and/or other countries: MICROSOFT, WINDOWS, ODBC

For more copyright & trademark information see ibm.com/legal/copytrade.phtml




Agenda

- *DB2 Managed FlashCopy Solution in V8*
 - ▶ *DB2 System Level Backup using Copy Pools*
 - ▶ *Backup System and Restore System Utilities*
 - ▶ *FlashCopy Preserve Mirror in DS8K 4.2*
- *What's New in DB2 9*
 - ▶ *Automatically manage Copy Pool backups to tapes*
 - ▶ *Allow table space recovery using System Level Backups*
 - ▶ *Incremental FlashCopy*
- *Possible Future Enhancements*



DB2 Managed FlashCopy Solution in V8

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





DB2 Managed FlashCopy Solution in V8 ...

- *new utilities in DB2 for z/OS V8:*
 - ▶ *BACKUP SYSTEM*
 - ▶ *RESTORE SYSTEM*

- *Takes system-level copies of data and logs*
 - ▶ *Exploits SMS Copy Pool*
 - ▶ *DB2 data and logs must be SMS-managed*
 - ▶ *Write log activity is NOT suspended*
 - ▶ *Suspends data set creation, deletion, rename, and extend operations*



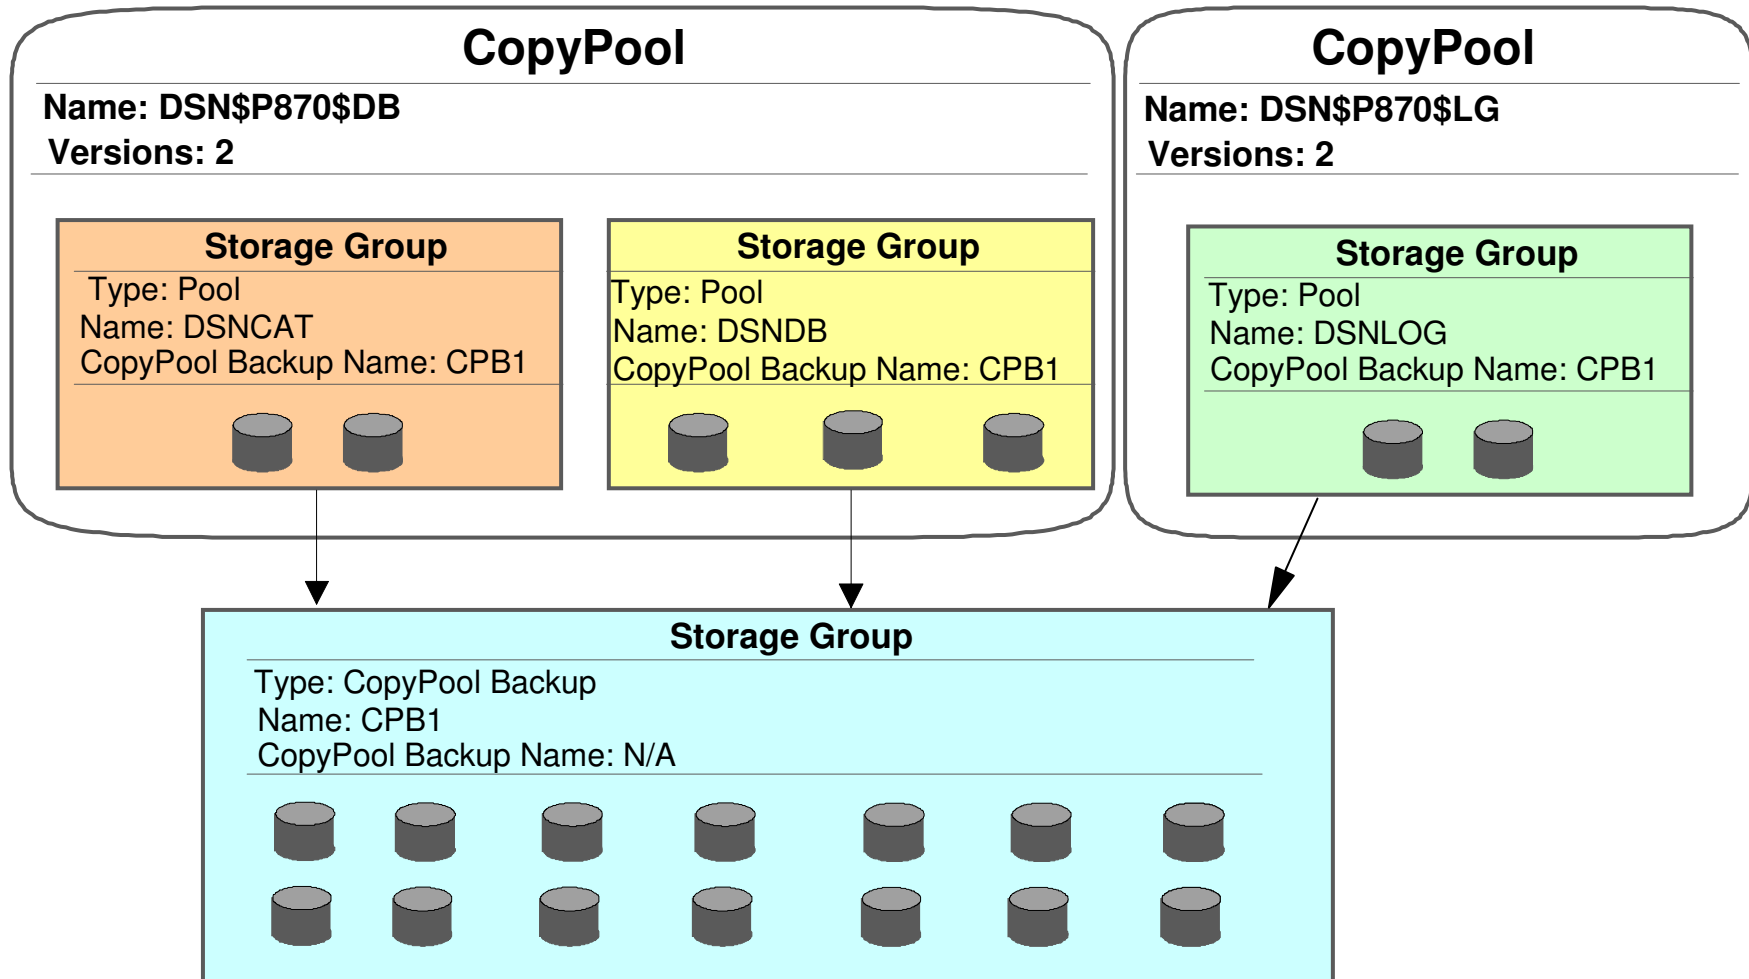
SMS COPYPOOL

- *SMS construct*
- *Set of SMS storage groups - maximum 256*
- *Has a VERSIONS attribute - maximum 85*
- *Each DB2 system has two SMS COPYPOOLS*
 - ▶ *DATA COPYPOOL (DSN\$location_name\$DB)*
 - ▶ *LOG COPYPOOL (DSN\$location_name\$LG)*

- *Copy Pool Backup*
 - ▶ *New storage group type*
 - ▶ *Used to hold volume copies of DASD defined in the COPYPOOL*



SMS COPYPOOL - Example





Preparation of the DB2 Environment

- *DB2 logs and BSDS must be separated from the rest of the DB2 data sets*
 - *Own pool of volumes*
 - *Own ICF user catalog*
 - *Defined in a separate SMS storage group*
- *Special care should be taken to ensure that the ICF user catalogs stay synchronized with the data*
 - *Must have separate ICF user catalogs for data and logs*
 - *ICF user catalog(s) for data must reside in the data copy pool*
 - *ICF user catalog for logs must reside in the log copy pool*
 - *Do not share ICF user catalog with non-DB2 data*



Backup System Process

- *No DB2 quiesce point is required, nothing stops as in SET LOG SUSPEND*
- *Invokes DFSMSHsm (FRBACKUP) to take fast volume copies of the DB2 data and/or logs*
- *Record the Recover Base Log Point (RBLP) in DBD01*
 - *Starting point for RESTORE SYSTEM to apply logs*
- *Two flavours:*
 - *BACKUP SYSTEM FULL*
 - *Allow recovery of the entire system in later stage*
 - *Have to define the database and log copy pools*
 - *Backup both data and then log (active logs and BSDS)*
 - *BACKUP SYSTEM DATA ONLY*
 - *Only database copy pool has to be defined for database backup*



Point-in-time Recovery of the Entire System

- *To recover the system to the time of the backup:*
 - *Use the DFSMSHsm FRRECOV command to restore the database and log copy pools*
`FRRECOV COPYPOOL(DSN$locn$DB) VERIFY(YES) TOKEN(token)`
`FRRECOV COPYPOOL(DSN$locn$LG) VERIFY(YES) TOKEN(token)`
 - *Start DB2 and inflight URs are backed out*
 - *If active logs are striped, need to use conditional restart to truncate logs using the “data complete lrsn” in BSDS or in DSNU1614I message (issued by the BACKUP SYSTEM utility)*
- *RESTORE SYSTEM allows you to recover the entire subsystem or data sharing group to an arbitrary point in time after the time of the backup (and up to the end of the logs)*



Restore System Utility

- Establish the 'PITR' conditional restart record
 - ▶ CRESTART CREATE SYSPITR=log-point
 - ▶ Allow to specify “FF...F” in V9
 - ▶ Each active data sharing member must have SYSPITR created
- Start DB2 with a PITR CRCR
 - ▶ Implicitly apply DEFER ALL, FWD=NO, Access(Maint)
 - ▶ Install SYSADM authority required
 - ▶ Can only run “Restore System” utility
- Write logs to rollback uncommitted changes
- Make sure ICFCTLG data sets are not opened
 - ▶ F CATALOG,OPEN to query status
 - ▶ F CATALOG,DEALLOCATE to close



Restore System Utility ...

- Run the RESTORE SYSTEM utility
 - ▶ Unallocate ICFCTLG - z/OS 1.11
 - ▶ Restore Data CopyPool backup
 - ▶ Apply logs from RBLP to SYSPITR value of CRCR

- Objects not recovered from “log no” events
 - ▶ LOAD LOG NO, REORG, CREATE INDEX
 - ▶ RECP for table spaces and “Copy Yes” indexes
 - ▶ RBDP for “Copy No” indexes

- If there are objects in RECP/RBDP, the utility will complete with Return Code = 4



PPRC/XRC Restrictions

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



FlashCopy Preserve Mirror

- *Require DS8K 4.2 – Remote Pair FlashCopy*
 - ▶ *Available on 4/24/2009*
- *Prreq z/OS 1.8 with DFSMSHsm APAR OA24814 and DFSMSdss APAR OA24811*
- *ALLOWPPRCP(NO, YES, PMNO, PMPREF, PMREQ)*
 - ▶ *A new keyword on FRBACKUP and FRRECOV*
 - ▶ *Can also specify ALLOWPPRCP(FRRECOV, YES, NO, PMNO, PMPREF, PMREQ) on FRBACKUP*
 - ▶ *Options specified on FRBACKUP are saved for all subsequent use except for CopyPools with Versions 0*
 - ▶ *Ignored when recovering from tapes*
 - ▶ *Move to SMS CopyPool definition in z/OS 1.11*



What is New in DB2 9

- ***BACKUP SYSTEM** and **RESTORE SYSTEM** utilities will manage Copy Pool backups to/from tapes*
 - ▶ *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*
 - ▶ *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*
 - ▶ *FROMDUMP - user can specify that they don't want to use the DASD version*
 - ▶ *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:*
 - ▶ *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)*
 - ▶ *If FlashCopy background copy is not complete*
 - ▶ *If the production volume is the source of PPRC/XRC*
 - *Use FlashCopy with Remote Pair FlashCopy on DS8K 4.2.*
 - ▶ *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*
 - ▶ *Support for data sets that have extended to new volumes*
 - ▶ *DB2 Recovery Expert V2 can alleviate this problem*
 - ▶ *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*

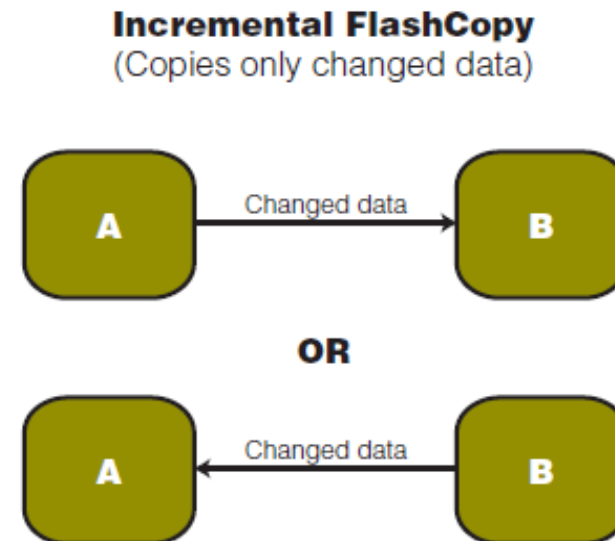
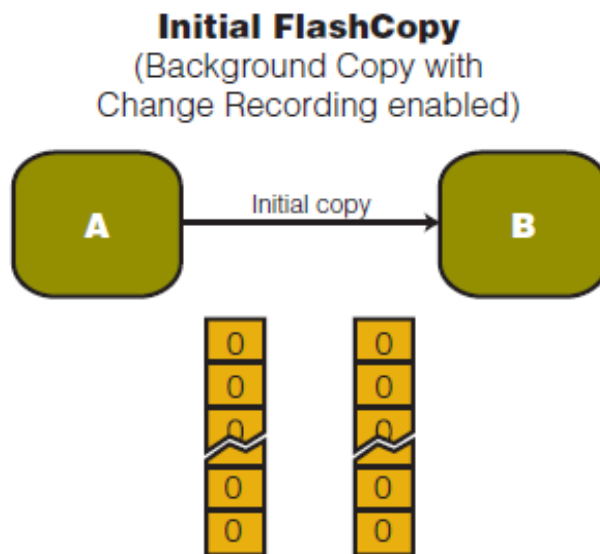
- *If the restore of datasets from DASD fails, then the recovery of the object will not proceed*
 - ▶ *Use the RECOVER RESTOREBEFORE option to direct the utility to use a recovery base prior to the system-level backup*

- *If FROMDUMP is specified:*
 - ▶ *Data sets are restored from tapes*



Incremental Flash Copy

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





RECOVER to PIT with consistency

Enhance DB2 RECOVER utility to:

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



Possible Future Enhancements

- *Consistent SHRLEVEL Change COPY*
 - ▶ *Will not quiesce applications*
 - ▶ *Need data set level FlashCopy capability*
- *Extend COPY and LOAD/REORG inline COPY to support backups using FlashCopy*
- *Extend RECOVER to support a point-in-time recovery via rollback using logs*
- *REPORT RECOVERY support for System-level backup*



Possible Future Enhancements ...

- *Allow RESTORE/RECOVER to use Sytem-level backups without waiting for FlashCopy background copy to complete*
 - ▶ *Use DS8K Fast Reverse Restore (i.e. FRR) feature*
 - ▶ *Need DS8K HW enhancement to allow FRR to work with Remote Pair FlashCopy*
- *Allow Backup System to use Space Efficient FlashCopy to keep backups on tapes*
- *Allows implicitly created table space to use TRACKMOD = NO (a new zparm parameter)*



Summary

- *A fast and non-disruptive backup solution using*
 - ▶ *FlashCopy and DB2 Backup System Utility*
 - ▶ *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