# What's New in DB2 9 for z/OS for Backup and Recovery

*Dr. Jim Teng*
*IBM Distinguished Engineer*

**DB2 for z/OS Technical Conference**
**Sept 05-07, 2007 - Shanghai, China**
**Sept 10-12, 2007 - Beijing, China**

## Disclaimers & Trademarks*

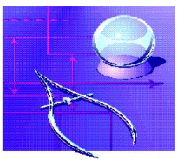

Information in this presentation about IBM's future plans reflect current thinking and is subject to change at IBM's business discretion. You should not rely on such information to make business plans. Any discussion of OEM products is based upon information which has been publicly available and is subject 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

DB2 for z/OS Technical Forum

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in the operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only. This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

# Agenda

- *DB2 Managed FlashCopy Solution in V8*
- *Copy Pool enhancements in V9*
  - *Automatically manage Copy Pool backups to tapes*
  - *Use Copy Pool backups as the source for DB2 table/index recovery*
  - *Incremental FlashCopy*
- *Restart Enhancements*
- *Extend RECOVER utility to support point-in-time recovery with transaction level consistency*
  - *Rollback uncommitted changes*

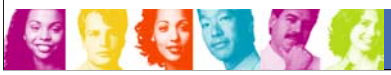

## DB2 Managed FlashCopy Solution in V8

- *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 DFSMShsm to support system level PIT recovery*

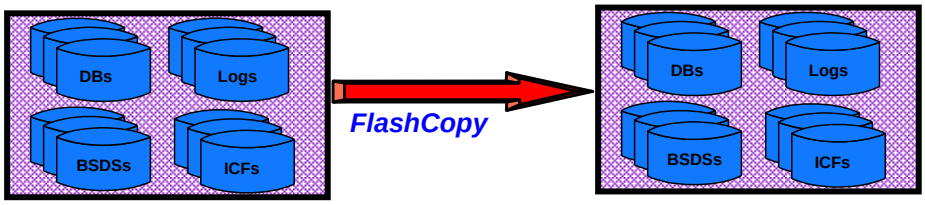

## DB2 Managed FlashCopy Solution in V8 ...

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

- Takes system-level copies of data and logs
  - Exploits SMS CopyPool in z/OS 1.5
  - DB2 data and logs must be SMS-managed
  - Write log activity is NOT suspended
  - Suspends data set creation, deletion, rename, and extend operations

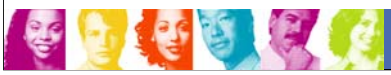

# FlashCopy

- *A relationship is established between a source volume and a target volume*
- *The copy is considered successful when the relationship is established.*

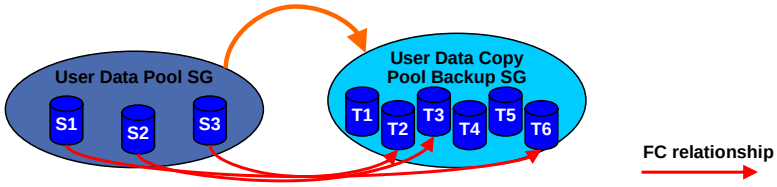


- *A background copy is then started which will result in a target volume that looks like the source volume when the relationship was established.*
- *The relationship goes away when the background copy completes.*

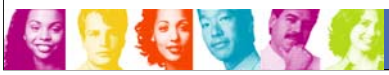

# COPYPOOL

- *SMS construct in z/OS 1.5*
- *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*

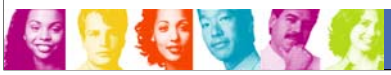

# CopyPool

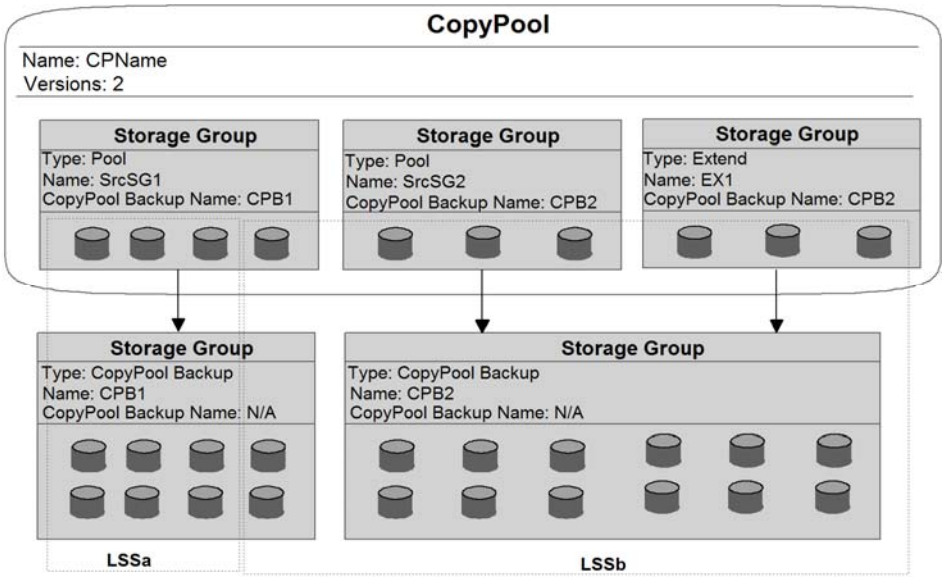
CopyPool
Name: CPName
Versions: 2
Storage Group
Type: Pool
Name: SrcSG1
CopyPool Backup Name: CPB1
Storage Group
Type: Pool
Name: SrcSG2
CopyPool Backup Name: CPB2
Storage Group
Type: Extend
Name: EX1
CopyPool Backup Name: CPB2
Storage Group
Type: CopyPool Backup
Name: CPB1
CopyPool Backup Name: N/A
Storage Group
Type: CopyPool Backup
Name: CPB2
CopyPool Backup Name: N/A
LSSa
LSSb

# BACKUP SYSTEM

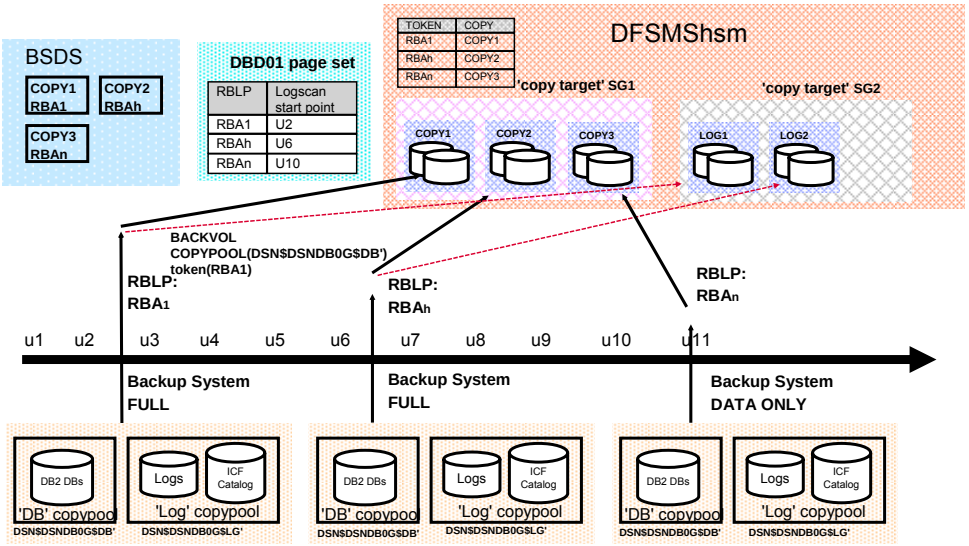

# RESTORE SYSTEM

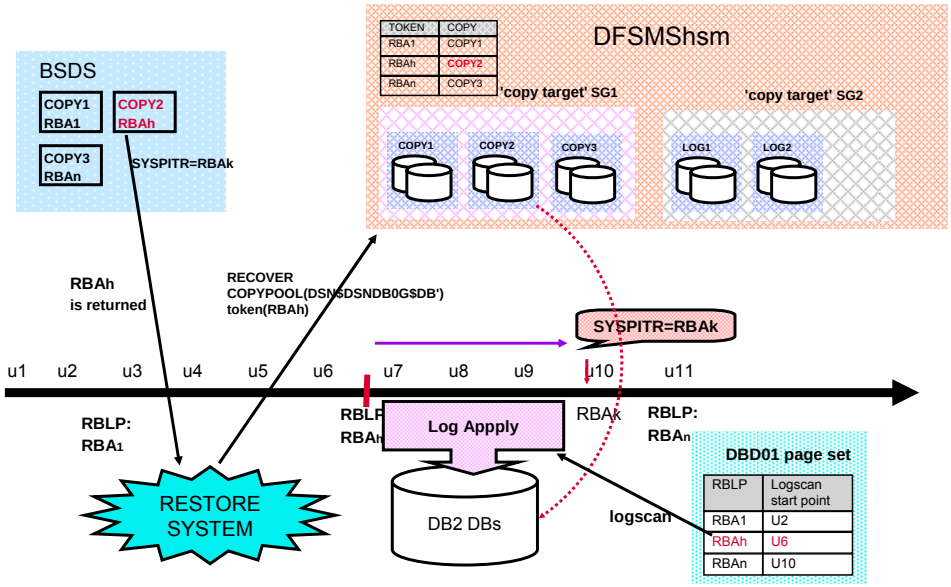

# Copy Pool Enhancements in Version 9

- *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*

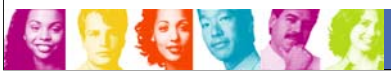

# CopyPool Backups on Tapes

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

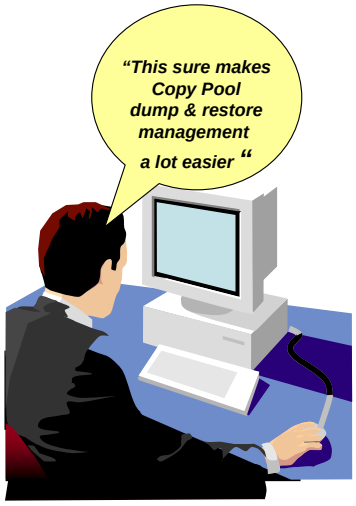


- *Integrated tape management between DB2 and DFSMShsm*
- *Retaining Copy Pool backups for long term use*
- *Providing a means of recovery from media failure*
- *Remote site recovery*

# BACKUP SYSTEM syntax for Tape Support

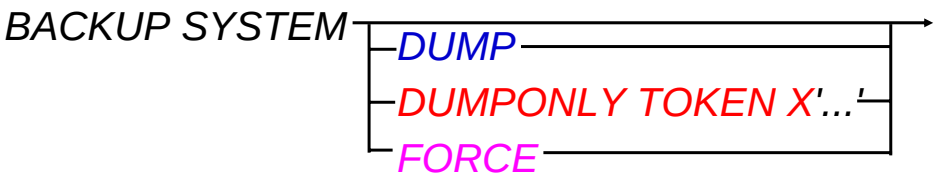


- *DUMP - initiate dump processing to tape when the copy pool(s) is logically complete*
- *DUMPONLY TOKEN – create a dump on tape for an existing Copy Pool backup on DASD or "restart" the dump processing if the dump has failed. TOKEN is optional – if not specified, dump the most recent Copy Pool backup to tapes.*
- *FORCE - allow to overlay the oldest Copy Pool backup whose required dumps have not completed*
- *DUMPCLASS sub option also available on both DUMP and DUMPONLY*

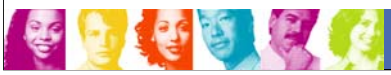

# *BACKUP SYSTEM Utility – Tape Support*

- *DUMP - Automatic dump to tapes when the copy pool backup to DASD are logically complete*
- *Up to five DFSMShsm dump classes may be specified.*
- *Optionally, invoke the BACKUP SYSTEM Utility twice:*
  - Once without the DUMP keyword to initiate the copy pool backup on DASD
  - Then a second time with the DUMPONLY keyword to initiate the dump processing to tape.

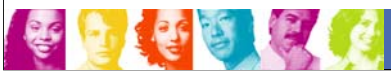

# RESTORE SYSTEM – Tape Support

- *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*
  - *DUMPCLASS (dc) - user can specify a specify HSM dump class to restore from*

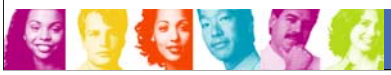

# RESTORE SYSTEM syntax

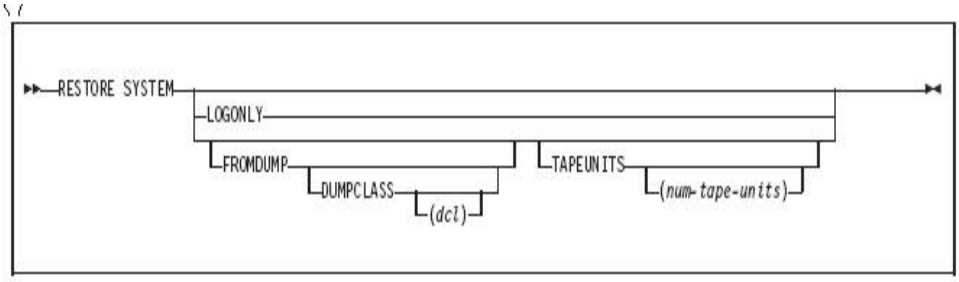


- *FROMDUMP – Restore only requires dumps on tape of the DB copy pool*
- *TAPEUNITS – Specifies the tape drives required during restore of DB copy pool*

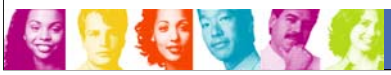

# RECOVER utility – use Copy Pool backup

- *RECOVER utility enhancements permit using a backup taken at the system-level as the recovery base for a subset of objects in the system*

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

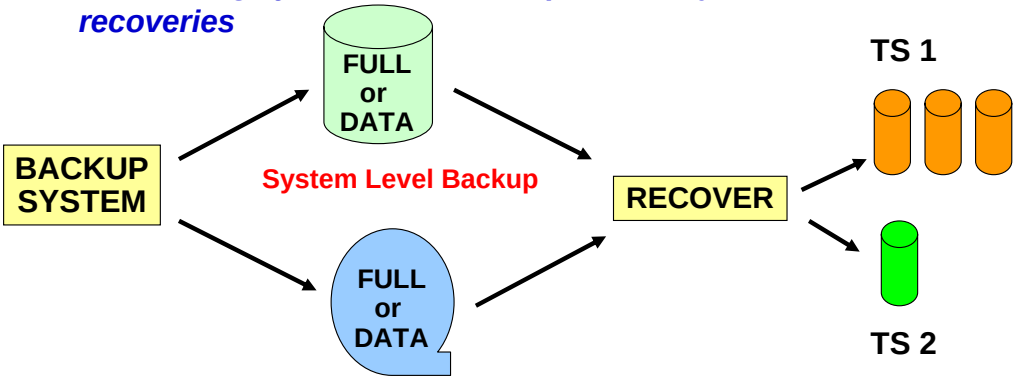

## *RECOVER utility – use Copy Pool backup ...*

- *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)*
- *If the FlashCopy background copy is not complete, normal I/O is used*
- *When restoring a list of objects, the restore process is done in parallel*

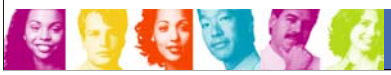

## RECOVER utility – use Copy Pool backup ...

- *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*
- *If the restore of datasets from DASD fails, then the recovery of the object will not proceed*
- *If FROMDUMP is specified:*
  - *Data sets are restored from tapes*

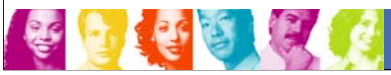



like the RECOVER

# RECOVER utility – use Copy Pool backup ...

- *PIT Recovery restrictions:*
  - *if tablespace/indexspace has "moved" since the copy pool backup, then the copy pool backup cannot be used as the recovery base*
  - *find/use the previous recovery base*
  - *utilities that can "move" objects*
    - *reorg, load replace, recover from image copy, rebuild*
    - *Recommend to use REUSE option to avoid “move” objects*

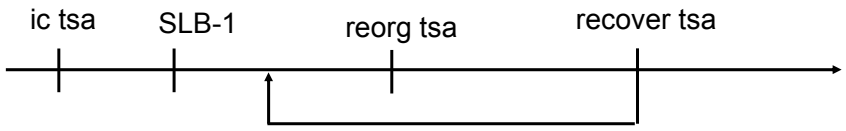


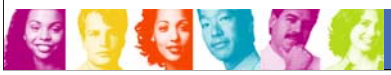

# Incremental FlashCopy

- *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*
- *Considerably reduces elapsed time of physical copy*

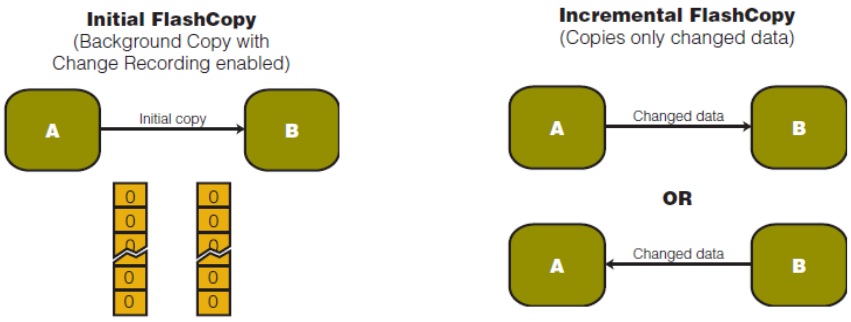

# Incremental FlashCopy - Notes

- *A persistent relationship is established between two DASD devices*
- *All tracks on the source volume are considered to be changed when the relationship is established so all tracks are copied.*
- *Subsequent incremental copies will copy only the tracks that have changed on the source volume since the last copy was taken*
- *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*

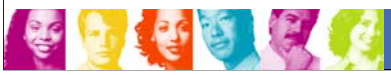

# BACKUP SYSTEM – New Keywords

- *ESTABLISH FCINCREMENTAL*
  - *Establish a persistent incremental FlashCopy relationship for a DB Copy Pool version and take a full copy*
  - *Use this keyword once for a DB copy pool*
  - *If the version being replaced is an incremental FlashCopy, only copy changed data on the source volume since the last copy was taken*
  - *The keyword is ignored if the version being replaced is an incremental FlashCopy*
  - *If the version being replaced is for a full FlashCopy, DSNU1630I with RC8 will be issued*
- *END FCINCREMENTAL*
  - *Take the last incremental copy and withdraw the FlashCopy relationship*

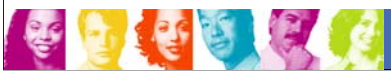

# Incremental FlashCopy with two generations

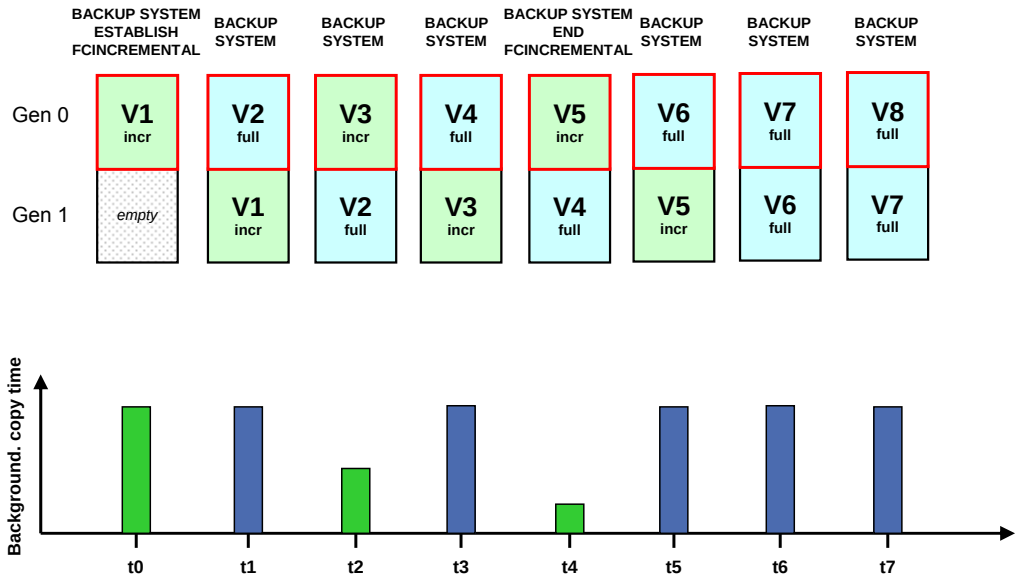

# DSNTIP6 – A new utilities install panel

**The entries on this panel affect the execution of DB2 Utilities.**

```
DSNTIP6          INSTALL DB2 - DB2 UTILITIES PARAMETERS
===>
Enter system-level backup options for RESTORE SYSTEM and RECOVER below:

| 1 SYSTEM-LEVEL BACKUPS  ===> NO       As a recovery base: YES or NO
| 2 RESTORE/RECOVER       ===> NO       From dump (NO or YES)
| 3 DUMP CLASS NAME       ===>          RESTORE/RECOVER from dump
| 4 MAXIMUM TAPE UNITS    ===> NOLIMIT  For RESTORE SYSTEM: NOLIMIT or 1-255

Enter other DB2 Utilities options below:
  5 TEMPORARY DS UNIT NAME===> SYSDA    Device for temporary utility data sets
  6 UTILITY CACHE OPTION    ===> NO     3990 storage for DB2 utility IO
  7 STATISTICS HISTORY      ===> NONE   Default for collection of stats history
  8 STATISTICS ROLLUP       ===> NO     Allow statistics aggregation. NO or YES
  9 UTILITY TIMEOUT         ===> 6      Utility wait time multiplier

PRESS:    ENTER to continue    RETURN to exit    HELP for more information
```

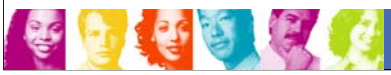

# Data Sharing Restart Enhancements

***Improve performance, usability and availability During DB2 restart.***

- *Performance*
  - *Avoid acquiring certain locks for GBP dependent objects*
  - *Open the objects involved in restart as early as possible.*
- *Usability*
  - *Automatically initiating the GBP Recovery (GRECP) of the GRECP objects at the end of restart.*
- *Availability*
  - *Supporting table level granularity of the retain locks for postponed abort Unit of Recovery*

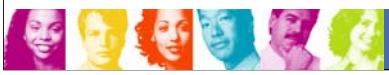

# PIT Recovery today

- *Recovering data to a prior time, and not to the present, are referred to as "Point in time recoveries"*
- *Today's options on the RECOVER utility include:*
  - *TOCOPY*
  - *TOLOGPOINT*
  - *TOLASTCOPY*
  - *TORBA*
  - *TOLASTFULLCOPY*

**Each terminates recovery at a Specific Point**

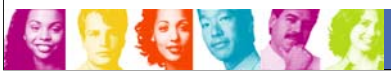

# Possible PIT problems faced today

- *PIT Recovery could cause a data inconsistency problem*
  - *– Point recovered to is NOT a consistent point*
  - *– Remember there is no backout process of the inflight URs*
- *Today we recommend taking QUIESCE points*
  - *– For later recovery to PIT with consistency*
- *Downside to running the QUIESCE utility*
  - *– Blocks applications in high volume system*
  - *– Deadlock between QUIESCE and SQL applications*
  - *– Unwanted overhead on production systems – from frequent running*
- *In reality, many PIT recoveries must be done to unplanned points in time*

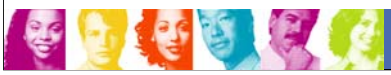

# RECOVER to PIT with consistency

*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*

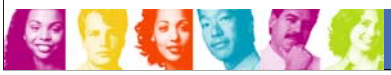

# RECOVER to PIT with consistency

- *RECOVER to PIT with consistency will become the behavior for:*
  - *RECOVER TORBA*
  - *RECOVER TOLOGPOINT .*

- *For RECOVER TOCOPY, TOLASTCOPY and TOLASTFULLCOPY using SHARELEVEL CHANGE copy:*
  - *The behavior is still as V8*
  - *NO consistency is ensured.*
    - *Since we can specify multiple objects in the recover list*
    - *Difficult to determine which time should be the consistency point .*

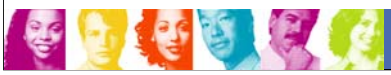

# New Log Analysis Phase - LOGCSR

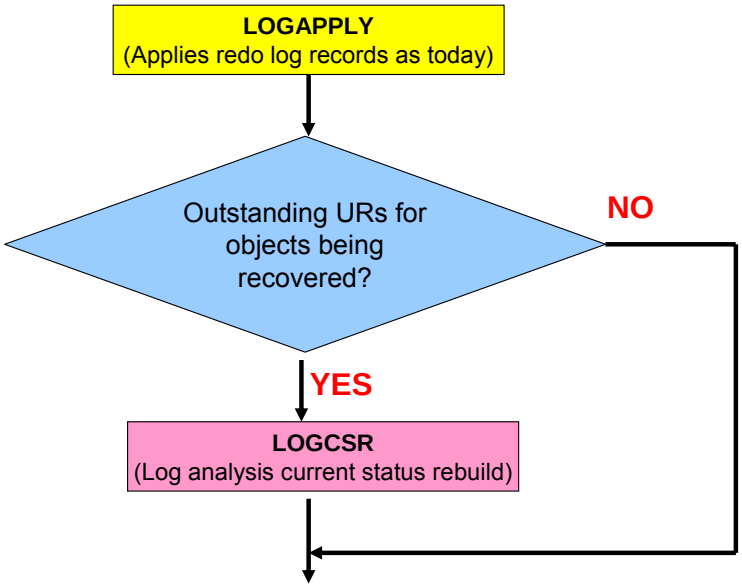

# LOGCSR processing

- ***During the LOGCSR phase, RECOVER identifies the URs that:***
  - *Were active( INFLIGHT, INABORT, INDOUBT or POSTPONED ABORT) during the recovery point.*
  - INDOUBT URs will be treated as INABORT
- ***Log analysis for each DB2 member***
  - DSNU1550I shows the member name whose log will be analyzed
  - DSNU1551I marks the end of log analysis for this member
  - DSNU1552I marking the end of LOGCSR phase

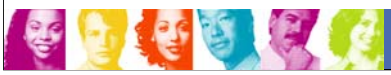

## *New LOGUNDO Phase*

- *LOGUNDO handles any UR whose changes to recovered objects need to be backed out.*
- *In a data sharing environment, RECOVER backs out the changes made on recovered objects for all members with active URs.*
- DSNU1554I marks the start of backout processing on each member
- DSNU1555I is issued periodically to report the progress of the RECOVER LOGUNDO phase

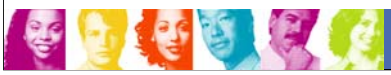

# Catalog changes

*SYSCOPY table PIT_RBA column will be used to store the finish time of share level change copy.*

- *Recover utility will use this value to locate the correct image copy to start recover if recover to certain timestamp .*

- *SYSCOPY table STYPE column will be used to remember if a recover to PIT was done with or without consistency.*
  - *ICTYPE is P, and STYPE is blank , means the recover to PIT was done without consistency*
  - *ICTYPE is P and STYPE is C, means recover to PIT was done with consistency .*

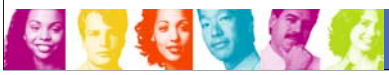

# MODIFY RECOVERY Utility

- *Delete SYSLGRNX records according to AGE or DATE criteria even if there are no SYSCOPY records to delete*
- *Insert a new SYSCOPY record ICTYPE=M, STYPE=R with START_RBA= highest RBA/LRSN of the SYSCOPY or SYSLGRNX records deleted*
- *Delete SYSCOPY records based on*
  - *RETAIN LAST n – keep last n image copies*
  - *RETAIN LOGLIMIT – based on archive logs in BSDS*
  - *RETAIN GDGLIMIT – based on the image copy GDG limit*

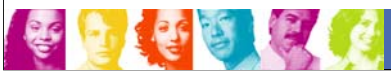

# Summary

- *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*
- *Automatic GRECP Recovery for Disaster Recovery*
- *Recover Utility can recover tables/indexes to PIT with transaction level consistency*

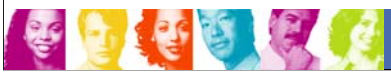

Session Title: What's New in DB2 9 for z/OS for Backup and Recovery

Session: 1272

*Dr. Jim Teng*
*IBM Distinguished Engineer*
*IBM Silicon Valley laboratory*
*jteng@us.ibm.com*