#SHAREorg

SHARE
Technology · Connections · Results

DB2 for z/OS System Level Backup Update

Judy Ruby-Brown
IBM Corp.
Advanced Technical Skills (ATS)

Wednesday, March 14, 2012
Session Number 10499

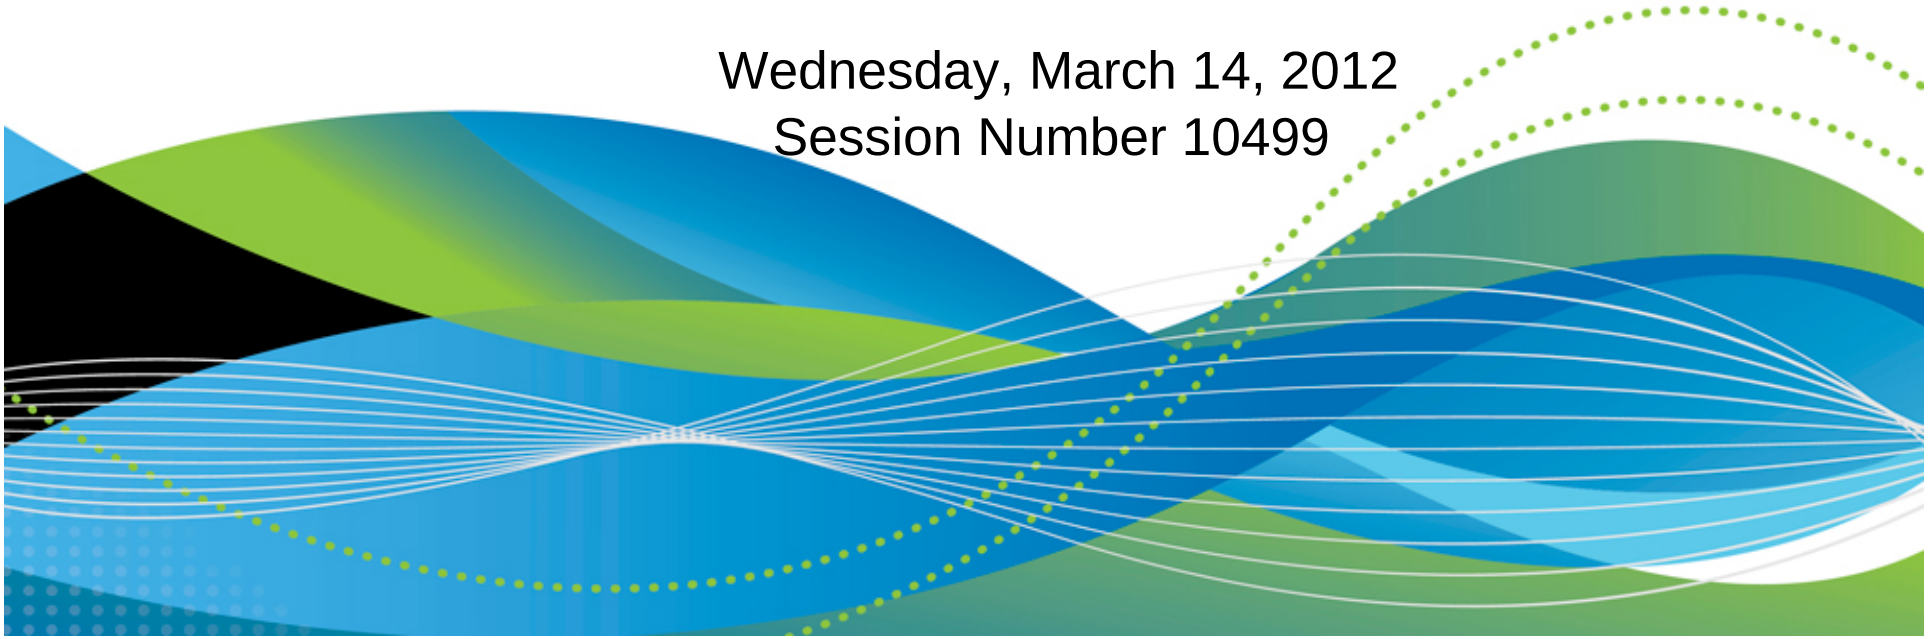

Notes

- System Level Backup, consisting of BACKUP SYSTEM and RESTORE SYSTEM utilities, was introduced in DB2 V8. Originally created for local recovery of an entire DB2 subsystem, it has become useful for offsite disaster recovery scenarios, as its advantages are ease of backup and restoration. We cover two of those.
- z/OS and DFSMShsm have added functionality through z/OS 1.12. The meaning of those enhancements for DB2 users are explained in this session in DB2 terms.
- We take a deeper dive into SLB and discuss some of the nuances.
- While SLB has attraction for customers using PTAM (“Pick up Truck Access Method”) to transport tape backups to the recovery site, enterprises now are embracing DASD mirroring. For some the driver is fear of local DASD failure. Those customers may use Peer to Peer Remote Copy (PPRC) Metro Mirror for all local DASD and invoke a Hyperswap capability that switches all DASD to secondary PPRC volumes within seconds in the same machine room.
- “Preserve mirror” capability allows SLB to be run at each site in parallel, within the site, not across the links between the sites. Duplexing of the links is retained, and there is no traffic on them as track updates are copied to the copy pool backup storage groups.
- Setting up the environment takes more time and is usually performed by the storage administrators. The requirements, how and where to enter the correct values are described in the Setup section, to provide DB2 administrators a brief exposure to them.
- Summary / final thoughts – should you abandon your image copy schedule in favor of SLB?

Disclaimer

Why Update?

Improvements to SLB functions since 2004 by

- DB2
- z/OS
- DFSMS
- storage

Agenda

1. Basic background of System Level Backup (SLB)
2. Good things z/OS® versions do for SLB
3. Setup
4. Summary

Slides titled "Notes" are for your reference and are not shown in the presentation

- This presentation is Part 2 of a two part presentation.
  - Its focus is on mirroring environments, especially those that have implemented PPRC Metro Mirror
- SLB is described in more detail in Part 1
  - DB2 for z/OS Disaster Recovery for the Rest of Us
    - Session 10501: Thursday, March 15, 6PM-7PM
  - Its focus is on SLB use for Offsite disaster recovery
- The SHARE attendee will benefit from downloading both parts

Terms used in this presentation

- SLB – System Level Backup (both Backup System/Restore System)
- FC – FlashCopy
- FCIC – FlashCopy image copies (DB2 10 only)
- FRR – Fast Reverse Restore
- SEFC – Space Efficient FlashCopy
- DFSMShsm – (HSM) Hierarchical storage manager
- DFSMSdss – (DSS) data set services
- FRRECOV – HSM command to restore the Copy Pools
- FRBACKUP – HSM command to back up the Copy Pools
- CP Backup SG (abbrev) – Copy Pool Backup Storage Group

More terms…

- DASD mirroring terminology
  - PPRC – Peer to Peer Remote Copy (DASD mirroring) (normally assumed Metro Mirror below)
  - PPRC Metro Mirror (MM) – synchronous PPRC
  - RPFC – Remote Pair FlashCopy (aka “preserve mirror”) applies to PPRC MM
  - FCTOPPRCP – FlashCopy to PPRC Primary
  - FRBACKUP keyword – Fast Replication Backup setting on ISMF Copy Pool panel for FCTOPPRCP
  - FRRECOV keyword – Fast Replication Recovery setting on ISMF Copy Pool panel for FCTOPPRCP
  - XRC – z/OS Global Mirror – eXtended Remote Copy
  - GM – Global Mirror
  - HyperSwap – switches all PPRC primary DASD to secondary in a few seconds – protects against control unit failure
  - TPC-R – Tivoli Productivity Center for Replication
  - GDPS – Geographically Dispersed Parallel Sysplex - IBM premier solution (services) using DASD mirroring to provide no data loss DR for enterprise. Supports HyperSwap and OEM DASD. Widely used among largest customers

References

- *DFSMShsm Fast Replication Technical Guide*, SG24-7069
- *IBM System Storage DS8000: Remote Pair FlashCopy (Preserve Mirror),* REDP-4504 (May, 2009)
- *z/OS V1R12*
  - *DFSMS Advanced Copy Services*, SC35-0428-18 (z/OS 1.12)
  - *DFSMSdss Storage Administration*, SC35-0423-14
  - *DFSMS Installation Exits*, SC26-7396-13
  - *DFSMSdfp Storage Administration*, SC26-7402-14
- DFSMS V1.12 Technical Update**,** SG24-7895-00
- *Using the Utilities Suite,* redbook, SG24-6289

Introduction to System Level Backup (SLB)

DB2 9 >

Background

- SLB introduced in DB2 V8 for local subsystem-wide recovery
  - BACKUP SYSTEM utility
  - RESTORE SYSTEM utility

- Replaced manual steps to obtain a consistent and “instant” backup of data, BSDS, active logs, catalog/directory taken while DB2 is active
  - -Set Log Suspend operator command
  - Manual submission of FlashCopy or OEM “instant copy” jobs
  - Tracking of completion of jobs
  - -Set Log Resume operator command

FlashCopy® - What it Does

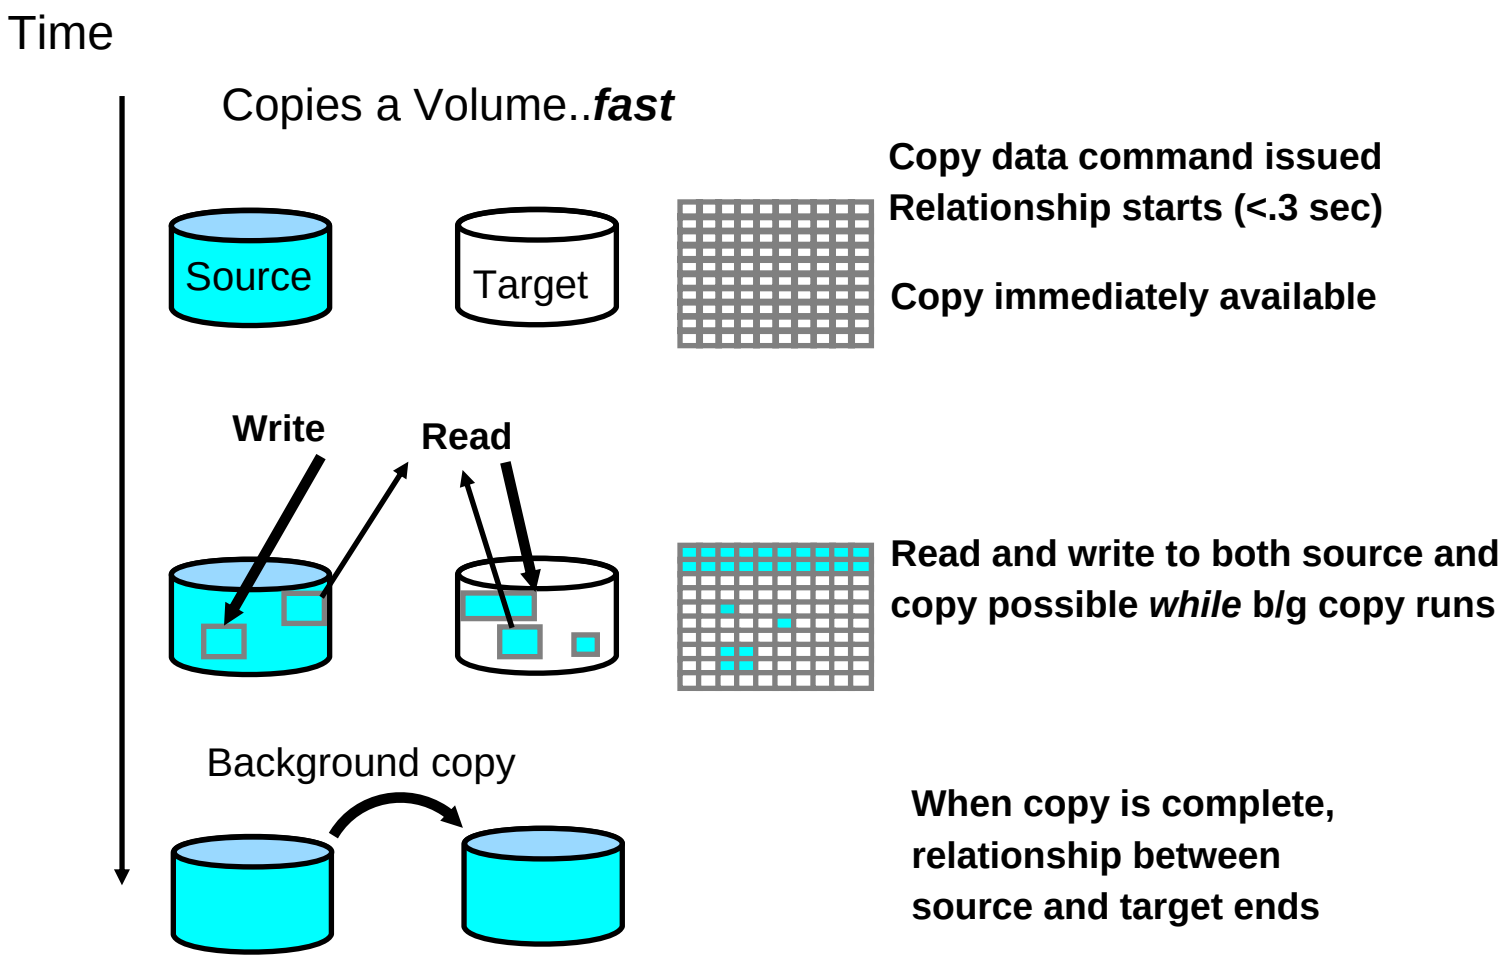

Notes – many DBAs do not know how FC works

- Produces ' instant copy' of a volume, as this example shows (can be data set)
- FlashCopy volumes can only be in a single FlashCopy relationship at any time
- Source and Target volumes require real disk space AND must be on the same DASD ssid
- Source and target volumes must be same track geometry
  Sample JCL for volume level FlashCopy

```
//COPYJOB JOB...
//INSTIMG EXEC PGM=ADRDSSU
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=V,OUTLIM=3000
//SYSIN DD *
 COPY FULL INDYNAM (src)) OUTDYNAM (tgt) DUMPCONDITIONING (SMS managed)
 COPY FULL INDYNAM (src)) O UTDYNAM (tgt) NOCOPYVOLID  ( non SMS managed) *
```

- DFSMSdss checks if source and target are eligible for FlashCopy. If not, DFSMSdss will do a normal copy.
- Once the FlashCopy "Logical Complete" occurs, the DFSMSdss job completes. It does not wait until the copy is physically complete (which is performed by the ESS/DS8K hardware)
- Once the copy is physically complete, the relationship between source and target is ended.
- NOCOPY reserves space (for copied tracks), but does not start the background copy task. Target is used as a cache for updated tracks only. Secondary Relationship stays till terminated or until all source tracks have been copied because they were all updated. For DR tape copies, you would explicitly withdraw it. See *z/OS 1.12 DFSMS Advanced Copy Services*, SC35-0428-18, Ch 26.

BACKUP SYSTEM utility

- Invokes DHSMShsm services to take fast, (few minutes) minimally disruptive volume copies of the DB2 data and / or logs
- No DB2 Quiesce point is required, nothing stops as in SET LOG SUSPEND (thus copy pool volumes are very fuzzy)
- Two flavors: FULL / DATA ONLY
  - Two sets of SMS Copy Pools (database and log)
  - BACKUP SYSTEM FULL copies both
  - BACKUP SYSTEM DATA ONLY copies database copy pool
- Make sure active log/BSDS has separate ICF catalog from data

BACKUP SYSTEM

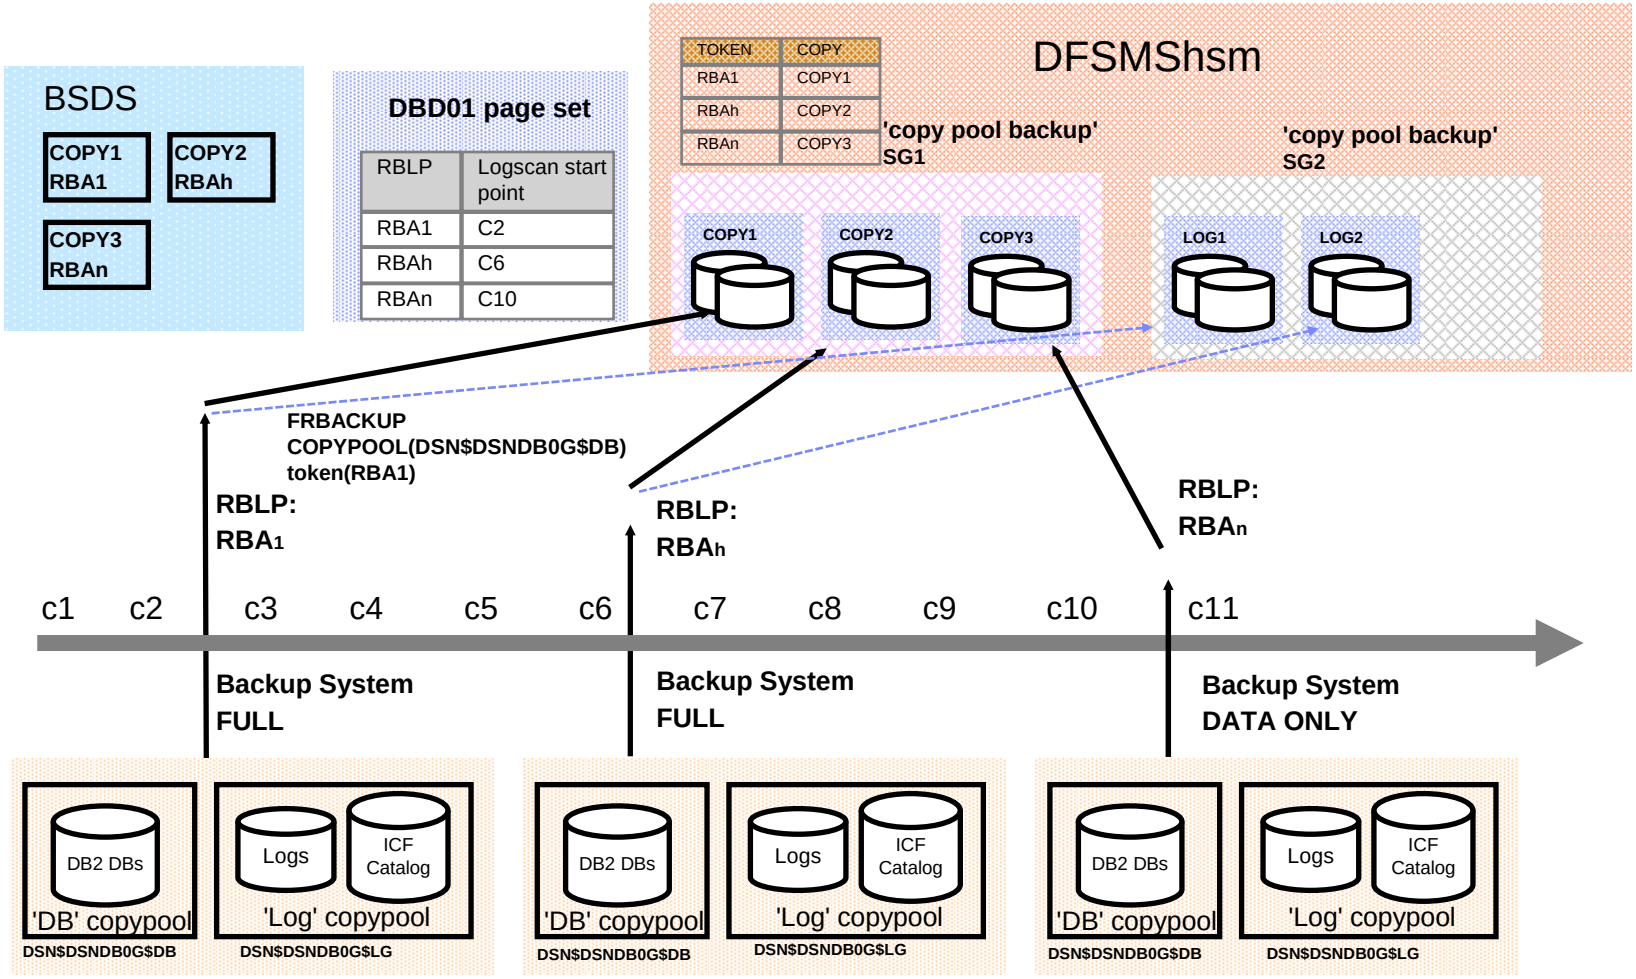

DB copypool also contains ICF catalogs for the DBs (not shown above)

Notes

- This picture shows generally three backups at different times: 2 of data and logs, one of data only. DFSMS stores and tracks them. Various parts necessary for identification are stored in the BSDS and for restart in the DBD01 header page.
  - Note that the volume copies are fuzzy with respect to the checkpoint values verbally. Use C2-C3 for example.
  - Due to limitations of the slide, the database copy pool does not show any ICF catalogs while the log copy pool does. Both copy pools contain at least one ICF catalog.
- RBLP value is the log scan starting point for log apply phase. Previous checkpoint value..similar to restart..
- BSDS RBAnn represents token information including the RBLP denoted by RBA1, 2, n.
- RBA1 represents the token that is comprised of the concatenation of SSID+Start of STCK+ RBLP
- Token is concatenation of SSID - START STCK - RBLP and is a hexadecimal representation of it.

Restore System Example

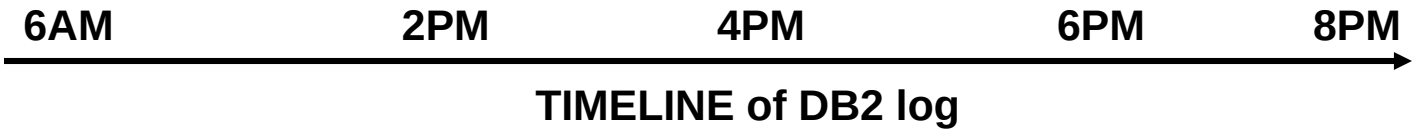

- Scenario
  - BACKUP SYSTEM is run at 6AM and 6PM daily for SAP system
  - Errors discovered at 8PM
  - Data was correct at 2PM

- Solution: Restore subsystem to 2PM using SLB from 6AM (-1 version)

- RESTORE SYSTEM does not restore active logs
- At local site logs are available, either as active or archive logs
- Prepare DB2 for 2PM recovery point

Prepare DB2 for start up

- Develop the CRESTART record using time format
  - CRESTART SYSPITRT=yyyydddhhmmsst
- CRESTART SYSPITRT=20111802000000
  - Assuming day 180 of 2011 (1400 is -6 hours from UTC if CST)
- Stop xxxxADMT (if it is used)
- Stop DB2
- Run Change Log Inventory utility (DSNJU003)
- Make sure LOGAPSTG=100 in DSNZPARM (V9)
  - Enables fast log apply for faster recovery

Notes

- This is SYSPITR to truncate the log. Do not code other values in the CRESTART record
  - If data sharing do for all non-dormant members
- LOGAPSTG =100 means DB2 uses Fast Log Apply, in this case 500M, since nothing else is active. In DB2 10 this DSNZPARM has been removed (you get it).
    - V9 100MB log section – 1000 tasks
    - V10 500MB log secion – 1000 tasks
- Using time is easier, but if you choose, you can determine PITR same old way - via time and translate to RBA or LRSN (if data sharing)
  - For data sharing, do not specify checkpoint frequency in terms of log records. A member used primarily for query will not advance the log much. It could set the RBLP used for log apply to a very old value.
- That becomes the logpoint (use EXACT RBA, no round up). If you are in data sharing, it is a LRSN.
  - If the value is at the end of an archive log, you subtract 1 from the LRSN to avoid spanned log records..

Recovery of DB2 subsystem

- **Start DB2**
  - DB2 system enters into System Recover Pending state (sets internally DEFER ALL and MAINT mode)
- **Submit RESTORE SYSTEM utility**
  - DB2 invokes DFSMShsm to restore the database copy pool
  - DB2 reads the RBLP and applies all log to 2PM – one pass
    - Automatically – no human hands
    - No “Recover Tablespace Logonly” statements – it is just like restart
      - *No complex order - Not single threaded*
- **When it completes, System Recover Pending mode is reset**

- **Recycle DB2**
- **Following restart, RECOVER or REBUILD PENDING objects**

Restore System Timeline

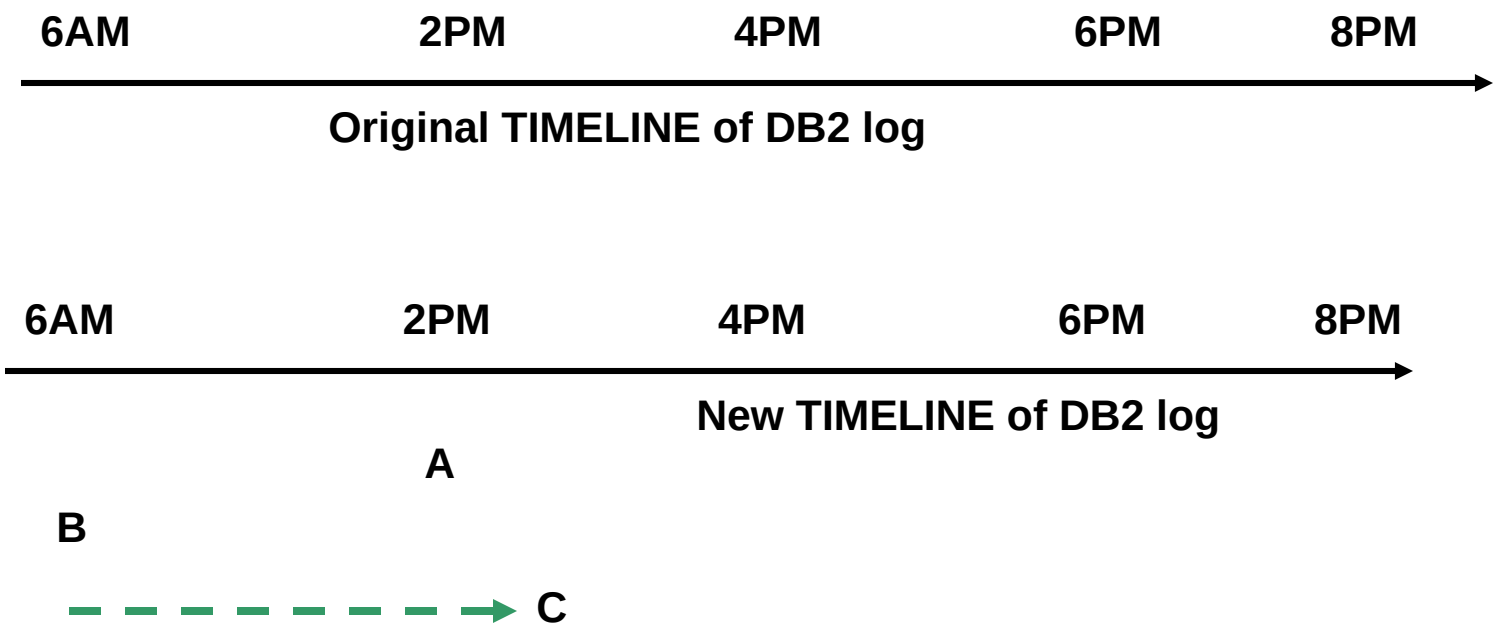

1. A – Log truncated via SYSPITR
2. B – SLB at 6AM chosen and restored. **Log Apply STARTS when Establish Phase completed – no wait for background copy!!**
3. C – Log applied to this point

DB2 Versions Summary

- DB2 V8 – introduced SLB
  - Copy pools to DASD only
    - Tools like Recovery Expert needed for Tape copies
- DB2 9
  - Copy pools to DASD and TAPE
  - Incremental FlashCopy supported for copy pool volumes
  - Recovery from SLB permitted ..but restrictions
  - Can specify SYSPITRT as yyyydddhhmmsst (UTC)
  - Can specify SYSPITR as x'ffffffffffff' (end of log)
- DB2 10
  - Recovery from SLB even if table space or partition moved via reorg to other volumes
    - Must have enough space on original volumes
    - DFSMShsm option to capture ICF catalog information also required

What z/OS (DFSMShsm) does for SLB

z/OS 1.13	**available 9/2011**	**nothing**
z/OS 1.12	**available 9/2010**	
z/OS 1.11	**available 9/2009**	

z/OS 1.11 capabilities

- **Unallocates internally ICF catalogs before RESTORE SYSTEM**
  - If user specified “CAPTURE CATALOG INFORMATION FOR DATA SET RECOVERY (**N**)” field of the copy pool definition for each catalog
  - Before that you had to use z/OS command ***F catalog,unallocate..***

- **RECOVER can use SLB to restore data sets that moved to other volumes**
  1. If “CAPTURE CATALOG INFORMATION FOR DATA SET RECOVERY (**R**)” field of the copy pool (R=Required)
     - “P”(P=Preferred) might work if catalog information was captured – no guarantees
  2. If enough space on the original volumes to restore them
  3. DB2 behavior
     - DB2 10 – even if reorged (I/J name switch)
     - DB2 9 - “moved” but not reorged

Capture Catalog considerations

- **Prereq for “Required” - Some ICF catalogs are not “clean”.**
  - Example: data sets cataloged that don’t exist on the volume
  - Not uncommon for catalogs in use for years
  - IDCAMS DIAGNOSE can identify many errors, but not all
  - IBM Support center “DFSMS Catalog” can work with you to make it clean
- **If you specify “Preferred”**
  - ICF catalogs will be unallocated when Restore system runs
  - Attempt to capture the named ICF catalogs. If yes, you can restore individual DB2 table spaces. If no, you cannot.
- **If you specify “None”**
  - ICF catalogs will be unallocated when Restore system runs
  - You will not be able to recover individual DB2 table spaces
- **If you leave it blank**
  - you must manually unallocate the ICF catalogs before Restore System runs

Notes

- We are restoring the ICF catalogs for the data sets in the copy pools. Locally they exist as of now and we will restore them to an earlier time.
- Capture Catalog options. There are three options: N, P, R.
  - N - Do not collect data set info (use if you only want HSM to unallocate the ICF catalogs before the SLB
  - P - Try to collect data set info, but proceed if unsuccessful
  - R - Collect data set info and fail SLB if unsuccessful (recommended)
- Even if unallocated at the start of Restore System, it does not prevent them from being allocated later on (Establish for DASD/DASD, entire Restore System for Tape)
- DFSMShsm expects a clean user catalog which contains only data set entries which point to data sets that reside on volumes within the managed Copy Pool. Catalog Search Interface does returns non zero code when the catalog contains orphan entries. A non-zero CSI return code prevents data set level recovery if the data set has moved or has been deleted at the time of the recovery. The condition of a clean catalog is very strict but necessary to successfully create volume copies of the managed copy pool volumes to allow successful data set recovery

Tape Recovery in z/OS 1.12

- DFSMS**hsm** supports recovery of copy pools on tape
- Now DFSMShsm parallelizes tape recovery for FRRECOV

- Makes restore procedure at DR site now *easy*
  - *N/A for GDPS* (pertains to tape based DR)

- Before
  - DFSMS**dss** used for FRRECOV was single threaded recovery
  - RESTORE SYSTEM needed for any kind of parallelism – at DR

Notes

If necessary to restore log copy pools a more complex (DR) was required. Following is procedure

- LIST COPYPOOL (DSN$location$LG) ALLVOLS(ALLVERS)
  - *It returns volser information for the log copy pool*
- Restore active log/BSDS on a volume by volume basis
  - FRRECOV TOVOLUME (volser) FROMDUMP

Fast Reverse Restore – z/OS 1.12

- Fast Reverse Restore (FRR)
  - Ability to restore a flash copy before the background copy is complete
  - Set in the copy pool and recorded for each backup
  - Establish phase must have completed
  - SLB only (volume level, not data set level)
  - What about SLB incremental FC? Yes: it is a persistent relationship
    - DFSMShsm automatically withdraws the relationship and restore volumes
      - *Best practice afterwards: create another full version for protection*
  - If FRR cannot be used, recovery fails (b/g copy must complete first)
  - FRR cannot be used in combination with Preserve Mirror operation (Remote Pair FlashCopy) PPRC

Speed Up Recovery

- Enables DFSMShsm to FlashBack for recovery even though background copy has not completed
- Enables recovery from NOCOPY versions
- Copy pool setting to enable FRR
- QUERY COPYPOOL command indicates percent complete
  - Determine if you should just wait for background copy to complete

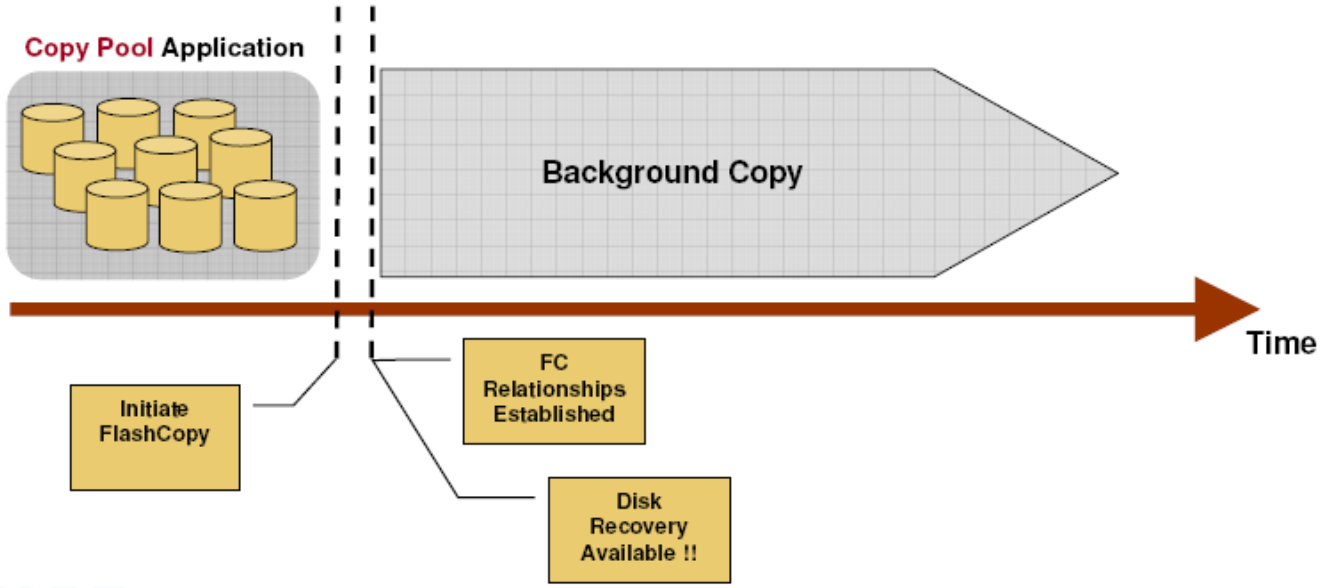

Space Efficient FlashCopy (z/OS 1.12)

DFSMShsm selects available Space Efficient target volumes when NOCOPY (VERSIONS=0) is selected (copy pool)

- Support is transparent to DB2 (any version)
- Volumes only

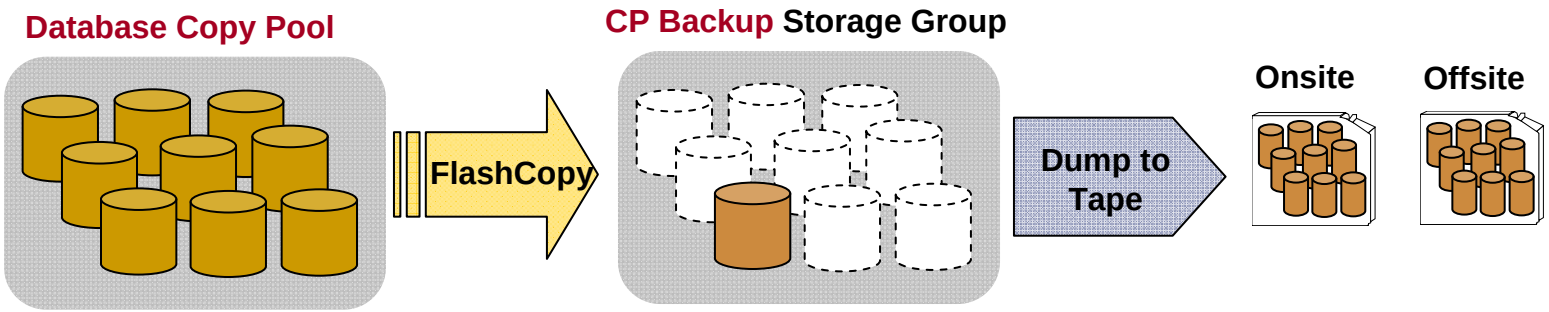

Fully Provisioned

- 1 TByte copy pool needs 1 TByte of copy space

Space Efficient Target Volumes

- ★1 TByte copy pool that changes <10% between copies only needs 100 GBytes of target space

Using FRR with Space Efficient FC (SEFC)

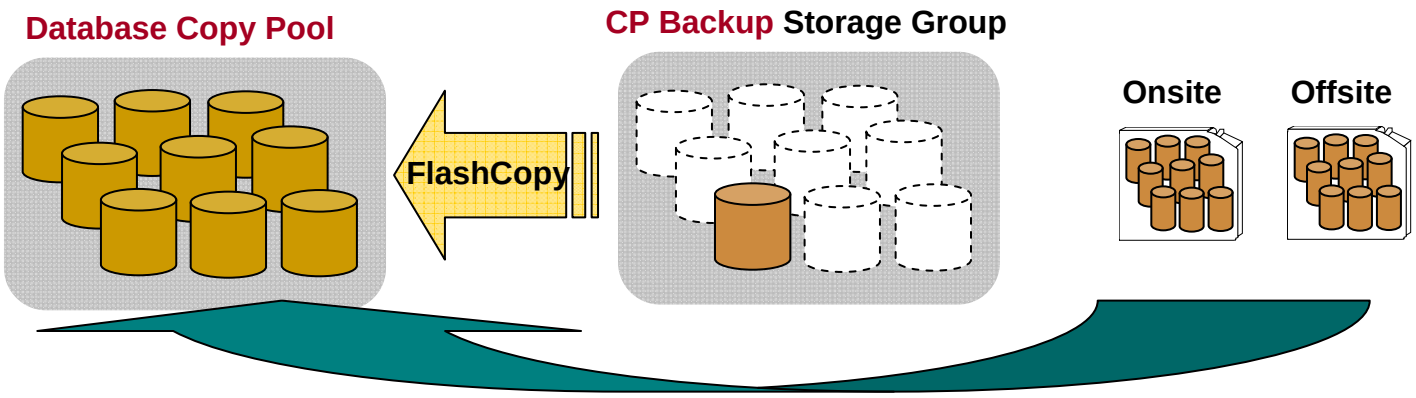

With Fast Reverse Restore…

- In addition to being able to recover from tape, a disk recovery can be performed!

- SEFC is special case of SLB for users that keep backups on tape

DASD mirroring users

- SLB not easily compatible with these mirroring solutions:
  - Global Mirror
  - XRC (z/OS Global Mirror)
- FlashCopy to a target (copy pool backup storage group) that is also a PPRC Primary volume sets the secondary volume "duplex pending".
  - Secondary site not restartable
  - HyperSwap fails - protection against local control unit failures
    - allows primary DASD to switch to a secondary local DASD w/i few seconds
  - Mirroring normally controlled by GDPS
- Customers typically mirror all DASD volumes at enterprise

- Concepts important to understand for later set up parameters

Notes (history)

- Before z/OS V1R8, any FlashCopy to a target that was also a primary PPRC would fail
  - User had to break the links for the duration and then reestablish them manually
- z/OS APAR OS23849 (7/08) allows FlashCopy to PPRC primary volume
  - But it puts the links of the secondary in “duplex pending” until all tracks copied
- z/OS APAR OA24814 (3/09) and OA24811 implemented Remote Pair FlashCopy (RPFC) that allows FlashCopy to execute in parallel between the primary and secondary pairs
  - Full duplexing of links is maintained
  - Needs microcode on DASD subsystem also
- All service is at least 2+ years old - the z/OS versions are out of support
- Neither XRC nor Global Mirror support FlashCopy to primary volumes
  - Backup System would not choose one
  - Restore System would fail

FlashCopy to PPRC Primary

- **Problem**: If you FlashCopy to a volume that is also a PPRC primary volume, the secondary target volume becomes duplex pending until the tracks are copied across the links from the primary to the secondary.
  - Bad for GDPS and Hyperswap – don't break 'full duplex' mode
    - Both fail

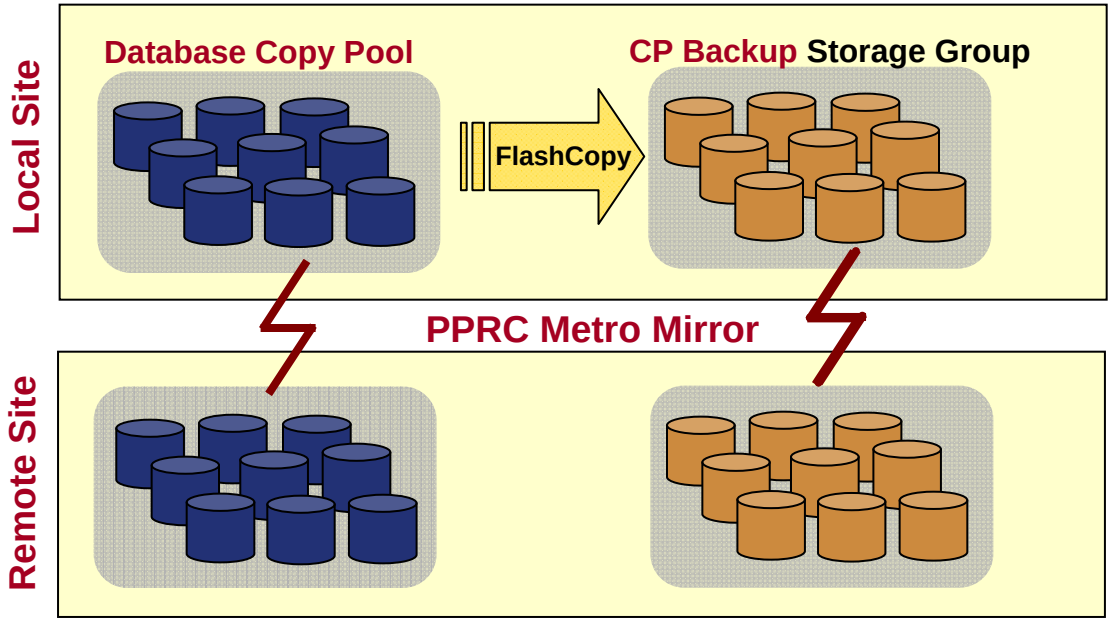

Remote Pair FlashCopy (RPFC) (Preserve Mirror)

- **Solution:** Preserve Mirror: FlashCopy performed at local site is also performed at secondary site via a command from the local site
  - No sending of changed tracks from local volume to secondary volume
- Secret Sauce is the FlashCopy inband connection
  - With Preserve Mirror, under Backup System, DFSMShsm sends the same command to the DASD controller at the secondary site.
- When Establish phase complete on both sites (minute+), DFSMShsm gets control back
  - Background copies complete on each site and not on the Metro Mirror PPRC links

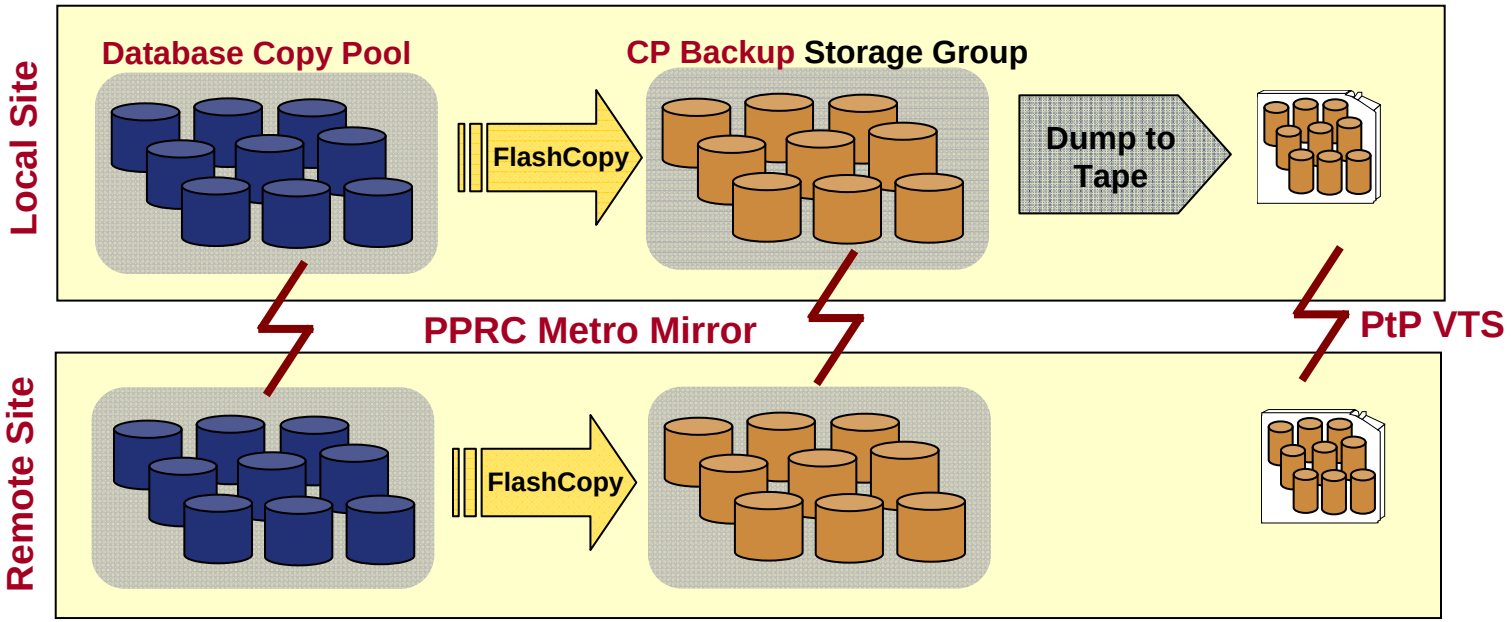

Restore System using Preserve Mirror/RPFC

1. Restore System invoked local site
2. DFSMShsm sends the same command to the DASD controller at the secondary site
3. When Establish phase is complete on both sites (minute), DFSMShsm gets control back
4. Log Apply starts immediately (b/g restore continues in the DASD subsystem)
   - Tracks restored are read from the source
   - Tracks not yet restored are read from the CP backup storage group
5. No link traffic!

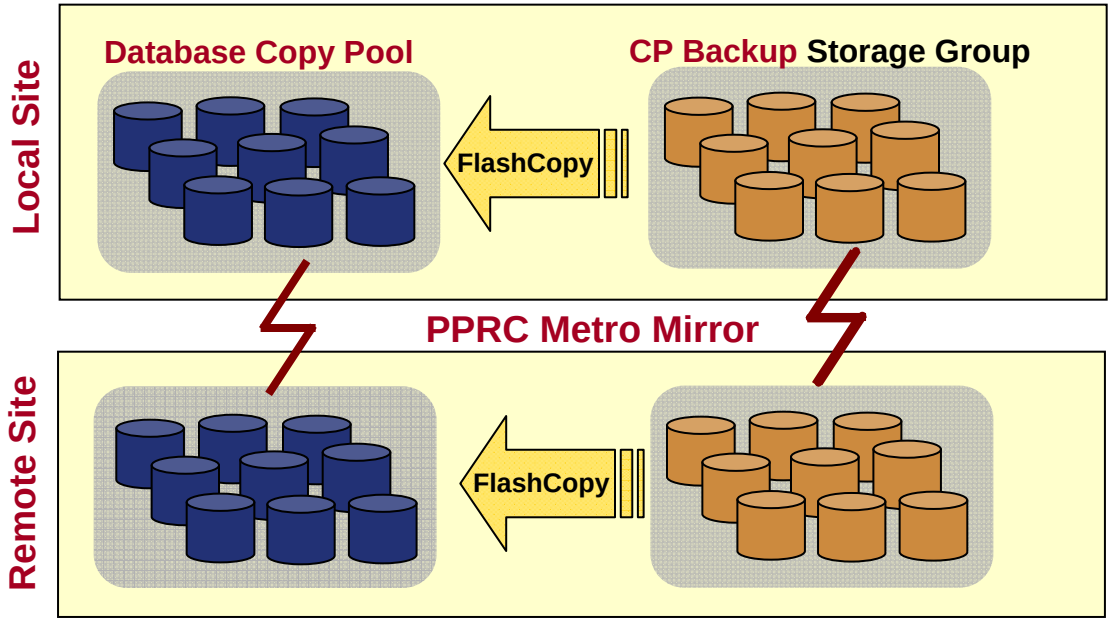

Error condition #1 - Restore System with HyperSwap

1. Restore System invoked locally (A > A')
2. Same command sent to secondary site (B > B')
3. When Establish phase is complete on both sites (minute+), DFSMShsm gets control back
4. HSM returns to DB2 – RC0 EOJ
5. **Hyperswap occurs – links suspended**
6. Local site (A) now connected to secondary DASD (B pair)
7. Restore System b/g copy finishes on its own (B pair)

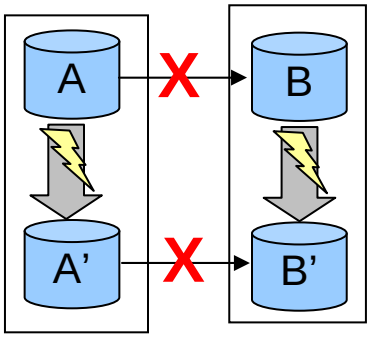

Error condition #2 - Restore System with HyperSwap

1. Restore System invoked locally (A > A')
2. Same command sent to secondary site (B > B')
3. Establish phase in progress
4. **Hyperswap occurs – links suspended**
5. Local site now connected to secondary DASD (B pair)
6. Restore System terminates (Establish incomplete)
7. Submit Restore System again
8. **Remote pair Flash Copy now allowed even if either or both pairs are suspended or duplex-pending**
   - **Enabled by DS8K Rel 6.2 with Preserve Mirror Required**
9. Establish phase completes on (B pair)
10. Background copy B > B' completes

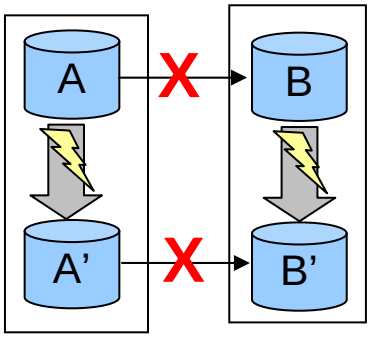

Preserve Mirror Summary

- Applies only to PPRC Metro Mirror
  - Cannot be used for Global Mirror or z/OS Global Mirror (XRC)
    - Or PPRC secondary that is Primary to one of the above in a 3-site configuration
- z/OS 24811 (F904)
- Options specified on the ISMF Copy Pool (ISPF panels)
  - Preserve Mirror Required (PMR)
- Microcode on IBM storage hardware DS8000 series
  - Include µcode DS8800 Rel 6.2 (Nov 2011) - FC now allowed, even if either or both pairs are suspended or duplex-pending AND if required to preserve mirror

- *Fast Reverse Restore and Space Efficient FlashCopy cannot be used*

- GDPS users must understand to set up RPFC for SLB correctly.
- *IBM System Storage DS8000: Remote Pair FlashCopy* REDP-4504 (May, 2009)

Setup

Notes: Prerequisites

- z/OS V1R8 or above for DFSMShsm copy pools on tape (out of service)
- DASD control units which support ESS Flashcopy APIs – vendor can purchase
- DB2 datasets reside on SMS-managed storage groups
- Whole subsystem restoration
- FlashCopy 2
  - Required for Backup System Incremental
  - Not for Backup System Full, but about 10x slower than FC2
  - Not for Tape dumps
- NFM not required
- Separate ICF catalogs for Active logs/BSDS from data (can have multiple ICF catalogs)
- DFSMShsm
  - z/OS V1.11 unallocates ICF catalogs before RESTORE SYSTEM
- DB2 10 for recovery of single object following reorg movement

Copy Pool Storage Group

- SMS construct, consists of SMS storage groups.
  - Versions attribute allow specification of the number of copy versions to be maintained on DASD (max is 85).
  - Each version is a complete set of the source DASD, so greater than 2 are unlikely – 1 is okay
  - Advanced capabilities of FlashCopy are specified on the ISMF Copy Pool backup storage group definitions.
- Each DB2 system or data sharing group has two HSM copy pool backup storage groups with prescribed DB2 naming convention (30 bytes in length)
  1. DATABASE COPYPOOL (DSN$location_name$DB)
  2. LOG COPYPOOL (DSN$location_name$LG)

BACKUP SYSTEM setup

- **Copy Pool - finest granularity is SMS Storage Group** (SG)
  - Several storage groups for **Database Copy Pool Backup SG**
    - User data
    - Catalog/directory
    - ICF catalogs containing only definitions for
      - *All user data*
      - *Catalog/directory*
    - Storage groups for additional extents if not included in above
  - Do not mix data from another DB2 in the storage groups
  - Do not add other non-DB2 data sets (VSAM, TSO etc)
  - Not required to include work DB (DSNDB07/WRK)
    - Catalog them in a separate ICF catalog

Notes:

- Not necessary to include the work database in the copy pool
- DB2 does nothing with them during either SYSPITR or conditional restart, since they are not needed at that time (only the logs are read)
- During normal restart, DB2 resets the work DB.
- For non mirroring DR you can define them at the recovery site and put them in their own ICF catalog ie, or include them in the database copy pool and avoid an extra step at the recovery site – your choice.
- If you do not plan to use the SLB for DR, there is no point to including them because they exist locally and the restart following Restore System will reset them.
- For GDPS users, the work database is part of the enterprise consistency group, so are mirrored to the recovery site.

BACKUP SYSTEM setup…

- **Log Copy pool backup storage group**
  - Active logs
  - BSDS
  - ICF catalogs containing only definitions for above
  - Volumes containing only additional extent definition for above
  - Do not mix data from another DB2 (unless data sharing) in the storage groups
  - Do not add other non-DB2 data sets (VSAM, TSO etc)
  - Do not put archive logs in Log copy pool – too many volumes to restore
- You will need more SMS storage groups when BACKUP SYSTEM is used
- Keep image copies in different ICF catalog / storage group from DB Copy Pool
  - Any inline copies needed for recovery following reorg/load would be lost when Database Copy Pool is restored

Specify Copy Pool Options

- **Fast Reverse Restore** – z/OS 1.12
  - Enable a fast replication (FlashCopy) recovery before b/g copy completes
    - *Enables a recovery immediately after the backup "establish"*
  - You can run Restore System before the Backup System background copy completes
- **Space Efficient FlashCopy** – z/OS 1.12
- **Preserve Mirror** – z/OS 1.9
- Note: **Incremental FlashCopy** (FCINCREMENTAL) affects the copy pool but is only specified in the BACKUP SYSTEM utility

Specification on ISMF Copy Pool

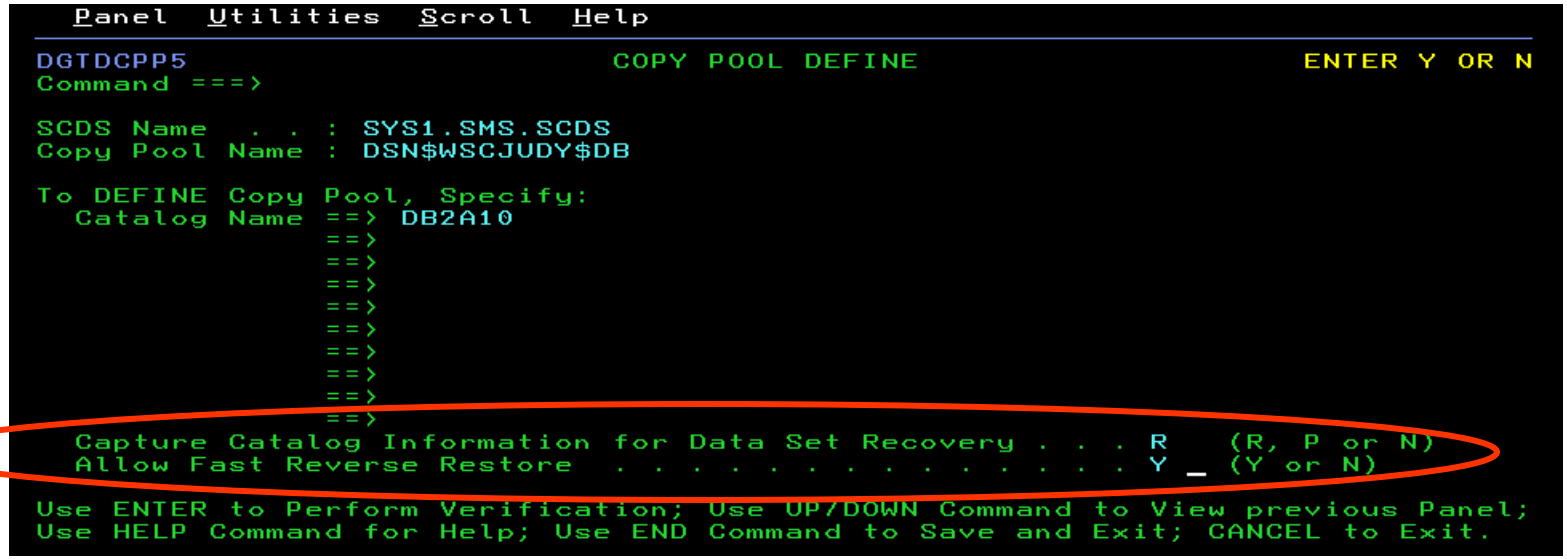

- Capture Catalog Information – R critical if you want to restore table spaces from SLB
  - If ICF catalog not “clean”, SLB fails
- Fast Reverse Restore (“N” if Preserve Mirror / RPFC)
- All ICF catalog names for the copy pool

Copy Pool Definition (non mirrored)

```
DGTDCPP1                          COPY POOL DEFINE                    Page 1 of 5
Command ===>

SCDS Name  . . : SYS1.SMS.SCDS
Copy Pool Name : DSN$WSCJUDY$DB

To DEFINE Copy Pool, Specify:
  Description ==> _
              ==>
  Auto Dump . . . N    (Y or N)          Dump Sys/Sys Group Name . . .
  Dump Class  . .                        Dump Class  . .
  Dump Class  . .                        Dump Class  . .
  Dump Class  . .

  Number of DASD Fast Replication Backup
   Versions with Background Copy . . . . . 2        (0 to 85 or blank)
  FRBACKUP to PPRC Primary Volumes allowed . . .     (NO, PN, PP, PR or blank)
  FRRECOV to PPRC Primary Volumes allowed  . . .     (NO, PN, PP, PR or blank)

Use ENTER to Perform Verification; Use DOWN Command to View next Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

- Options FRRBACKUP and FRRECOV – for users who do not mirror:
  - Blank – same as NO – PPRC Primary volumes cannot be targeted
- **Options ignored if target is not PPRC Primary**
- Versions – if 0 means Space Efficient FlashCopy if SEFC volumes available, otherwise it is standard NOCOPY using tape

Copy Pool Definition for Preserve Mirror

```
 Panel  Utilities  Scroll  Help
DGTDCPP1                          COPY POOL DEFINE                    Page 1 of 5
Command ===>

SCDS Name  . . : SYS1.SMS.SCDS
Copy Pool Name : DSN$WSCJUDY$DB

To DEFINE Copy Pool, Specify:
  Description ==> DATABASE COPY POOL FOR JUDY SUBSYSTEM
              ==>
  Auto Dump . . . N   (Y or N)          Dump Sys/Sys Group Name . . .
  Dump Class  . .                       Dump Class  . .
  Dump Class  . .                       Dump Class  . .
  Dump Class  . .

  Number of DASD Fast Replication Backup
   Versions with Background Copy . . . . . 2          (0 to 85 or blank)
  FRBACKUP to PPRC Primary Volumes allowed . . . PR  (NO, PN, PP, PR or blank)
  FRRECOV to PPRC Primary Volumes allowed  . . . PR  (NO, PN, PP, PR or blank)

Use ENTER to Perform Verification; Use DOWN Command to View next Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

- Options FRRBACKUP and FRRECOV – for Remote Pair FlashCopy users:
  - Preserve Mirror Required (PR) – Preserve mirror between primary and secondary sites – volumes must not go “Duplex pending”. Fail the FlashCopy if that would occur
  - Preserve Mirror Preferred (PP) - Attempt to preserve mirror, but if impossible, take the FlashCopy and allow the volumes to become duplex pending
  - Preserve Mirror None (PN) – PPRC Primary can be target. Volume then is duplex pending
  - NO – do not FlashCopy to PPRC Primary Volumes. Fail FlashCopy if that would occur
  - Blank – same as NO – PPRC Primary volumes cannot be targeted
- Versions – if 0 means Space Efficient FlashCopy if SEFC volumes available (n/a RPFC)

dss keywords: FASTREPLICATION and FCTOPPRCPrimary

- **Both** used for DFSMSdss, which actually copies the volumes
- Set values on Copy Pool before Backup System – can't be done later.
- **DFSMShsm does not support slow DSS copy using standard I/O**
  - DFSMShsm in SLB passes to DSS: FASTREPLICATION REQUIRED

- For PPRC with Preserve Mirror copy pool panel specify:
  - FRBACKUP/FRRECOV settings for SLB – **PR** means FCTOPPRCPrimary

- If you use GDPS with RPFC –
  - SLB uses both FASTREPLICATION(REQUIRED) and FCTOPPRCP(PMR)
  - Guarantees that SLB runs or any h/w error condition causes it to FAIL
  - Assures links remain full duplexed and HyperSwap enabled

Rationale for Preserve Mirror Required (PR) for RPFC

- **Prereq: DS8800 Release 6.2 microcode ("µcode")**
- PR (Preserve Mirror Required) means if the mirrors cannot be preserved, fail the FlashCopy
- If HyperSwap occurs while RESTORE SYSTEM is running
  - If Establish not completed, rerun RESTORE SYSTEM
  - If Establish completed, the b/g copy will finish as usual.
- DS8K Release 6.2 (Nov 2011) allows Restore System even if mirrors suspended
- Do not choose PP
  - could put the mirrors in duplex pending for normal Restore System

Copy Pool Definition for PPRC Metro Mirror

```
DGTDCPP1                          COPY POOL DEFINE                    Page 1 of 5
Command ===> _

SCDS Name  . . : SYS1.SMS.SCDS
Copy Pool Name : DSN$WSCJUDY$DB

To DEFINE Copy Pool, Specify:
  Description ==>
              ==>
  Auto Dump . . . N    (Y or N)          Dump Sys/Sys Group Name . . .
  Dump Class  . .                        Dump Class  . .
  Dump Class  . .                        Dump Class  . .
  Dump Class  . .

  Number of DASD Fast Replication Backup
   Versions with Background Copy . . . . . 2        (0 to 85 or blank)
  FRBACKUP to PPRC Primary Volumes allowed . . . NO  (NO, PN, PP, PR or blank)
  FRRECOV to PPRC Primary Volumes allowed  . . . PN  (NO, PN, PP, PR or blank)

Use ENTER to Perform Verification; Use DOWN Command to View next Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.
```

- For users **without** Remote Pair FlashCopy. CP backup SG outside the mirror:
  - FRBACKUP: NO or 'blank' - Do not FlashCopy to a PPRC Primary target
  - FRRECOV: Preserve Mirror None (PN) – PPRC Primary can be target. Volume becomes duplex pending if chosen
    - Must have PN on FRRECOV
    - Restore System fails if either NO or blank
- **Options ignored if target is not PPRC Primary**
- Versions – if 0 means Space Efficient FlashCopy

Validation of configuration

- If you change configuration of your source copy pools or copy pool backup storage groups, validate the backup environment using

  HSENDCMD FRBACKUP COPYPOOL(DSNlocDB) PREPARE
  HSENDCMD FRBACKUP COPYPOOL(DSNlocLG) PREPARE

- If volume removed from a copy pool backup storage group, and you need it, RESTORE SYSTEM will fail.

- If you add volumes to your source for new/changed tables, make sure there are enough volumes in the copy pool backup storage group or BACKUP SYSTEM will fail.

Alternate configuration for PPRC (no RPFC)

- **CP backup storage group outside mirror – no secondary volumes**
  - BACKUP SYSTEM runs w/o intervention
  - Dump CP backup storage group to tape and transmit to secondary site; is available (Peer-to-Peer VTS is shown)
  - RESTORE SYSTEM
    - Local user must break duplexing (Freeze secondary) for duration of run, including b/g copy
    - Or manually restore each volume of CP backup storage group (hours) – duplexing maintained

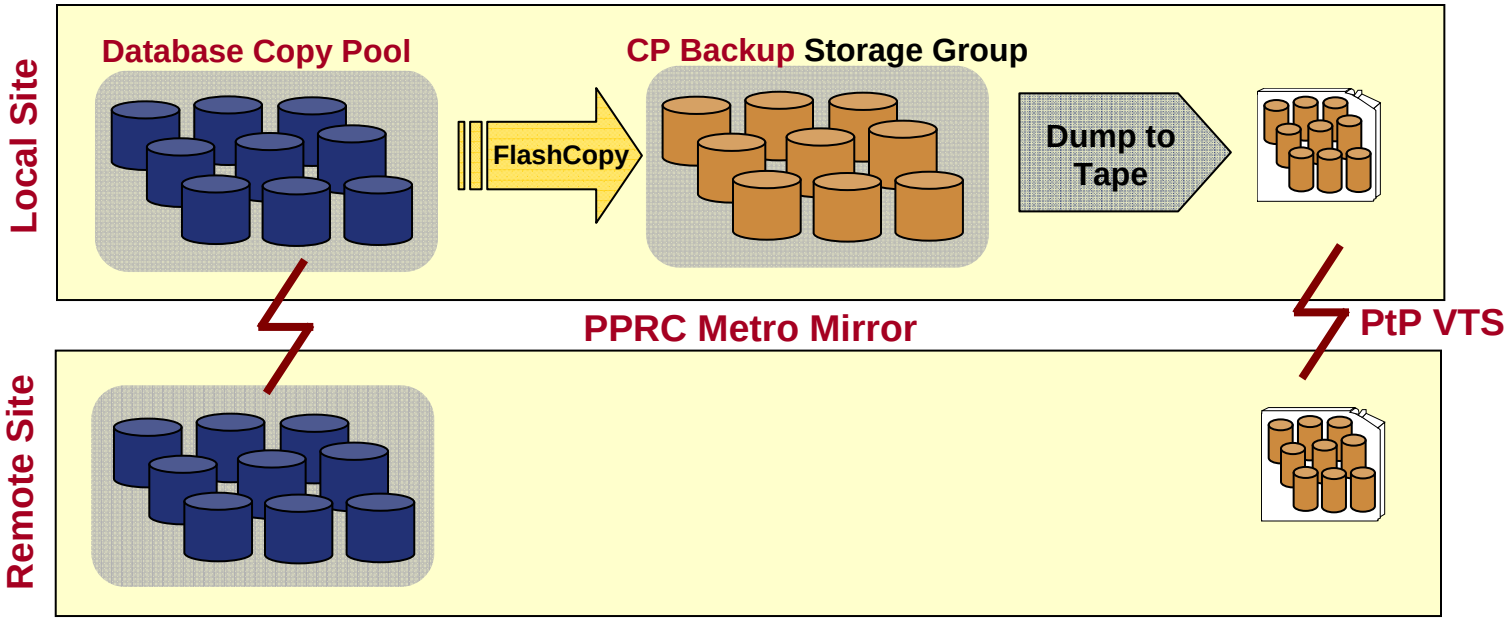

RESTORE SYSTEM for PPRC (no RPFC)

FRBACKUP=NO and FRRECOV=PN (preserve none – allow FC to PPRC Primary):

1. Set DSNZPARM **RESTORE_ RECOVER_FROMDUMP=YES** as safety – tape only use
2. Source SAP system is within enterprise mirror
3. CP Storage Pool backup outside the enterprise mirror

4. Set DSNZPARM **RESTORE_ RECOVER_FROMDUMP=NO** (allows DASD use)
5. GDPS STOP SECONDARY to preserve the secondary consistency point
6. RESTORE SYSTEM FC from CP Storage Pool Backup to Source
7. After ESTABLISH phase is complete….
8. GDPS START SECONDARY
9. B/G copy completes later to secondary (resync complete – full duplex restored)

- 1. What is the status of the links at #5? **Suspended**
- 2. What is the status of the links after #8 executes? **Duplex Pending**
- 3. What is the status of the links after #9? **Full Duplex**
- **HyperSwap disabled from #5 - #9**

Summary

- Setup is not a trivial task – nor fast
- Most work performed by storage team
  - Much movement of data to align correctly in storage groups
    - Once done, much easier to maintain
- Standards may have to be changed
  - If user has SMS-managed storage but only one storage group; almost the same as no SMS!
  - SLB requires separation of DB2 data from log data – 2 storage groups/SSID or DS Group minimum
  - Move miscellaneous data sets off the storage groups
- DB2 staff has to create another ICF catalog for logs, copy them from old volumes, and update BSDS.
- Once setup performed, execution is easy for each utility

Summary

Summary

- Ease of use by DB2 personnel
  - Difficult setup for storage personnel
- Fast restoration of DB2 system from DASD
- SLB can be used for recovery of table spaces (better in DB2 10)
- Can NOT be used for recovery of an application or subset of data within the subsystem
- Should you use SLB for your only backup, eliminating image copies other than inline?
  - When you need to recover, the utility must be rock solid.
  - SLB not as mature as RECOVER. Even in 10, restrictions can prevent
  - If you choose this, Reorg/image copy weekly – insure enough log for traditional RECOVER utility, if necessary

Recovery Expert capabilities

- Creates its own SLB (does not invoke DB2 Backup System/Restore System
- DASD can be multi vendor
- Can recover an object from RE SLB when DB2 cannot:
  - Even if unclean ICF catalog
  - Even if object moved even as a result of reorg
  - Even if not enough space on original volume (must be enough space in Storage Group)
- Can create a DB2 image copy from RE SLB
  - Can register it in SYSIBM.SYSCOPY
- You can use that image copy as input to DB2 Recover
  - You must have registered in SYSCOPY first

Summary

- SLB implementation practical for
  - PPRC Metro Mirror with Remote Pair FlashCopy (RPFC)
  - PPRC Metro Mirror without RPFC
  - No mirror
- Any other mirroring scenario
  - Lacks FlashCopy capability in most conditions
  - Requires specialized solution

Recommendations for PPRC Preserve Mirror / RPFC

- For SLB DFSMShsm passes FASTREPLICATION (REQUIRED) to DSS
- If you use GDPS –
  - SLB uses both FASTREPLCICATION(REQUIRED) and FCTOPPRCPrimary(option)
  - Set Copy Pool values to PR for FRBACKUP and FRRECOV
    - sets FCTOPPRCP(PMR) (preserve mirror required)
  - Guarantees that SLB will run or any error condition causes it to FAIL
  - Assures links remain full duplexed and HyperSwap enabled
- Insure DS8K µcode Release 6.2 is on
- Cannot specify Fast Reverse Restore
- Cannot specify Space Efficient FlashCopy (SEFC) (versions=0)

PPRC Metro Mirror without Preserve Mirror

- Set Copy Pool values PN for FRRECOV
  - Allows Restore System FlashCopy to PPRC Primary volumes
- Insure copy pool backup storage group is outside the mirrored DASD
- Toggle RESTORE_SYSTEM_FROMDUMP=YES for normal ops, to NO when the business needs to RESTORE SYSTEM from DASD
- Plan how you will restore if necessary
  1. Recovery Expert or other tool might simplify a manual recovery of volumes or
  2. Suspend duplexing while Restore System runs and then resync many volumes