

DB2 for z/OS: Disaster Recovery for the Rest of Us

Judy Ruby-Brown
IBM

Advanced Technical Skills (ATS)

Thursday, March 15, 2012
Session Number 10501



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. .
- 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.
- This presentation assumes you are NOT mirroring and you want SLB primarily for offsite DR.
- 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.



Disclaimer

© Copyright IBM Corporation 2012. 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, FlashCopy, 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

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

- This presentation is Part 1 of a two part presentation.
- The attendee will benefit by downloading
- Share presentation: session 10499
 - DB2 for z/OS System Level Backup Update

Agenda

1. Basic background of System Level Backup (SLB)
2. Deeper dive into SLB
3. Setup
4. Use of SLB for DR recovery (z/OS 1.12)
5. Summary

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

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
- Copy Pool Backup Storage Group – CP Backup SG (abbrev)

References

- *DFSMSshm 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

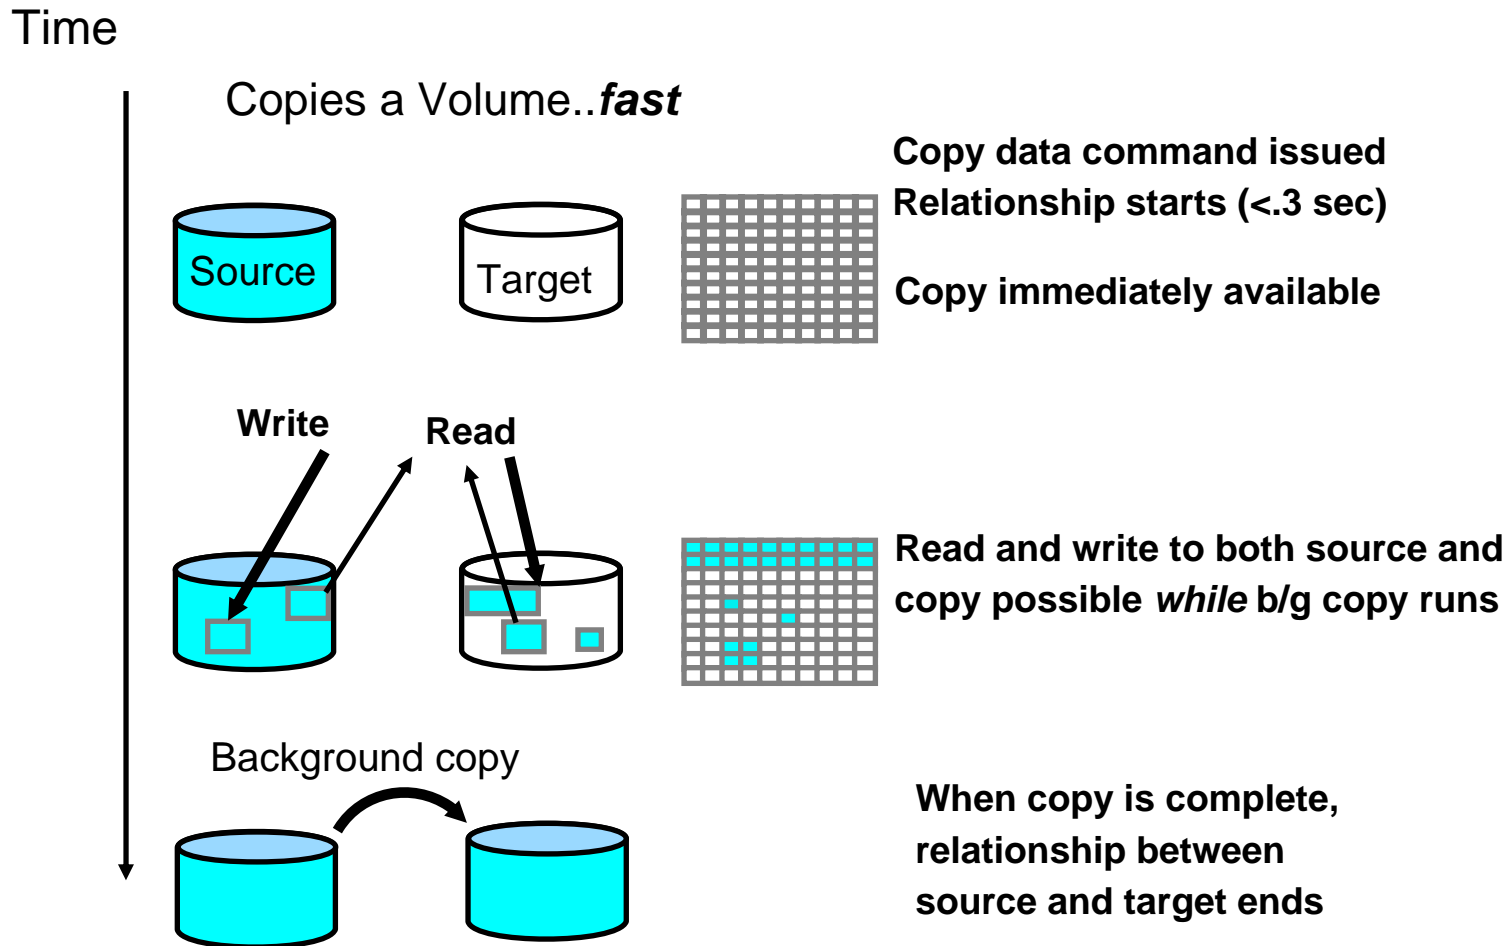
Introduction to System Level Backup (SLB)

DB2 9 >

Background

- SLB introduced in DB2 V8 for local subsystem-wide recovery for SAP
 - 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



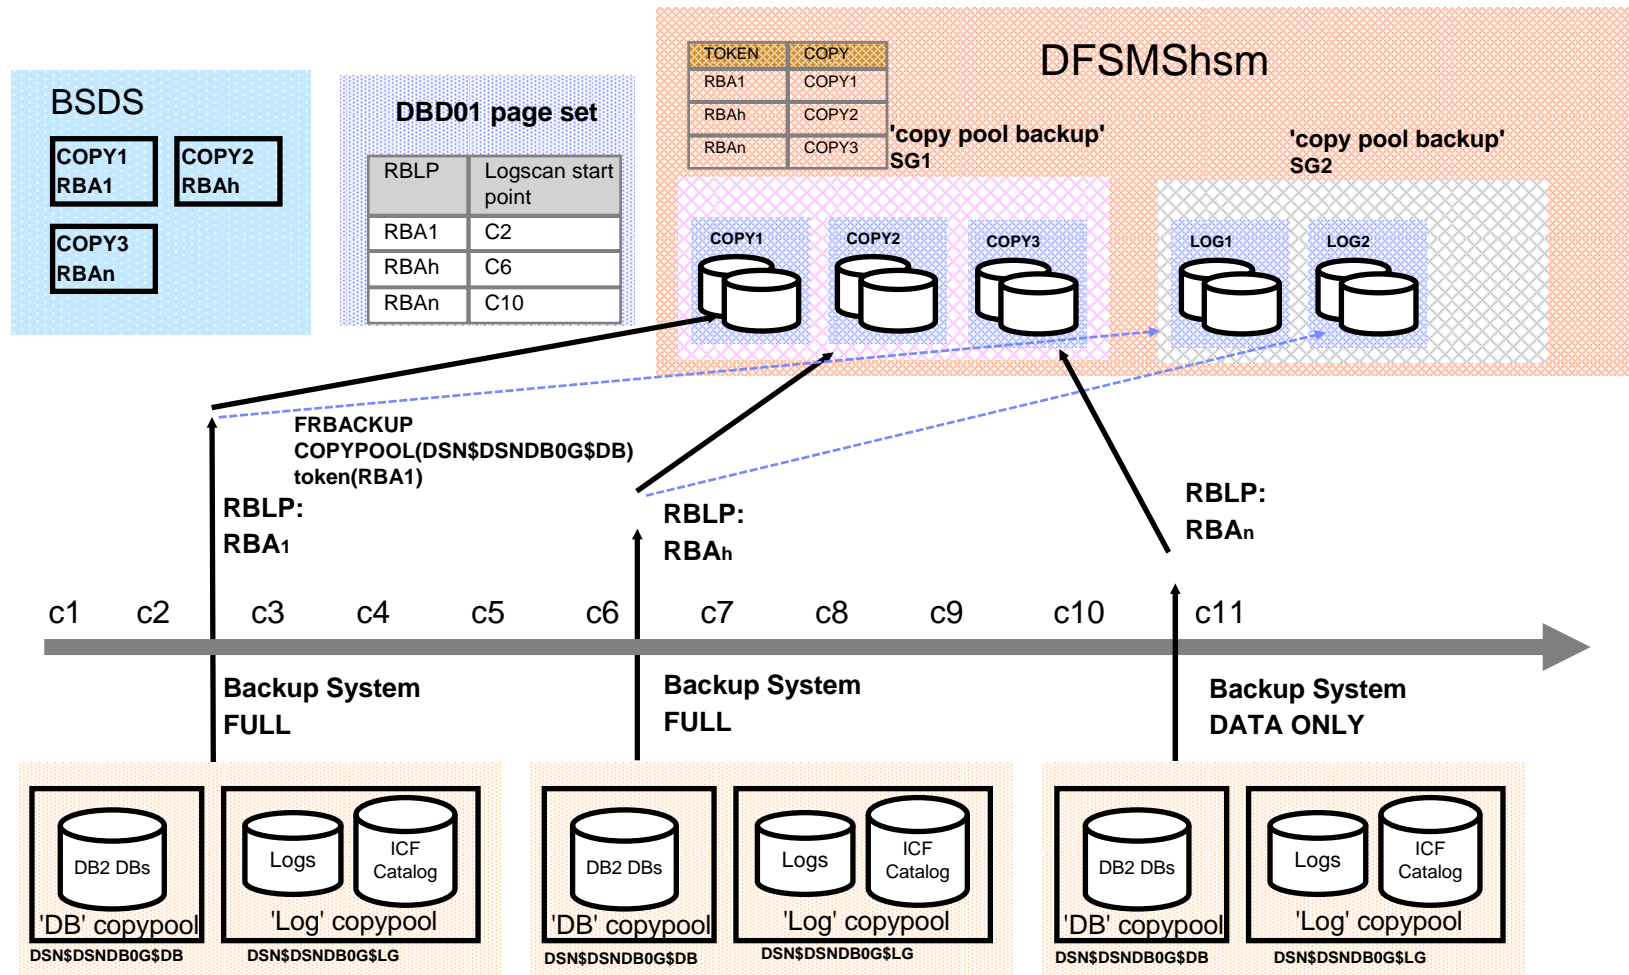
Notes – many DBAs do not know how FC works

- 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 DHSMSshm 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 cypool 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 (locally)



- 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
- RESTORE SYSTEM does not restore active logs
- Locally all 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)
  - Enable 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 section – 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**

## 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
  - Can specify SYSPITRT as yyyydddhhmmsst (UTC)
  - Can specify SYSPITR as x'ffffffffffff' (end of log)
  - Individual table space Recovery from SLB permitted ..but restrictions
- DB2 10
  - Individual table space 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

## Deeper Dive into SLB

Backup System  
Restore System

## BACKUP SYSTEM...

1. **Records the Recover Based Log Point (RBLP) in DBD01**
    - **VIP** - Starting point for DB2 log apply in the RESTORE utility and is critical for DR scenarios where additional log is applied
  2. **Invokes DFSMSHsm for DASD-DASD copies of “DB” copy pool**
  3. **Calls DFSMSHsm for DASD-DASD copies of “log” copy pool**
  4. **DFSMSHsm returns control to DB2 once “logical” copies have completed (several minutes)**
    - Backups performed in parallel – volumes in pool are fuzzy/unordered
  5. **Issues DSN1614I with Data Complete LRSN – use for later recovery**
  6. **Resumes quiesced activities**
- 
- **Less disruptive than SET LOG SUSPEND**
    - Allows most updates during backup, but none to ICF Catalog
    - No delete, define, rename, extend allowed

## Notes

- Since it is a checkpoint value (lowest of all members), make sure CHKFREQ is at least time based for data sharing users
- 15 DFSMSHsm tasks each invoking DFSMSDss with 24 volumes to process in parallel achieved (in 2004)
  - 650 3390-3 volumes in THREE minutes (2100 cylinder data) - <0.3sec each
- Customer (non-benchmark) 500 Mod 54s took 5-7 minutes to establish
- Log MUST be applied as the copy pools can be very fuzzy – each volume at a different point in time and log records may be written for a volume after it is copied. Logs are dumped following data for that reason. BSDS Data Complete LRSN is used for “end of the log”
  - Data sharing logs are NOT in sync with each other. Later we see why it matters.
- DB2 data can have more than one ICF catalog and they reside in the data copy pool.
- Logs must have separate ICF catalog as well as they are backed up sequentially. Prior to V8 customers did not usually separate them.

## Notes

- DB2 updates the DBD01 header page with the recovery base log point (RBLP) RBA and writes the page to DASD. This field is used by the RESTORE SYSTEM utility to determine the starting point for the log scan for system PIT recovery.
- DB2 outputs last RBA/LRSN which should be used for later restart – it is safe as it does not depend on whether or not the logs are striped

DSNU1614I - *csect-name* BACKUP SYSTEM UTILITY FOR *object* COMPLETED SUCCESSFULLY, COPY POOL = *copy-pool-name*, TOKEN = X'*token*',  
**DATA COMPLETE LRSN = X'*rba-or-lrsn*', ELAPSED TIME = *hh:mm:s***

- The message occurs after the backup of the data copy pool. The Data Complete LRSN is the RBA or LRSN of the last log record written to the active log data sets after the backup of the database copy pool was complete. It is reported in DSNU1614I for the data and logs, as well as recorded in the BSDS. Use it as the log truncation point for the conditional restart of DB2 for system recovery to a point in time if the log copy pool is restored.
  - For the database copy pool ('DATA'), message is issued to the Utility job output and to the system console log.
  - For the log copy pool ('LOGS'), it is only issued to the Utility job output.


## Example of FC incremental copy

- Every version registered to DFSMSHsm requires a full set of DASD for its copy pools – more than 2 is unlikely
- Assume
  - 2 versions of log and database copy pools to be registered with DFSMSHsm
  - DB2 SSID with one data source volume, DB2A, and one log source volume, DB2L
  - DFSMSHsm database copy pool backup storage group associated with target volumes DB2A1-4
  - DFSMSHsm log copy pool backup storage group associated with target volumes DB2L1-4

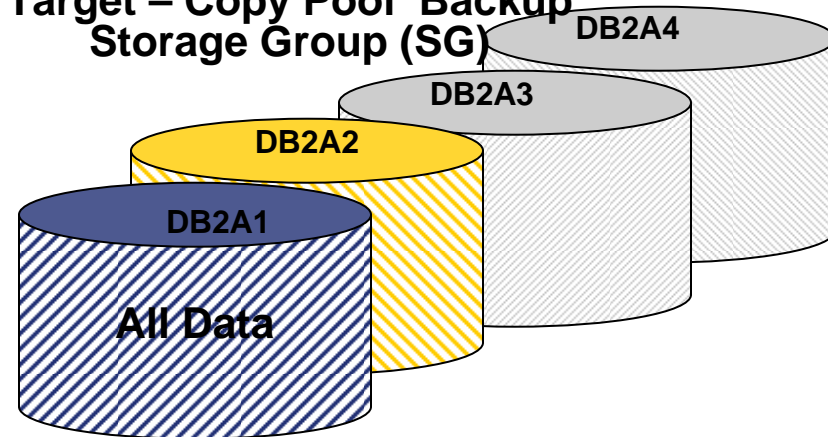


# Day 1 – Begin Incremental Version 1

Source – Database Copy Pool



Target – Copy Pool Backup Storage Group (SG)



- Day 1 – V1
  - BACKUP SYSTEM FULL ESTABLISH FCINCREMENTAL
    - Copies data source DB2A to candidate volume in the database CP backup SG, DB2A1, establishes the FC Incremental relationship
      - *All tracks on DB2A are copied to DB2A1*
    - Copies logs on DB2L to candidate volume in the log copy pool, DB2L1 (not shown)


## Day 2

- Day 2 – V2
  - BACKUP SYSTEM FULL
    - New version created using a different volume in database CP backup SG, DB2A2
    - Copies data source volume DB2A to DB2A2 as full copy
    - FC Incremental exists - only one version in copy pool can have an incremental relationship
    - Copies log copy volume DB2L in entirety to DB2L2

## Day 3

Source –  
Database Copy  
Pool

Target – Copy Pool  
Backup SG



**BACKUP SYSTEM FULL control statemet**  
**Only tracks updated since Day 1 are copied to DB2A1**

## Day 3 – V3

- Day 3 – V3
  - BACKUP SYSTEM
    - The incremental relationship still exists – while the version may be deleted, the volumes containing the database CP Backup SG (DB2A1) are not
    - All the tracks of volume DB2A that changed since Day 1 (V1) are copied to volume DB2A1 and registered as V3
    - Since 2 versions exist, DFSMSHsm rolls off V1, the oldest
    - Dumps log copy volume DB2L in entirety to DB2L1 and registered as V3
  - Tape backup of V3 for DR, using BACKUP SYSTEM DUMPONLY DUMPCLASS(DR), contains all tracks of DB2A1 and all the tracks of log copy pool, DB2L1
    - **Warning:** If current version of database copy pool was taken later than log copy pool, there is a mismatch! (Beware of Backup System DATA ONLY)

## Notes

- The number of versions don't have to be the same between logs and data. DFSMSHsm is unaware of any relationships among copy pools. All it keeps track of is the number of versions.
- Incrementals apply only to the database copy pool backup storage group
- For Day 3, DFSMSHsm knows that the version that is being replaced is an incremental copy, so it will issue a DFSMSdss command that looks like:  
COPY IDY(VOLA) ODY(VOLA1) DUMPCOND FR(REQ) PUR  
FCINCREMENTAL
  - If it didn't have the FCINCREMENTAL parameter, a full copy would have been made.
- You can have 1 version registered and still use incrementals

# Backup System Options (default FULL) Summary

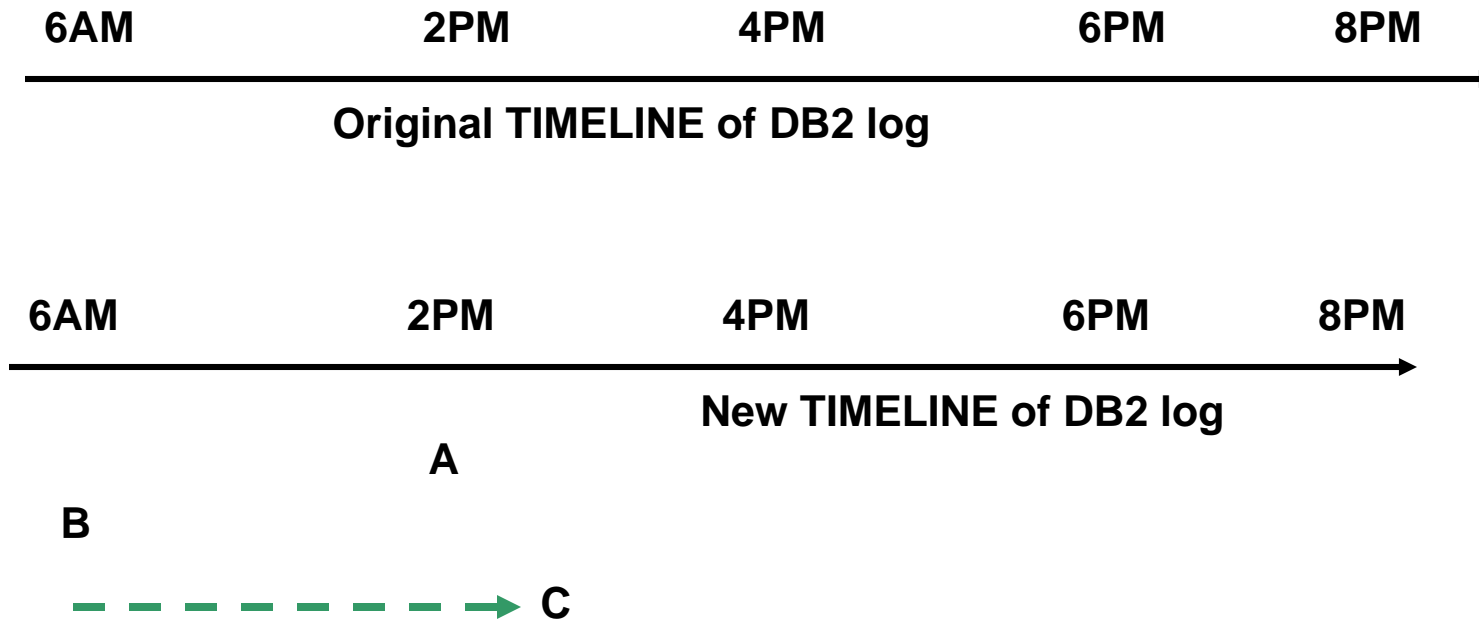


- Backup SYSTEM
  - FC Source to Target Copy Pools on DASD – data first, then logs
- Backup SYSTEM DUMP
  - FC to Copy Pools on DASD – then to tape (after establish for Copy Pools)
- Backup SYSTEM DUMPONLY
  - Copy the most current version of DASD Copy Pools to Tape
- Backup SYSTEM DUMPONLY TOKEN('xxxx')
  - Copies the version of DASD Copy Pool Backup SG to tape
  - No effect if only one version of Copy Pool Backup SG

## Restore System notes

- When copy pool backup storage groups are on DASD (local), DB2 does not have to wait for the b/g copy to complete – starts log apply when establish phase is complete – faster than volume restores.
  - For DR when CP backup SGs are on tape and are restored to source DASD, b/g copy must complete first.
- When DB2 is in System Recover Pending state
  - Only RESTORE SYSTEM utility is allowed
  - START DATABASE command is not allowed
  - TERM UTIL command is not allowed
  - DISPLAY UTIL command will display only the status of RESTORE SYSTEM utility
  - SQL operation is not allowed
    - Claim request on any DB2 objects will be rejected with a -904 SQL code (reason code of 00C20269)
- If you need to run Recover, you must recycle DB2 first.

# Restore System deeper dive



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



## Log apply notes

- **Handles table space**
  - CREATEs - will define data sets (DB2-managed)
  - DROPs – will delete data sets (DB2-managed)
  - LOG NO events (Load or Reorg)
    - If LOG NO, the associated object is entered into RECP or RBDP state.  
Table spaces and indexes with COPY YES attribute marked RECP.  
Indexes with COPY NO marked RBDP
  - Any objects you "forgot" are placed in LPL when accessed
- **Uses fast log apply (FLA) to recover objects in parallel (6x faster)**
  - "Restart logic" applies log to all objects in one pass

## Notes

- FLA (zparm LOGAPSTG > 0 V9) at least 6x faster than without it
  - DB2 9 – 100MB FLA buffer and 1000 tasks
  - DB2 10 – 500MB FLA buffer and 1000 tasks
- At the end of log apply phase
  - Issues informational message if any object is marked RECP, RBDP or LPL during the log apply phase
    - Will occur for new objects created during this phase
  - Releases the PITR lock and PITR state of each member
- Log apply phase takes periodic checkpoints
  - Writes changed pages to DASD
  - Triggers system checkpoint
  - Updates the DBD01 header page with the new RBLP value
- RESTORE SYSTEM utility is restartable in the log phase
  - In data sharing, only that member which issued the original RESTORE SYSTEM can issue the restart request

## Restore System w/o log truncation

- Designed for specialized *local* recovery where active log exists
  - Think DASD control unit failure with 1000 of objects, BUT active log set intact
  - Restore System probably faster and much easier to do vs. separate Recovers
- You tell DB2 you want to recover to currency (to end of active log):
  - CRESTART CREATE,SYSPITR=FFFFFFFFFFFFFF
    - No need to run log print utilities – ease of use
- Can be used for data sharing, BUT logs for each member must represent the same point
  - Not usually true for Data Sharing
- Do not use if active logs are striped

# General Restore Systems options summary

- In all cases you run Change Log Inventory to set the SYSPITR/T and you make available necessary range of log.

## 1. Restore System

- DFSMSHsm restores DASD database copy pool backup to source DB2 volumes (local)

## 2. Restore System FROMDUMP

- Restores database copy pool on tape to source DB2 volumes
  - (not to DASD CP backup storage group)
- Note: As long as tape data sets are cataloged and w/i 50 saved Backup Systems, you can go back as far as you have log range
  - *Even if one Copy Pool version, you can go back to the -4 tapes*

## 3. Restore System LOGONLY

- You have already restored source DB2 volumes manually.
- When restarted, DB2 starts at RBLP and processes to “current”
- This flavor does not require Backup System

# 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 DR you can just 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.

## 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 storage groups when BACKUP SYSTEM is used
- Place image copies in different ICF catalog and storage group from Database Copy Pool
  - Any inline copies needed for recovery following reorg/load would be lost when Database Copy Pool is restored

# Specification on ISMF Copy Pool

```

Panel Utilities Scroll Help
DGTDCPP5 COPY POOL DEFINE ENTER Y OR N
Command ==>

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

To DEFINE Copy Pool, Specify:
 Catalog Name ==> DB2A10
 ==>
 ==>
 ==>
 ==>
 ==>
 ==>
 ==>
 ==>
 Capture Catalog Information for Data Set Recovery . . . R (R, P or N)
 Allow Fast Reverse Restore Y _ (Y or N)

Use ENTER to Perform Verification; Use UP/DOWN Command to View previous Panel;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.

```

- Capture Catalog Information – R critical if you want to restore table spaces from SLB, but N(one) fine for DR purposes
- Fast Reverse Restore – can be used only for local recovery
- ICF catalog names for the copy pool

# Capture Catalog Information

## 1. Used to unallocate ICF catalogs before Restore System

- ICF catalogs that are part of the copy pool backup storage group also exist locally
- The local version must be unallocated so the backup one can be restored through Restore System
- Before z/OS 1.11, part of the Restore System process required a manual step to unallocate all the ICF catalogs that exist on the copy pool backup storage group.
- Manual step unnecessary if any option entered other than blank

## 2. Also used to allow recovery of an individual table space/partition from an SLB backup, specifying R/P/N options

- Described in SLB Update presentation, as it does not apply to offsite DR scenario

# 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 . . . 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.

```

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

## Validation of configuration

- If you change configuration of your source copy pools or copy pool backup storage groups, validate the backup environment using  
HSEND CMD FRBACKUP COPYPOOL(DSN\$loc\$DB) PREPARE  
HSEND CMD FRBACKUP COPYPOOL(DSN\$loc\$LG) 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.

## Summary

- Setup is not a trivial task – nor fast
- Tools such as IBM Recovery Expert V2.2 can make sure all data sets are included and reduce the labor involved
- 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

## Offsite Disaster Recovery Procedure using z/OS 1.12

- Scenario 1: SLB is restored / restarted
- Scenario 2: SLB restored – apply more archive log w/RESTORE SYSTEM



## SLB used for Offsite Disaster Recovery

- **Backup SYSTEM output can be directed to tape – Useful for DR**
  - BACKUP SYSTEM FULL DUMP DUMPCLASS(DR)
    - Dumps copy pools to DASD and also to tape
  - BACKUP SYSTEM FULL DUMPONLY DUMPCLASS(DR)
    - Dumps existing copy pools on DASD to tape
    - (BACKUP already performed)
  - Even if incremental versions are used, DUMPCLASS dumps the complete version (see next slides)
    - Time to dump is that of a FULL
- **Take also the DFSMSHsm CDS after the tape dump is complete**
- **Transport tapes offsite**

## Notes

- CDS positioning is so it “knows” about the tape dump..otherwise the FRR commands can’t be used
- DUMPCLASS specified on DSNZPARM
- DUMPCLASS name at remote site (DSNTIP6)
- Assume database copy pool was 1TB. If incremental used it just copies updated tracks to the original “1T” copy pools..dumping to tape dumps full 1T every time. This is not like incremental image copies. You save time only on backup system.
- Also incrementals are not applicable to Restore side, you always restore the full 1TB. And you can only restore system to the latest incremental, no way to go back to the -1 incremental.

## Use **BACKUP SYSTEM FULL**

- If you mix **DATA ONLY** with **FULL**, and you later issue **DUMPOONLY**, you might get a mismatch with more current database and less current log copy pool. For DR we want the most current versions and want both to match:
  - **BACKUP SYSTEM FULL (Day 1)**
  - **BACKUP SYSTEM DATA ONLY (Day 2)**
  - (Later) **BACKUP SYSTEM DUMPOONLY DUMPCLASS(DR)**
  - Now there is a mismatch with Day 2 data and Day 1 log – **BAD!**
- Usually DB2 will not start

## DR preparation

- May be better to use DUMPONLY
  - BACKUP SYSTEM stresses source I/O subsystem while FC runs
    - Incremental dumps many fewer tracks > reduces stress
  - DUMPONLY is performed after the BACKUP SYSTEM
    - No contention on source volumes
    - Perform at a more convenient time (tape drive availability)
  - Time to dump is same as FULL, even if incremental BACKUP SYSTEM
- Take CDS so that registration occurs
- Transport any needed image copies and/or archive logs as before (automated tape libraries)
- DB2 libraries
  - Libraries must be available or DB2 will not start – z/OS staff takes them to the DR site
- Do not allow xxxxADMT to start (if used) – set DSNZPARM ADMT=,

## Notes

- To run FRRECOV or RESTORE SYSTEM, you need the CDS taken after the tape dump completes, space for the production volumes, and both sets of copy pools on tape. You do not need any copy pools on DASD, nor do you need the existing SSID data sets on DASD. DB2 libraries are necessary and are normally taken by z/OS staff.
- In this scenario both copy pools were produced at the local site using BACKUP SYSTEM FULL, so the database and logs reflect approximately the same time
- For data sharing GRECP recovery is performed automatically during restart and ends with DSNR056I? For each member
- LPL recovery does not occur automatically during restart. You may see a REASON=RESTART during DB2 startup.
- Following restart, issue start commands for any objects still in LPL.

## Active log issues for striping (simple restart)

- We are restoring the log copy pool (and restarting DB2)
  - There is a risk that a newer log CI was written and backed up before an older log CI (for the same active log dataset).
  - DB2 restart cannot handle ...especially if a log record spans multiple CIs.
- For non-data sharing, the FFFFFFFFFFFFF cannot be safely be used for SYSPITR/T
- The only safe point is the Data Complete LRSN stored in the BSDS of the system on which Backup System was run.
  - If it is an LRSN, it is common to all members' active logs
  - No calculations and searching through Print Log Maps unnecessary!

## Tape Recovery in z/OS 1.12

- DFSMShsm supports recovery of copy pools on tape
- ➔ Now DFSMShsm parallelizes tape recovery for FRRECOV
- Makes restore procedure at DR site now easy
- Before
  - DFSMSdss used for FRRECOV was single threaded recovery
  - RESTORE SYSTEM necessary 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



## Notes

- If active logs are striped, use only the Data Complete LRSN (which is an RBA if data sharing is not used) – set it as SYSPITR on all members.
- For SYSPITRT (before Data Complete LRSN) you would have to ensure that the timestamp is covered by the active logs for each member and that there are no "gaps" (some log CIs not written/backed up) prior to the timestamp.
- How many transactions will be lost? Can't know so choose a non fuzzy point. It depends on
  - how much logging activity is occurring on each member
  - how many FC volumes are in the log copy pool
  - which FC volumes were FCed first.

## Restoring the SLB for DR (DB2 9/10) – z/OS 1.12

- **Scenario 1:** Simplified scenario for DR when you want to use the last SLB and restart
  - DFSMSHsm supports parallel recovery of copy pools on TAPE
- hsm FRRECOV COPYPOOL(DSN\$LOCATION\$LG) FROMDUMP
- hsm FRRECOV COPYPOOL(DSN\$LOCATION\$DDB) FROMDUMP
  
- Conditional restart using Data Complete LRSN in BSDS
  - CRESTART CREATE,ENDLRSN=DataCompleteLRSN
    - Truncates the log
- Start DB2

## More about conditional restart with copy pools

- You can use conditional restart locally (data sharing or not)
- Before SLB, using manual **Set Log Suspend – FlashCopy volumes, Set Log Resume**, all volumes represent the same point
  - When volumes restored, DB2 start made transactions consistent
- Under SLB, dump of volumes within the copy pools unordered, as DFSMSHsm parallelizes copy of the volumes
  - “fuzzy” since copied at different points in time – simple restore insufficient
- DB2 compensates by preventing certain activities (updating of the restart/redo LRSN during checkpoints, pseudo-close, physical close)
  - A database copy pool restoration followed by a DB2 ENDRBA (not SYSPITR)
    - with a log truncation point of the Data Complete LRSN of the SLB
    - ensures all objects in the system have transaction consistency

## With z/OS 1.12

## Restore and Restart

## Without z/OS 1.12

- hsm FRRECOV  
COPYPOOL(DSN\$LOCATION\$LG)  
FROMDUMP
- hsm FRRECOV  
COPYPOOL(DSN\$LOCATION\$DB)  
FROMDUMP
- Conditional restart using Data  
Complete LRSN in BSDS
  - CRESTART  
CREATE,ENDLRSN=DataCompleteLRSN
    - Truncates the log
- Start DB2
- Define work DB objects
- Recover any objects RECP/RBDP

- Restore the log copy pool by VOLSER
  - hsm LIST COPYPOOL  
(DSN\$location\$LG)  
ALLVOLS(ALLVERS)
    - Gets volsers for the log copy  
pool
  - Restore active log/BSDS by volume  
(DSS)
    - FRRECOV TOVOLUME  
(volser) FROMDUMP
- Create a SYSPITR/T using the “Data  
Complete LRSN” in the BSDS
  - CRESTART  
CREATE,SYSPITR=DataCompleteLRS  
N
- Start DB2
- Submit Restore System Utility  
(FROMDUMP)
- Stop and start DB2 for normal activity
- Define work DB objects
- Recover any objects RECP/RBDP

## Notes

- When the log copy pool was FRRECOVed (prior slide), you get the current BSDS / logs.
- Set up for conditional restart like traditional DR, documented since 1992 in the *DB2 for z/OS Administration Guide* for your version.
- If RESTORE SYSTEM is used, change DSNZPARM to ADMT=, so that the xxxxADMT ASID does not start with DB2. Else ICF catalog will be allocated and RESTORE SYSTEM fails.
- GRECP
  - Not set for SYSPITR startup because we know RESTORE SYSTEM will resolve them by applying log.
  - If conditional restart, resolved by DB2 during/following restart
- After stop/start of DB2 for normal activities
  - Term any stopped utilities
  - Dis DB(\*) sp(\*) restrict(\*) limit(\*)
  - Any objects from Log NO events during log apply processing RECP

## Scenario 2 - RESTORE SYSTEM (apply log)

Combine copy pools with archive logs sent to recovery site

- Using a remote automated tape library
- Exporting ICF catalogs frequently

### 1. Restore Log Copy Pool:

```
hsm FRRECOV COPYPOOL(DSN$LOCATION$LG) FROMDUMP
```

### 2. Set DSNZPARM RESTORE\_RECOVER\_FROMDUMP=YES

### 3. Prepare to use the archive log for more current recovery

- Repr BSDS from last archive
- Delete/Newlog add active logs without ranges
- Create a SYSPITR/T at the end of the archive
- Run Change Log Inventory

## Scenario 2 - RESTORE SYSTEM (apply log)

4. Start DB2
  5. Submit Restore System Utility (FROMDUMP)
  6. Stop and start DB2 for normal operation
  7. Define work DB objects
  8. Recover or rebuild objects in RECP/RBDP
- More complex to implement, longer RTO, but more current RPO

## Notes

- When the log copy pool was FRRECOVed, you get the current BSDS / logs.
- Creation of SYSPITR same as traditional DR (end of log common to all members)
- Change DSNZPARM to ADMT=, so that the xxxxADMT ASID does not start with DB2. Else ICF catalog will be allocated and RESTORE SYSTEM fails.
- GRECP isn't set for SYSPITR startup because RESTORE SYSTEM resolves them by applying log.
- After stop/start of DB2 for normal activities
  - Term any stopped utilities
  - Dis DB(\*) sp(\*) restrict(\*) limit(\*)
  - Any objects from Log NO events during log apply processing may be RECP
- There are likely more objects in exception states the more log that is processed for this scenario. For that reason I recommend daily Backup System processing vs. every 3 days to a week.
  - For the latter, RTO is increased substantially. Additionally, there are more objects that need recovery from image copies for a week than for a day.



## Notes

- **RESTORE\_RECOVER\_FROMDUMP**

Is a DSNZPARM designed for use in a disaster recovery situation. After the production DFSMSHsn environment is restored, a DB2 RESTORE SYSTEM is done. The only SLB copies are on dump tapes. The parameter forces a recovery from tape even if the RESTORE job is set up to use DASD.

Without the parameter, a restore system will flash back the targets even though they have no data. In other words, the source volumes would end up with junk and RESTORE SYSTEM would finish with a RC0.

Do not use this parameter in a local production environment.

# Summary

## Summary

- Ease of use by DB2 personnel
  - Difficult setup for storage personnel
  - Can take months
  - Recovery Expert and other tools can aid
- Faster restoration of DB2 copy pool backup storage group from tape now parallelized with one FRRECOV command
- Complex DR procedure becomes simple for Scenario #1.
- No need to recover individual table spaces, rebuild indexes, as in the past PTAM methods
- No outage to produce SLB as in Set Log Suspend – Resume scenarios
- Good solution between tape recovery and mirroring scenarios
  - See SLB Update presentation for latter

## Summary

- Scenario 1 - Easy DR with z/OS 1.12
- Scenario 2 - For DR with log apply, run Backup System daily
  - The less frequently run, the more log you must have on hand, and the more objects will be in RECP or other exception states due to LOAD/REORG
  - You will need inline image copies at the recovery site (automated tape libraries)

## Scenario #1 compared to Tape based DR

- RPO is <24 hours behind in both cases
  - SLB tapes and image copies/archives log transported daily
- No existing computing facility infrastructure is assumed in either case
- Traditional D/R RTO is 24 hours to be recovered and accepting new work.
- This scenario takes only FRRECOV restore time (depending on number of volumes). New work is accepted after restart.
  - May reduce RTO by 12x (if 2 hours to FRRECOV)
- It eliminates labor costs of manual procedures.
- **It offloads work from DB2 staff to storage systems staff!**
- **May be a good improvement in your processes without mirroring.**