



DB2 for z/OS: [Some] Backup and Recovery Utility Enhancements in V8/9/10

Bart Steegmans
DB2 for z/OS L2 Performance
March 2011

Disclaimer

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, or shall have the effect of:

- Creating any warranty or representation from IBM (or its affiliates or its or their suppliers and/or licensors); or
- Altering the terms and conditions of the applicable license agreement governing the use of IBM software.


Many thanks to Haakon Roberts and Jim Teng for allowing me to ‘re-use’ some of their material

Agenda

- **DB2 for z/OS V8 Recap**
 - DB2 Managed FlashCopy Solution in V8
 - Backup System and Restore System Utilities
 - FlashCopy Preserve Mirror in DS8K 4.2
 - Miscellaneous
- **DB2 9 enhancements**
 - Automatically manage Copy Pool backups to tapes
 - Allow table space recovery using System Level Backups
 - Incremental FlashCopy
- **DB2 10 utility enhancements**
 - Flashcopy image copies and inline copies
- **Summary**

DB2 Managed FlashCopy Solution in V8

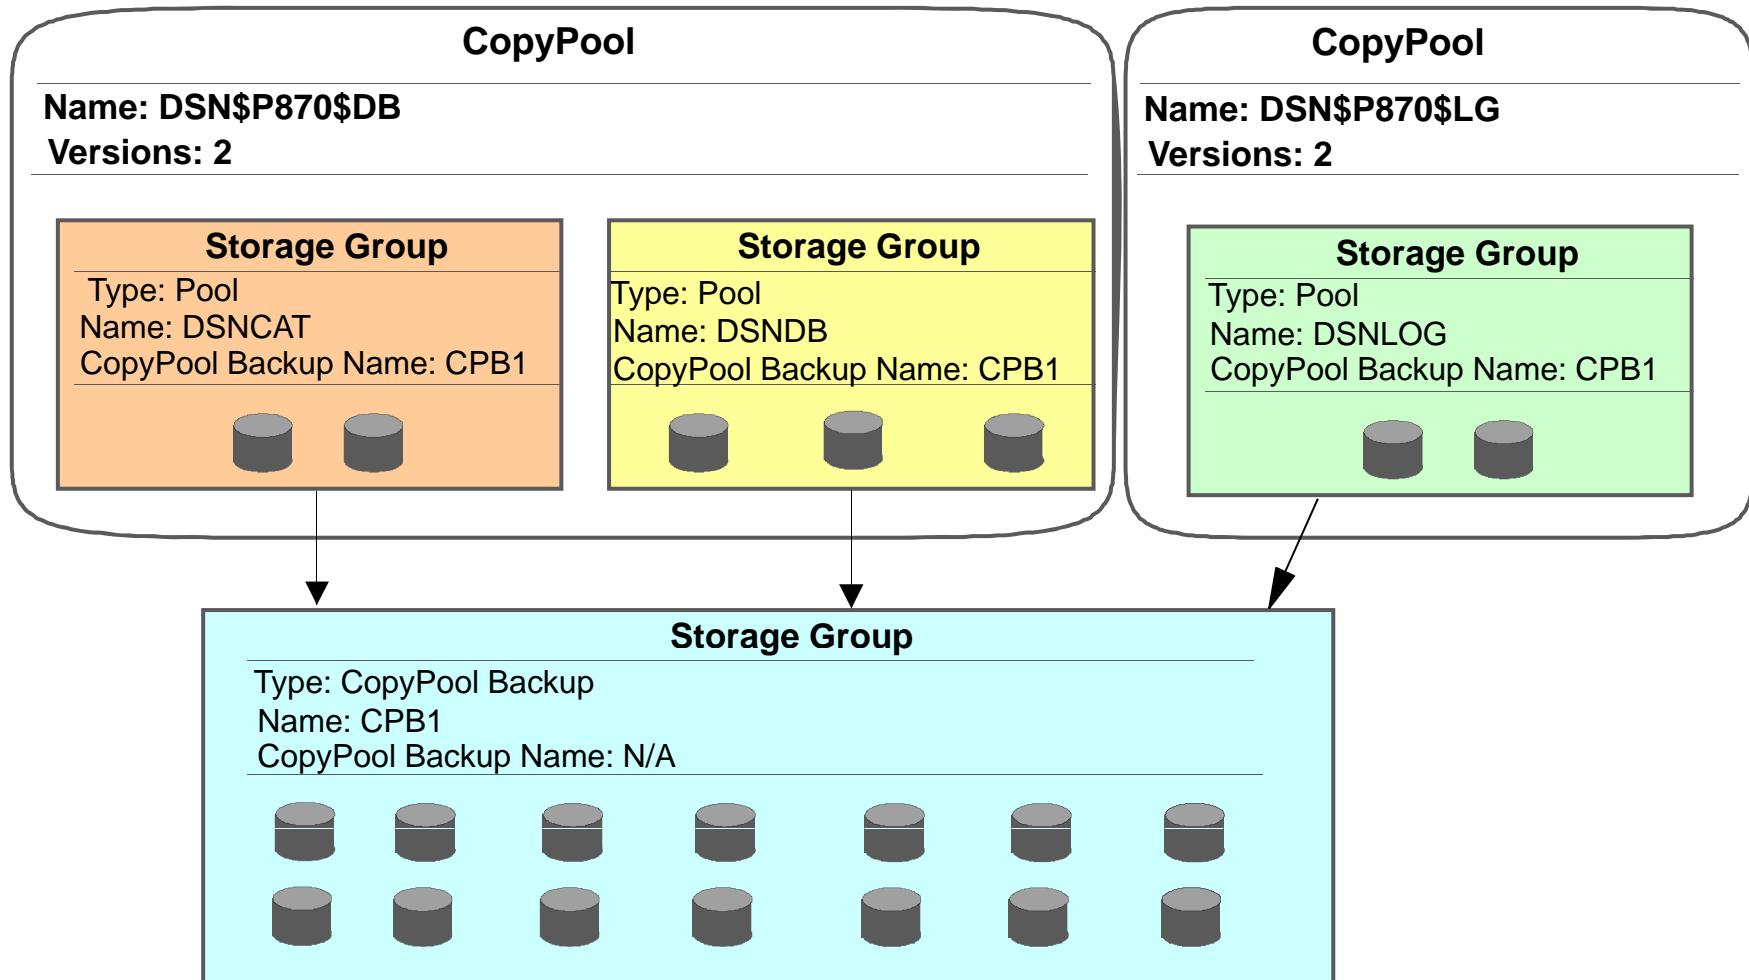
- Provide an easier and less disruptive way for fast volume-level backup and recovery
- Use FlashCopy technology to backup DB2 data and logs
- No longer need to suspend DB2 logging
- Backups are managed by DB2 and DFSMSHsm to support system level PIT recovery, Disaster Recovery and Cloning



DB2 Managed FlashCopy Solution in V8

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

SMS COPYPOOL - Example



Preparation of the DB2 Environment

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

Backup System Process

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

Point-in-time Recovery of the Entire System

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

Restore System Utility

- **To Establish the 'PITR' conditional restart record**
 - CRESTART CREATE SYSPITR=log-point
- **Allow to specify “FF...F” in V9 (PK51979)**
 - Each data sharing member must have SYSPITR created
- **Start DB2 with a PITR CRCR**
 - Implicitly apply DEFER ALL, FWD=NO, Access(Maint)
 - Install SYSADM authority required
 - Can only run “Restore System” utility
- **Write logs to rollback uncommitted changes**
- **Make sure ICFCTLG data sets are not opened**
 - F CATALOG,OPEN to query status
 - F CATALOG,DEALLOCATE to close

Restore System Utility

- **Run the RESTORE SYSTEM utility**
 - Unallocate ICFCTLG - z/OS 1.11
 - Restore Data CopyPool backup
 - Apply logs from RBLP to SYSPITR value of CRCR
- **Objects not recovered from “log no” events**
 - LOAD LOG NO, REORG, CREATE INDEX
 - RECP for table spaces and “Copy Yes” indexes
 - RBDP for “Copy No” indexes
- **If there are objects in RECP/RBDP, the utility will complete with Return Code = 4**

PPRC/XRC Restrictions

- **RESTORE SYSTEM [& RECOVER from SLB in V9]** fails if target is a primary in a PPRC relationship, or part of an XRC relationship
 - DFSMSHsm APAR OA23849 permits FlashCopy to a PPRC primary
 - But volumes will be in “duplex pending” state until background copy completes
 - GDPS / Hyperswap failover fails for volumes in Duplex Pending
 - Resolve this issue by shipping function rather than data through PPRC (FlashCopy Preserve Mirror solution – next foil)
 - No XRC solution in the short term

FlashCopy Preserve Mirror

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

V8 Miscellaneous

- **COPY SYSTEMPAGES YES option**
 - Useful when using UNLOAD from image copy
- **RECOVER CURRENTCOPYONLY**
 - Performance enhancement when using CONCURRENT copy
 - No 'fallback' to prior copy in case of a failure with most recent one
- **REBUILD INDEX SCOPE PENDING**
 - Usability enhancement
- **CHECK INDEX SHRLEVEL CHANGE**
 - PQ96956
- **COPY performance with large LISTDEF lists**
 - PK78865 (V8 & V9)
 - Reduce writes to SYSUTILX

V9 Enhancements

- **BACKUP SYSTEM & RESTORE SYSTEM enhancements**
 - Support for tape (up to 5 tape copies)
 - Support for incremental FlashCopy
- **Object-level recovery from system-level backup**
- **Prereq. DFSMSHsm and z/OS 1.8**
- **RECOVER to any point in time with consistency**
 - Avoid QUIESCE impact on applications
- **TEMPLATE switching**
 - Dynamic switching between tape and disk
- **COPY**
 - Always uses CHECKPAGE
 - Use BM MRU
 - No longer sets COPY pending flag to avoid unavailability
- **REBUILD INDEX SHRLEVEL CHANGE**

CopyPool Backups on Tapes

- **New BACKUP SYSTEM options**
 - DUMP – First to disk then to tape
 - DUMPONLY – Only copy (a disk version) to tape
- **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 (Disaster recovery)**


RESTORE SYSTEM – Tape Support

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

RECOVER utility – use Copy Pool backup

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

Connecting system-level backups with object level recoveries



RECOVER utility – use Copy Pool backup ...

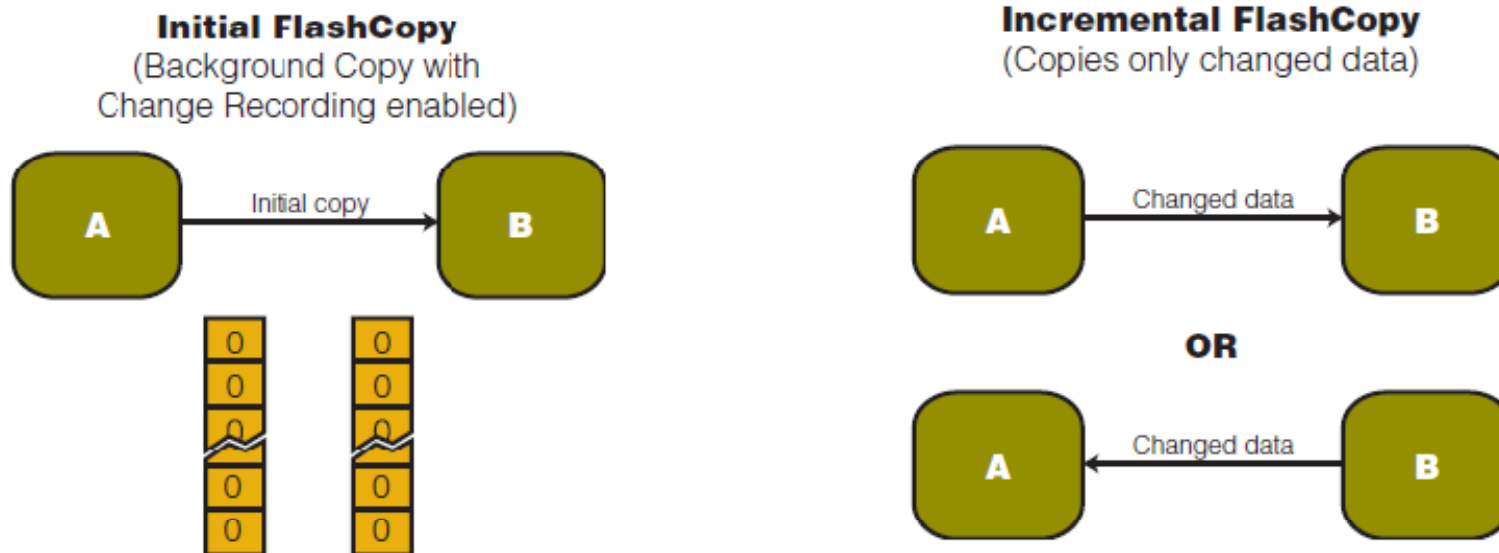
- **Most recent recovery base (prior to the recovery point) is chosen:**
 - could be image copy, concurrent copy, log yes event, or Copy Pool backup
- **Takes sub-second to restore a data set if the backup is on DASD (independent of its size)**
- **Use normal I/O (i.e. not FlashCopy)**
 - If FlashCopy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

RECOVER utility – use Copy Pool backup ...

- **Data set must be cataloged and allocated on the same volumes that it resided on at the time of the backup**
 - Support for data sets that have extended to new volumes
 - DB2 Recovery Expert V2 can alleviate this problem
 - DFSMS support is in z/OS 1.11
 - SMS option to capture ICFCTLG for Copy Pool
 - Allow recovery for moved/deleted data sets
 - Still need to have sufficient free space on the original DASD volume
 - DB2 support in V10
- **If the restore of datasets from DASD fails, then the recovery of the object will not proceed**
 - Use the RECOVER RESTOREBEFORE option to direct the utility to use a recovery base prior to the system-level backup
- **If FROMDUMP is specified:**
 - Data sets are restored from tapes


Incremental Flash Copy

- **Introduced by DFSMSHsm in z/OS 1.8 (+ OA17314)**
 - 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**



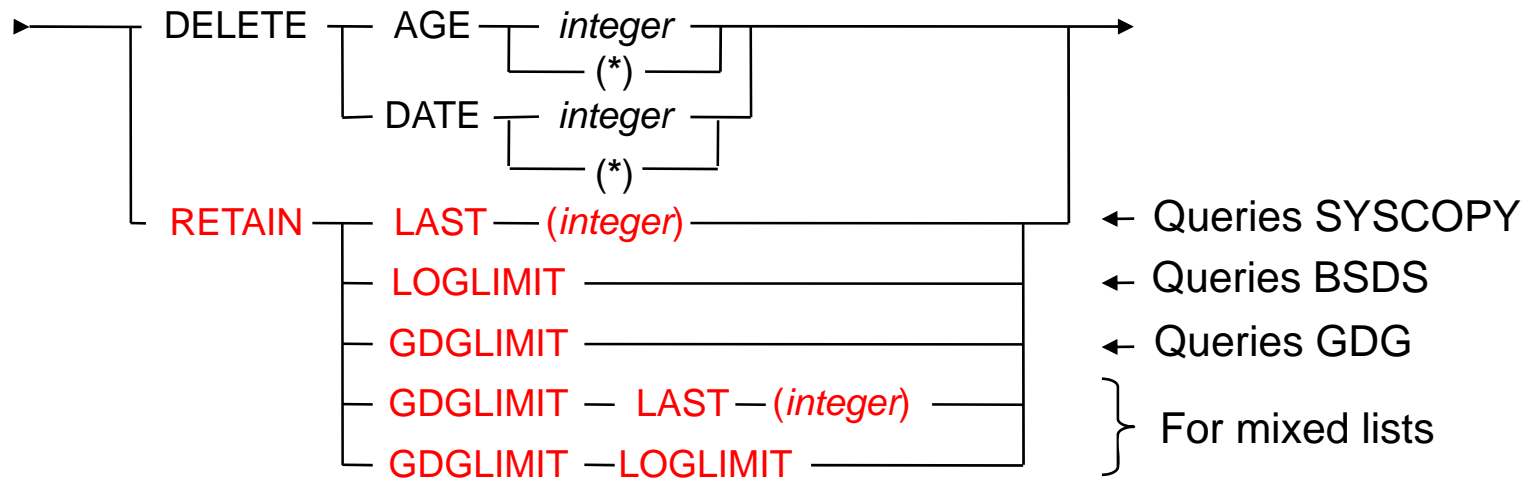
RECOVER to PIT with consistency

- **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.
- **May reduce the need to regularly run the QUIESCE utility**
 - Reduces disruption to DB2 users and applications



MODIFY RECOVERY

- **MODIFY RECOVERY simplification & safety**



- **Template switching for COPY utility**

- E.g. copy to disk if small, to tape if large

```

TEMPLATE LRG DSN &DB..&TS..D&DA..T&TI. UNIT=TAPE
TEMPLATE SML DSN &DB..&TS..D&DA..T&TI. UNIT=SYSALLDA LIMIT(20 CYL, LRG)
COPY TABLESPACE SMALL.TS COPYDDN(SML)
COPY TABLESPACE LARGE.TS COPYDDN(SML)
  
```

V10 Enhancements

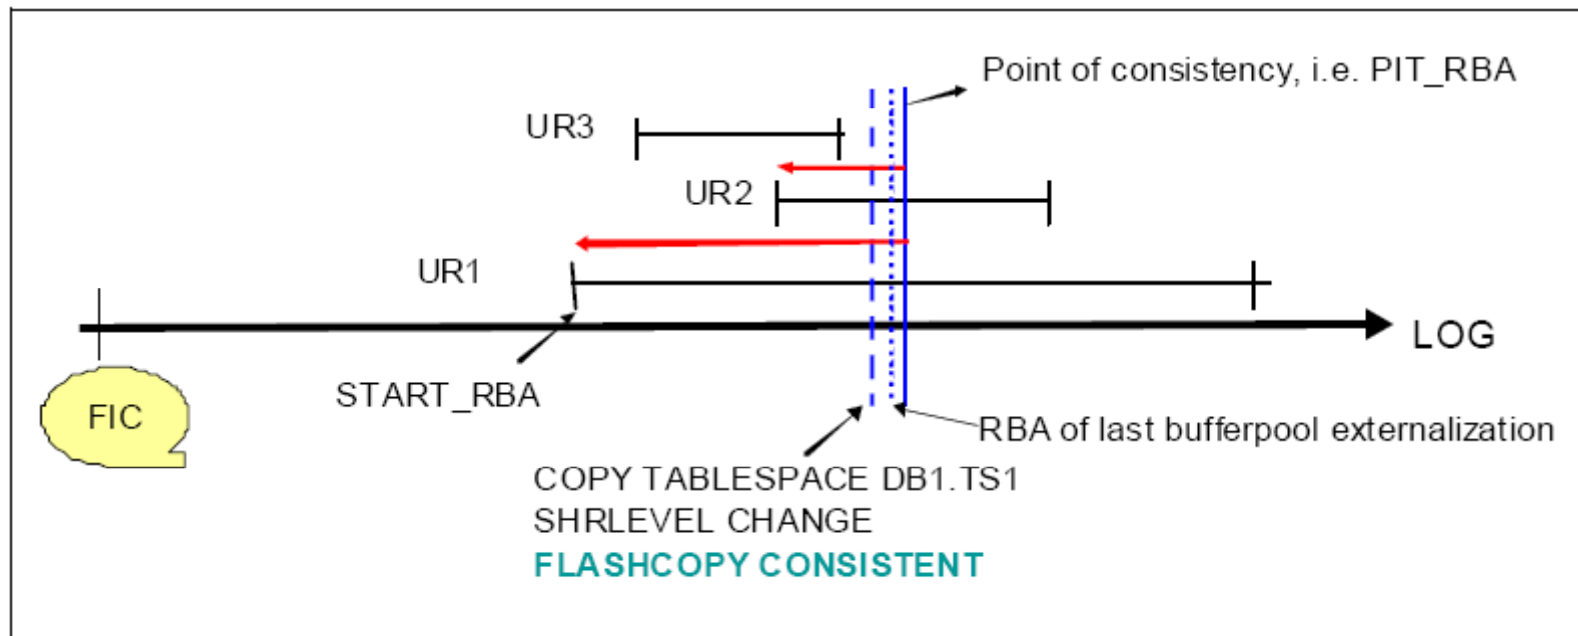
- **Extend COPY, LOAD, and REORG to take backups using data set level FlashCopy**
- **Allow RECOVER utility to use data set level FlashCopy during restore**
- **Data set level FlashCopy backups as input to**
 - DSN1COPY, DSN1PRNT
- **Consistent COPY with Shrlevel Change using FlashCopy**
- **Extend RECOVER to support a point-in-time recovery via “rollback” using logs only**
- **DB2 system checkpoint can be time based and # log records (whichever occurs first) CHKTYPE=BOTH**

COPY

- **Dataset-level FlashCopy support**
 - COPY, RECOVER, REORG, LOAD, REBUILD INDEX, REORG INDEX
 - New zparms & utility parms to govern
 - Dramatic CPU & elapsed time reduction
 - Can make 'sequential copies' from the FlashCopies
- **Improved dataset management & performance**
 - CHANGELIMIT will not allocate copy dataset unless copy taken
 - Incremental copy will not allocate copy dataset unless pages changed
 - Insert dummy SYSCOPY record for incremental copy even though no pages changed

Consistent SHRLEVEL CHANGE image copies

- **Create consistent image copies from SHRLEVEL CHANGE**
 - Via FLASHCOPY CONSISTENT keywords
 - Will not quiesce applications
 - Requires data set level FlashCopy capability



RECOVER

- **New BACKOUT YES option for point in time recovery**
 - True rollback, not run of generated SQL undo statements
 - Requires COPY YES for indexes
- **VERIFYSET option to fail PIT recovery if entire set not included**
 - Base, LOB, XML, history – not RI
- **ENFORCE option to avoid CHKP/ACHKP when subset of set recovered**
 - Applies to Base, LOB, XML and RI
 - Improved performance due to avoidance of set checking (RI, aux)

CHECK

- **CHECK utilities will no longer set CHKP/ACHKP by default**
 - New ZPARM CHECK_SETCHKP to override
- **CHECK SHRLEVEL CHANGE default changed to fail if Flashcopy not available**
 - ZPARM CHECK_FASTREPLICATION (PM19034)
- **CHECK DATA enhanced for XML support**
 - Document validation
 - Schema validation
- **Automated exception table processing for XML documents**
 - INCLUDE XML TABLESPACE option

Other

- **Removed UTSERIAL lock for greater utility concurrency**
- **SQL SELECT on SYSLGRNX (work in progress)**
- **LISTDEF & TEMPLATE enhancements**
 - LISTDEF support for CHECK DATA
 - LISTDEF support for multiple part ranges on REORG
 - LISTDEF support for DEFINED YES|NO|ALL
 - Improved utility performance since unnecessary to build & then discard structures for undefined objects
- **REPORT RECOVERY support for SLBs**
- **Allows implicitly created table space to use TRACKMOD = NO (a new opaque zparm parameter IMPTKMOD)**

HSM Enhancements in z/OS 1.12

- **Allow RESTORE to use Sytem-level backups without waiting for FlashCopy background copy to complete**
 - Use DS8K Fast Reverse Restore (i.e. FRR) feature
- **Allow Backup System to use Space Efficient FlashCopy to keep backups on tapes**

Summary

- **Allow a fast and non-disruptive backup solution using**
 - FlashCopy (incl. Incremental Flash) and DB2 Backup System Utility
- **Copy Pool backups can be used as the source for DB2 table/index recovery**
- **Restore System Utility can recover DB2 system from DASD or Tapes**
- **Recover Utility can recover table space and index to PIT with transaction level consistency**
- **Tight collaboration between DB2 and Storage teams**
- **Ensure utilities are non-disruptive**
- **Continued delivery of performance improvements & features of real value**
- **Toleration, support & exploitation of new features from Day #1**

Upcoming Event April 7 2011 IBM Forum Brussels

- **Morning DB2 for z/OS sessions**
 - *DB2 10 Technical Overview* - Roger Miller, DB2 for z/OS® Technical Evangelist, Designer and Architect, IBM Software
 - *DB2 10 for z/OS Migration Planning* - Roger Miller
 - *How to get the DB2 9 and 10 for z/OS Performance benefits and Reduce Costs ?* - Cristian Molaro - Independent Consultant - IBM Information Champion
- **Afternoon**
 - *Parallel sessions on data governance and analytics*
- **Short link:**
 - <http://bit.ly/hT1nB3>

Information Integration & Governance Forum 2011

Delivering Trusted Information for Smarter
Business Decisions.



Questions ?



Thank you !

Bart_steegmans@be.ibm.com