

# Effective Backups: Selecting the Right Backup Mechanism to Match the Business Requirement

*Session 16723*

*Chris Taylor, IBM Corporation*

*ctaylor1@us.ibm.com*



SHARE is an independent volunteer-run information technology association that provides education, professional networking and industry influence.



## Session Abstract

Backups are a core part of any availability management and disaster recovery strategy. They are also crucial for long-term retention of critical data to conform with regulatory requirements that your company or installation may have. During this session the speaker will discuss different backup options available and their suitability to fulfill the various requirements to ensure a complete and compliant solution.

# Agenda

- Why backup?
- Who is responsible?
- Types of backup
- DFSMShsm
- DFSMSdss
- Other options
- Auditing backups

# Why Backup?

- Availability Management
  - Data has been created and backup is required
  - Data has changed and backup is needed to capture changes
- Disaster Purposes
  - The need to access data at another location
- Long-term retention
  - Ensure compliance for data availability for regulatory purposes

## Chris Taylor “Pearl of Wisdom”

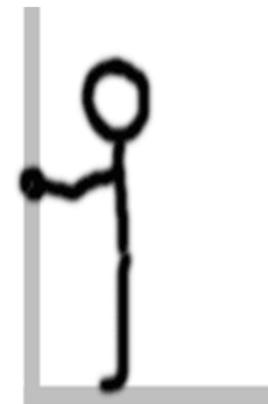
- Data is not important
- Until you need it!!



- We tend to focus on the backup requirements and often neglect recovery considerations

## Monday morning phone call to Storage Admin

- “Job ABC abended over the weekend and we need to recover the data. We thought we had a backup but job XYZ did not run and so we are hoping that you have one in HSM”
- Have you ever received such a call?



## Question #1: Who is responsible?

- Applications or Centralized Storage management?
- Application groups are often responsible for their own backups
  - They have a better understanding of their data
  - They have control over the JCL and choice of backup utility
- Storage Admins control HSM or other centralized backups
  - Ensure backups are taken of critical data
  - Retained for correct amount of time

## Question #2: What kind of backup is needed?

- Point-in-time backup
  - Executed as part of batch process
  - Has to run at a specific point in the cycle
  - Easy recoverability for restart/rerun purposes
- Often under the control of application group
- Can also be used for long-term retention
  - Make sure that correct data is being retained
  - Daily production backups may not suffice for long-term retention

## What kind of backup needed? (cont.)

- Availability Backup
  - Backup can be taken at any time
    - As long as a backup is taken
  - Often performed with DFSMSHsm, FDRABR, CA-Disk, etc.
  - Performed during production “slow” time
- Controlled by Storage Administration group
  - DFSMS constructs ensure centralized policy management
    - Correct retention policies applied
- SMS and HSM are the policy enforcers
  - Business determines the retention policies

## What kind of backup needed? (cont.)

- Long-term Retention
  - Taken for regulatory purposes
    - 7 or 10 years retention requirement not unusual
  - Onsite or offsite storage
    - Will define whether largely utility-based or not
  - If offsite, encryption requirements
- How do you know that the required backup is being retained correctly?
- It is possible to retain data too long!
  - “We never delete anything” probably means you’re not compliant with published retention policies

## What kind of backup needed? (cont.)

- DR Backup
  - Full volume backup approach
  - Flashcopy and subsequent dump
  - Encryption
    - If using physical tape or offsite disk storage
    - Cloud??

# Using DFSMShsm

# Backup methods using DFSMSHsm

- Data set/Application backups
  - Automatic Backup
  - Data set level backup
  - ABARs
- Full volume backups
  - Automatic Volume Dump
  - Fast Replication/Continuous Data Protection (CDP)

# Automatic Backup

- Controlled by storage admin group
- Usually runs every day
- Time determined by other activities
- Can be run from different HSM hosts within the HSM plex
  - Storage groups can be managed from different systems
    - Test vs production
- Managed by rules defined in management class and storage group

# Benefits of Automatic Backup

- Centrally controlled policy
  - The business defines the policy
  - SMS and HSM enforce the policy
- Number of backup versions and length of time retained determined by business requirements
  - Data retained too long can be a compliance violation
- Once setup, takes the responsibility away from applications
- Activities and results can be monitored by tools
  - Activity logs, SMF records (FSR), etc.

# Automatic Backup concerns during backup

- Activity time may not be appropriate for all applications
  - Nowadays more 7X24 availability necessary
  - Backup may interfere with batch processing
- Data sets could always be in use
  - Standard HSM approach is using an ARCCMDxx parm and applies to all backups
    - Serialization(preferred) could end up in a fuzzy backup
  - ARCBDEXT exit can be used for a more granular approach
- Without the right tools/reporting, backup errors may not be noticed

# Automatic Backup concerns during recovery

- Data set may not be there!
  - Error during backup or in use
- Backup was taken before latest update
- Difficult to recover
  - HRECOVER command syntax?
  - Are end-users allowed to recover data sets?
- If used at DR, application recovery may take a long time
  - Data located on multiple tapes
  - Conversely, tape contention when trying to recover multiple applications



# DFSMShsm data set backups

- HBACKDS command
- Allows data sets to be processed explicitly or using patterns
- Can be setup as part of application process
  - Controlled by application user
- Optional parameters can also be supplied
  - CHANGEDONLY, CC, SPHERE, etc.

## HBACKDS examples

- HBACKDS 'CHRIS.APPL1.\*\*' WAIT
  - Backup data sets using mask and wait for backup to complete
- HBACKDS 'CHRIS.APPL1.\*\*' CHANGEDONLY NOWAIT
  - Send data sets backup command to HSM for changed data sets but do not wait for completion
  - Backup may or may not be successful
- HBACKDS 'CHRIS.APPL1.\*\*' TARGET(TAPE) CC(PREFERRED)
  - Backup to tape with (virtual) concurrent copy if available

## Data set backups (cont.)

- Inline backup
- Also known as ARCINBAK
- Backups can be performed in the middle of a job
- Overcomes some of the problems seen with IKJEFT01 approach
  - Relative generation number not allowed for GDG
  - Enqueue issues with data sets created within same job
    - ARCINBAK will process as unserialized

# ARCINBAK example

```

File Edit Edit_Settings Menu Utilities Compilers Test Help
-----
EDIT      P390.JCL.CNTL (ARCINBK3) - 01.01          Columns 00001 00072
Command ==> _____ Scroll ==> CSR
***** Top of Data *****
000001 //P390I JOB (OCOP), 'CHRIS',
000002 // NOTIFY=P390, REGION=6M, MSGCLASS=X, CLASS=A
000003 //*
000004 //GEN1      EXEC PGM=IEBGENER, REGION=4M
000005 //SYSPRINT  DD SYSOUT=*
000006 //SYSUT1    DD DISP=SHR, DSN=P390.LOG.MISC
000007 //SYSUT2    DD DISP=(,CATLG), DSN=P390.TESTGDG(+1),
000008 // UNIT=SYSDA, SPACE=(TRK, (1,1))
000009 //SYSIN     DD DUMMY
000010 //*
000011 //STEP2    EXEC PGM=ARCINBAK, PARM=('RETAIN DAYS (00001),
000012 //          TARGET (TAPE), CC=(PREFERRED, PHYSICALEND) ')
000013 //ARCPRINT DD  SYSOUT=*
000014 //ARCSNAP DD  SYSOUT=*
000015 //BACK01 DD  DSN=*.GEN1.SYSUT2, DISP=SHR
000016 //*
***** Bottom of Data *****

```

\*DSLIST

## ARCINBAK example comments

- DDs for backup in format BACKxxxx
  - //BACK01 DD DSN=\*.GEN1.SYSUT2,DISP=SHR
- Concurrent copy parms can also be included
  - CC=(PREFERRED,PHYSICALEND)
  - TARGET can be specified
    - TARGET(TAPE)
- Use of TARGET, NEWNAME and RETAINDDAYS can be protected with RACF (or other security system)

## Long-term data retention for data sets using HSM

- Normal HSM backup processing limited to max. 100 versions
  - Greater number rolls off
- Copy data set to new name and migrate for x years
  - Allows both migration and backup
  - GDS, datestamp?

# RETAINDDAYS for long-term retention

Data set backup was enhanced in z/OS DFSMS V1R11 to enable a specific retention period to be assigned to a backup copy

- Specified Retention value overrides Management Class and SETSYS settings for retention
- Retention value can be used to keep a backup copy for a *shorter or longer* than normal period of time

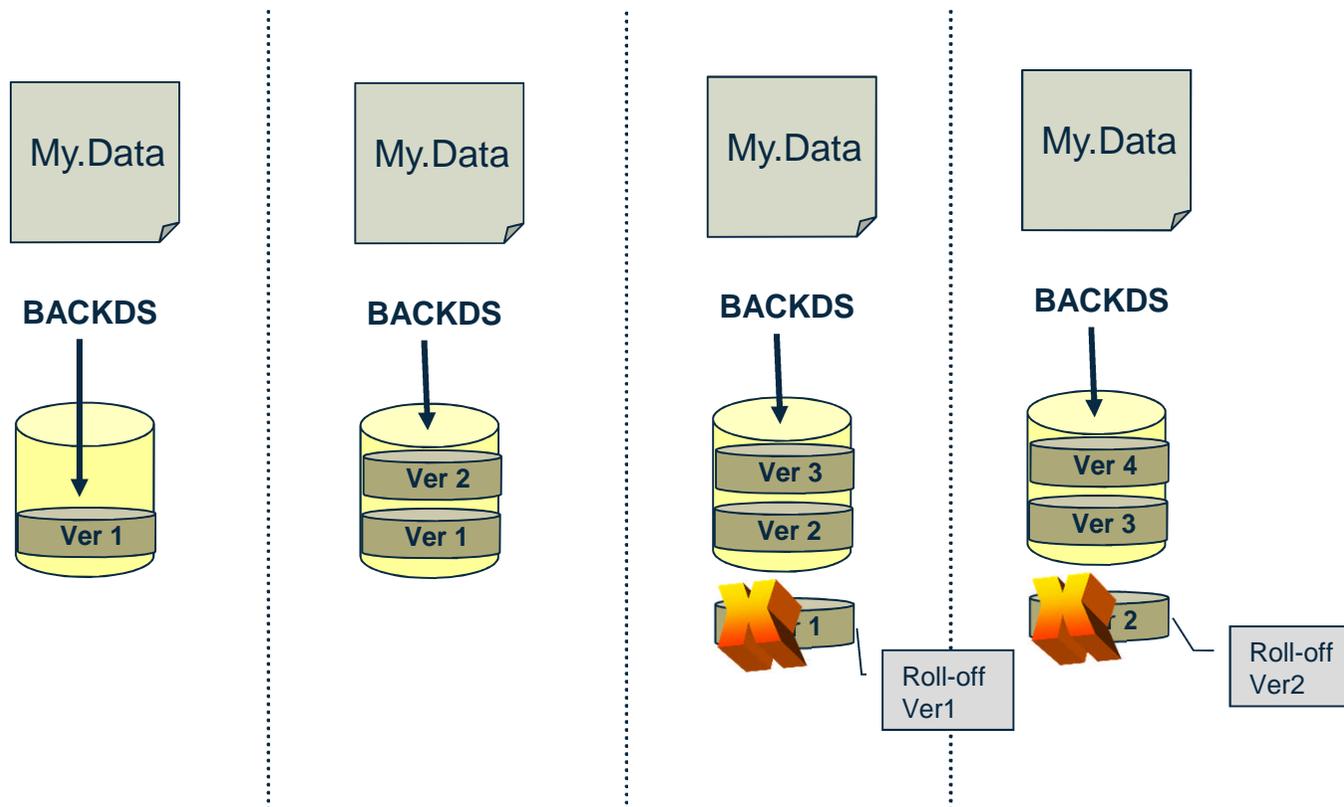
## Data Set Backup Terminology

- **Active Copy:** A backup version that is within the maximum number of backup copies specified by the Management Class or SETSYS value
- **Management Class Retention Period:** The maximum number of days to maintain a backup copy
- **RETAINDDAYS (new):** The minimum number of days to maintain a backup copy. (This value takes precedence).
- **Retained Copy:** A backup copy that has rolled-off from being an active copy, but has not yet met its RETAINDDAYS value

# DS Backup RETAINDDAYS

## How it works normally

- Maximum 2 Backup Copies

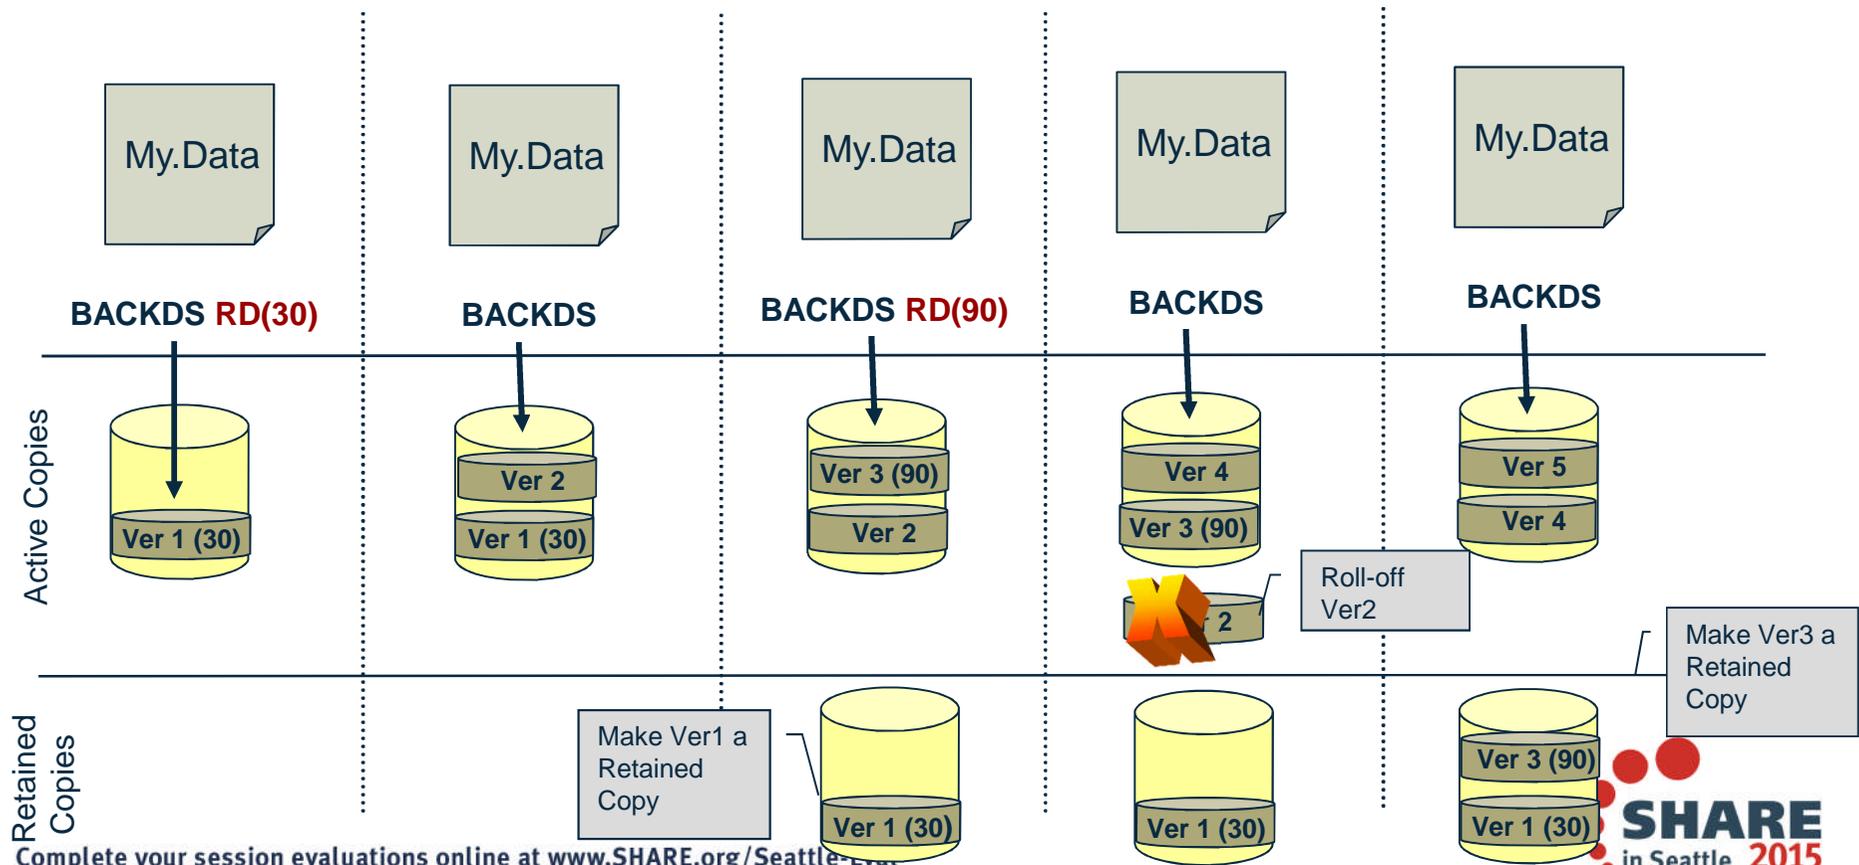


# DS Backup RETAINDDAYS



## How it works with RETAINDDAYS

- Maximum 2 Backup Copies
- Keep Retained Backup Copies *longer* than normal



Complete your session evaluations online at [www.SHARE.org/Seattle](http://www.SHARE.org/Seattle)

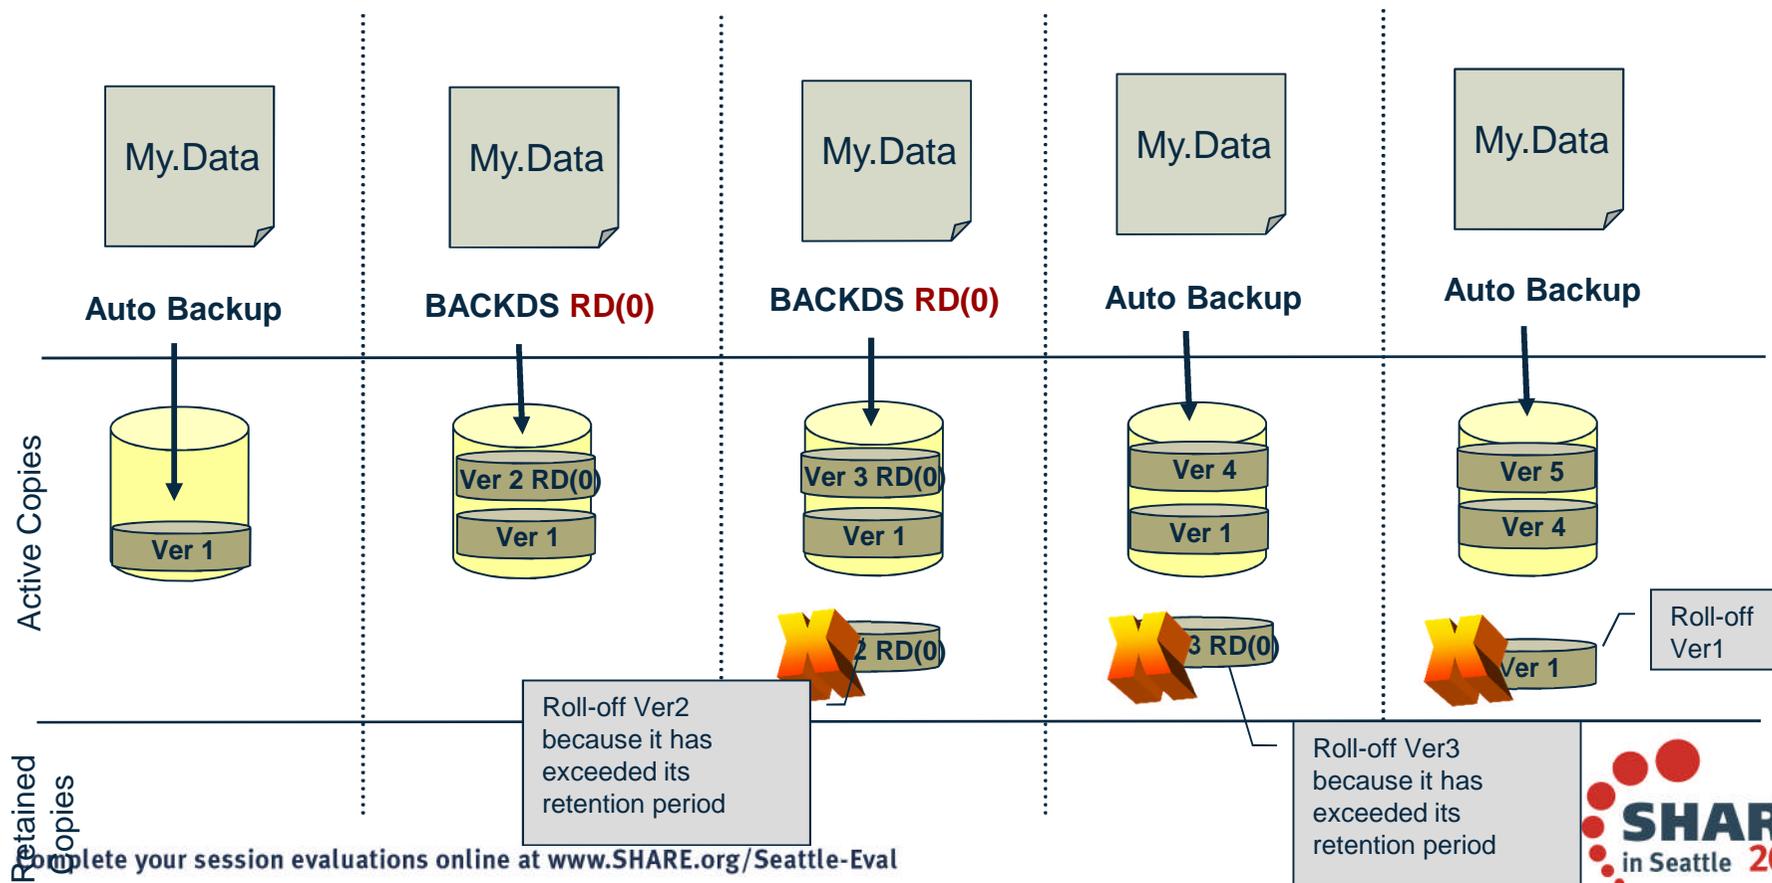
Session 16273: Effective Backups



# DS Backup RETAINDDAYS

## How it works with RETAINDDAYS

- Maximum 2 Backup Copies
- Keep Retained Backup Copies *shorter* than normal



# DS Backup RETAINDDAYS



## RETAINDDAYS are examined during Roll-off and EXPIREBV

- **Roll-Off processing**

- Roll-off processing occurs when the creation of a new backup copy causes the maximum number of 'active copies' to be exceeded
- First, all 'active copies' (except the one that was just created) are examined to determine if any of them have met their RETAINDDAYS value
  - If so, they are deleted
- If the maximum number of 'active copies' has still been exceeded, then all excess versions are examined to determine if they have an unmet RETAINDDAYS value
  - Versions with unmet RETAINDDAYS values are converted to 'retained copies'
    - Backup copy is no longer tracked by version number and is managed via another record internally to HSM
  - Otherwise, the excess versions rolls-off

- ★ **EXPIREBV**

- EXPIREBV *must be run* to expire 'retained copies' that have met their RETAINDDAYS value

# DS Backup RETAINDDAYS

## DFSMShsm can now maintain a *nearly unlimited* number of backup copies

- 'Active Copies' are still limited to 100 per data set
- 'Retained Copies' are nearly unlimited per data set
- New 'MCBR' record – Mirrors MCB record

## RETAINDDAYS(*nnnnn*) can be specified on all Data Set Backup commands

- BACKDS
  - HBACKDS
  - ARCINBAK
  - ARCHBACK
- Can also be applied during Automatic Backup using ARCBDEXT exit

## RETAINDDAYS Values

- Valid values: 0-50000 or 99999
- 99999 means 'Never Expire'

# DS Backup RETAINDDAYS



## ‘Retained Copies’ can only be referenced by Date and Time

- Version number would not be unique
- Version and Generation not listed in LIST output
- Data Set Recover and BDELETE commands updated to accept TIME

## LIST

- By default, both ‘active’ and ‘retained’ backup copies will be listed
- SELECT option enables only ‘active’ or only versions with a RETAINDDAYS value to be listed
  - Using SELECT(RETAINDDAYS) enables you to view all of the backup copies for which a RETAINDDAYS value has been specified

## New Facility Class Profiles:

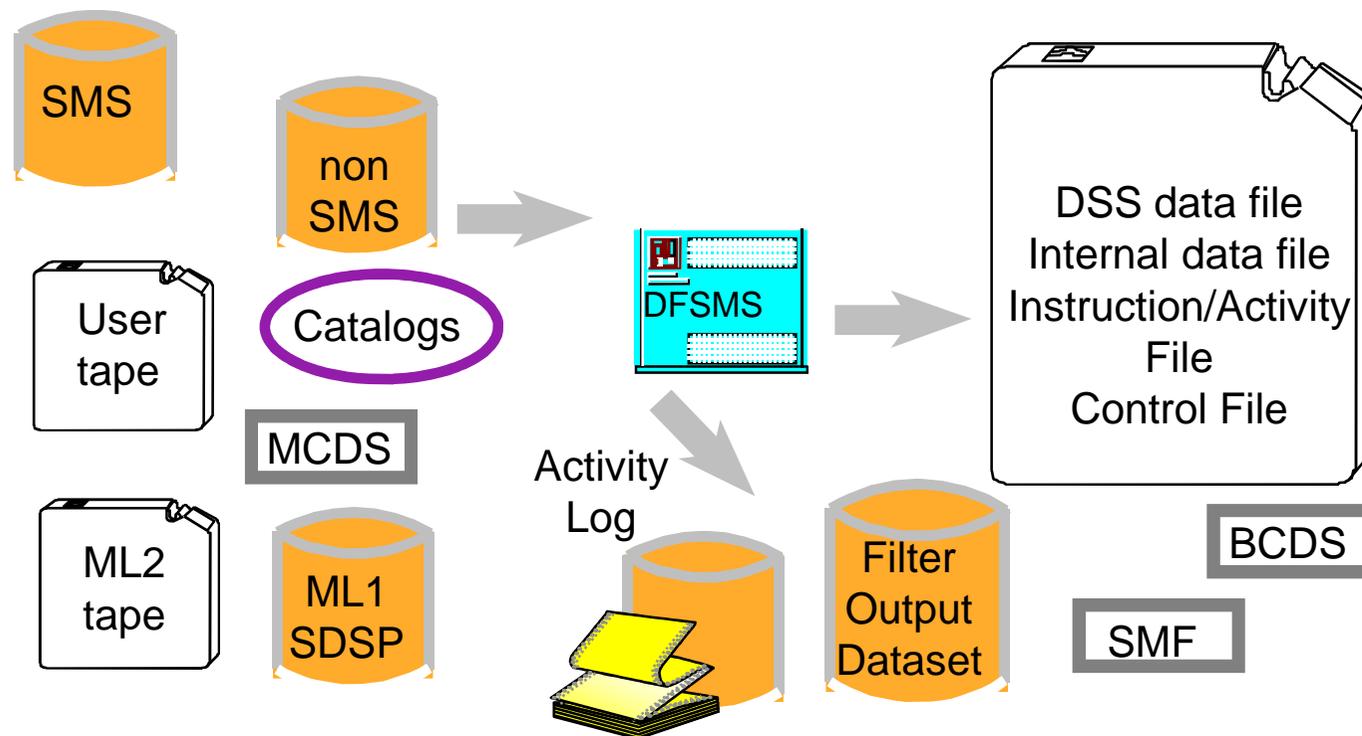
- STGADMIN.ARC.ENDUSER.HBACKDS.RETAINDDAYS
- STGADMIN.ARC.BACKDS.RETAINDDAYS

# ABARS

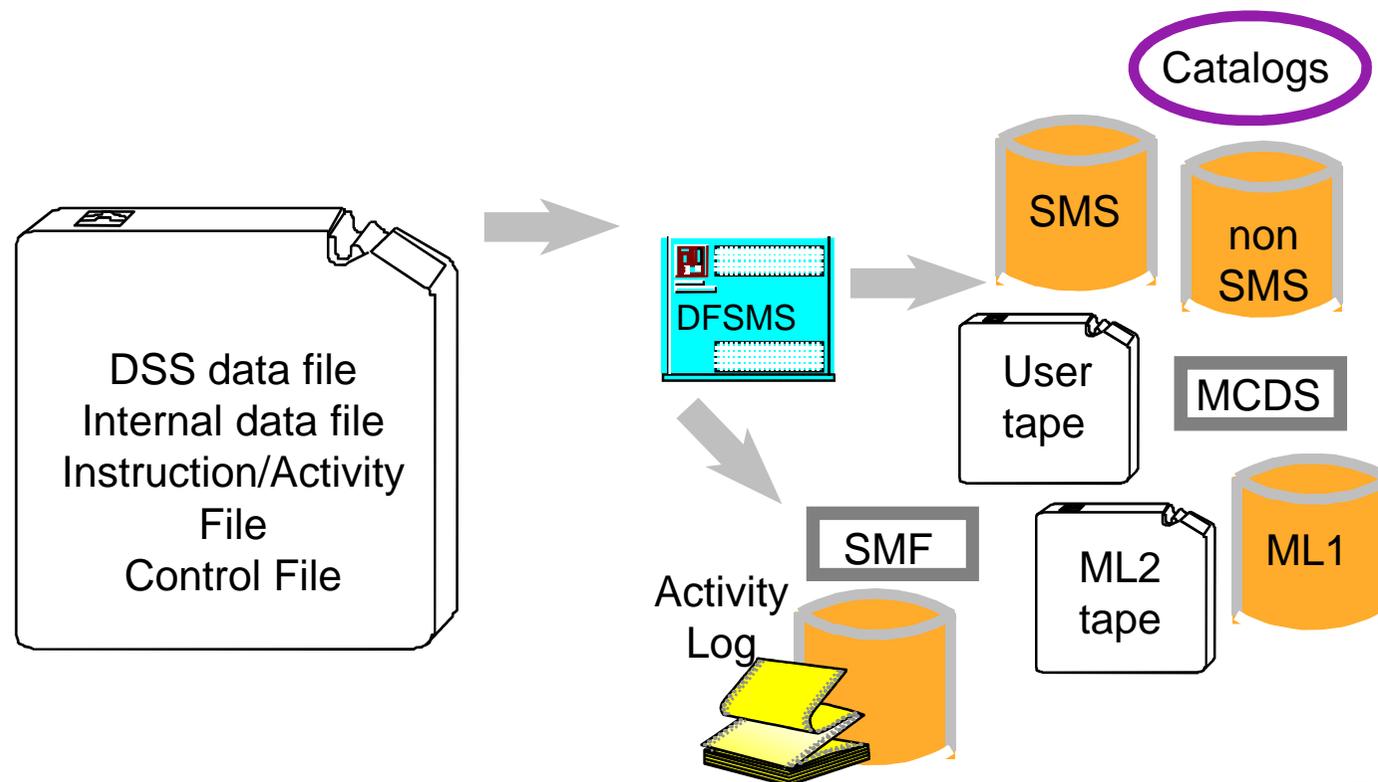


- Allows data sets to be grouped together during backup
  - Easier recovery of application (aggregate)
- Data can be backed up from primary disk, ML1, ML2 and tape
  - Migrated data is processed directly from disk or tape
  - Not recalled to primary disk
- HSM data is not immediately accessible at recovery site
  - ARECOVER needs to run first
- Single data set recovery is possible
- ABARS is best used with an ABARS management tool
  - E.g. Tivoli Advanced Backup and Recovery for z/OS

# The ABACKUP process



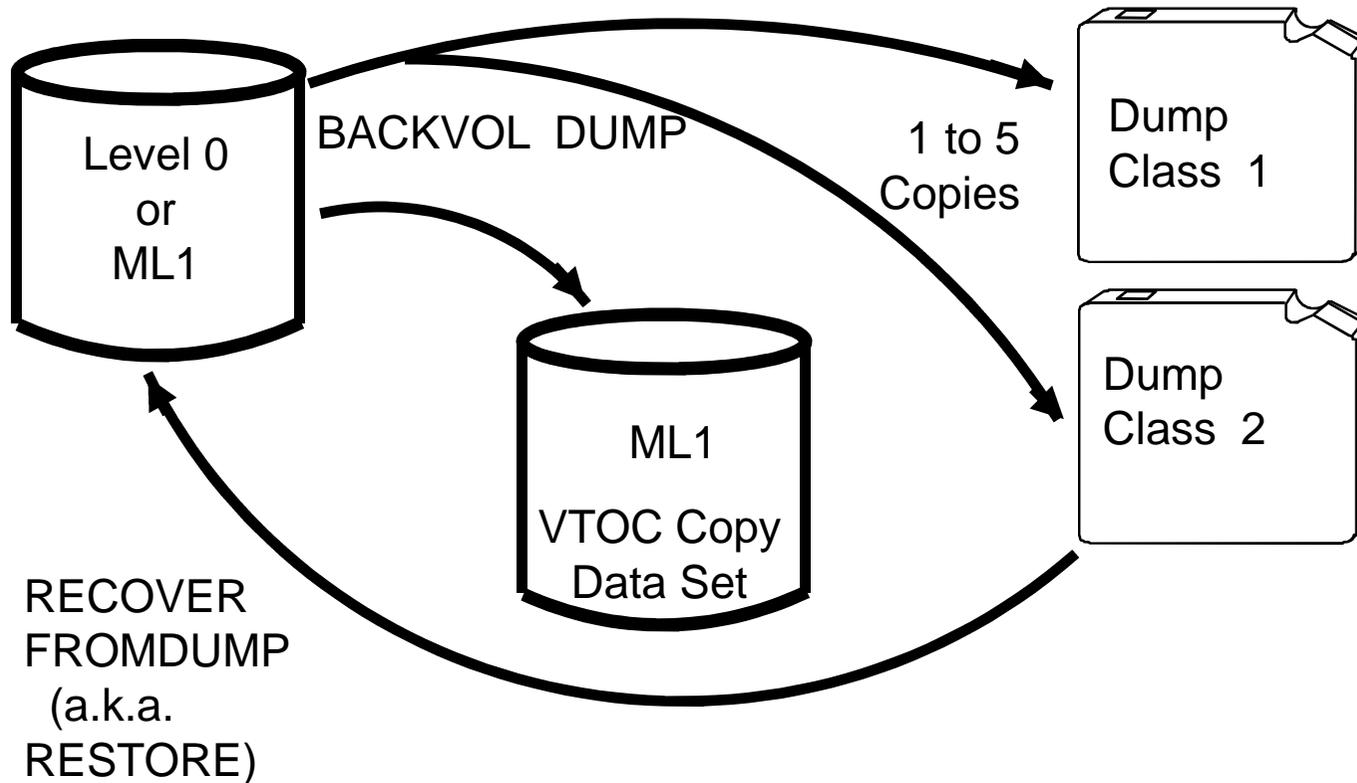
# The ARECOVER process



# Full Volume Dump

- Why Volume Availability is necessary
  - Disk Volume device unavailable
- Full Volume Dump intended for
  - System Packs required for IPL at cold site
  - Migration Level 1 DASD (ML1)
  - Non-SMS volumes with single-volume data sets
  - SMS volumes with read-only or slow-changing data

# Dump Data Flow

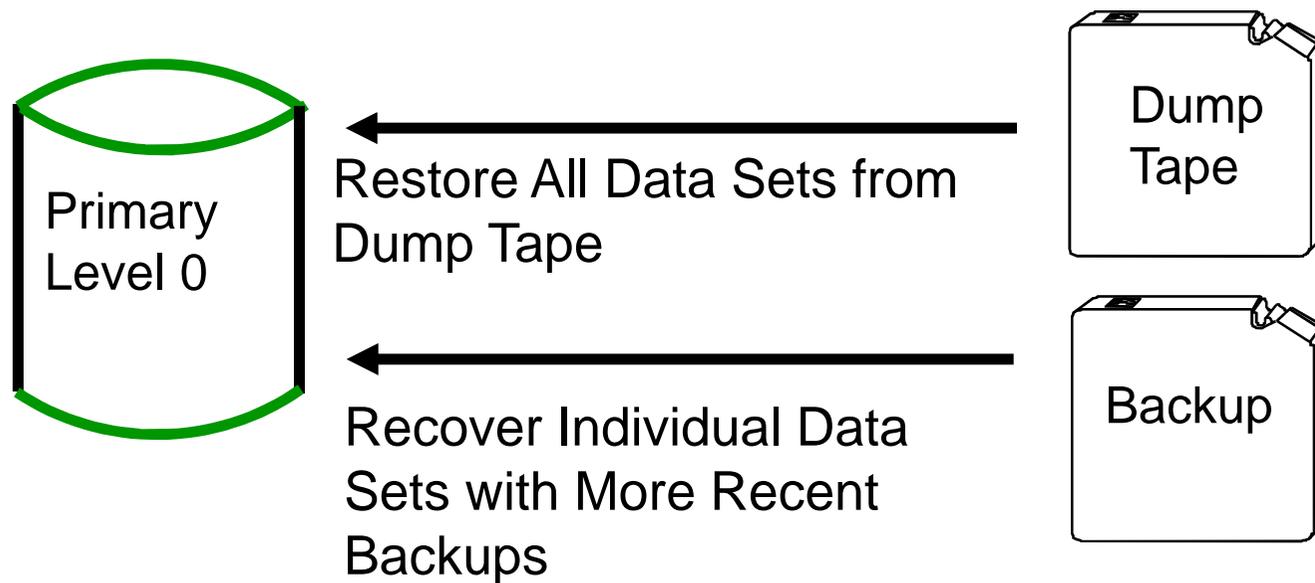


# Types of Full Volume Dump

- Automatic Dump
  - Expire Dump Tapes and Return to Scratch pool
  - Dump Level 0 and ML1 DASD volumes
  - Roll-off Excess Dump VTOC Copy data sets created on ML1 DASD
- Command Volume Dump
  - Specify Volume list or Storage groups
  - Stacking determined by Dumpclass or BACKVOL parameter override

## Scenario - Lost Volume

RECOVER \* TOVOLUME(PRIM37) UNIT(3390)  
FROMDUMP(APPLYINCREMENTAL)



# Business Continuity and Continuous Data Protection

# Business Continuity Overview

## *Business Continuity*

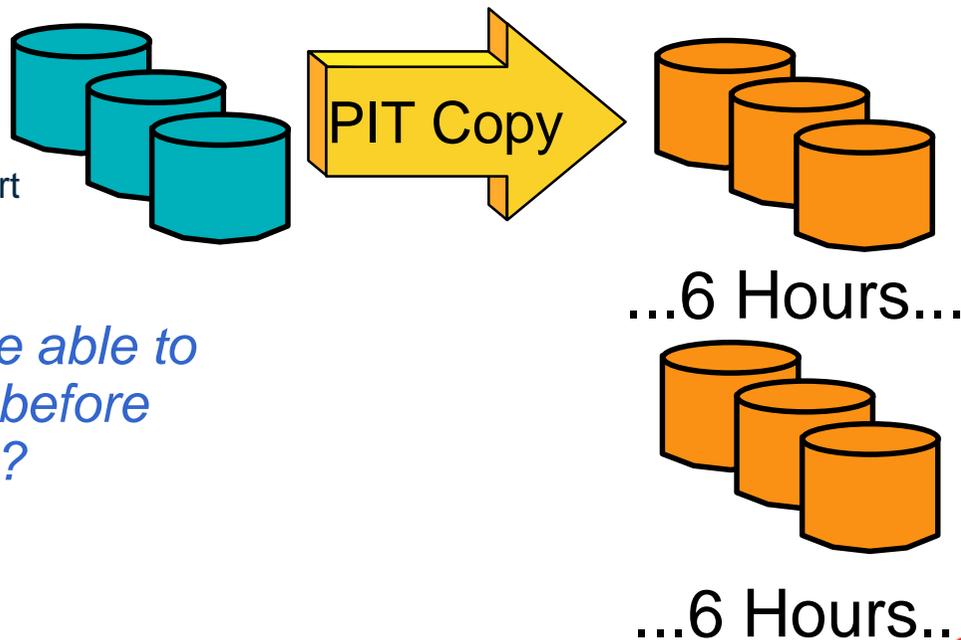
*Maintaining business operations in the event of an outage – with processes and infrastructure that are responsive, highly available and scalable*

### Three key characteristics

- ✓ Recovery Time Objective (RTO)
  - The *acceptable* amount of time you can afford to be without your data
- ✓ Recovery Point Objective (RPO)
  - The amount of data that can be *acceptably* recreated
- ✓ Backup Window Objective (BWO)
  - The *acceptable* amount of time operations can be quiesced to create a copy

# Introduction to Continuous Data Protection

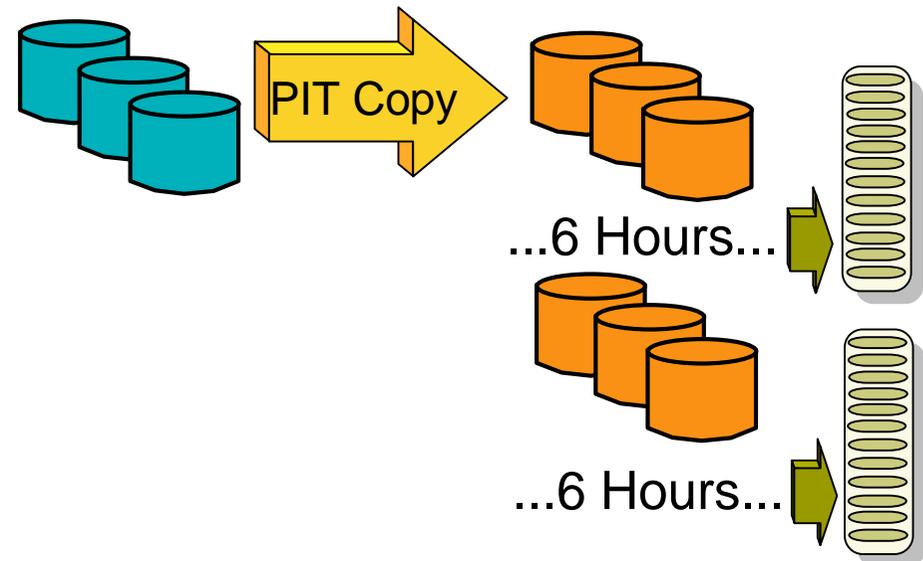
- Traditional Point-in-Time Backup
  - Taken at specific time or data points
  - Only captures data at the point of the backup
  - Long RPO
  - RTO requirements vary
    - Disk – Short
    - Tape - Long
  - BWO requirements vary
    - Point-in-Time Copy – Short
    - Standard I/O – Long



- *Wouldn't it be nice to be able to recover to a point right before the data was corrupted?*

# Introduction to Continuous Data Protection *(continued)*

- **Continuous Data Protection (CDP):**
  - ★ Continuously captures all changes
    - Journaling combined with Point-in-Time copies
  - ★ Eliminates backup window
    - Short/Transparent BWO
  - ★ Short RPO
  - ★ *Generally* short RTO
    - Long from tape



# Types of Continuous Data Protection

- Block Based
  - Capture done at *storage level*
  - Time-ordered capture of every block write
  - Capturing process does not '*understand*' the data
    - Post processing may be required for a *data consistent* recovery
  - *True CDP*
- Software Based
  - Specific *software* journals every update
  - Recovery is tightly integrated with the software
  - Enables data consistent recovery
  - *True CDP*
- File Based
  - Runs on application host (Linux, AIX, Windows, etc)
  - Backup created when file is written to disk
  - Policies can be based on needs of various file types
  - *Near CDP*

Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)

Session 16273: Effective Backups

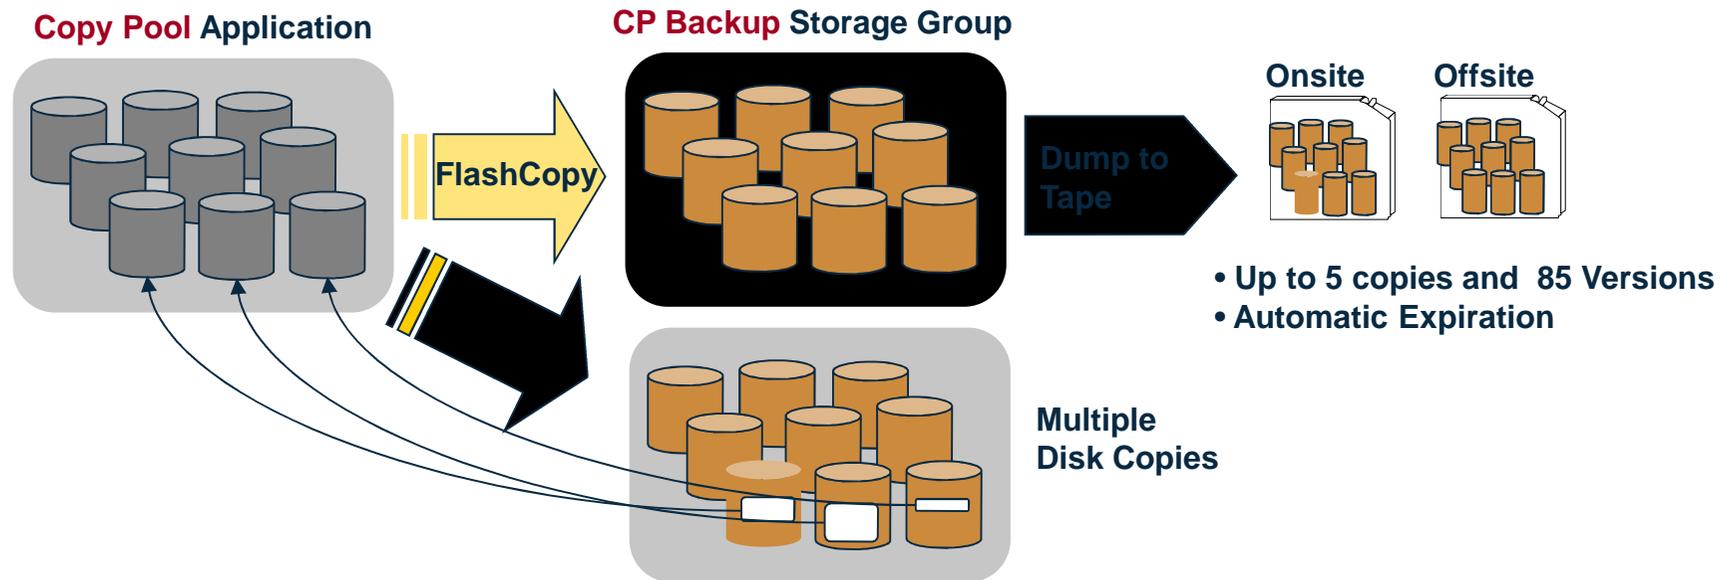
# zCDP for DB2

- Software based CDP for DB2 on System z
  - Joint solution between DFSMS and DB2
- Solution based on Point-in-Time (PIT) backups combined with DB2 logging
  - ★ Eliminates need for DB2 Log Suspend
    - Only Object-level creates, extends, renames and deletes are suspended
    - Hundreds of volumes backed up in a matter of minutes
  - ★ Managed tape copies created from PIT copies
  - ★ Recovery at the System or Tablespace level
- Base Support: DB2 V8, z/OS V1R5 (2003)
- Enhanced Support: Current z/OS and DB2 releases

# zCDP for DB2

## HSM function that manages Point-in-Time copies

- Combined with DB2 BACKUP SYSTEM, provides non-disruptive backup and recovery to any point in time for DB2 databases and subsystems (SAP)

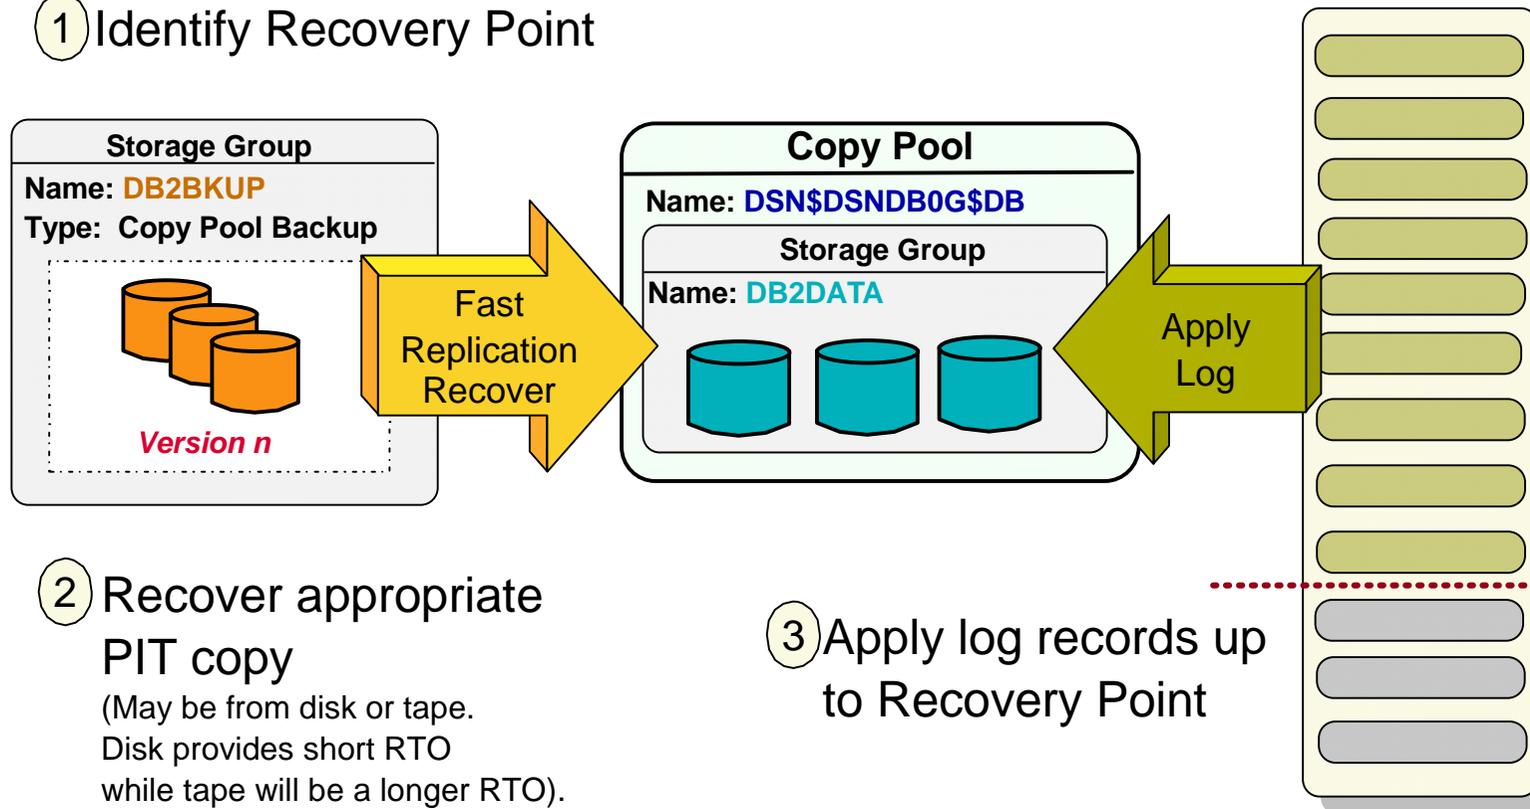


### ★ Recovery at all levels from either disk or tape!

- Entire copy pool, individual volumes and ...
- Individual data sets

# DB2 RESTORE SYSTEM

## 1 Identify Recovery Point

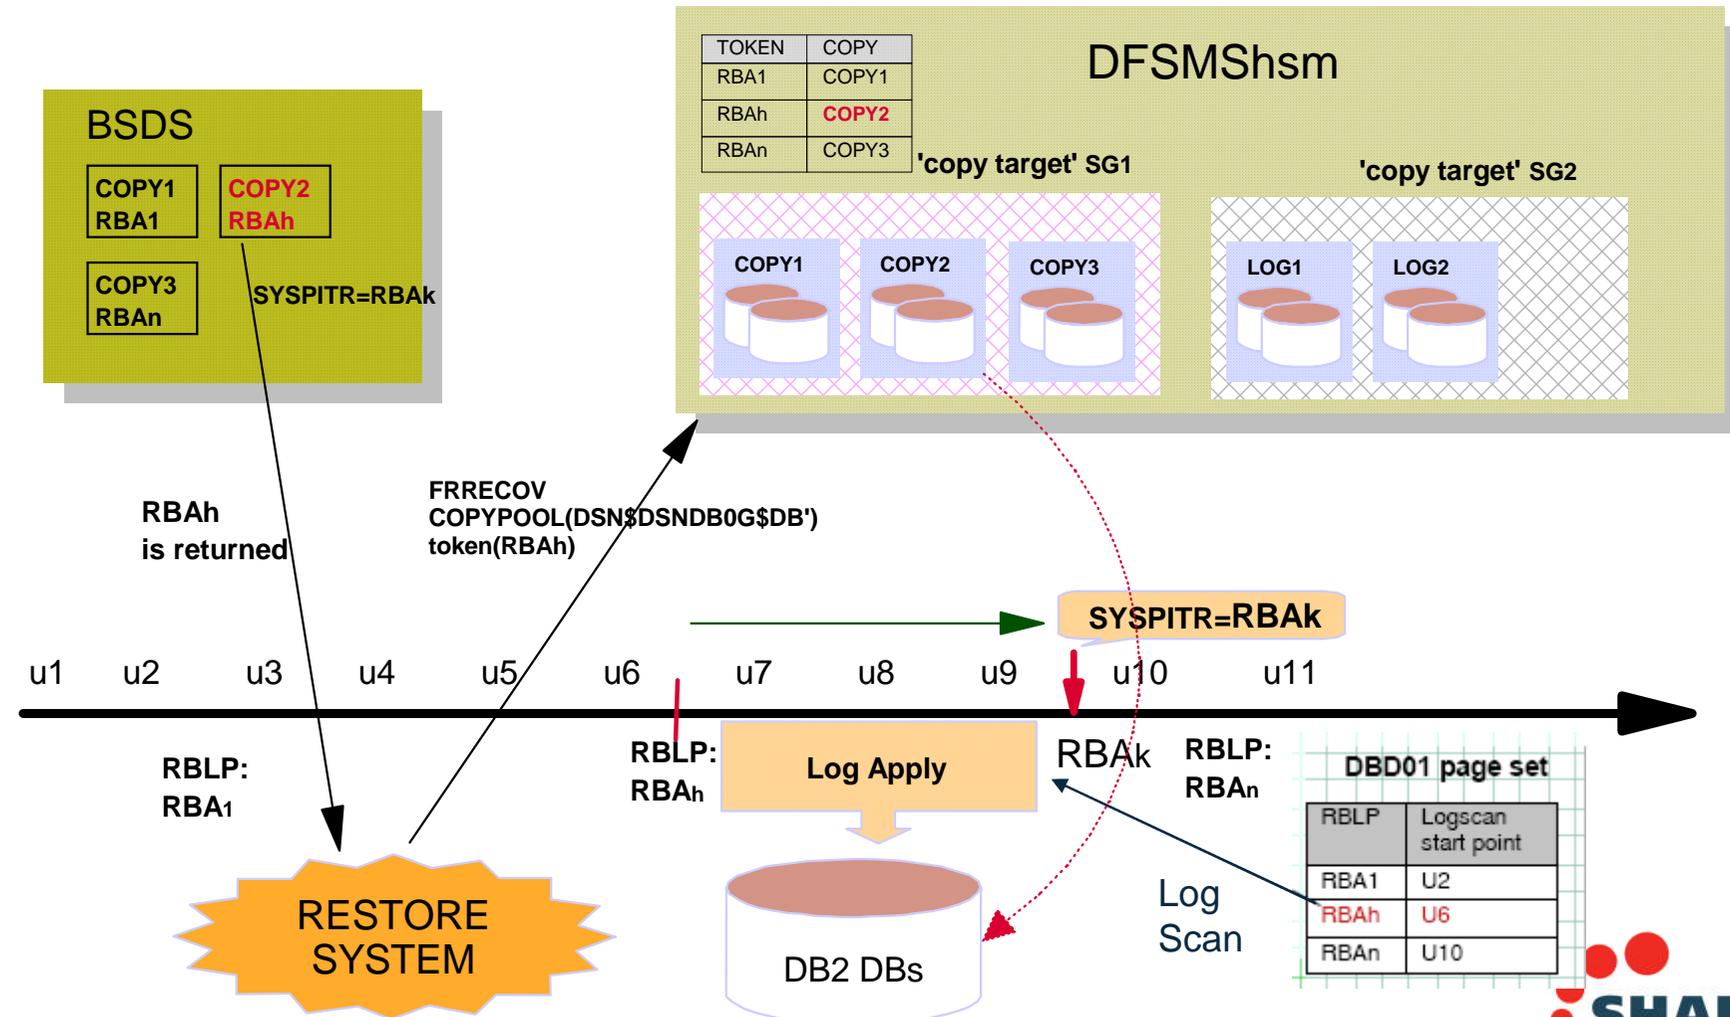


## 2 Recover appropriate PIT copy

(May be from disk or tape.  
Disk provides short RTO  
while tape will be a longer RTO).

## 3 Apply log records up to Recovery Point

# DB2 RESTORE SYSTEM



Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)

Session 16273: Effective Backups

# Using DFSMSdss

# DFSMdss Terms

- **Physical Processing**
  - Volume view for task
- **Logical Processing**
  - Dataset view for task
- **Data Set Filtering**
  - INCLUDE list consists of what data sets you want processed
    - Specific data set names (fully qualified), wildcards (\*, \*\*, %)
  - EXCLUDE list contains data sets you don't want processed
    - Specific data set names (fully qualified), wildcards (\*, \*\*, %)
  - BY Filtering
    - Check data set type, size, catalog status, SMS class, etc.
  - FILTERDD
    - DD Statement points to data set containing INCLUDE/EXCLUDE/BY filtering

# DFSMdss Dump

- Used to backup data to tape or disk
  - One or more input datasets will be written out to a single sequential output file
- Supports data set name and attribute filters
- Process all associated VSAM components using SPHERE keyword
- Will utilize best available technique, Concurrent Copy, SnapShot, Flashcopy

# Dump Keywords (abbreviated list)

- CANCELERROR: Specifies that a dump is to be terminated if an input I/O error is encountered
- CONCURRENT: Specifies that the dump is to use concurrent copy if possible
  - Snapshot may be used depending on input device
- DYNALLOC: Specifies that dynamic allocation is to be used for serialization instead of enqueue
- LOGINDYNAM: Specifies that volumes containing the input dataset names are to be dynamically allocated
- SELECTMULTI: Determines how multivolume datasets are to be processed
  - All: requires that the volume list contain all volumes which contain part of the input dataset being processed (default)
  - Any: requires that the volume list contain any volume which contains part of the input dataset being processed
  - First: requires that the volume list contain the volume which contains the first part of the input dataset being processed
  - Input list to be determined by LOGINDDNAME, LOGINDYNAM or STORGRP
- OUTDDNAME: Specifies the output dataset to be used for the dump
- RESET: Specifies that the change flag be reset for datasets successfully processed
- SHARE: Specifies that dataset is available for read processing during dump operation
- SPHERE: Specifies that all components of a VSAM dataset are to be processed including associated AIX clusters and paths
- TOLERATE: Specifies that processing should continue even if a shared or exclusive access fails

# Filter attributes

- Further filtering supported by data set attribute (BY)
  - Allocation unit
  - Catalog status
  - Backup status
  - Create, Last reference or Expiration date
  - Single or multi-volume data set
  - Data set organization
  - SMS contract name
  - Extent or track utilization
  - Use with EQ, NE, GT, LT, GE, LE

## Filter BY Attribute Examples

- This is a subset of support attribute filters:
  - **CREDIT** Creation date (absolute or relative)
  - **DSCHA** Whether the data-set-changed flag is on or off
  - **DSORG** Data set organization (SAM, PAM, PDS, PDSE, BDAM, HFS, EXCP, ISAM, VSAM or zFS)
  - **REFDT** Last-referenced date (absolute or relative)
  - **DATACLAS** Data class for SMS
  - **MGMTCLAS** Management class for SMS
  - **STORCLAS** Storage class for SMS
- Use with OP keyword
  - **BY((ALLOC EQ CYL) (CATLG EQ YES))** you receive all cataloged data sets with cylinder allocation.
  - **BY(FSIZE GE 100)** you receive all data sets whose size is greater than or equal to 100 tracks.
  - **BY(DSORG EQ (PAM,SAM))** DFSMSDss selects all partitioned and sequential data sets.

# Dump Examples ...

## Example 1: Dumping Data Sets Constantly in Use

```
//JOB1    JOB    accounting information,REGION=nnnnK
//STEP1   EXEC   PGM=ADRDSSU
//SYSPRINT DD   SYSOUT=A
//DASD1   DD     UNIT=3380,VOL=(PRIVATE,SER=111111),DISP=OLD
//DASD2   DD     UNIT=3380,VOL=(PRIVATE,SER=222222),DISP=OLD
//TAPE    DD     UNIT=3480,VOL=SER=TAPE02,
// LABEL=(1,SL),DISP=(NEW,CATLG),DSNAME=USER2.BACKUP
//SYSIN   DD     *
DUMP LOGINDDNAME(DASD1) OUTDDNAME(TAPE) -
      DATASET(INCLUDE(**)) TOL(ENQF) WAIT(0,0)
/*
```

DFSMSDss does not wait (WAIT(0,0)) if a data set is in use. Instead, it processes the data set without serialization or enqueueing (TOL(ENQF)).

# Dump Examples ...

## Example 3: Logical Data Set Dump Operation with Catalog Filtering

```
//JOB3    JOB    accounting information,REGION=nnnnK
//STEP1   EXEC   PGM=ADDRSSU
//TAPE    DD     UNIT=3480,VOL=SER=TAPE04,
// LABEL=(1,SL),DISP=(NEW,CATLG),DSN=USER3.BACKUP
//SYSPRINT DD   SYSOUT=A
//SYSIN   DD     *
DUMP OUTDD(TAPE) -
      DS(INCL(USER1.**))
/*
```

All data sets cataloged in the standard search order whose first-level qualifier is USER1 are to be dumped. Because some of these data sets are multivolume, source DASD volumes are not specified, resulting in data set selection by catalog.

Example 3 can be modified as follows to dump only data sets changed since the last backup. In addition, data sets that end with a qualifier of LISTING are not to be dumped (EXCL(\*\*.LISTING)).

```
//SYSIN   DD     *
DUMP OUTDD(TAPE) -
      DS(INCL(USER1.**)) -
      EXCL(**.LISTING) -
      BY((DSCHA EQ 1))
/*
```

Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)

Session 16273: Effective Backups

# VSAM Backups

- IDCAMS is the standard utility for VSAM processing
- DFSMSdss can also be used to perform logical backups
  - Trying to change options during restore can be “challenging”
- Several VSAM and Catalog products available from vendors
  - Allow additional selection criteria during backup
  - VSAM attributes can be modified during restore
  - Care should be taken when backing up open VSAM data sets
    - “Fuzzy backup” may not be appropriate for the application

# Auditing and monitoring backups

- Consider investing in a backup auditing tool
- It should:
  - Track backups to identify critical data sets that do not have a backup
    - Tracking performed regardless of backup mechanism/tool
      - IEBGENER, DFSMSHsm, SORT, DFSMSDss, FDR, etc.
  - Identify data sets that have duplicate backups
  - Assist with recovery from local, unplanned outages and facilitate remote recovery at DR site
  - Provide reporting capabilities to show backed up data sets
    - Both successful and exceptions

## Reference material

- z/OS DFSMSHsm Managing Your Own Data
  - SC23-6870-00 (V2R1)
  - SC35-0420-10 (V1R13)
  
- z/OS DFSMSHsm Storage Administration
  - SC23-6871-02 (V2R1)
  - SC35-0421-13 (V1R13)
  
- z/OS DFSMSdss Storage Administration
  - SC23-6868-02 (V2R1)
  - SC35-0423-17 (V1R13)

## What was discussed in this session

- Why backup?
- Who is responsible?
- Types of backup
- DFSMShsm
- DFSMSdss
- Other options
- Auditing backups

# Questions?



Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)  
Session 16273: Effective Backups

# Legal Disclaimer

## NOTICES AND DISCLAIMERS

Copyright © 2008 by International Business Machines Corporation.

No part of this document may be reproduced or transmitted in any form without written permission from IBM Corporation.

Product information and data has been reviewed for accuracy as of the date of initial publication. Product information and data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or programs(s) described herein at any time without notice.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Consult your local IBM representative or IBM Business Partner for information about the product and services available in your area.

Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectually property rights, may be used instead. It is the user's responsibility to evaluate and verify the operation of any non-IBM product, program or service.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein.

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

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

Complete your session evaluations online at [www.SHARE.org/Seattle-Eval](http://www.SHARE.org/Seattle-Eval)

Session 16273: Effective Backups

# Trademarks

The following are trademarks of the *International Business Machines Corporation*:

**IBM, DFSMS/MVS, DFSMSHsm, DFSMSrmm, DFSMSdss, DFSMSopt, DFSMS Optimizer, z/OS, eServer, zSeries, MVS, FlashCopy®**

The information contained in this presentation is distributed on an 'AS IS' basis without any warranty either expressed or implied, including, but not limited to, the implied warranties of merchantability or fitness for a particular purpose. The use of this information is a customer responsibility and depends on the customer's ability to evaluate and integrate it into the customer's operational environment.

# Effective Backups: Selecting the Right Backup Mechanism to Match the Business Requirement

*Session 16723*

*Chris Taylor, IBM Corporation*



SHARE is an independent volunteer-run information technology association that provides education, professional networking and industry influence.

