



The Life and Times of a Data Set: You Wouldn't Want Your Relatives Hanging Around, Why Your Data ?

Chris Taylor
ctaylor1@us.ibm.com
IBM Corporation
August 16, 2013
Session Number 13772



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 intellectual 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 sessions evaluation online at SHARE.org/BostonEval



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.

Session Abstract

- DFSMS has been around since the early 1990s and has been implemented by a large number of z/OS customers. However, many sites have not revisited the SMS constructs and routines since they were first implemented. The speaker will review features in DFSMS which you may have missed and suggest updated data set allocation practices. In addition, he will describe how these simple changes can optimize your use of DFSMSHsm.

Agenda

- Birth
- Adolescence
- Adulthood
- Retirement




Agenda

- Birth
 - Allocation and Placement
- Adolescence
 - Usage
- Adulthood
 - Migration
- Retirement
 - Expiration and Deletion



Birth – Allocation and Placement



Birth – Data Class

- Data class ACS is called for both SMS and non-SMS allocations
 - Decision for SMS management is made in storage class ACS
- Data class originally designed as a “fill-in” construct
 - Envisaged as completing JCL attributes
 - Data class used instead of model DCB statements
 - Largely used to assign attributes requiring SMS management
 - Compressed
 - Extended format

Birth – Data Class – Overrides

- Can now be used to enforce some allocation attributes
 - Available since z/OS DFSMS V1R10
- Enforce primary space and secondary allocation quantities
- **Override Space** set to Y
 - If set to Y, then space information is overridden for nonSMS as well as SMS allocations
- **System Determined Blocksize**
 - Set **System Determined Blocksize** to “Yes”

Birth – Data Class – Volume Count

- **Volume Count** has been available in data class since early releases
 - 59 volumes maximum for disk
 - 255 volumes maximum for tape
- **Volume count** assigns candidate volumes in the catalog
 - Seen as VOL001 * * * *
 - Increases catalog space usage, depending on implementation

Birth – Data Class – Dynamic Volume Count

- **Dynamic Volume Count** introduced in z/OS DFSMS V1R3
 - Up to 59 volumes
 - Does not increase the number of candidate volumes in the catalog
- We do not recommend specifying *all* data sets up to 59 volumes
 - Space is reserved in the TIOT
 - For example, if TEMP data sets are included, S878 abends can occur
- More information in Technote
 - <http://www-01.ibm.com/support/docview.wss?uid=isg3T1000065>

Birth – Data Class – Space Constraint Relief

- **Space Constraint Relief**
 - Specifying “Y” removes the traditional 5 primary extents limit on a volume
- **Reduce Space Up to** setting will perform a one-time reduction attempt
 - E.G. 200 cyls would attempt 100 cyl allocation when 50 is specified
 - Using 0 still allows extent limits to be removed
- Use of a spaceabend prevention product offers more functionality

Birth – Data Class – KSDS CA Reclaim

- Available in z/OS DFSMS V1R12
- Reduces need to run offline reorgs of VSAM KSDS data sets
- Activated in IGDSMSxx member in PARMLIB or operator command
 - CA_RECLAIM(DATACLAS|DATACLASS)
 - SETSMS CA_RECLAIM(DATACLAS|DATACLASS)
- Note that **CA Reclaim** default in data class is “Y”

Birth – Data Class as an SMS Decision point?

- SMS data sets are not required to have a data class
- After data set is created, data class cannot be altered
 - Using “official” methods.....
- Using data class as decision point tends to lead to more coding & exclusions later in the ACS process
- Data class ACS is not invoked during DFSMSHsm Recall and Recover
 - Also not for DFSMSdss Restore

Birth – Storage Class

- Determines whether data set is SMS or non-SMS
 - All SMS-managed data sets have a storage class
- Performance characteristics of data sets
- Can be used to supply a specific volume for allocation
- Should be used as main decision point for SMS management

Birth – Storage Class – SMS or Non-SMS

- Specify storage class for SMS-management
- SET **&STORCLAS=**” for non-SMS
- If data class was assigned by ACS routine:
 - Retained for SMS-managed data set
 - Applied but discarded for non-SMS data sets
- For non-SMS:
 - UNIT parameter will be used
 - STORAGE mounted volumes
 - An allocation product can also be used at this point
- Storage class name is placed in both BCS and VVDS

Birth – Storage Class – Performance

- Storage class is a list of storage objectives and requirements
- Provides the criteria to help determine an appropriate location to place a data set
- Attempts to seek a location that meets or exceeds requirements
- If multiple applicable devices, then usually volume with the most available space is selected
- Create a storage class for each level of service

Birth – Storage Class – Availability & Accessibility



- **Availability**
 - Determines whether a data set should reside on a volume with dual copy active
 - Array DASD
 - Replicated volume
- **Accessibility**
 - Allow access to a data set on a volume during copy or backup operations
 - Concurrent Copy/Virtual Concurrent Copy
- **Direct Bias/Sequential Bias**
 - Specify whether a data set is mainly referenced in read or write mode

Birth – Storage Class – Guaranteed Space

- **Guaranteed Space**
 - Used to allocate data sets to a specific volume
- Also used to preallocate primary space on all specified volumes of a multivolume data set
- Use sparingly and limit users allowed to specify this
- Goes against the concept of a pooling product
 - Allocation normally determines data set placement

Birth – Storage Class – Multi-Tiered SG

- Example: SET &STORGRP = 'SG1', 'SG2', 'SG3'
- When **Multi-Tiered SG=Y**
 - SMS uses the order specified in the storage group ACS
 - SG1 will be chosen before SG2, etc.
- When **Multi-Tiered SG=N** (default)
 - Allocation could take place in any of the supplied storage groups

Birth – Storage Group

- Various data types
 - Temporary data sets
 - Permanent
 - Batch
 - TSO
 - Database
 - *CICS/VSAM*
 - *IMS*
 - *DB2*
 - *Other database formats*
- Tape mount management candidates
- Tape Volumes & Objects

Birth – Storage Group – Size-based Storage Groups



- Data set size-based storage groups help deal with free space fragmentation
- Reduce abends due to space constraints
- Reduce need to perform DEFRAIG functions
- When data set is deleted, it can be replaced by one of a similar size
- Use **&MAXSIZE** in ACS routine to determine where data set is allocated
 - Non-VSAM calculation = **&SIZE** + (Secondary*15 extents)
 - VSAM calculation = primary + (123*secondary*volume count)
 - If **Space Constraint Relief** = Y **&SIZE** +(Sec*7257 extents)

Birth – Storage Group – Size-based Storage Groups

- Sample ACS rules can be found in SYS1.SAMPLIB
 - Member DFPSSSGR

```
WHEN (&MAXSIZE => 285MB)                /* Large data sets go to */
DO                                          /* their own storage group */
    SET &STORGRP = 'LARGE90', 'LARGE80'
    EXIT
END

OTHERWISE                                  /* Normal and medium-sized */
DO                                          /* data sets go to the */
    SET &STORGRP = 'PRIME80', 'PRIME90'    /* PRIME storage group */
    EXIT
END
```


Birth – Storage Group - Overflow

- **Overflow**
 - Used to handle periods of high demand for initial primary space allocations
- Still needs to be specified in storage group ACS routine
 - SET **&STORGRP** = 'SG1', 'OVERFLOW'
- Volumes in the Overflow storage group are less preferred than a volume from the primary storage group
- More preferred than quiesced volumes and storage groups

Birth – Storage Group - Extend

- **Extend SG Name**
 - Used to extend data sets when there is insufficient space in the primary storage group
- **Overflow** storage group can also be used as an extend storage group
- Only one extend storage group per primary storage group
 - Same extend storage group can be the target of multiple primary storage groups
- Extend storage groups will not be used for initial allocation unless specified in storage group ACS routine

Adolescence



Adolescence – Management Class - Backup

```

Panel Utilities Scroll Help
MANAGEMENT CLASS ALTER Page 3 of 5
Command ==>
SCDS Name . . . . . : USER.SCDS
Management Class Name : MCDEF

To ALTER Management Class, Specify:
Backup Attributes
Backup Frequency . . . . . 0 (0 to 9999 or blank)
Number of Backup Vers . . . . . 5 (1 to 100 or blank)
(Data Set Exists)
Number of Backup Vers . . . . . 1 (0 to 100 or blank)
(Data Set Deleted)
Retain days only Backup Ver . . . 365 (1 to 9999, NOLIMIT or blank)
(Data Set Deleted)
Retain days extra Backup Vers . . 30 (1 to 9999, NOLIMIT or blank)
Admin or User command Backup . . BOTH (BOTH, ADMIN or NONE)
Auto Backup . . . . . Y (Y or N)
Backup Copy Technique . . . . . S (P, R, S, VP, VR, CP or CR)

Use ENTER to Perform Verification; Use UP/DOWN Command to View other Panels;
Use HELP Command for Help; Use END Command to Save and Exit; Cancel to Exit.

```

Adolescence – Management Class - Backup



- **Number of backup versions (Data Set Exists)**
 - Number of backup versions to keep while data set is still cataloged
 - Data set could exist as migrated data set
- **Number of backup versions (Data Set Deleted)**
 - Number of versions to keep after data set has been deleted
 - HSM detects that data set has been uncataloged during EXPIREBV
- **Retain days only backup version**
 - Days to retain last backup version
- **Retain days extra backup versions**
 - Days to retain extra backup versions (other than last)
- **Frequency**
 - 0 every time Autobackup runs
 - 1 at most once per calendar day
 - 2 at most every other day

Adolescence – Management Class - Backup

- Review the number of backups that are really needed
 - Based on business requirements
 - If number of **backup versions** is exceeded, then older versions will be rolled-off immediately
- Consider backup using RETAIN_DAYS if data set versions need to be kept for regulatory reasons
 - 100 active backup versions can be tracked in BCDS
 - > 600,000 versions can be kept using Retaindays
- Most migrated data sets should also have a backup copy
 - Exceptions: TMM, Dump data sets, etc.
- Ensure that EXPIREBV is run on a regular basis

Adolescence - Backups - NonSMS

- Non-SMS backups using HSM have some significant limitations
 - Lack of granularity
 - SMS & Management Classes relieve these limitations
- Multi-volume data sets are not supported
- Versions based retention approach is usual approach
 - Variations from default values require explicit data set name
- Review EXPIREBV NONSMS options
 - DELETEIFBACKEDUP, CATALOGEDDATA, UNCATALOGEDDATA

Adulthood



Adulthood – Management Class - Migration

```

Panel Utilities Scroll Help
MANAGEMENT CLASS ALTER Page 2 of 5
Command ==>
SCDS Name . . . . . : USER.SCDS
Management Class Name : MCDEF
To ALTER Management Class, Specify:
Partial Release . . . . . N (Y, C, YI, CI or N)
Migration Attributes
Primary Days Non-usage . . . . 21 (0 to 9999 or blank)
Level 1 Days Non-usage . . . . 60 (0 to 9999, NOLIMIT or blank)
Command or Auto Migrate . . . . BOTH (BOTH, COMMAND or NONE)
GDG Management Attributes
# GDG Elements on Primary . . . . 1 (0 to 255 or blank)
Rolled-off GDS Action . . . . . MIGRATE (MIGRATE, EXPIRE or blank)
Use ENTER to Perform Verification; Use UP/DOWN Command to View other Panels;
Use HELP Command for Help; Use END Command to Save and Exit; CANCEL to Exit.

```


Adulthood – Management Class - Migration

- Review **Primary Days Non-usage** and **Level 1 Days Date/Days** values
 - **Level 1 Days Date/Days** needs to be larger than primary
 - **Level 1 Days Date/Days** = 0 goes straight to tape
- When using Level 1, make sure that there is enough space on ML1 volumes
 - Look for ARC1237I (RC=37) messages during space management
 - SETSYS ML1OVERFLOW for larger ML1 data sets

Adulthood – Management Class - Migration

- # GDG Elements on Primary allows early migration of GDGs from primary disk
 - Specifies the number of versions to be retained on primary disk for active processing
 - Does not guarantee that additional versions will be migrated
- Review recall activity to ensure that GDSs are not being recalled excessively
 - Most common cause of thrashing
 - For example, batch jobs where complete GDG is referenced
- Non-scratch GDSs can be processed using Rolled-off GDS Action
 - Expire or migrate

Adolescence/Adulthood – MC – Partial Release

- **Partial Release**
 - Used to release allocated but unused space in a data set
- Can be used against extended format VSAM or Non-VSAM data sets (extended or standard)
 - GDSs are generally good candidates
- Conditional space release allows space to be released against data sets only if they have secondary allocation
- Space can be released at close time or during HSM space management
- Use caution when releasing space on partitioned data sets
 - Do not release space on production change management products that use PDSs

Retirement



Retirement – Management Class – Expiration/Retention Limits

- 1st panel of ISMF Management Class Display

```

Panel Utilities Scroll Help
MANAGEMENT CLASS ALTER Page 1 of 5
Command ==>
SCDS Name . . . . . : USER.SCDS
Management Class Name : MCDB2LOG
To ALTER Management Class, Specify:
Description ==> DB2 ARCHIVE LOGS
==>
Expiration Attributes
  Expire after Days Non-usage . . 45 (1 to 9999 or NOLIMIT)
  Expire after Date/Days . . . . . 45 (0 to 9999, yyyy/mm/dd or
  NOLIMIT)
Retention Limit . . . . . 0 (0 to 9999 or NOLIMIT)

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.

```

Retirement – Management Class – Expiration/Retention Limits

- **Expire after Days Non-usage**
 - Data set *eligible* to expire xxx days after last reference
- **Expire after Date/Days**
 - Data set *eligible* to expire xxx days after creation
 - If Expire after Days Non-usage specifies NOLIMIT
 - If Expire after Days Non-usage specifies a value, then that will be also taken into consideration
- If both fields specify NOLIMIT, then the data set will not automatically expire
- If both fields are specified, then the later date is used

UNLESS.....

Retirement – Management Class – Expiration/Retention Limits

- If explicit expiration date is specified, then management class expiration attributes will not be used
 - Expiration date in VTOC or catalog entry
 - Common reason why data sets are kept longer than expected
- Check DFSMSHsm EXPIREDDATASETS(SCRATCH) in ARCCMDxx member
 - Make sure that you do not have data sets with expiration dates in the past
 - These will be deleted by Space Management
 - Implement with caution!!

Retirement – Management Class – Allocated but not opened

- Data sets that are allocated by a Clist, for example
- Needed by ISPF-based product to function
 - Even if not actually used
- Reference date could indicate data set has not been used in many years
- Ensure these are in a non-expire management class

Retirement – Management Class – Reasons that Data Sets are not expired

- Explicit expiration date not reached
 - Use Retention Limit in management class to remove or limit RETPD or EXPDT
- Data sets with unknown DSORG not processed
 - Assign “default” data class as last part of data class ACS routine
 - Use allocation product to initialize empty data sets
- Volume containing data set not being processed by space management
 - Storage group **Auto Migrate** set to N?
 - Ensure low threshold is set correctly

Agenda

- Birth
 - Allocation and Placement
- Adolescence
 - Usage
- Adulthood
 - Migration
- Retirement
 - Expiration and Deletion

Questions?





The Life and Times of a Data Set: You Wouldn't Want Your Relatives Hanging Around, Why Your Data ?

Chris Taylor
ctaylor1@us.ibm.com
IBM Corporation
August 16, 2013
Session Number 13772

