



What's New in DFSMSHsm

Glenn Wilcock
IBM

August 15, 2013
Session Number 14137




Agenda

- **General**
 - QUERY ACTIVE
 - RLS Server Error
 - UCB Storage Relief
 - Storage Group Priority
 - Recover w/ NEWNAME
- **Migration**
 - Throughput Improvements
- **Tape**
 - >40 Tapes
 - RECYCLE in place of TAPECOPY
 - Automatic restart of RECYCLE
- **Fast Replication**
 - Consistency Groups
 - Recover Data Set Enhancements
 - Performance
 - Usability
 - UCB Refresh
- **Storage Tiering**
- **Miscellaneous**

General QUERY ACTIVE

- **Today:** You have a problem with a tape drive and need to cancel the associated HSM task, but how do you know which one to cancel?



- **V2R1:** QUERY ACTIVE(TCBADDRESS) includes **tape volser**, **device address** and **task name**.

```
ARC0161I MIGRATING VOLUME TPR18B FOR USER **AUTO**, REQUEST  
NONE, TASK=MVOL1, TCB=x'00747690'; VOL=TPR18B, ADDR=0F12;  
VOL=T00001, ADDR=0B30; VOL=NONE, ADDR=NONE
```

General SMSVSAM Server (RLS) Error Handling

- **Today:** When DFSMSHsm accesses the CDSes in RLS mode and an SMSVSAM server error occurs:
 - DFSMSHsm takes a fatal abend and shuts down
 - After DFSMSHsm is restarted, all DFSMSHsm requests that were in progress at the time of the SMSVSAM server error are lost and must be reissued to complete.



General SMSVSAM Server (RLS) Error Handling

- **V2R1:** When an SMSVSAM server error occurs:
 - DFSMSHsm will detect the error, and quiesce all CDS I/O activity
 - Once the SMSVSAM server initializes, DFSMSHsm will automatically close and reopen the CDSes and continue processing
 - I/Os that failed will be retried
 - Some DFSMSHsm requests that are in progress at the time of the server error may fail, depending on the timing of the error with respect to a series of I/O operations required to complete the request



General Virtual Storage Constraint Relief

- **Today:** DFSMSHsm's large customers face periodic 878-Abends.
 - These abends have repetitively been shown to be caused by DFSMSHsm capturing a high number **UCBs below-the-line**. Historically, DFSMSHsm was required to capture UCBs below-the-line due to certain called services requiring UCBs to reside below-the-line.
- **V2R1:** DFSMSHsm will no longer capture UCBs below-the-line. All services have been updated to be able to process UCBs in 31-bit addressing mode.



UCB: Unit Control Block that describes a device

General Storage Group Priority

- The new Storage Group **Processing Priority** specifies the relative order in which storage groups should be processed during Automatic Space Management and Backup Processing
 - Storage Groups will be processed in the order of their priority
 - A higher value means a higher priority
 - Valid values are 1-100
 - Default value is 50



General Storage Group Priority



```
DRIVER [24 x 80]
File Edit View Communication Actions Window Help
Host: tucvmc6.storage.tucson.i Port: 23 LU Name: Disconnect
Panel Utilities Scroll Help
POOL STORAGE GROUP DEFINE Page 2 of 2
Command ==>
SCDS Name . . . . . : IBMUSER.TEMP.SCDs
Storage Group Name : SG1TEST
To DEFINE Storage Group, Specify:
Allocation/migration Threshold : High 85 (1-100) Low . . 1 (0-99)
Alloc/Migr Threshold Track-Managed: High 85 (1-100) Low . . 1 (0-99)
Guaranteed Backup Frequency . . . . . (1 to 9999 or NOLIMIT)
BreakPointValue . . . . . (0-65520 or blank)
Processing Priority . . . . . 50 (1-100)
Use ENTER to Perform Selection; Use UP Command to View previous Page;
MA D 15/044
Connected to remote server/host tucvmc6.storage.tucson.ibm.com using port 23
```



General

Recover VSAM data set with a NEWNAME

- **Today:** No support for renaming of a VSAM data set during recovery from dump.
- **V2R1:** Add Support
 - Existing optional parameter NEWNAME(newdsname) can now be specified for VSAM data sets when recovering with FROMDUMP
 - DFSMSDss RENAMEU and REPLACEU enhancements for this support
 - Physical data set COPY and RESTORE will be enhanced to support renaming of VSAM data sets when RENAMEUnconditional is specified. Renaming of AIX will remain unsupported.
 - Physical data set RESTORE will be enhanced to support REPLACEUnconditional.
 - Caveat – Data set can only be cataloged under certain conditions

General ARC0936I

- DFSMSHsm can issue the ARC0936I message during processing of SMS-managed volumes. Today, when its return code is RC12, the SUBREAS 4-byte field is always zero.
- With this enhancement, when the FAST VVDS ACCESS interface returns RC12, REAS98, DFSMSHsm will use the information in FVPPROB field to put it into SUBREAS field of the ARC0936I message. This will help in diagnosing SMS VTOC/VVDS errors.

ARC0936I ERROR RETRIEVING SMS VTOC/VVDS DATA, FUNC = (ACCESS | FREE), RC=retcode, REASON=reas1, SUBREAS=reas2

The explanation for RC = 12, REASON = 98 will be added to the message manual.

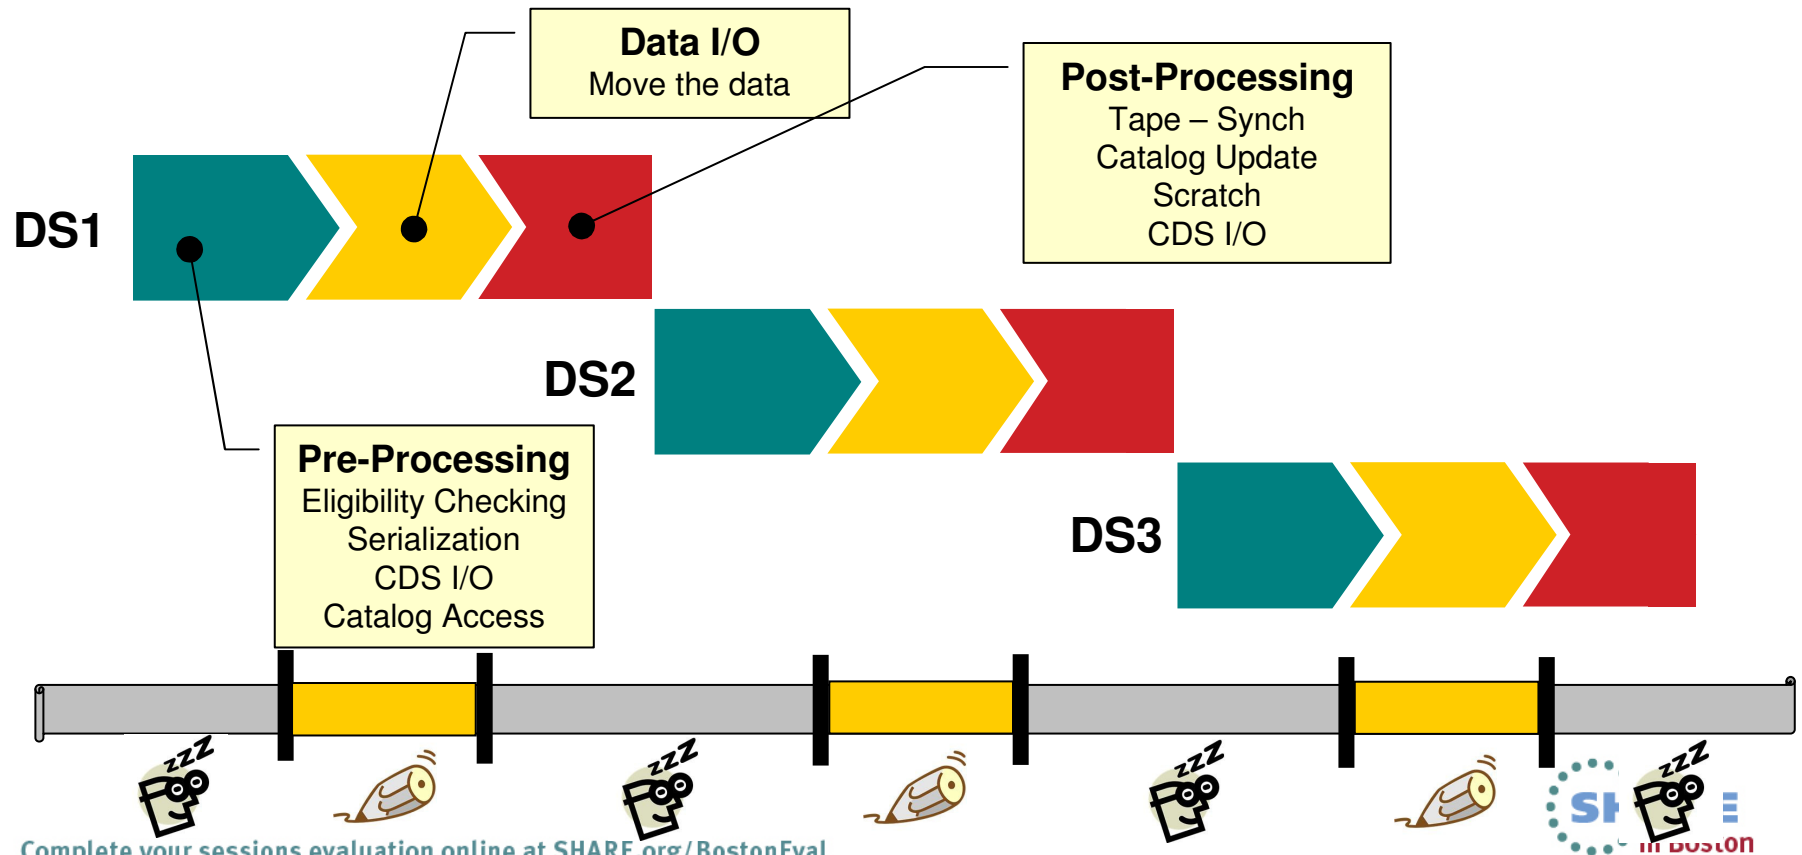
Code: RC = 12 REASON = 98-CATALOGENTRYMGR SUBREAS = Byte 1-Internal RC, Byte 2- Internal REASON, Bytes 3,4- catalog module identifier. Meaning: Error accessing CATENTRY.

Agenda

- **General**
 - QUERY ACTIVE
 - RLS Server Error
 - UCB Storage Relief
 - Storage Group Priority
 - Recover w/ NEWNAME
- **Migration**
 - Throughput Improvements
- **Tape**
 - >40 Tapes
 - RECYCLE in place of TAPECOPY
 - Automatic restart of RECYCLE
- **Fast Replication**
 - Consistency Groups
 - Recover Data Set Enhancements
 - Performance
 - Usability
 - UCB Refresh
- **Storage Tiering**
- **Miscellaneous**

Migration Throughput Improvement

- **Today:**
 - For small data sets, the overhead of pre- and post- migration processing accounts for as much as $2/3^{\text{rds}}$ of the processing time.



Migration Throughput Improvement

- **V2R1:**
 - For each Data Movement Task, run multiple subtasks that overlap pre- and post-processing with data movement.



Migration Throughput Improvement

- Migration subtasking can be enabled for Primary Space Management, On Demand Migration and Interval Migration.

SETSYS:

```

> _____>
  |                                     |
  |   _MIGRATIONSUBTASKS_ ( _YES_ ) |
  |   MIGSBTSKS _____ |   _NO_ |
  
```

- MIGRATIONSUBTASKS(YES) enables the function and can only be specified on startup.
- The default number of subtasks is a *patchable* value which is initially set to 5.
 - Internal testing found this to be an ideal balance of throughput, CPU usage and memory consumption.

Agenda

- **General**
 - QUERY ACTIVE
 - RLS Server Error
 - UCB Storage Relief
 - Storage Group Priority
 - Recover w/ NEWNAME
- **Migration**
 - Throughput Improvements
- **Tape**
 - >40 Tapes
 - RECYCLE in place of TAPECOPY
 - Automatic restart of RECYCLE
- **Fast Replication**
 - Consistency Groups
 - Recover Data Set Enhancements
 - Performance
 - Usability
 - UCB Refresh
- **Storage Tiering**
- **Miscellaneous**

Tape

Span datasets over 40 Tapes

- **Today:** Individual data sets can only span 40 tapes.
 - Not a problem for 'real' tape
 - Prevents DFSMSHsm from managing larger data sets in a virtual tape environment.

Example: $40_{\text{tapes}} * 4\text{GB}_{\text{logical tape}} * 2.5_{\text{Compaction}} = 400\text{GB}$

- **V2R1:** For migration and backup, individual data sets can span up to 254 tape volumes.

Example: $254 * 4\text{GB} * 2.5 = 2.5\text{TB}$


Tape Span datasets over 40 Tapes

- FSR records have been extended from a maximum of 1260 bytes to 6396 bytes in order to contain up to 508 volume serials (254 input and 254 output).
- FSRs that list more than 144 tape volumes will be truncated when written to the DFSMSHsm LOG data sets.
 - The existing fixed length 2048 LRECL prevents FSRs with more than 144 volumes from being written to the DFSMSHsm log data sets in their entirety.
- Truncated FSRs will affect ARCPRLOG and ARCPEDIT output
 - ARCPRLOG output will display only the portion of each FSR that was written to the log.
 - When RECYCLE volumes are truncated, ARCPRLOG and ARCPEDIT output will display "TOVOL=*****".

Tape

Use RECYCLE in place of TAPECOPY

- **Today:** When an alternate tape of a duplex copy pair fails to be created, DFSMSHsm generates an automatic request to create a TAPECOPY of the original.
 - TAPECOPY is single tasked and does not have a resume capability.
 - A Recall request can take-away the tape from tape copy, causing tape copy to have to start from the beginning.



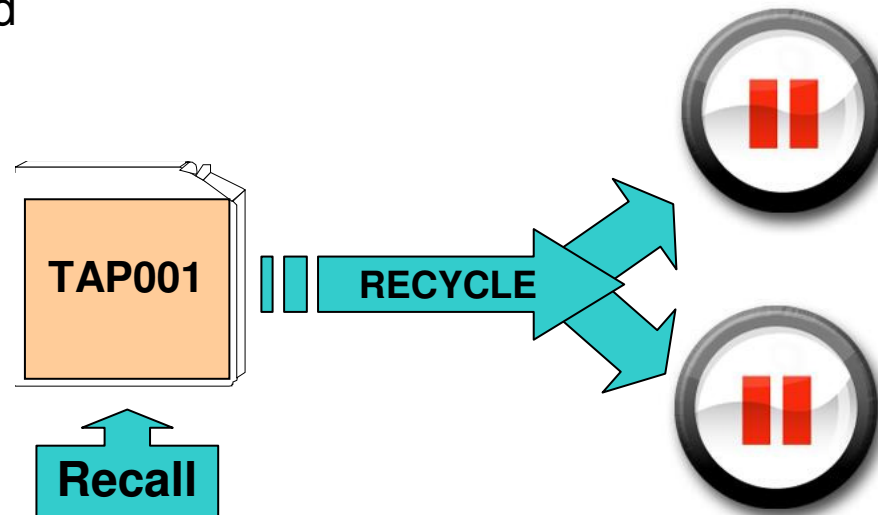
Tape

Use RECYCLE in place of TAPECOPY

- **V2R1:** A new option to create the alternate via RECYCLE

SETSYS DUPLEX ... ERRORALTERNATE (RECYCLE|TAPECOPY)

- Both Backup and Migration
- During Secondary Space Management, RECYCLE is used to replace the single original tape with a new original and alternate
 - RECYCLE can be multi-tasked and doesn't have to start over from beginning if it is interrupted



Tape Automatically restart RECYCLE

- **Today:** Tape take-away terminates Recycle processing
 - Recycle must be manually restarted, potentially multiple times
- **V2R1:** When a tape take-away occurs, a new Recycle command for the original tape will be automatically generated

SETSYS RECYCLETAKEAWAYRETRY(YES MAXRETRYATTEMPTS(*nn*) DELAY(*mmmm*) | NO)

- The default is NO
 - *nn*: The value can be 1 – 99. The default value is 12.
 - *mmmm*: The value can be 1 – 9999. The default value is 300 (5min).
 - When max is reached: **ARC0836I MAXIMUM NUMBER OF RECYCLE RETRY ATTEMPTS FOR TAPE <VOLSER> HAS BEEN EXCEEDED**
- To identify tapes that are in a Retry status or that have exceeded the maximum retries:

LIST TTOC SELECT (RECYCLETAKEAWAY)

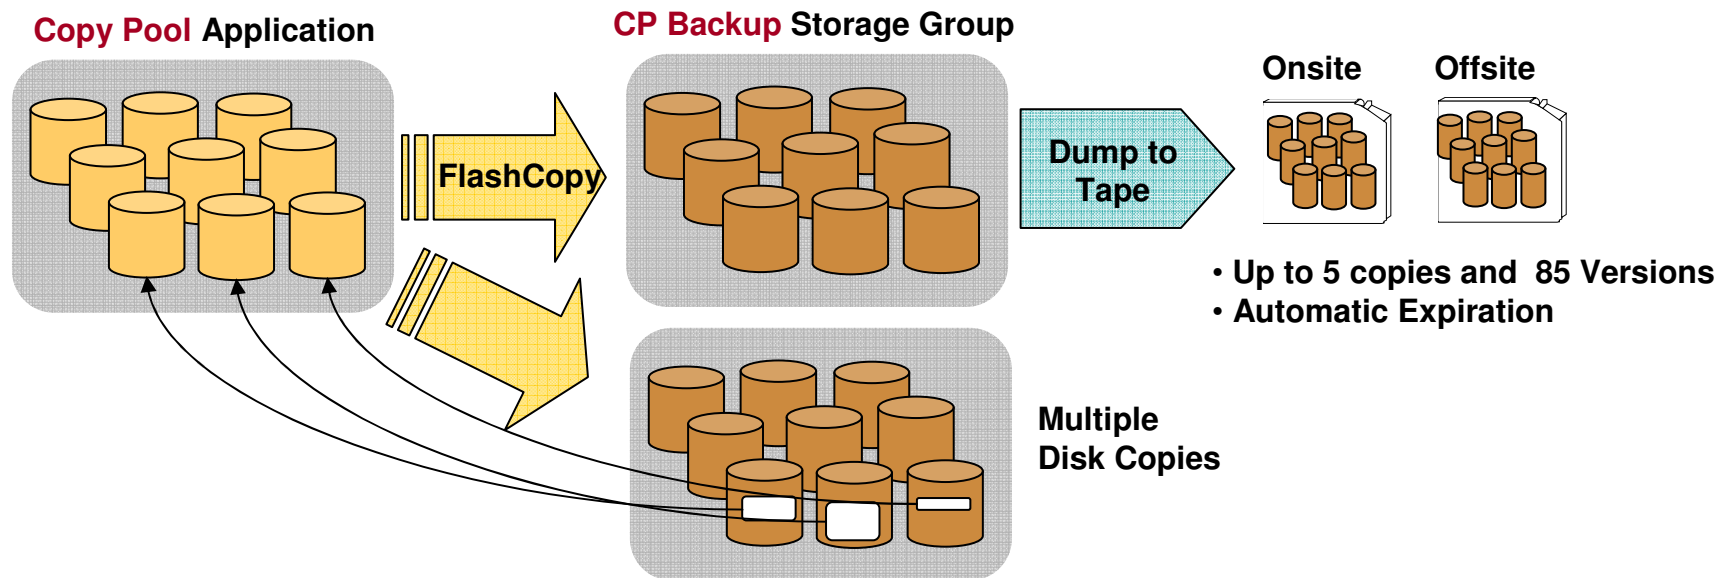
Agenda

- **General**
 - QUERY ACTIVE
 - RLS Server Error
 - UCB Storage Relief
 - Storage Group Priority
 - Recover w/ NEWNAME
- **Migration**
 - Throughput Improvements
- **Tape**
 - >40 Tapes
 - RECYCLE in place of TAPECOPY
 - Automatic restart of RECYCLE
- **Fast Replication**
 - Consistency Groups
 - Recover Data Set Enhancements
 - Performance
 - Usability
 - UCB Refresh
- **Storage Tiering**
- **Miscellaneous**

Fast Replication Overview

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)

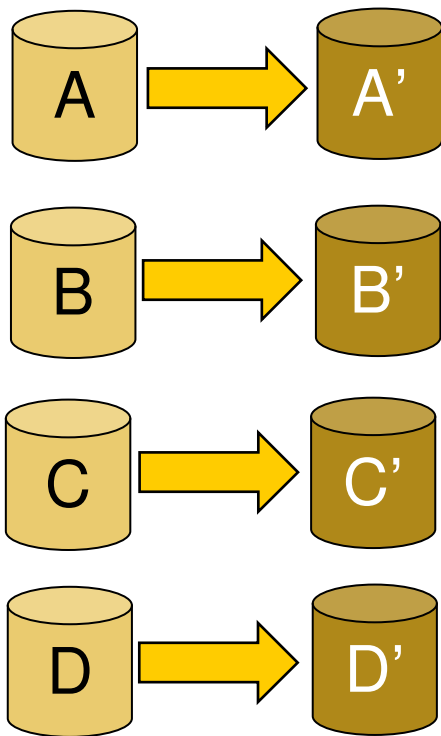


★ Recovery at all levels from either disk or tape!

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

Fast Replication Consistency Groups

- Standard FlashCopy – ‘Fuzzy’ Backup
- DB2 ensures data consistency during recovery




- **FC SrcA to TrgA'**
 - Writes held to SrcA during establish
 - Write activity on B, C, D
- **FC SrcB to TrgB'**
 - Writes held to SrcB during establish
 - Write activity on A, C, D
- **FC SrcC to TrgC'**
 - Writes held to SrcC during establish
 - Write activity on A, B, D
- **FC SrcD to TrgD'**
 - Writes held to SrcD during establish
 - Write activity on A, B, C

Fast Replication Consistency Groups

- **Freeze FlashCopy Consistency Group**
 - I/O activity to a volume in a consistency group is held in an extended long busy (ELB)
 - A window can be created during which dependent write updates will not occur
 - I/O activity resumes when a FlashCopy Consistency Group is created
- **Consistency Group support is intended for DB2 Logs**
 - Default ELB is two minutes
 - Create a consistent copy of the Log copy pool such that a conditional restart *is not required*

Fast Replication Consistency Groups

- **Consistency Group FlashCopy – Consistent copy**



- **FC SrcA to TrgA' / CGFREEZE**
 - Writes held to SrcA by ELB
 - Non Dependent write activity on B, C, D
- **FC SrcB to TrgB'**
 - Writes held to SrcA, B by ELB
 - Non Dependent write activity on C, D
- **FC SrcC to TrgC'**
 - Writes held to SrcA, B, C by ELB
 - Non Dependent write activity on D
- **FC SrcD to TrgD'**
 - TrgA – D contain a consistent copy
- **Consistency Group Created / CGTHAW**

Fast Replication Consistency Groups

- **Gotchas**
 - The Consistency Group THAW is an LSS level operation
 - The control unit does not enable logical volume connected Freezes
 - If two copy pools with volumes in the same LSS are being backed up concurrently, the copy pool to first issue the THAW will invalidate the backup of the other copy pool
 - FRBACKUP requests are single threaded on an HSM host, so the exposure is when FRBACKUP requests are issued from multiple HSM hosts
 - HSM recognizes the issue and will fail the backup, which will have to be reissued
 - If the time to establish the FlashCopies for the volumes being processed as a part of the consistency group exceeds the CG timer (2 min default), then the control unit will release the ELBs on all of the volumes and the backup will fail

Fast Replication Consistency Groups



```
Panel Utilities Scroll Help
----- DGTDGPP1 COPY POOL DEFINE/ALTER Page 1 of 5
Command ==>
SCDS Name . . . : USER6.MYSCDS
Copy Pool Name : COPYPL1
To DEFINE/ALTER Copy Pool, Specify:
Description ==> _____
                ==> _____
Auto Dump . . . N (Y or N) Dump Sys/Sys Group Name . . . _____
Dump Class . . _____ Dump Class . . _____
Dump Class . . _____ Dump Class . . _____
Dump Class . . _____
Number of DASD Fast Replication Backup
Versions with Background Copy . . . . . 2 (0 to 85 or blank)
FRBACKUP to PPRC Primary Volumes allowed . . . __ (NO, PN, PP, PR or blank)
FRRECOV to PPRC Primary Volumes allowed . . . __ (NO, PN, PP, PR or blank)
FlashCopy Consistency Group . . . . . N (Y or N)

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.
```

- **LIST COPYPOOL will have an indicator of the setting**



Fast Replication

Recover data set to any volume

- **Today**
 - FRRECOV DSNNAME requires that a data set be recovered to the volume(s) that the data resided on at the time of the backup
 - If there is not enough space, then the recovery is failed and enough free space must be regained to perform the recovery
- **V2R1**
 - If there is not enough space on a volume to perform the recovery, HSM will select the volume in the same Storage Group that has the most free space

Fast Replication NEWNAME

- **Today**
 - A data set must be recovered using the same name
- **V2R1**
 - A data set may be recovered to a new name
 - Recovery may from either Disk or Tape

```
FRRECOV DSNAME(ds.oldname) NEWNAME(ds.newname)  
FRRECOV DSNAME(ds.oldname) NEWNAME(ds.newname)  
REPLACE
```

- If there isn't enough space on the original volumes, then new volumes will be selected

Fast Replication Data Set Recovery Performance

- **Today**
 - DSS Scans the entire VVDS to locate the VVR for the data set being processed
 - When many data sets are being processed, huge performance hit
 - A customer was recovering thousands of data sets.
 - Impact when the VVDS is large
- **V2R1**
 - HSM will capture the VVR RBA along with other catalog data at time of backup
 - Requires Catalog Capture option
 - HSM passes this to DSS during FRRECOV processing so that DSS can directly access the VVR

Fast Replication D-Type APARS

OA39738

- For NOCOPY, volumes are always initialized after WITHDRAW, regardless of what the status of the volume is
- Some customers were reporting that target volumes were not getting initialized and being left with a broken VTOC

OA38533

- Increase Maximum Dump Tasks from 32 to 64

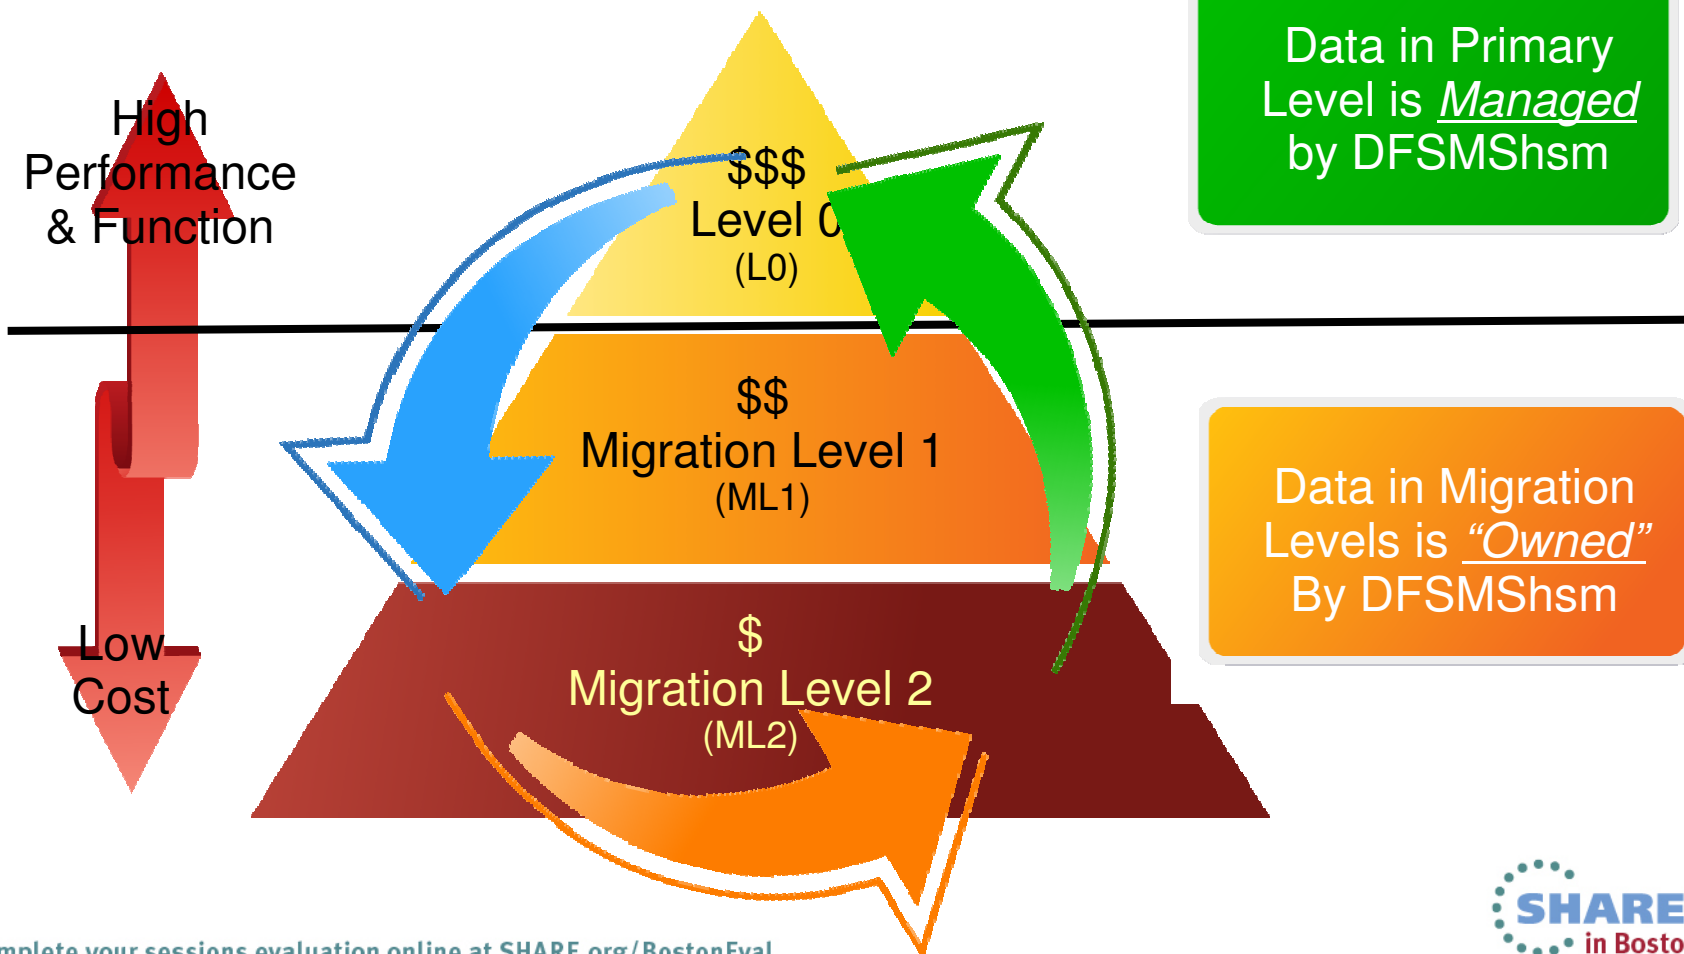
Agenda

- **General**
 - QUERY ACTIVE
 - RLS Server Error
 - UCB Storage Relief
 - Storage Group Priority
 - Recover w/ NEWNAME
- **Migration**
 - Throughput Improvements
- **Tape**
 - >40 Tapes
 - RECYCLE in place of TAPECOPY
 - Automatic restart of RECYCLE
- **Fast Replication**
 - Consistency Groups
 - Recover Data Set Enhancements
 - Performance
 - Usability
 - UCB Refresh
- **Storage Tiering**
- **Miscellaneous**

Storage Tiers: V2R1 Background



The *classic* DFSMS storage hierarchy (35 years old)



Storage Tiers: V2R1 Background

- **The Space Management Environment is evolving**
 - Typical configurations have changed to leave data on Level 0 longer and then migrate directly to ML2
 - ★ When ML2 is a VTS, the VTS disk cache replaces the ML1 tier
 - ★ *Eliminates MIPS required for software compression to ML1*
 - ★ *Eliminates DFSMSHsm ML1->ML2 processing*
 - ★ Tape storage controller functionality replacing DFSMSHsm duplexing
- **Shortcomings of today's DFSMS functionality**
 - ✓ No policy-based automation for moving data within the Primary Storage Hierarchy (Level 0)
 - ✓ No policy-based management of Active (open) data



Storage Tiers: V2R1 Overview



z/OS V2R1 DFSMS provides the Storage Tiers Solution...

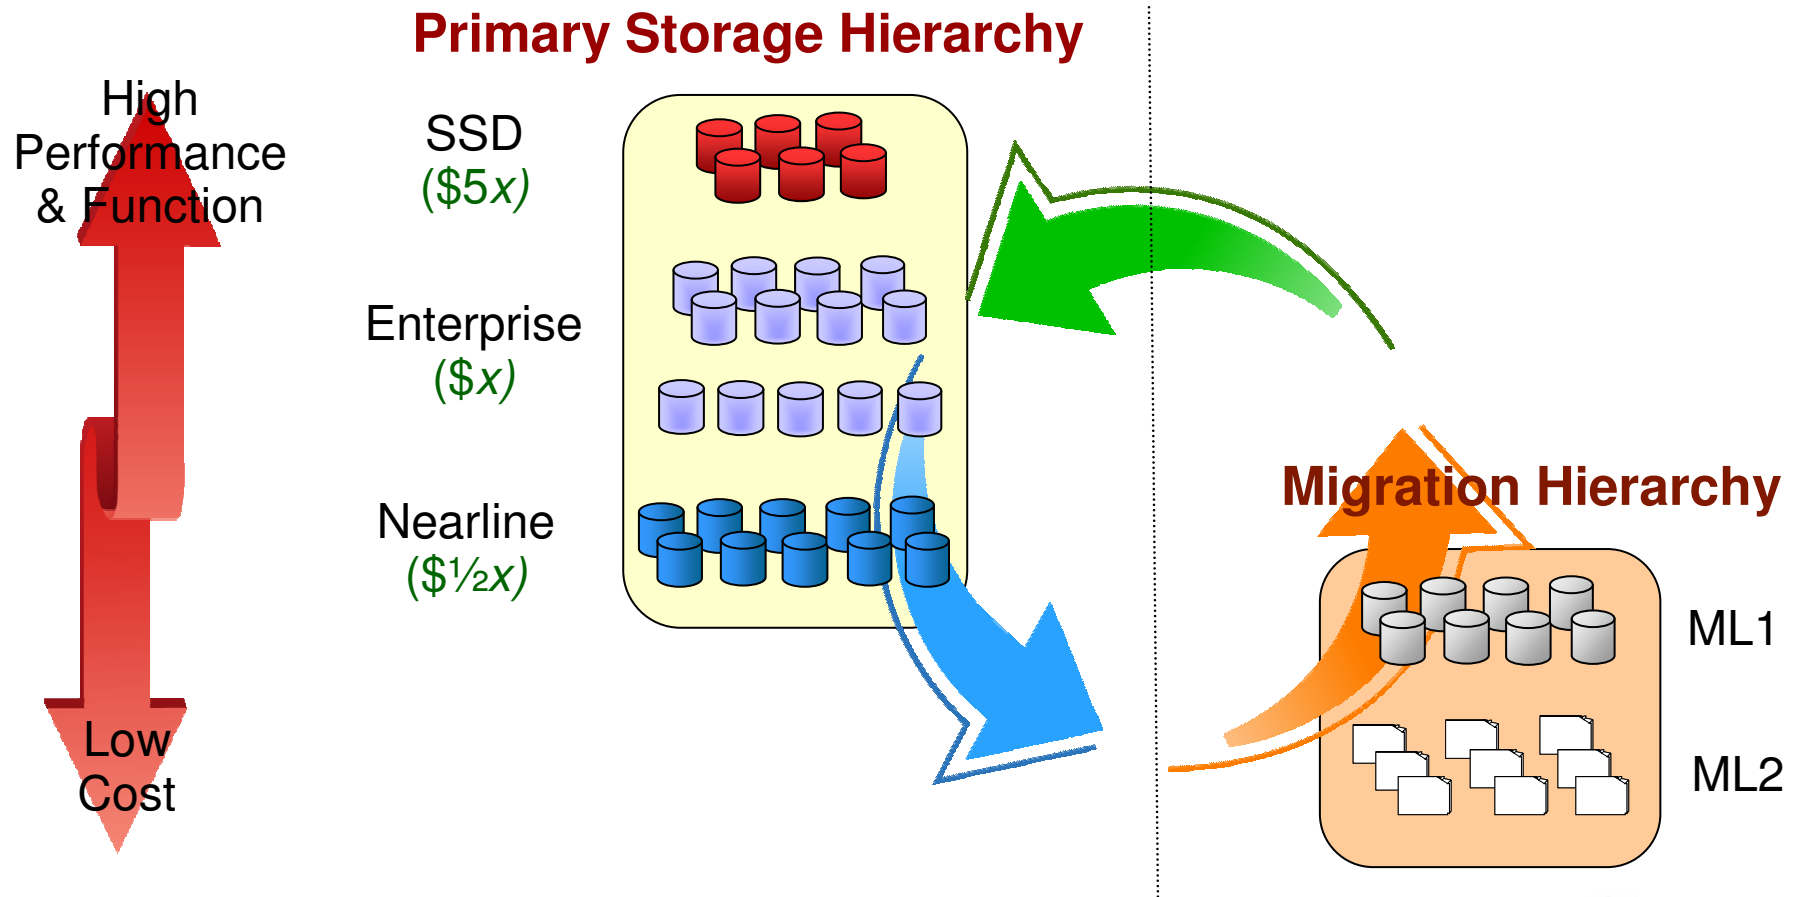
Automated, policy-based space management that moves data from tier to tier within the Primary (Level 0) Hierarchy

- ✓ Automated movement provided via the existing DFSMSshm Space Management function
 - Movement is referred to as a 'Class Transition'
 - Data remains in its original format and can be immediately accessed after the movement is complete
- ✓ Policies implemented via the existing Class Transition policies and updated Management Class policies
- ✓ Enhanced support for DB2, CICS and zFS data
 - *Open data temporarily closed to enable movement*

Storage Tiers: V2R1

Overview

The *classic* DFSMS storage hierarchy is modified to represent *distinct* Primary and Migration Hierarchies...



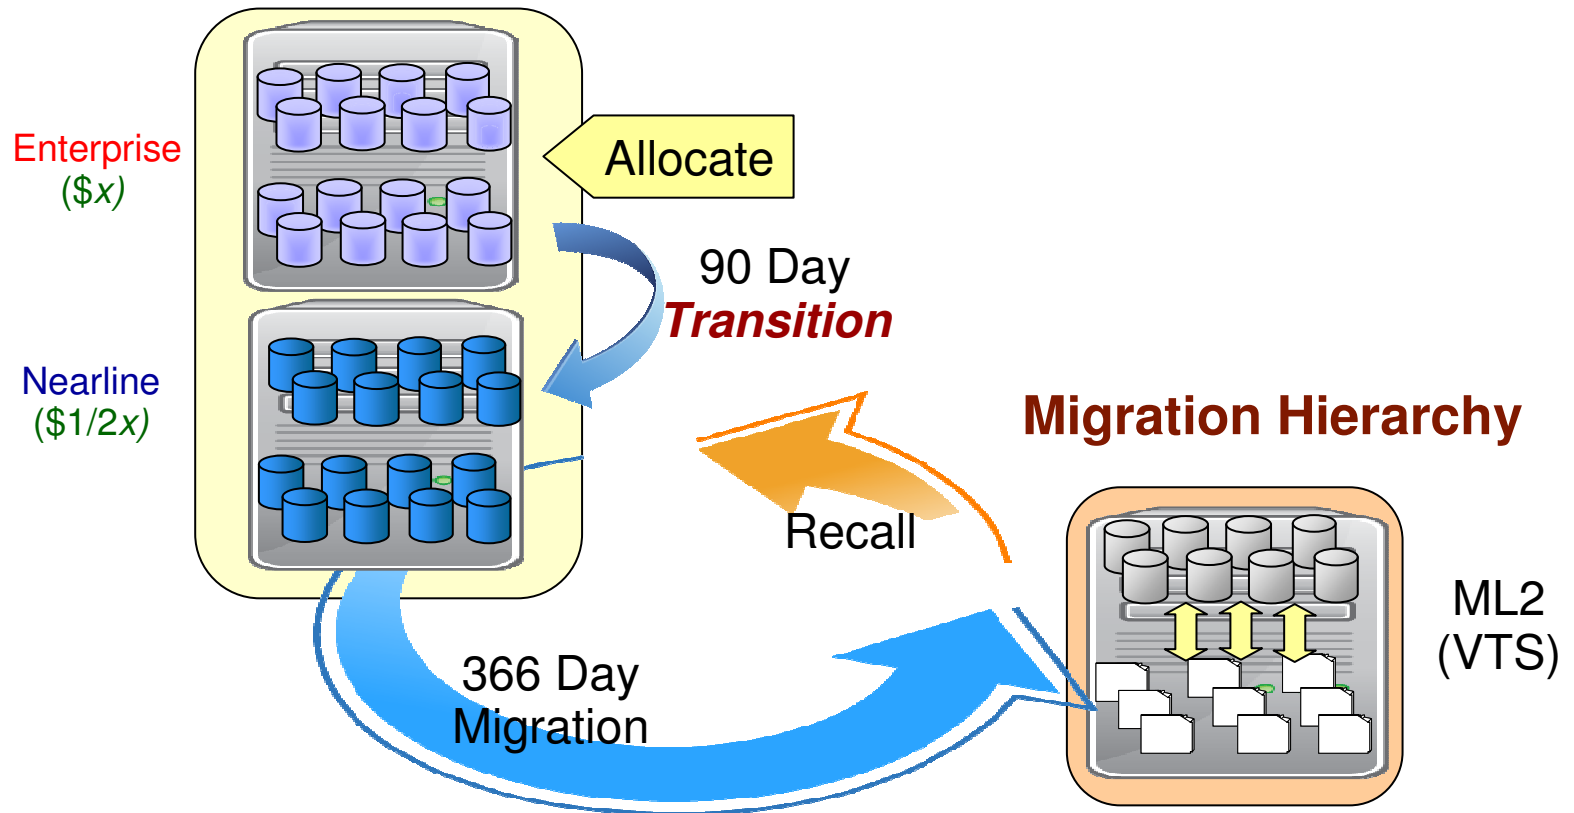
Storage Tiers: V2R1

Overview



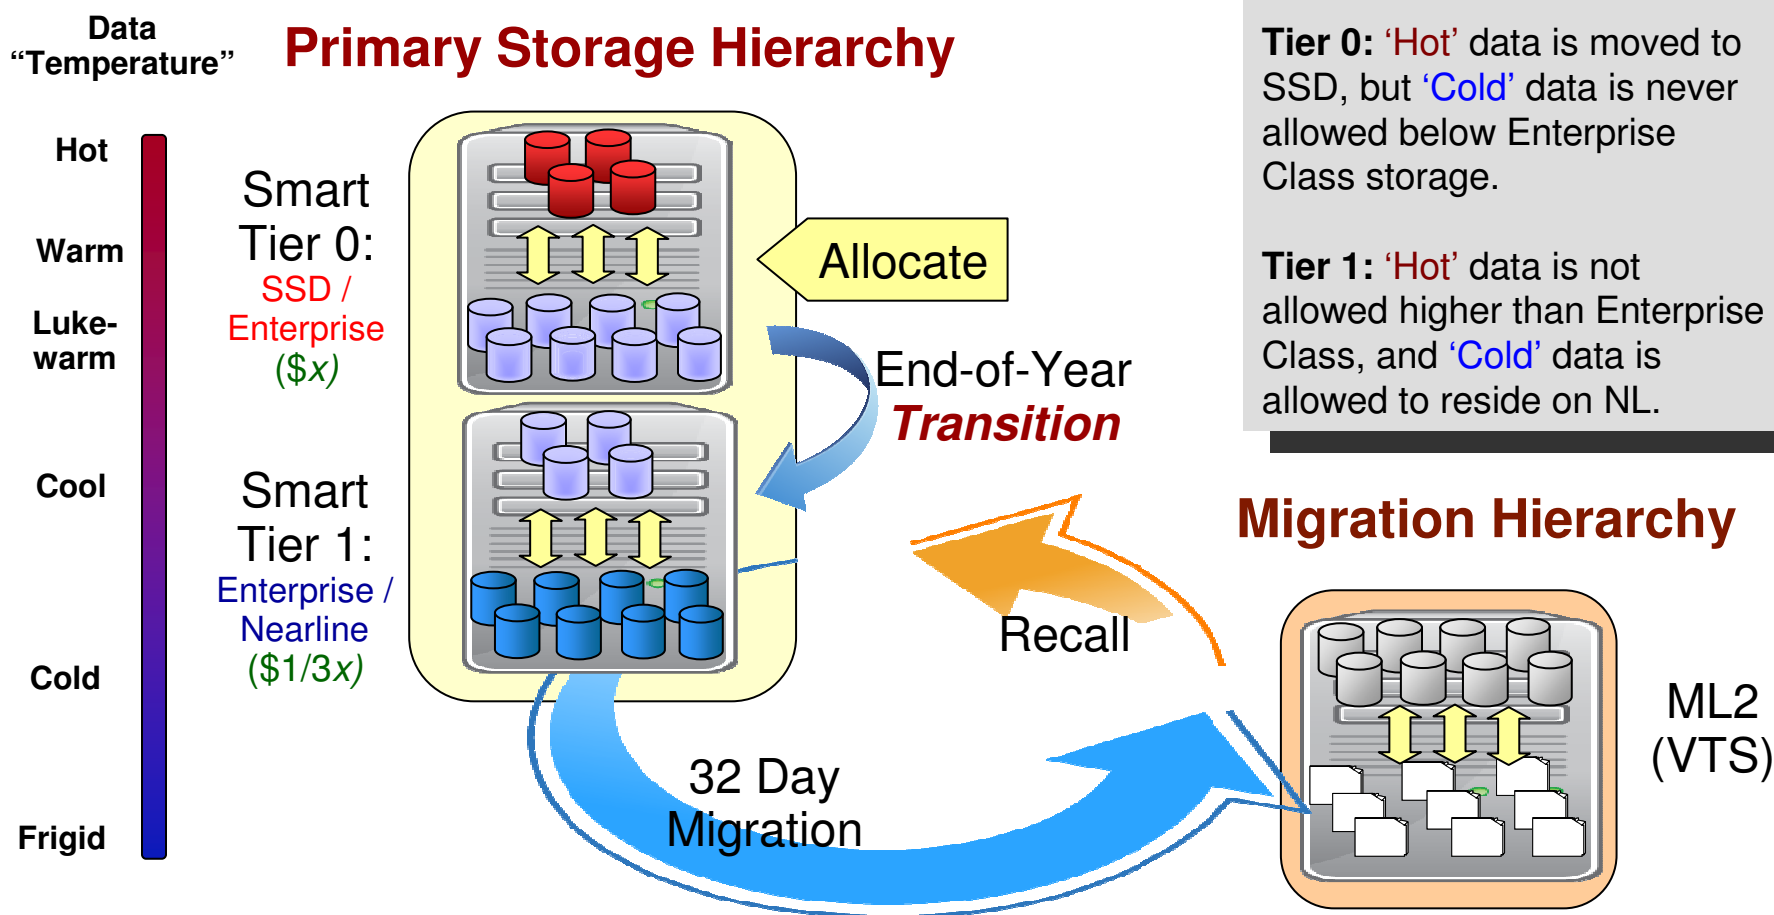
This example environment shows data being “transitioned” from Enterprise Class storage to Nearline (SAS/SATA) Class storage after 90 days, and then becoming eligible for migration after 366 days of inactivity.

Primary Storage Hierarchy



Storage Tiers: V2R1

This example environment shows data being “transitioned” from ‘Smart Tier 0’ to ‘Smart Tier 1’ at the end-of-year, and then becoming eligible for migration after 32 days of inactivity.



Tier 0: ‘Hot’ data is moved to SSD, but ‘Cold’ data is never allowed below Enterprise Class storage.

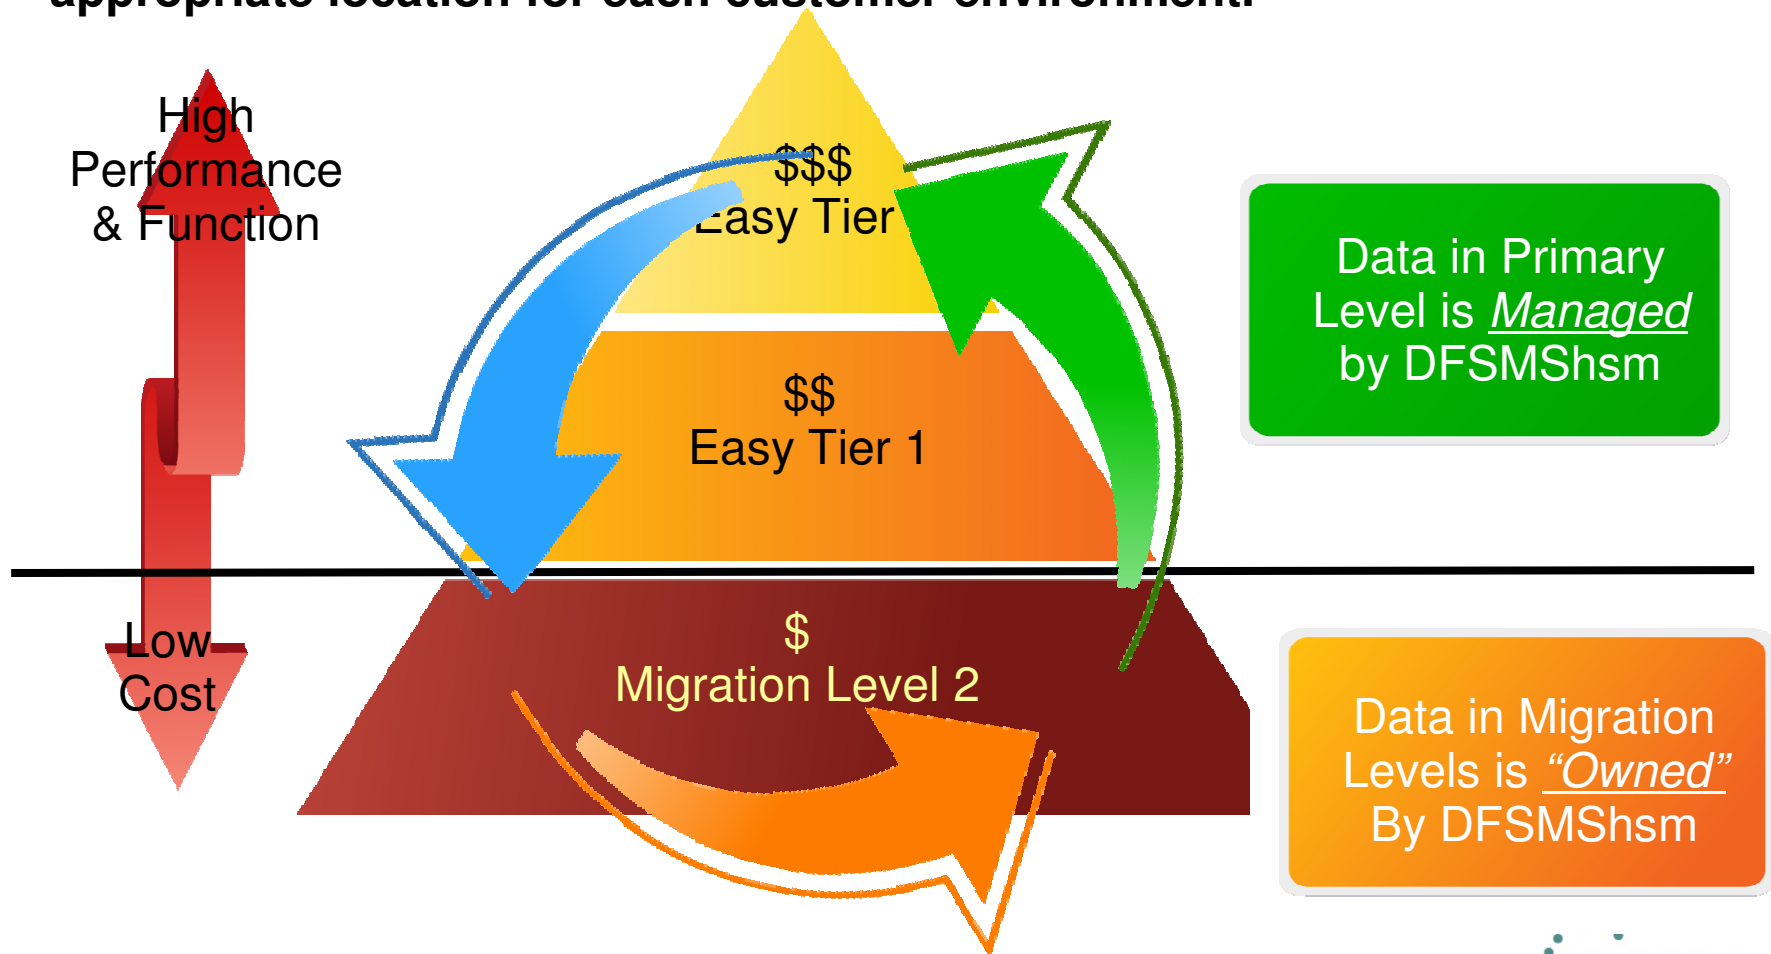
Tier 1: ‘Hot’ data is not allowed higher than Enterprise Class, and ‘Cold’ data is allowed to reside on NL.

Storage Tiers: V2R1

Overview




The *enhanced* DFSMS storage hierarchy enables “the bar” to be in the most appropriate location for each customer environment.



DFSMS Storage Tiers z/OS V2R1

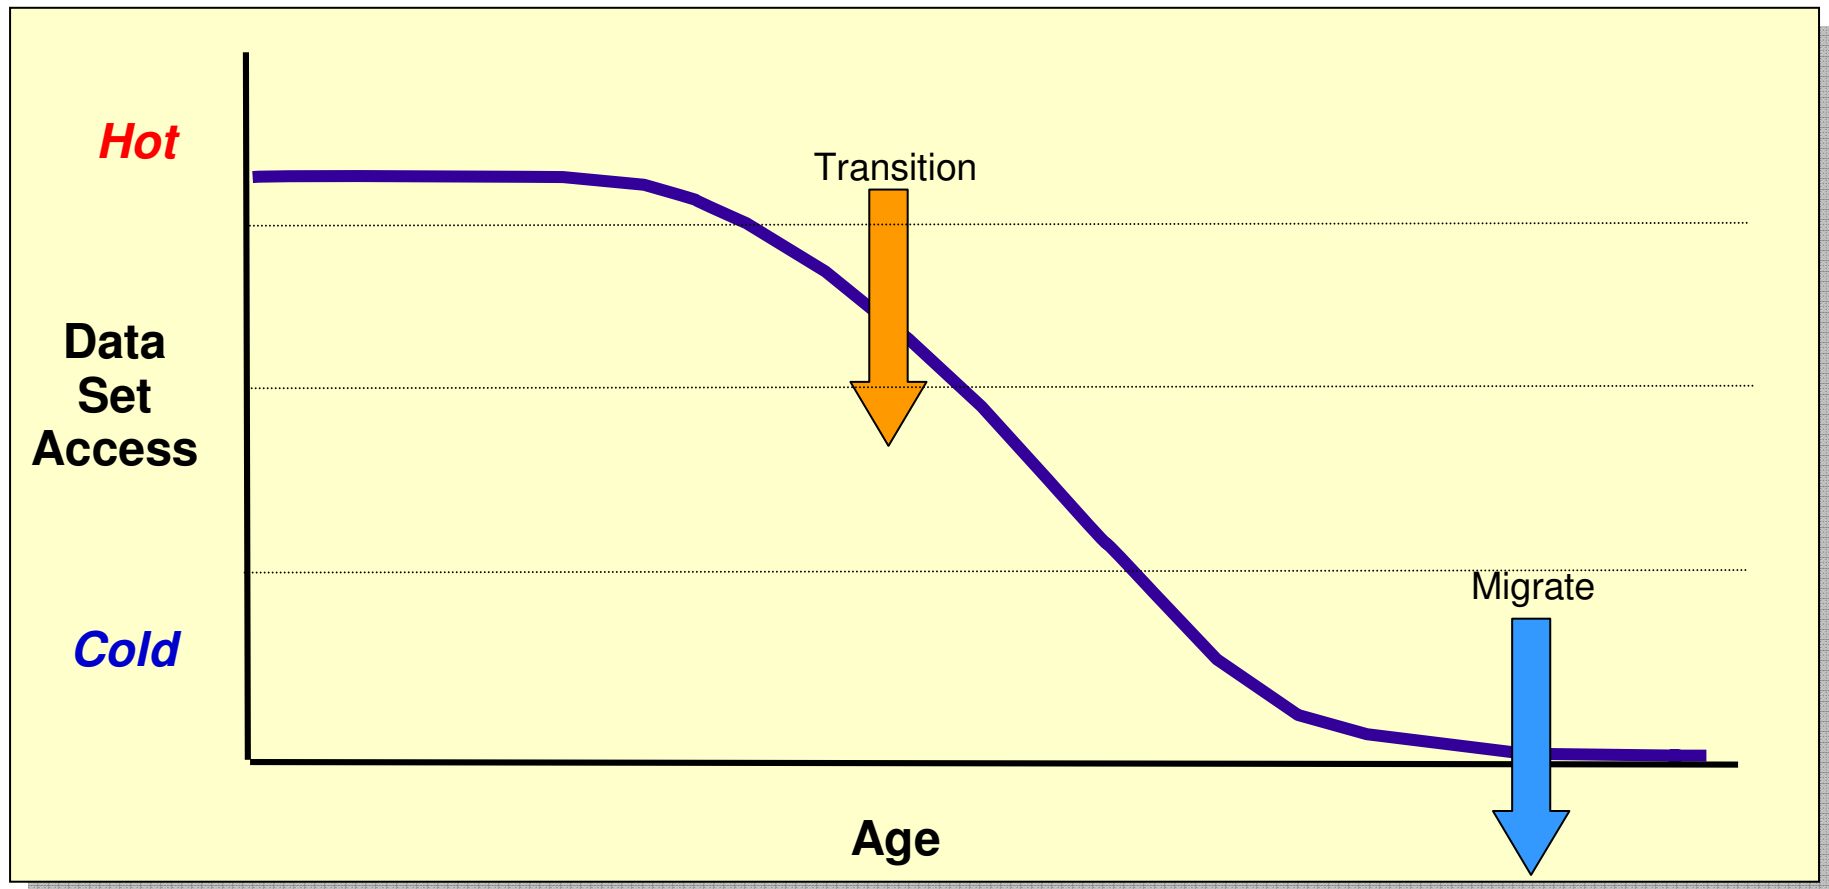


As a data set ages and goes through cycles of activity and inactivity, it can go through many migration and recall iterations.



Storage Tiers: V2R1 Usage

The migration / recall iterations can be replaced with a single class transition and potentially single migration.



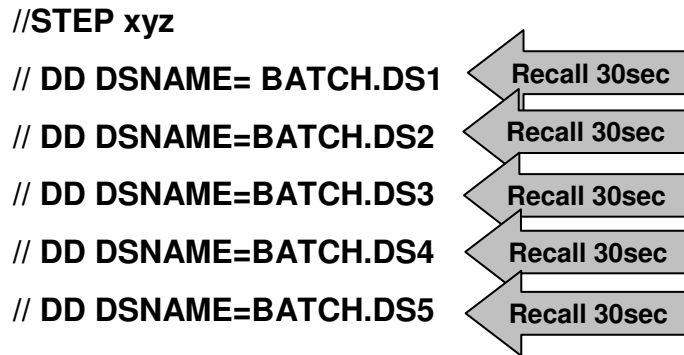
Agenda

- **General**
 - QUERY ACTIVE
 - RLS Server Error
 - UCB Storage Relief
 - Storage Group Priority
 - Recover w/ NEWNAME
- **Migration**
 - Throughput Improvements
- **Tape**
 - >40 Tapes
 - RECYCLE in place of TAPECOPY
 - Automatic restart of RECYCLE
- **Fast Replication**
 - Consistency Groups
 - Recover Data Set Enhancements
 - Performance
 - Usability
 - UCB Refresh
- **Storage Tiering**
- **Miscellaneous**


Miscellaneous

Allocation updated to enable you to specify that migrated data sets be allocated by a batch job be recalled in *parallel*, before each job step starts

**ALLOCxx Parmlib: BATCH_RCLMIGDS(PARALLEL)
SETALLOC: SETALLOC SYSTEM,BATCH_RCLMIGDS=PARALLEL**



Elapsed Recall Time = 150 seconds



Elapsed Recall Time = 30 seconds



What's New in DFSMSHsm

Glenn Wilcock
IBM

August 15, 2013
Session Number 14137



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.