#SHAREorg

SHARE

What's New in DFSMShsm

Glenn Wilcock
IBM

March 11, 2014
Session Number 15076

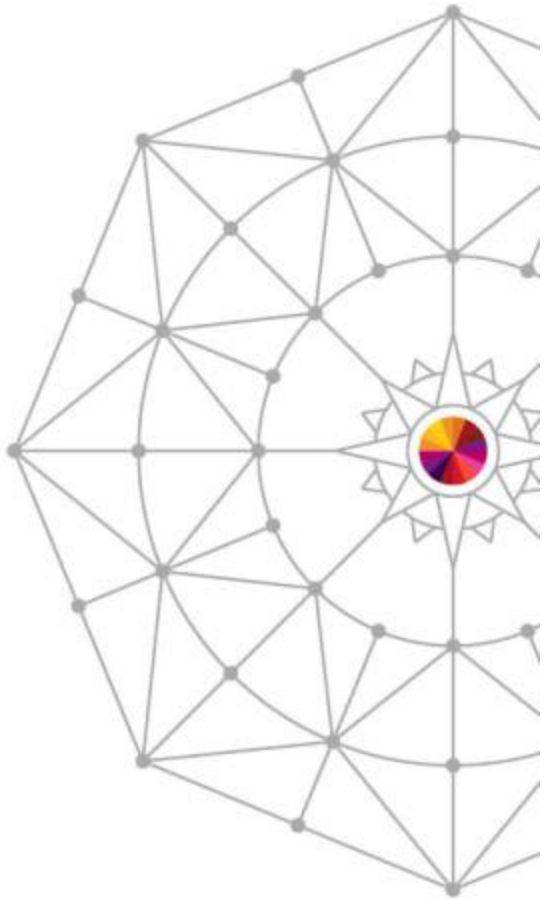

SHARE in Anaheim

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
- **Storage Tiering**
- **Fast Replication**
  - Consistency Groups
  - Recover Data Set Enhancements
  - Performance
  - Usability
  - UCB Refresh
- **Miscellaneous**

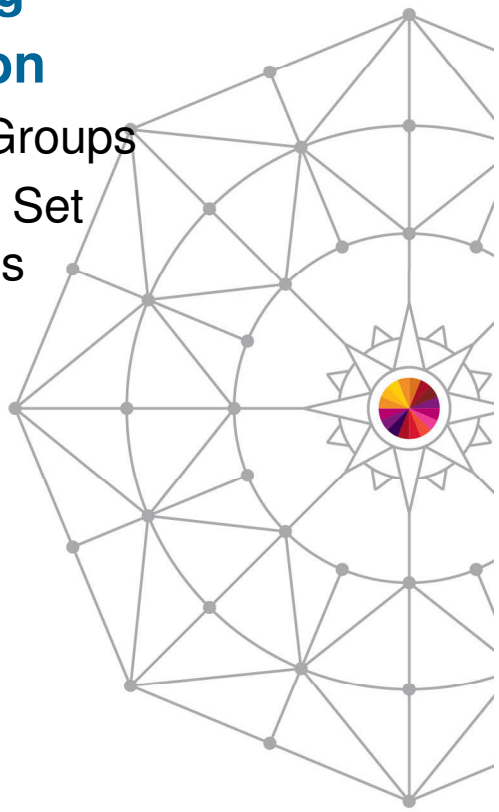

General
QUERY ACTIVE

- **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?

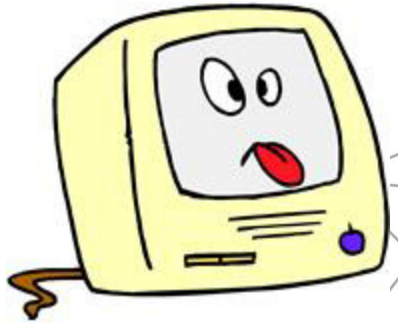

- **V2R1:** QUERY ACTIVE(TCBADDRESS) includes tape volser, device address and task name.

```
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

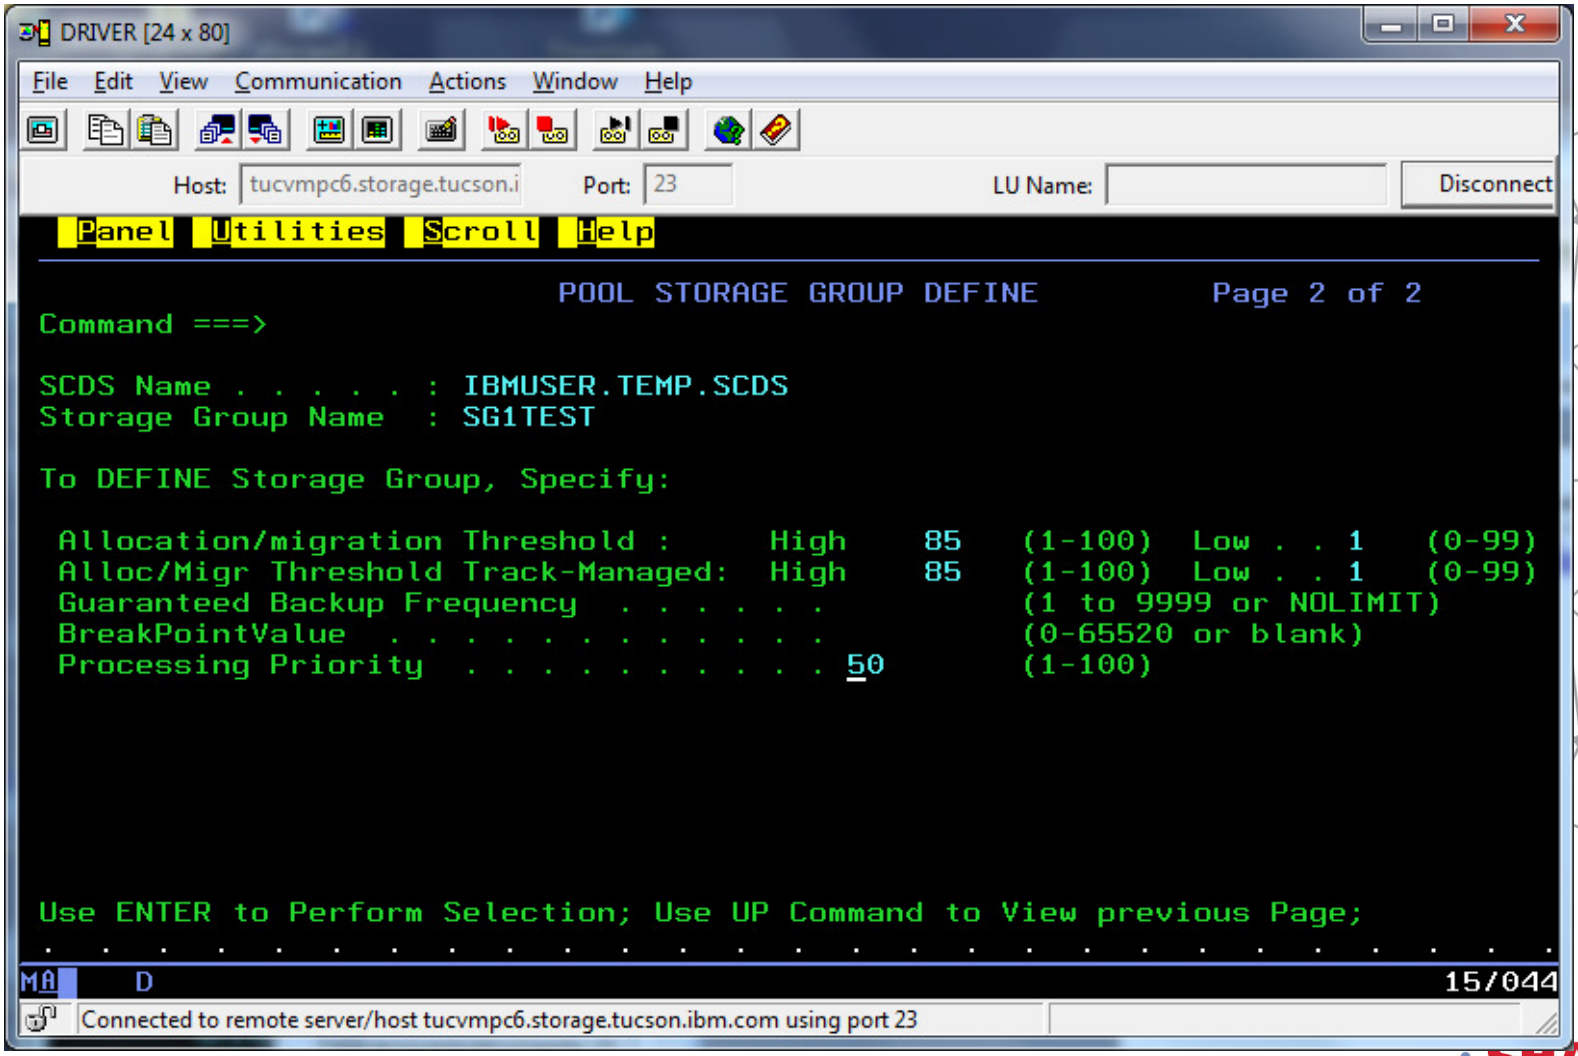

General
Recover VSAM data set with a NEWNAME

- **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

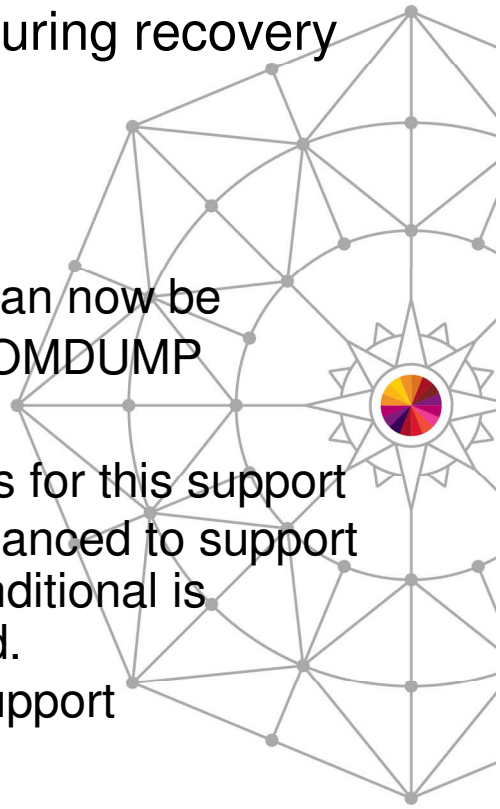

Agenda

- 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
- Storage Tiering
- Fast Replication
  - Consistency Groups
  - Recover Data Set Enhancements
  - Performance
  - Usability
  - UCB Refresh
- Miscellaneous

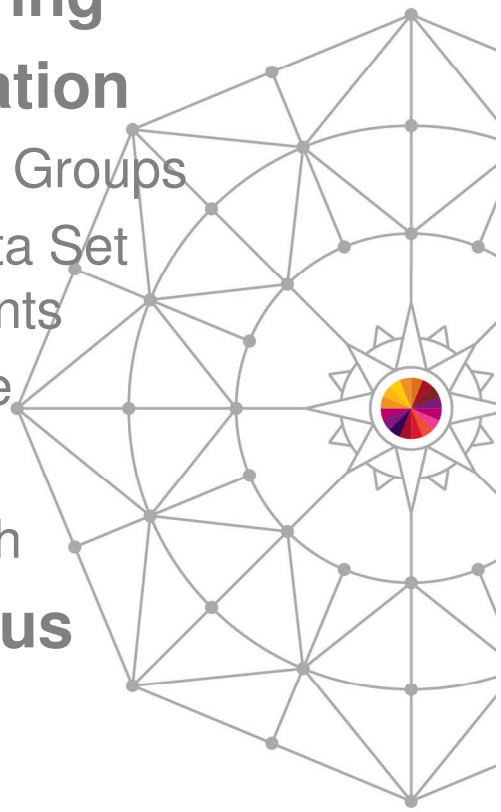

Migration
Throughput Improvement

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

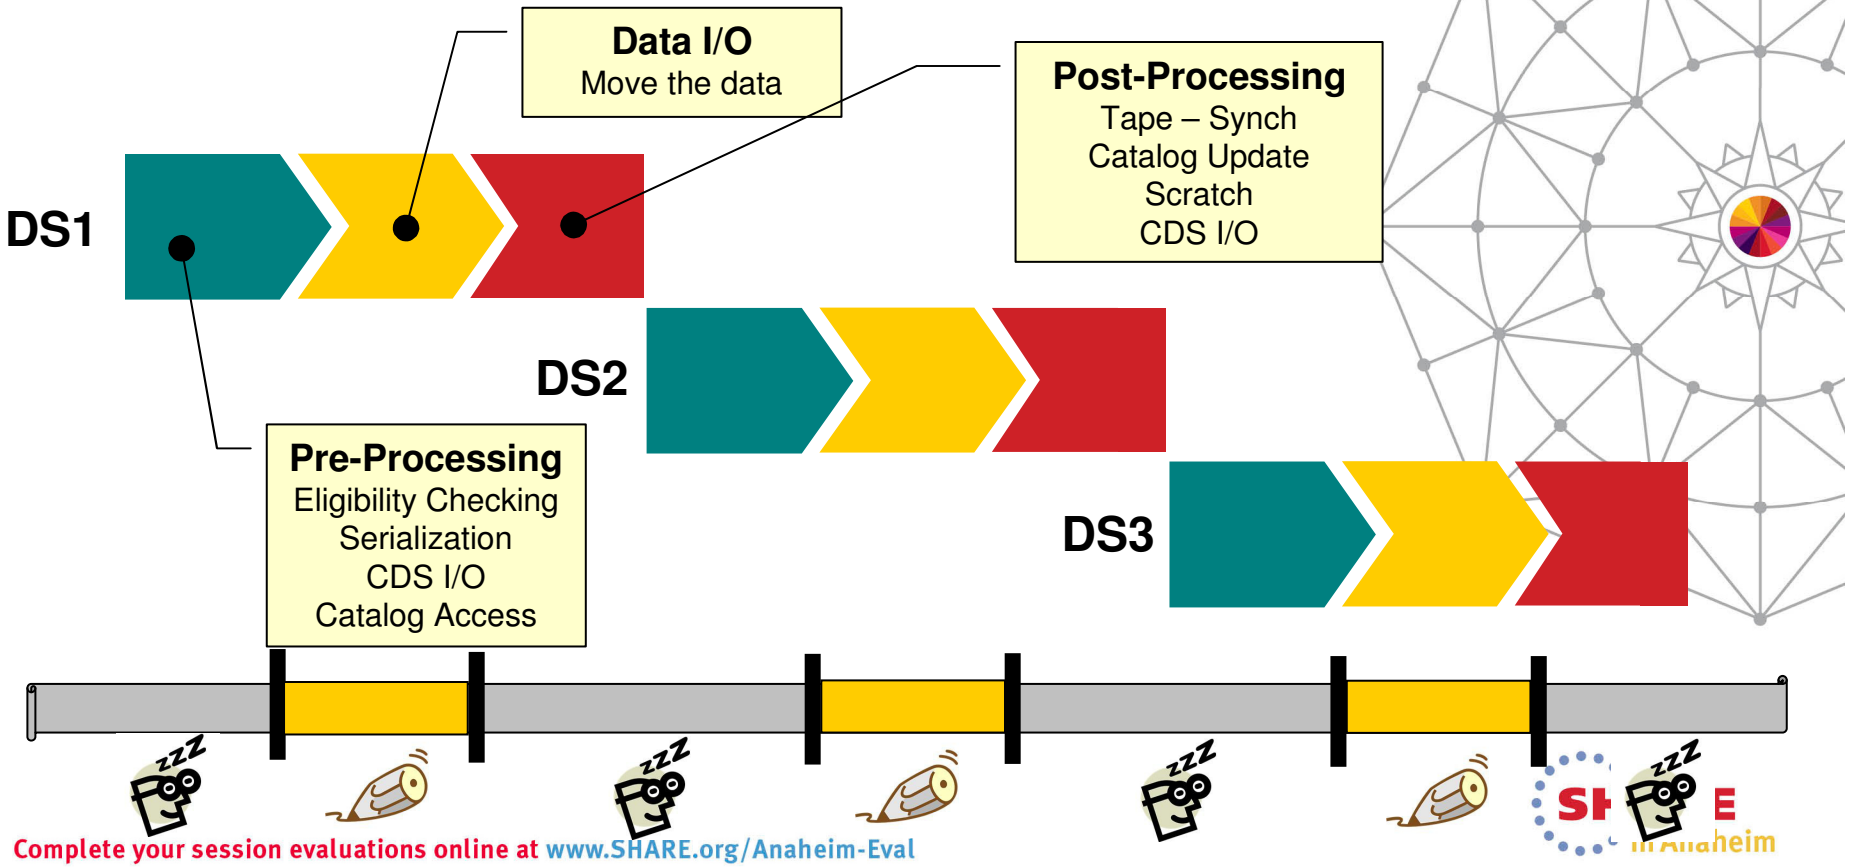

Migration
Throughput Improvement

- **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:
        >__ ___________________________________________ ____>
            |                                    _YES_      |
            |_ _ MIGRATIONSUBTASKS_ _(__|_____|__)_|
              |_ 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.
  - Maximum allowed subtasks is 105

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
- Storage Tiering
- Fast Replication
  - Consistency Groups
  - Recover Data Set Enhancements
  - Performance
  - Usability
  - UCB Refresh
- Miscellaneous

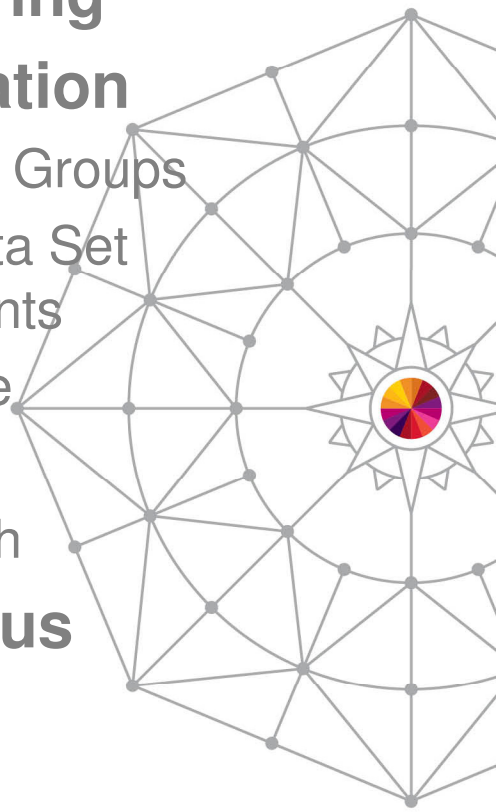

Tape
Span datasets over 40 Tapes

- **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_{tapes} * $4GB_{logical\ tape}$ * $2.5_{Compaction}$ = 400GB

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

  Example: 254 * 4GB * 2.5 = 2.5TB

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.

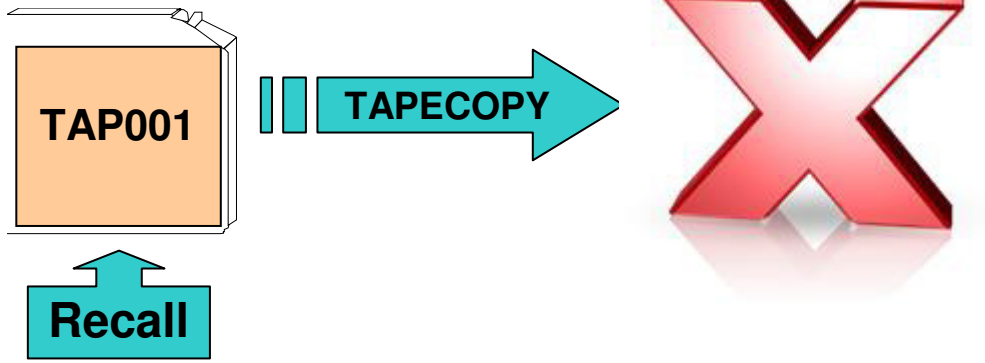

Tape
Use RECYCLE in place of TAPECOPY

- **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
  - RECYLE can be multi-tasked and doesn't have to start over from beginning if it is interrupted

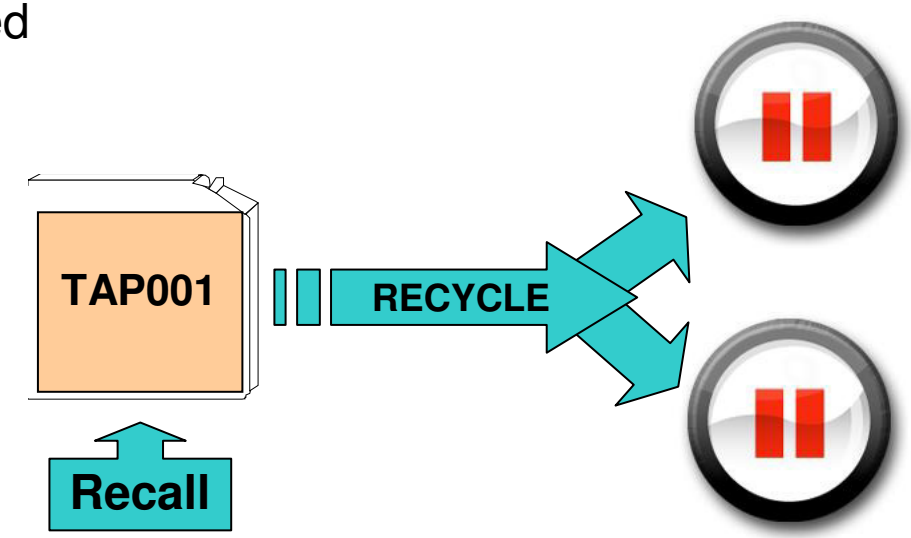

Tape
Automatically restart RECYCLE

- **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 ATTEMPS 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
- Storage Tiering
- Fast Replication
  - Consistency Groups
  - Recover Data Set Enhancements
  - Performance
  - Usability
  - UCB Refresh
- Miscellaneous

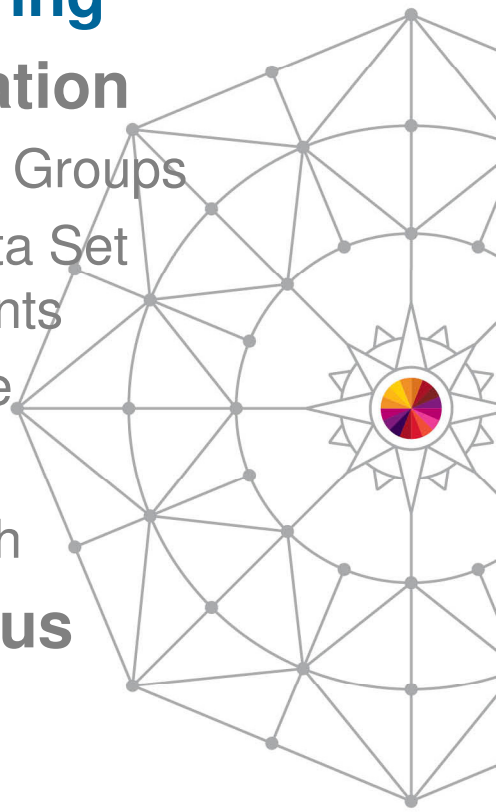

Storage Tiers Background

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

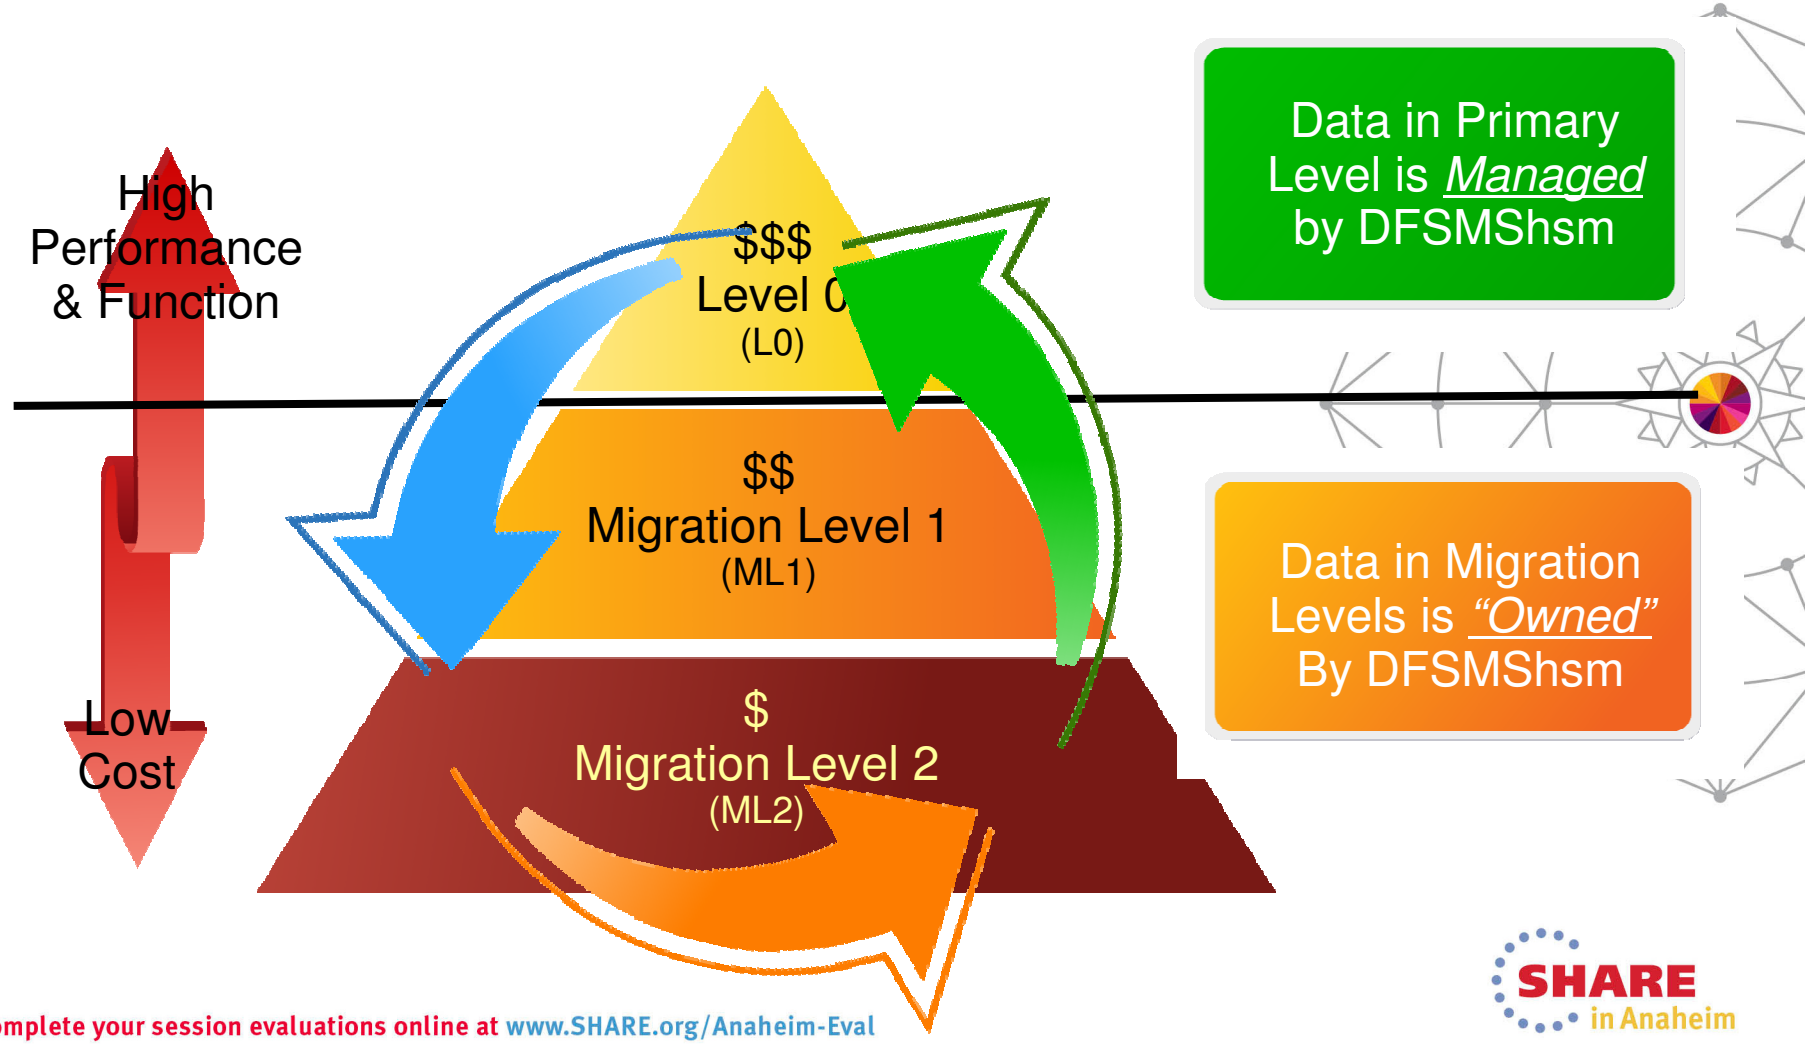

Storage Tiers Background

- **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 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 DFSMShsm 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*

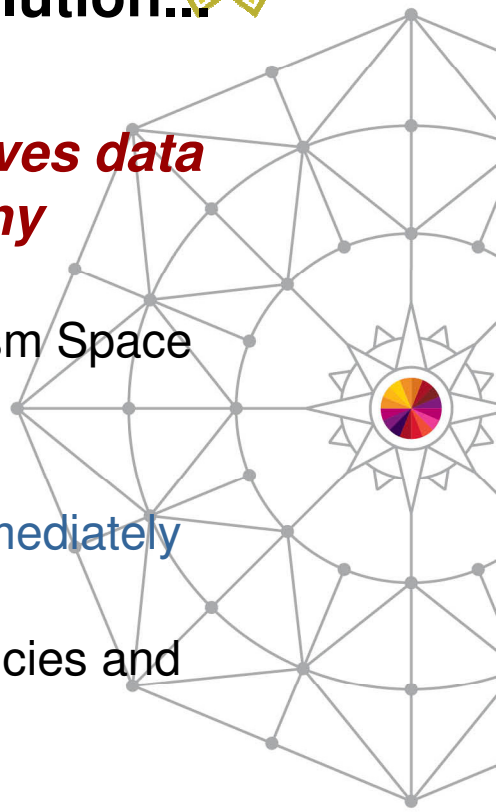

Storage Tiers Overview

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

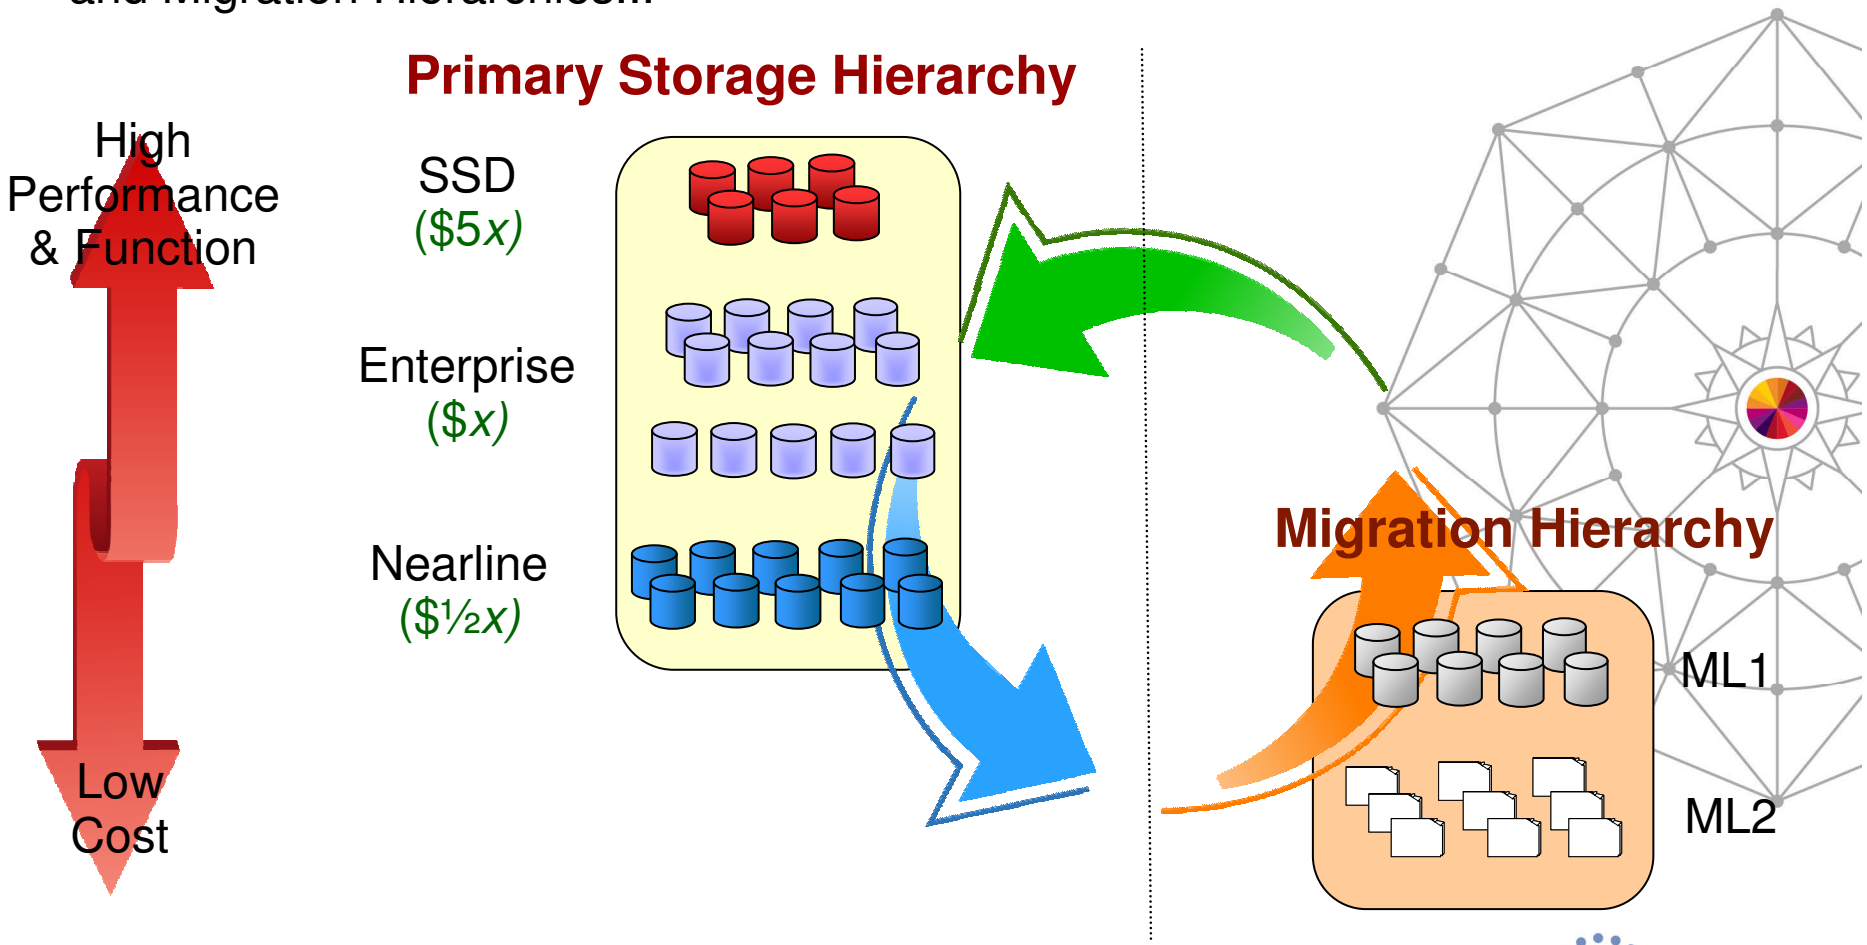

Storage Tiers Overview

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.

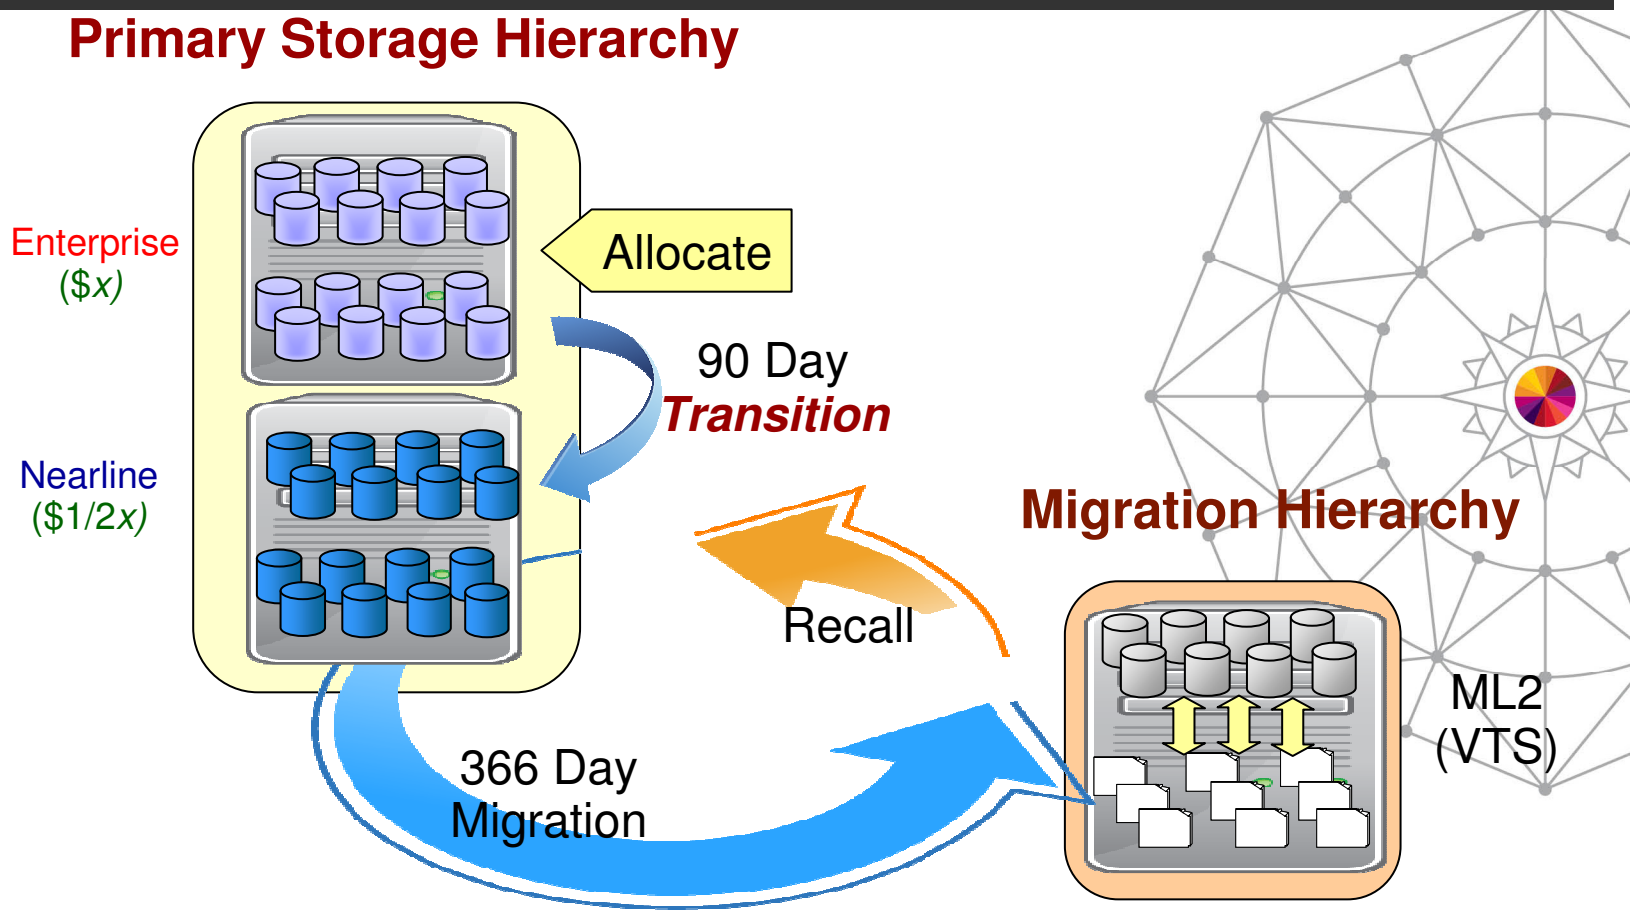

Storage Tiers

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.

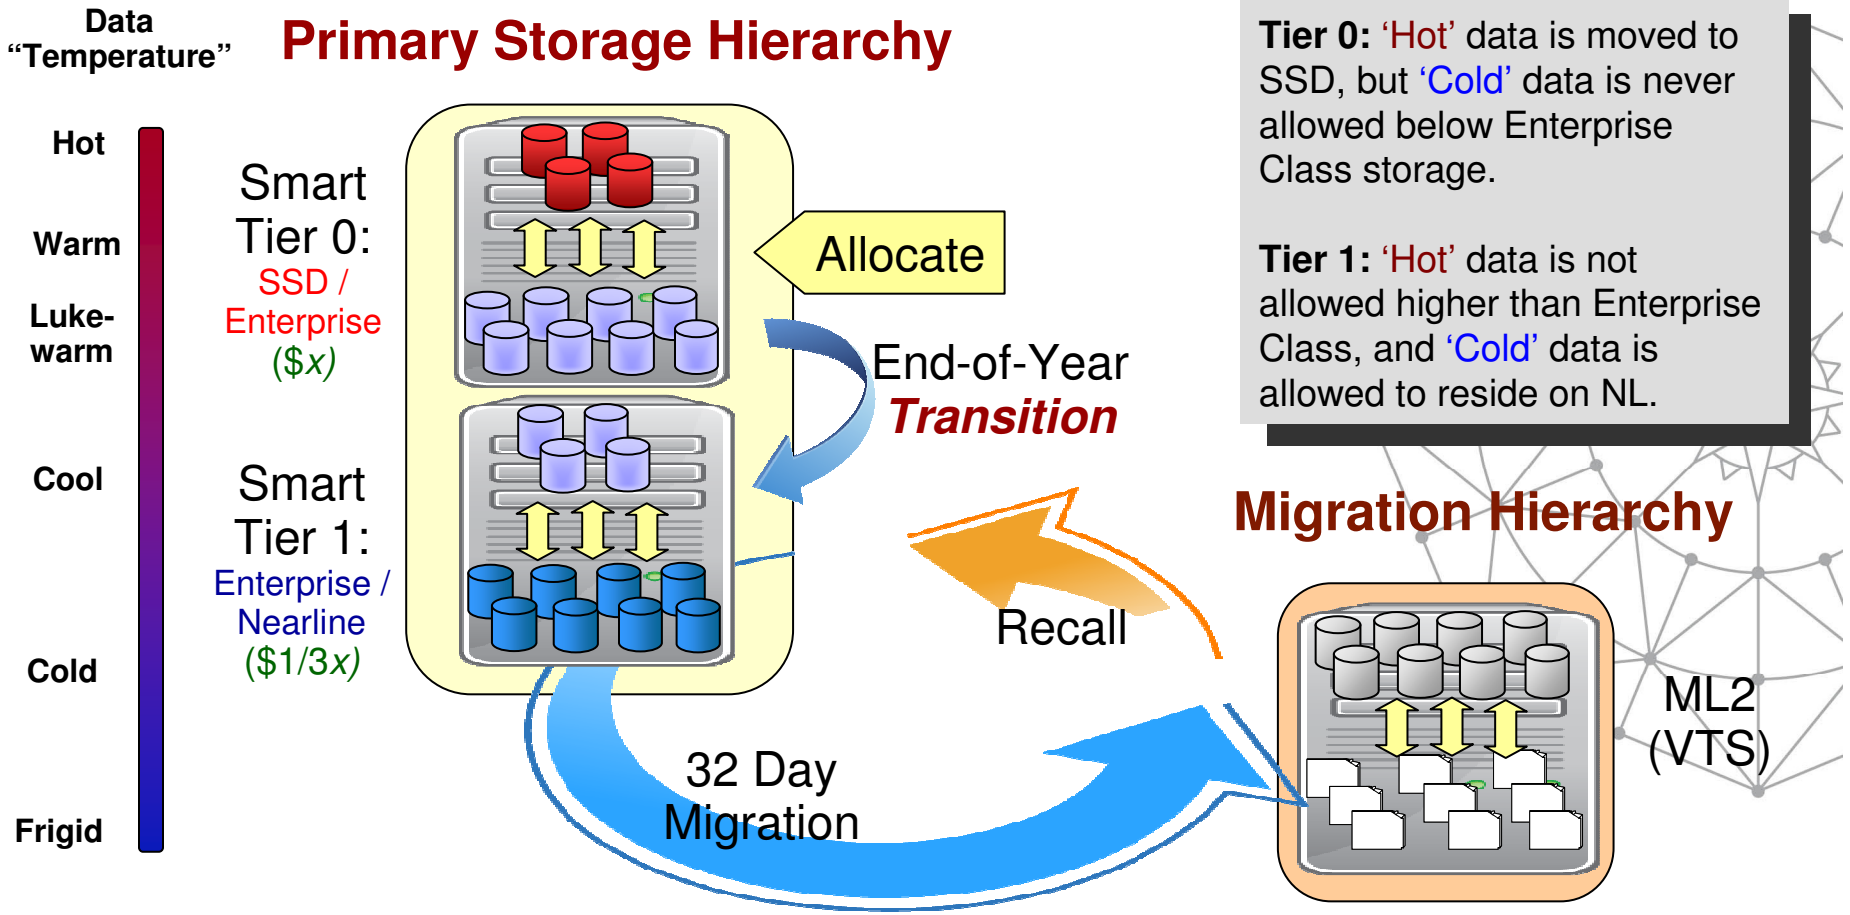

Tier 0: 'Hot' data is moved to SSD, but 'Cold' data is never allowed below Enterprise Class storage.

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

Storage Tiers

- *Contrasting DFSMS Tiering with Storage Controller Tiering*

	DFSMS Tiering	Controller Tiering
Movement Boundary	Data Set Level	Physical Extent Level
Level of Management	Policy Based	Data Temperature Based
Scope	Sysplex (across controllers)	Intra-controller
Access	Closed Data Only	Open and Closed Data
Impact	Data must be quiesced	Transparent
Cost	Host based MIPS	No host based MIPS

Storage Tiers: V2R1 Overview

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

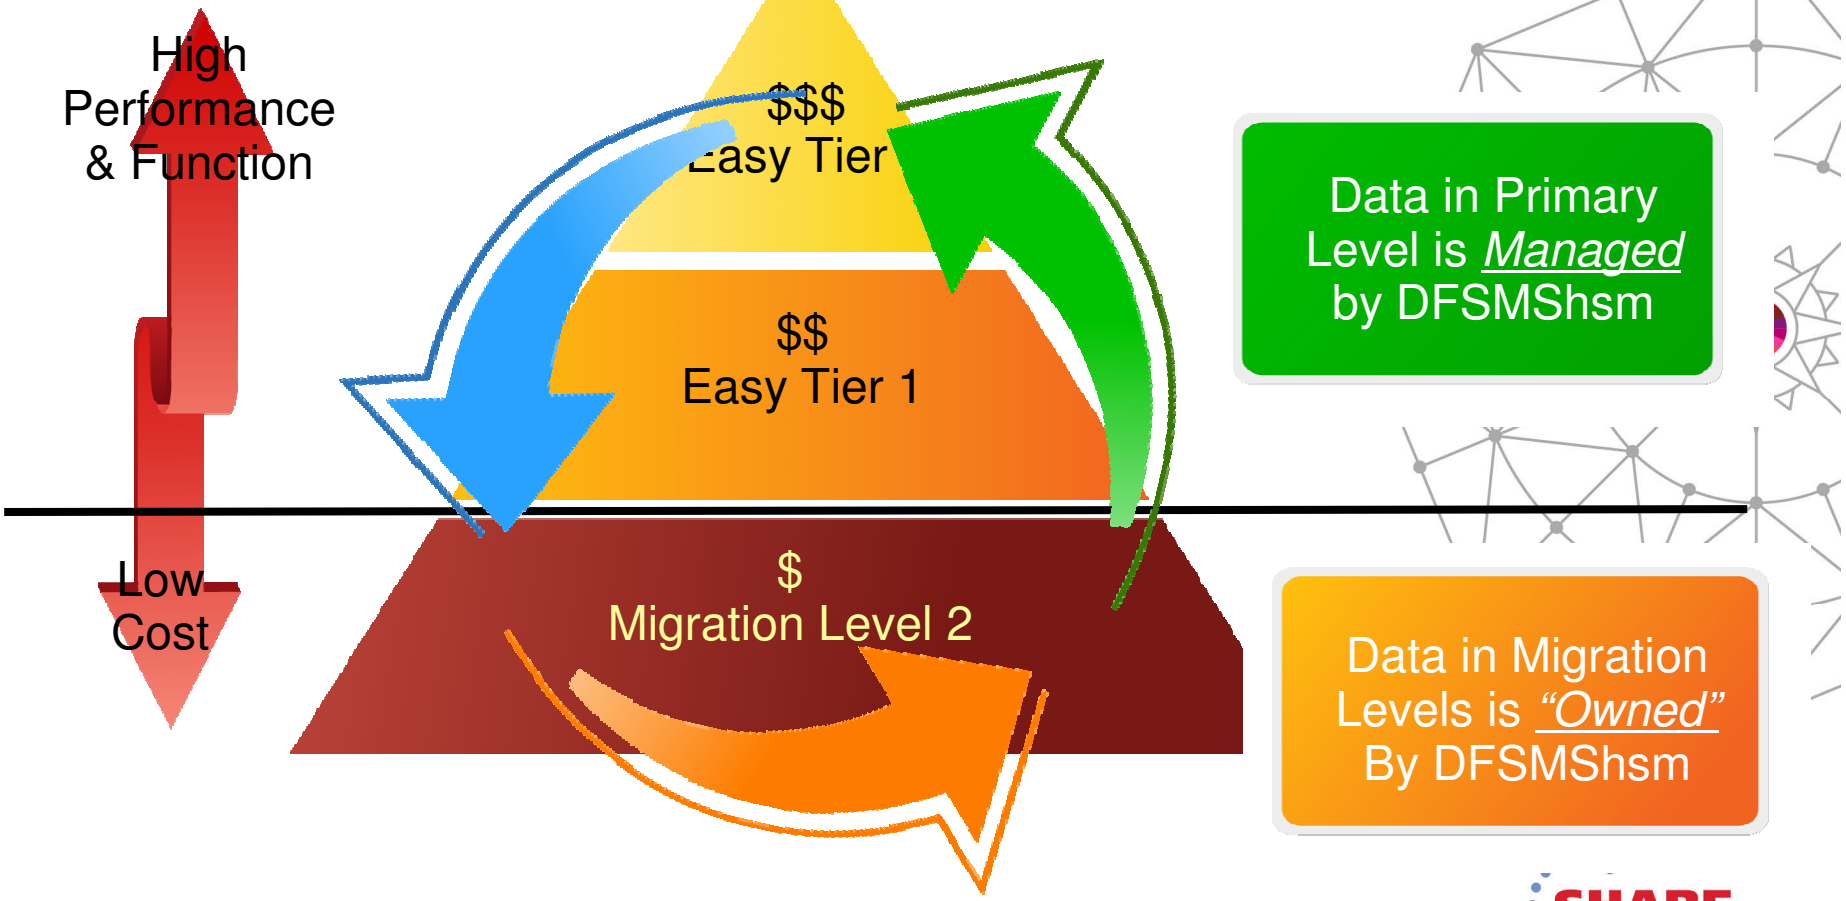

Storage Tiers
Overview

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

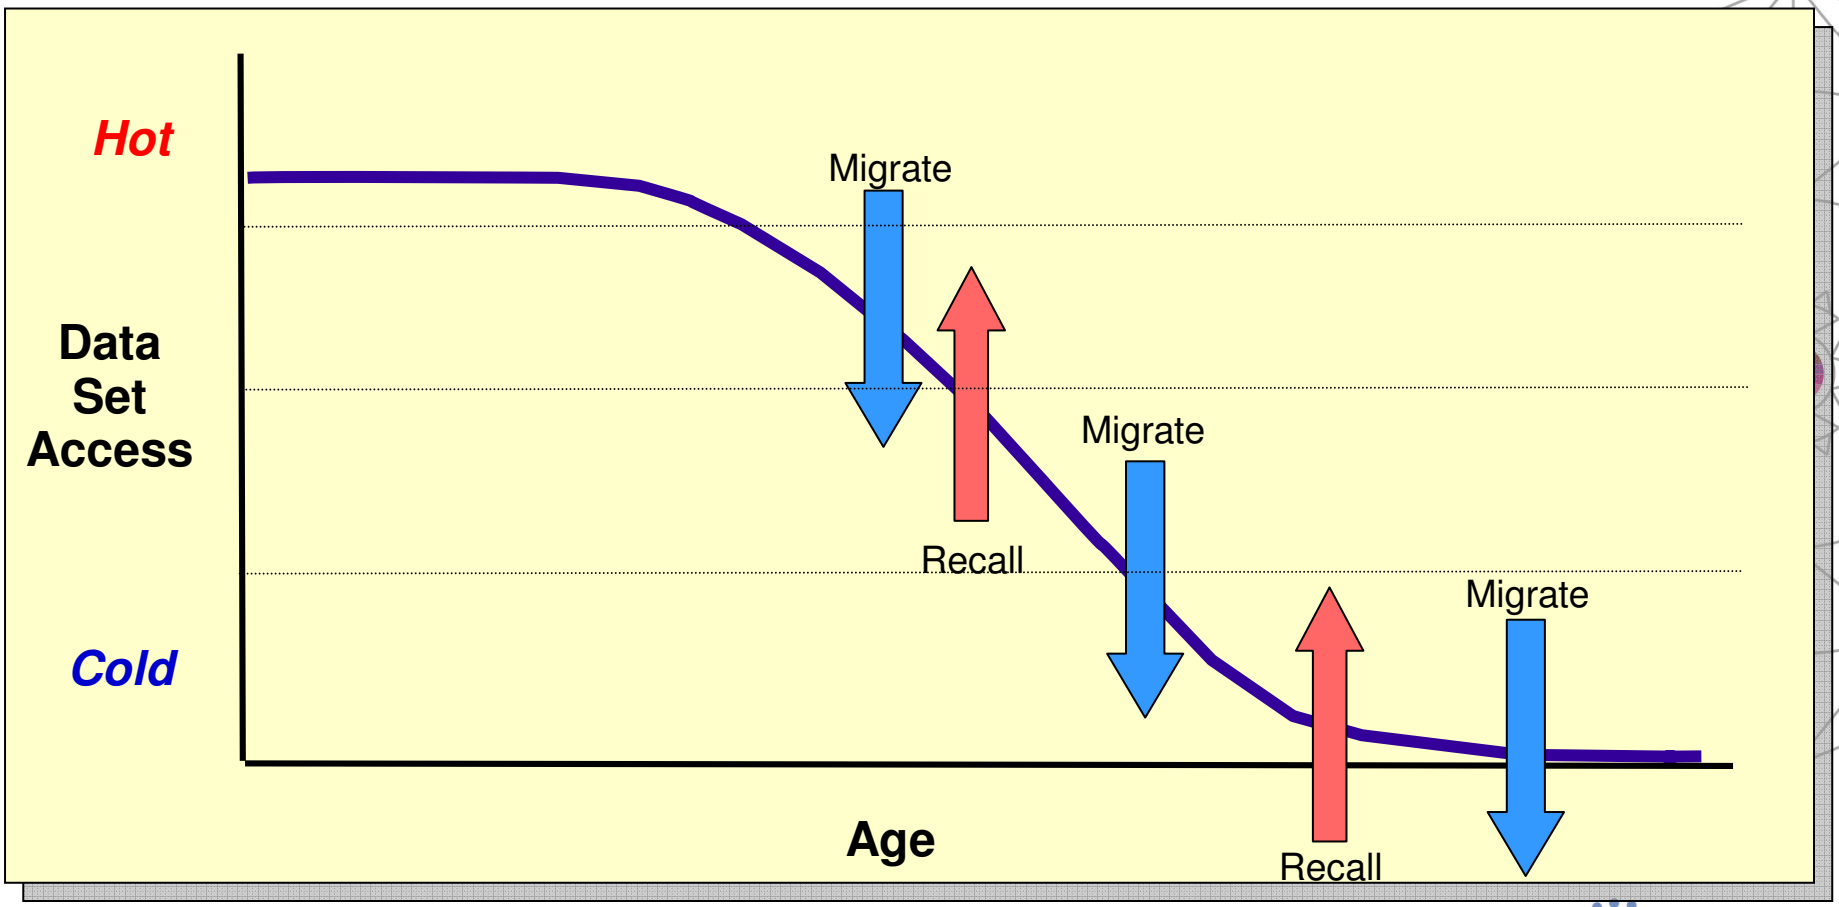

Storage Tiers
Overview

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

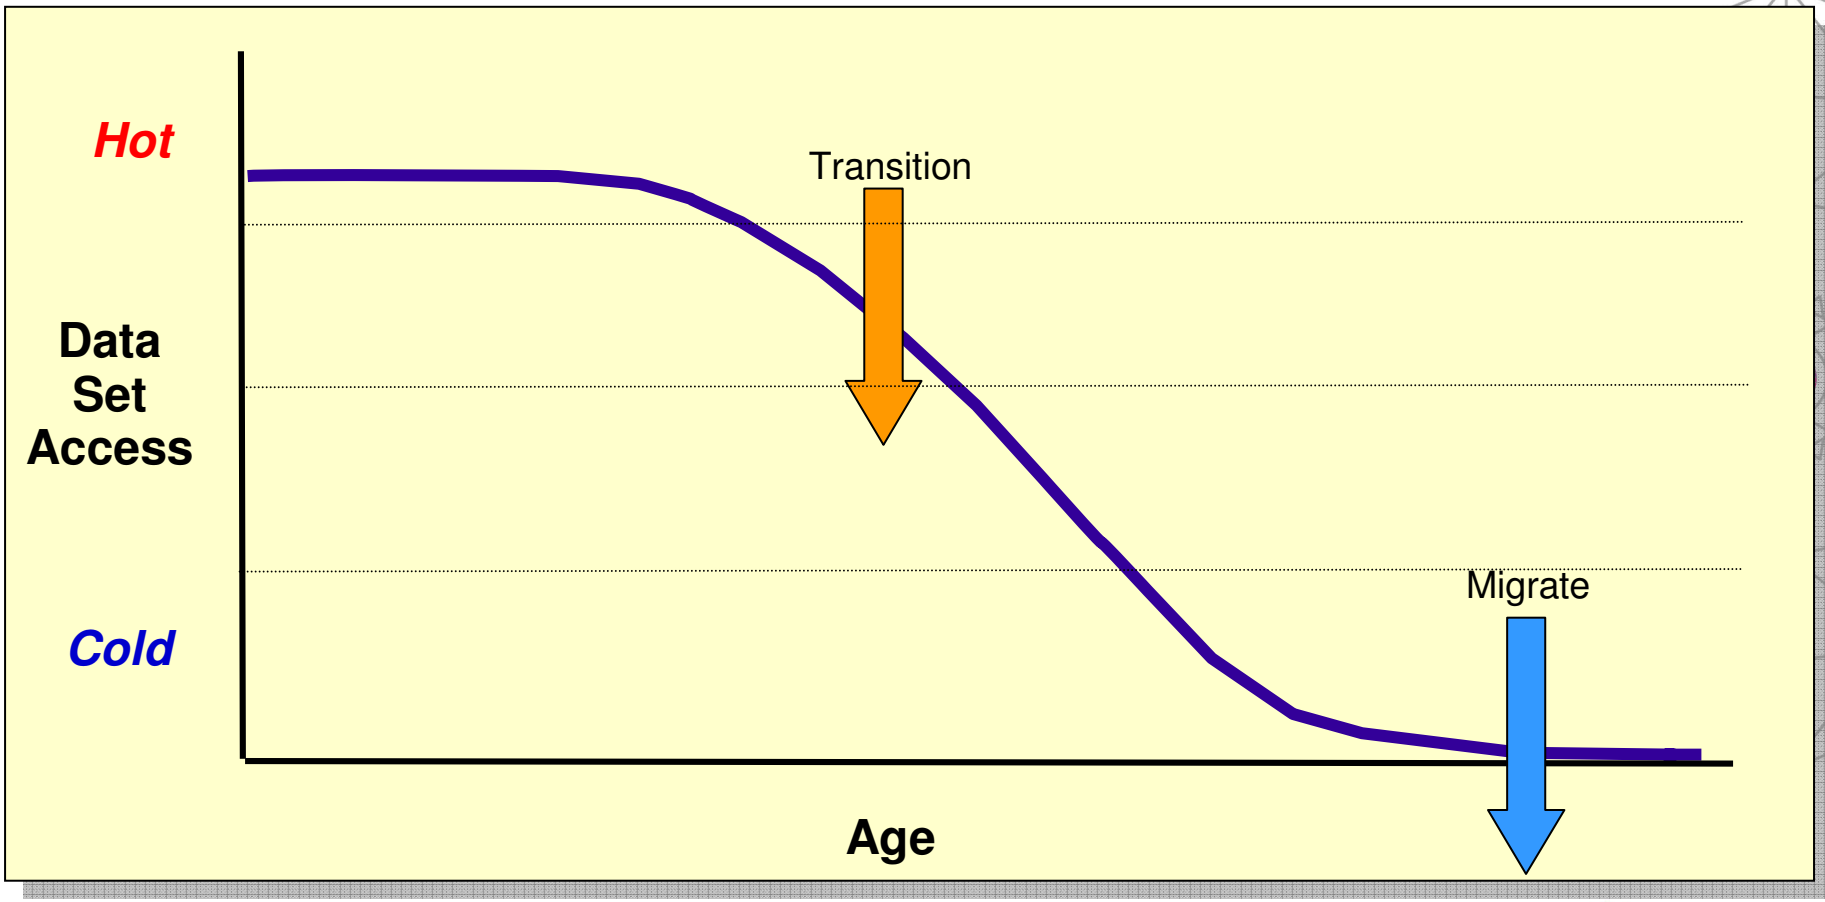

Storage Tiers Overview

Contrast the Space Management processing required when ...

- A data set is migrated to ML1, then ML2, recalled and remigrated just once

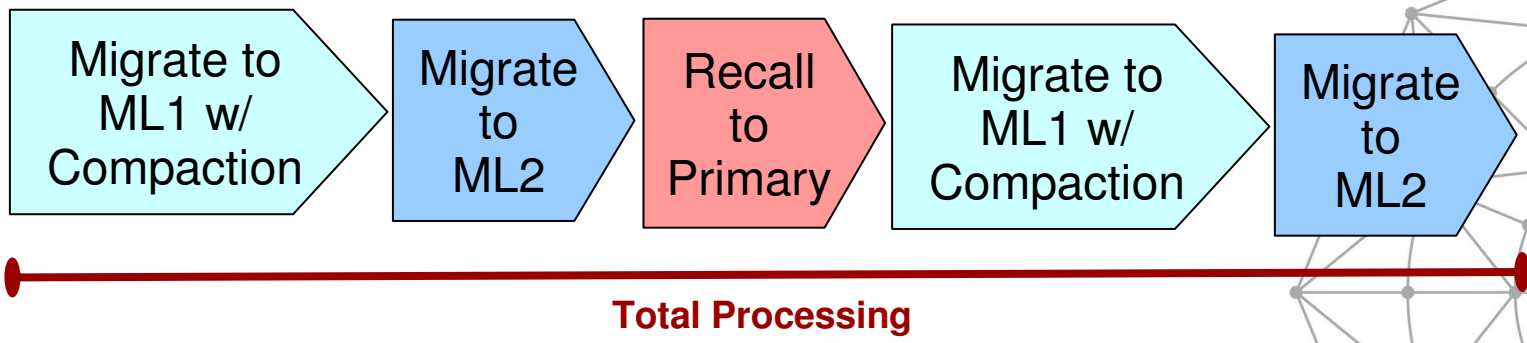

- A data set is transitioned using FlashCopy and then migrated directly to ML2

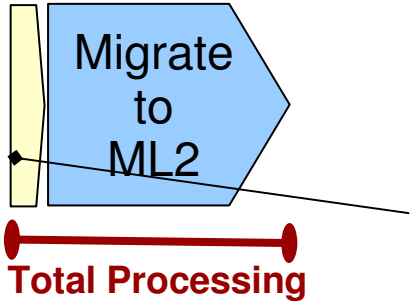

- **Transition** with FlashCopy uses minimal host MIPS. Data movement is done within the storage controller.

Storage Tiers
Usage

- **What are examples of data sets that would benefit from this solution?**

★Data sets currently not eligible for migration because they always need to be immediately accessible
- Recall delay is unacceptable
- Data sets could be allocated on a particular class of storage and then later *transitioned* to a less expensive class of storage for permanent retention

★Data sets that are eligible for migration today, but there would be a benefit to keep them online for a longer period of time.
- Convert the migration of data sets to *transition* to a lower cost storage.
- Increase the number of days that the data set must be unreferenced before migrating directly to ML2 (PRIMARY DAYS NON-USAGE)

Storage Tiers
Usage

- **There is a distinction between using Migrate/Recall and Class Transitions**
  - When a data set is recalled, it will be returned to the class of storage as directed by the ACS routines, which would typically be higher than where a data set would reside after a transition
  - When a data set transitions to a lower class of storage, it will remain there until it is transitioned again or until it migrates

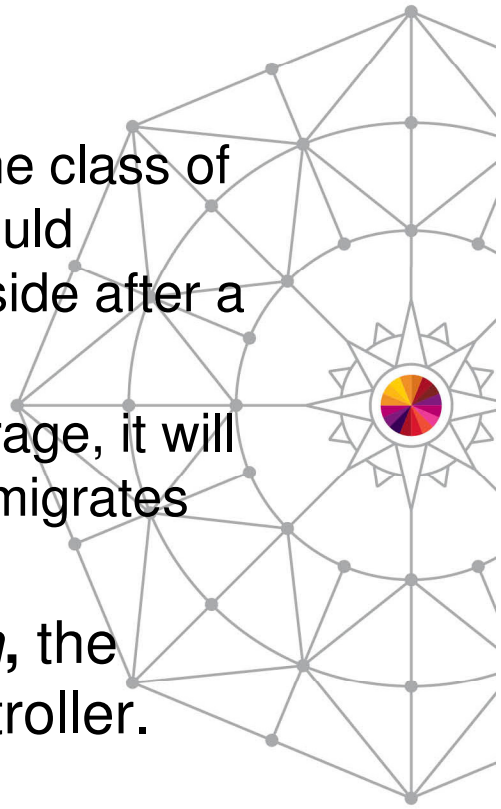

- *In order for FlashCopy to be used for a transition*, the movement must be within the same storage controller. This may be difficult to achieve.
  - ✓ FlashCopy will be discussed in detail later in the presentation

Storage Tiers
Invocation

- The function of *DFSMShsm Space Management* processing is to use *policy-based automation* to ensure that volumes within the Primary Storage Hierarchy have enough *free space* for new data and to ensure that data is stored at the *lowest acceptable tier in the Storage Hierarchy*
  - This is accomplished through
    - Data set expiration
    - Migration of unreferenced data to the Migration Hierarchy
    - *"Class Transitions" within the Primary Hierarchy*

- "Class Transition" processing is integrated into the existing DFSMShsm *Space Management* functions
  - Primary Space Management
  - On-Demand Migration
    - *New function introduced in V1R13. Performs space management on a volume as soon as it goes over its high threshold. Replacement for on-the-hour Interval Migration processing.*
  - Interval Migration

Storage Tiers
Invocation

- The SMS Management Class provides the Class Transition policies:
  - **Class Transition Criteria:** If and when a data set should be transitioned
  - **Serialization Error Exit:** Indicates what type of special processing should occur if the data set cannot be serialized
  - **Transition Copy Technique:** Which copy technique should be used to move the data set

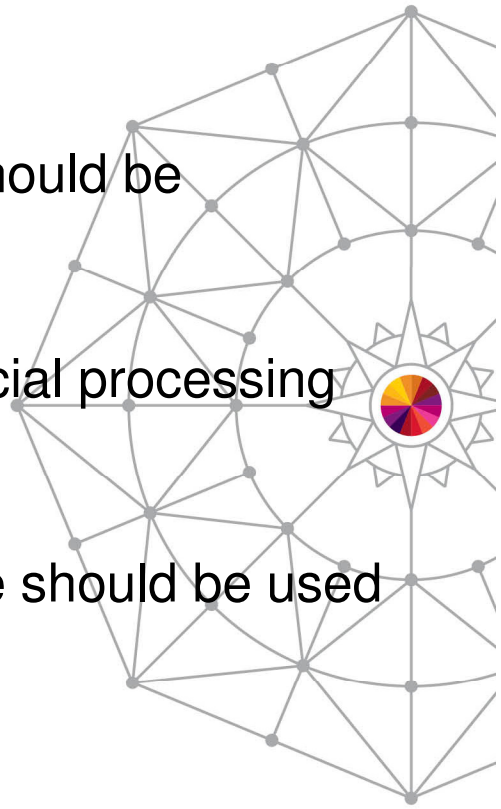

Storage Tiers
Invocation

- The Class Transition Criteria specifies if and when a data set should be transitioned.
  - **Default:** Class transitions are *not* performed
  - **Time since Creation:** Data set is eligible for a transition on or after this time.
    - This is a *subjective* setting. It indicates that regardless of the usage of the data set, it should be transitioned.
  - **Time since Last Use:** Data set is eligible for a transition on or after this time.
    - This is an *objective* setting. It indicates that a data set should not be transitioned until it has not been referenced for a certain period of time.
  - **Periodic:** Data set is eligible for a transition on a specific date.
    - This is a *subjective* setting. It indicates that regardless of the usage of the data set, it should be transitioned.
- Only one criteria may be specified.

Storage Tiers
Invocation

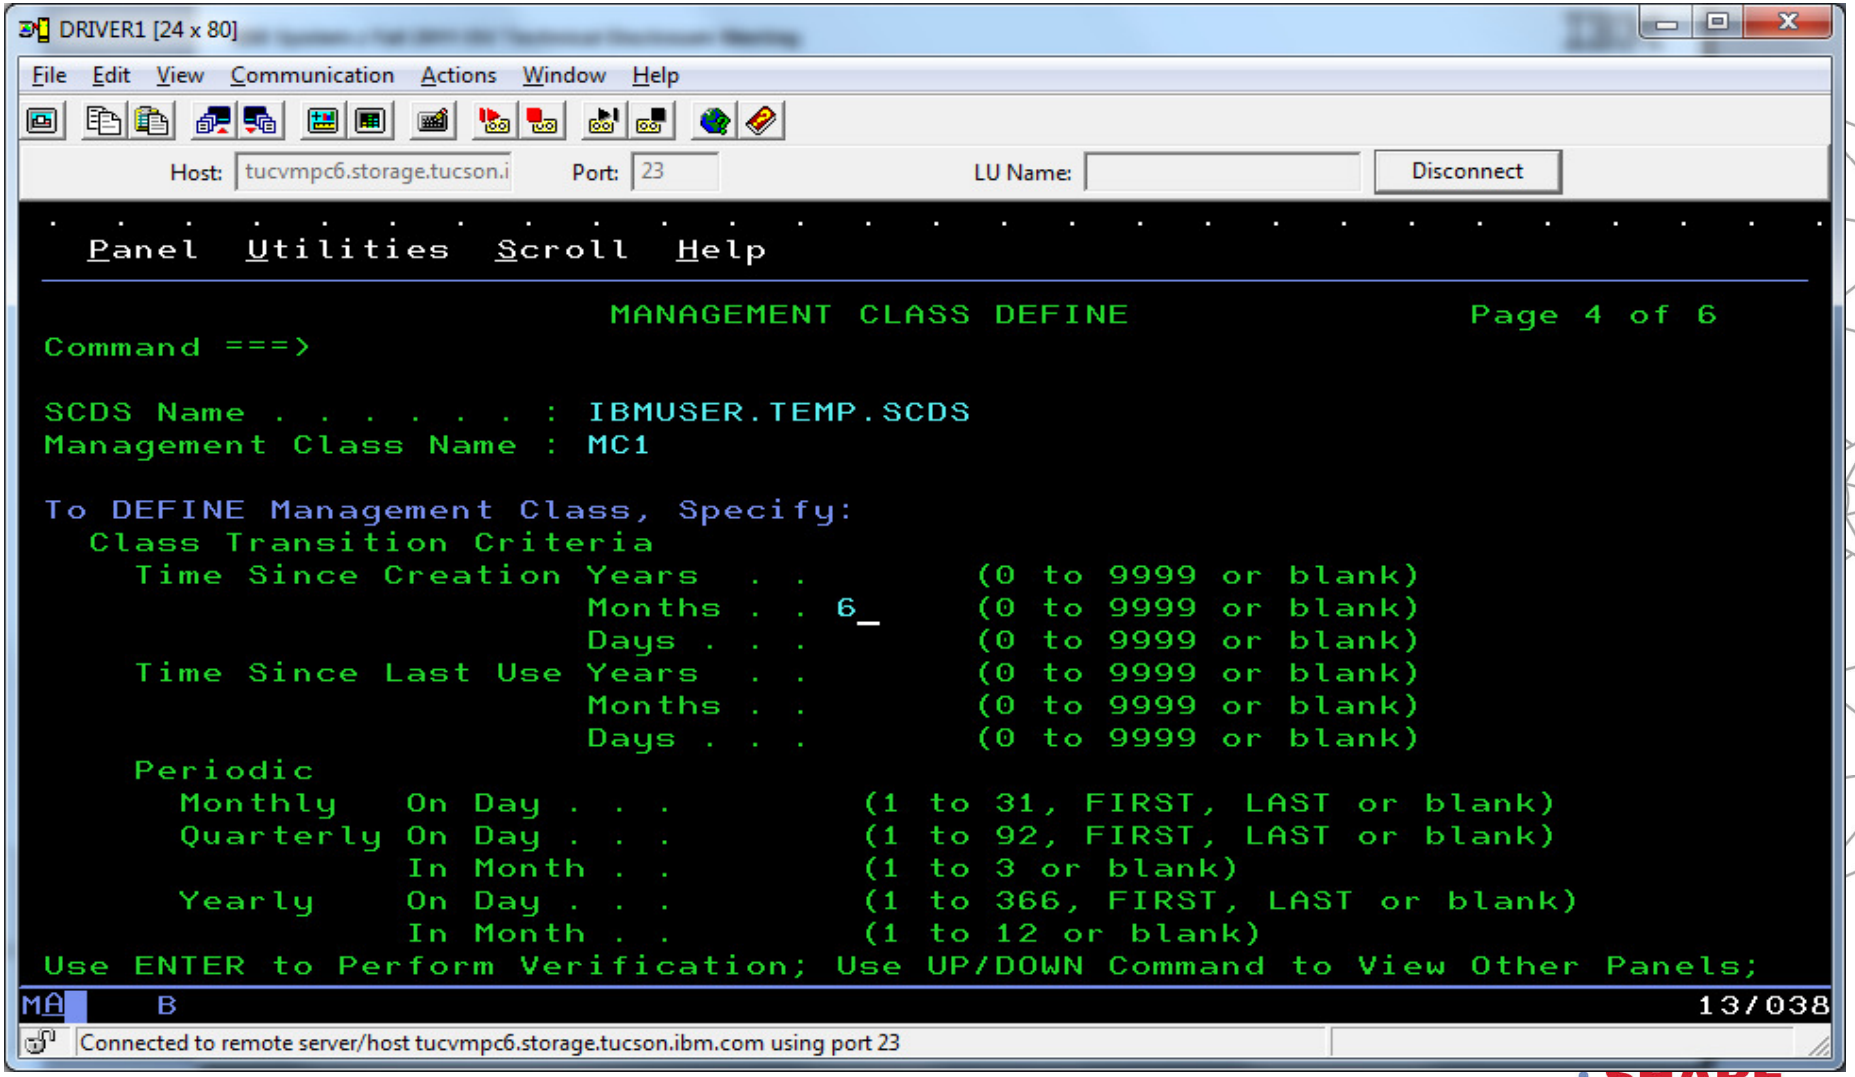

Storage Tiers
Invocation

- The Serialization Error Exit indicates what type of special processing should occur if the data set cannot be exclusively serialized
  - Since **'Time since Creation'** and **'Periodic'** may specify that a transition occur during a period of time when a data set is being accessed, this setting specifies what to do if a data set cannot be exclusively serialized for data movement
  - For database data, it may be expected that the data is always open and special processing must be done to transition the data at any time

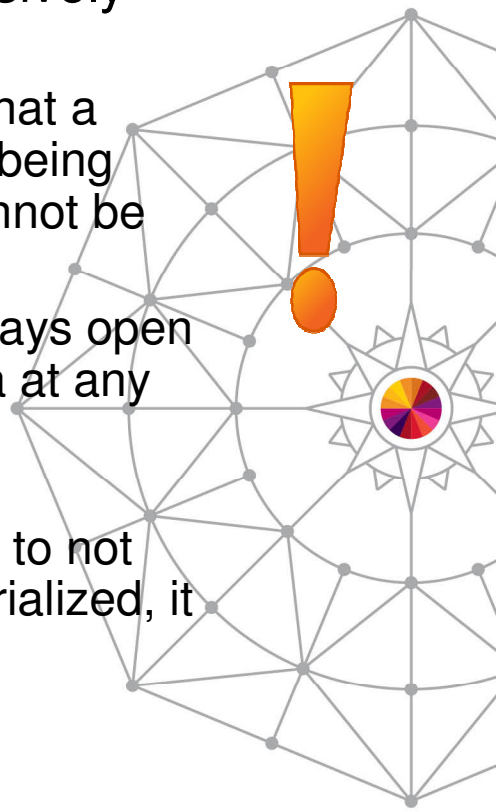

- Since it is expected that data sets may be open, the default is to not issue an error message if a data set cannot be exclusively serialized, it is just skipped *(similar to migration processing)*
  - To issue a DFSMShsm message for this condition, issue PATCH .MGCB.+EF BITS(. . .1 . . . .)
  - An FSR record is created to enable reporting for this condition

Storage Tiers
Invocation

- Serialization Error Exit settings
  - **NONE:** (Default) If a data set cannot be exclusively serialized, the data set is not transitioned
  - **DB2:** Data assigned to this management class are DB2 objects
    - Invoke DB2 to close and unallocate the object. If this is successful, the object is serialized and moved and DB2 is invoked to reopen the object
  - **CICS:** Data assigned to this management class are CICS objects
    - Invoke CICS to take the object offline. If this is successful, the object is serialized and moved and CICS is invoked to reopen the object
  - **zFS:** Data assigned to this management class are zFS data sets
    - Invoke zFS to unmount the data set. If this is successful, the data set is serialized and moved and zFS is invoked to remount the data set
  - **EXIT:** User exit is invoked to preprocess and post-process the data set
    - Enables users / ISVs to provide an exit that will be invoked before and after transitioning an allocated data set. The data set is only transitioned if serialization is obtained after the first invocation of the exit
    - Documented in *DFSMS Installation Exits (SC26-7396)*

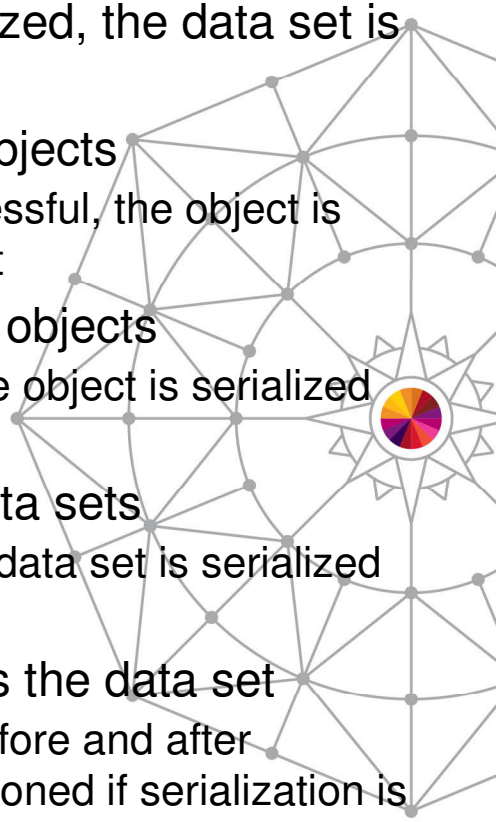

Storage Tiers
Invocation

- The Transition Copy Technique indicates what type of copy technique should be used to move the data set
  - **Standard:** (Default) Use standard I/O
  - **Fast Replication Preferred:** Prefer Fast Replication. If it cannot be used, then use standard I/O.
  - **Fast Replication Required:** Require Fast Replication. If it cannot be used, fail the data movement.
    - Requires the target volume to be in the same storage controller.
  - **Preserve Mirror Preferred:** Prefer Preserve Mirror. This indicates that a Metro Mirror primary volume is allowed to become a FlashCopy target volume. If Preserve Mirror cannot be used, FlashCopy or standard I/O may be used.
  - **Preserve Mirror Required:** Require Preserve Mirror. The transition is only performed if the Metro Mirror primary target volume will not go duplex pending. This parameter has no affect if the target volume is not a Metro Mirror primary volume.

Storage Tiers
Invocation

- Transition Copy Technique
  - If a copy technique *other than* 'Standard' is specified, then *a valid backup copy must exist and the DS Change Indicator OFF* before the data set is transitioned
    - This is required because DFSMShsm receives control immediately after the FlashCopy relationship is established.
    - DFSMShsm is not notified if there is a physical error within the storage controller during the background physical copy
    - While very unlikely, this ensures that if any physical error occurs while transitioning the data, that a backup copy is available to recover the data set
      - *This is critical because the data set is Deleted after the logical completion*

Storage Tiers
Invocation

- The new Storage Group Processing Priority specifies the relative order in which storage groups should be processed during Primary Space Management
  - *Class Transition processing will move data from one storage group to another*
  - In order to help ensure that the 'receiving' storage groups have enough space for the data sets that will be moved to them, a new storage group Processing Priority is provided
    - These storage groups should be assigned a higher priority
  - 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

Storage Tiers
Invocation

- Once DFSMShsm determines that a data set has met the Class Transition criteria specified by the Management Class, it invokes the ACS routines to determine what the transition should be
  - ACS Routines are invoked with new ACS environment (&ACSENVIR) of SPMGCLTR, for 'space management class transition'
  - The following routines are invoked *(in this order)*
    - Storage Class
    - Management Class
    - Storage Group
  - Any or all can be *transitioned*

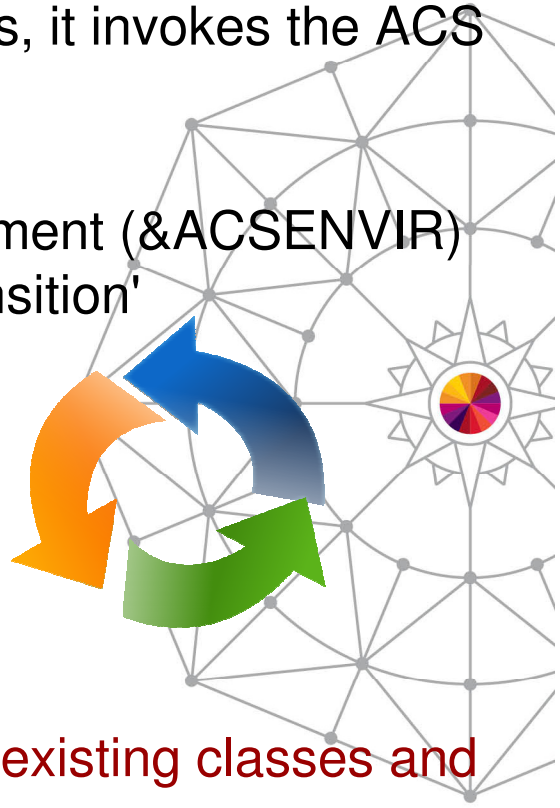

✓ If the classes and storage group returned match the existing classes and storage group, *then no transition occurs*

Storage Tiers
Invocation

- Storage Class
  - Storage Class indicates the *'preferred'* class of storage to which the data set should be allocated
  - If storage class changes, but storage group remains the same, *and a device matching the new storage class attributes cannot be selected*, the data set is *not* moved

```
IF &ACSENVIR = 'SPMGCLTR' THEN
   SELECT (&STORCLAS)
       WHEN ('SSD')        SET &STORCLAS = 'EASYTIER'
       OTHERWISE           SET &STORCLAS = &STORCLAS
   END
 ELSE ...
```

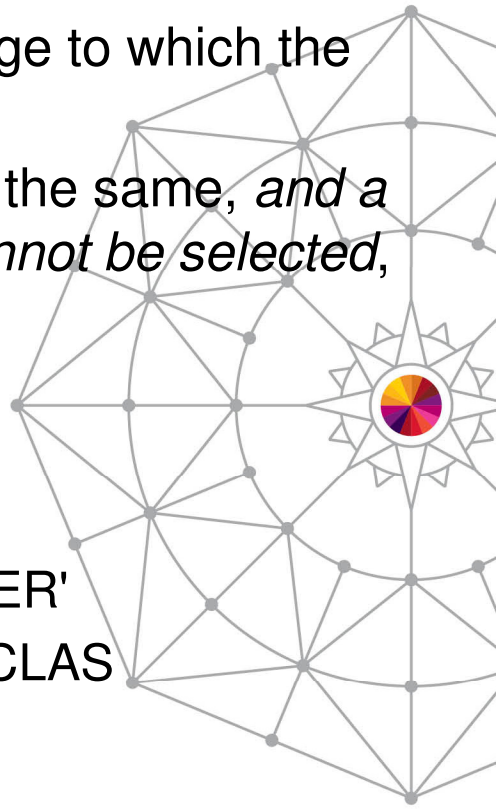

Storage Tiers
Invocation

- Management Class
  - When a new management class is assigned, DFSMShsm will begin using the newly assigned policies to manage the data set
  - If only the management class changes, then the data set is altered to assign it to the new management class and *no data movement is performed*
  - *Example*
- Upon creation, a data set is assigned to a management class for which the data set is only eligible to migrate to ML1 (not ML2) and 2 backup copies are maintained
  - After 120 days from creation, the data set is transitioned to a different management class for which the data set is eligible to migrate to ML2 and only 1 backup copy is maintained.

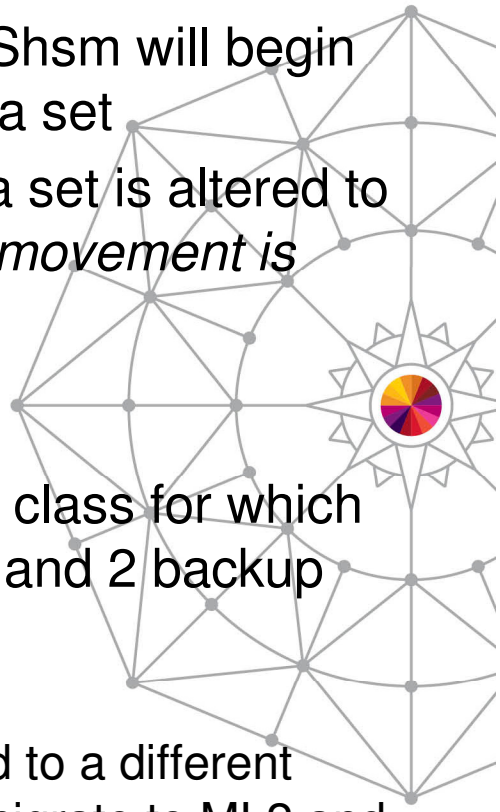

Storage Tiers
Invocation

- Management Class

```
IF &ACSENVIR = 'SPMGCLTR' THEN
/*  SPACE MANAGEMENT CLASS TRANSITION  */

   SELECT (&MGMTCLAS)
      WHEN ('NOML2')    SET &MGMTCLAS = 'ML2OK'
      WHEN ('DB2NEW') SET &MGMTCLAS = 'DB2AGED'
      OTHERWISE         SET &MGMTCLAS = &MGMTCLAS
      END
ELSE
   ...
```

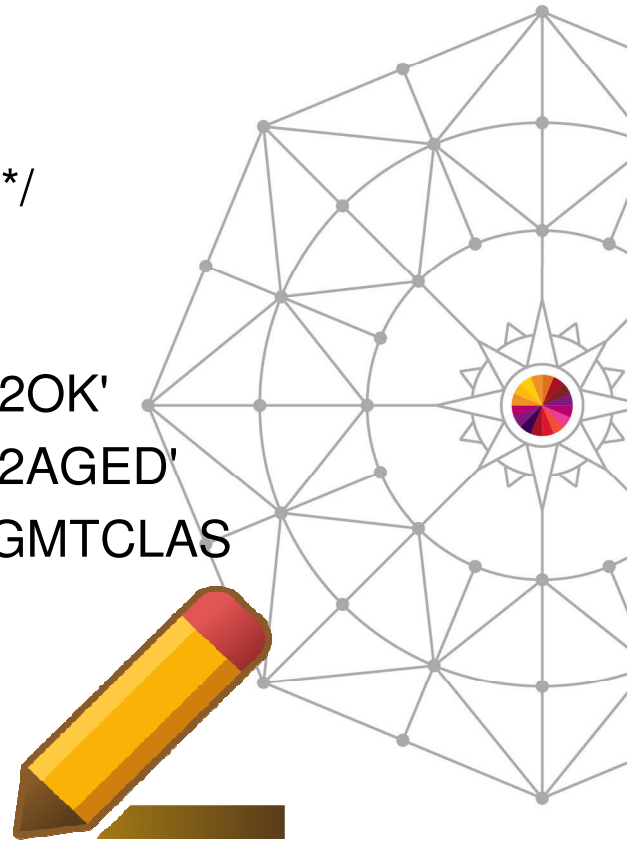

Storage Tiers
Invocation

- Storage Group
  - From 1 to 15 storage groups may be returned
    - It is the administrator's responsibility to ensure that a different storage group name provides a meaningful transition

```
IF &ACSENVIR = 'SPMGCLTR' THEN
   SELECT (&STORCLAS)
       WHEN ('SSD')          SET &STORGRP = 'ESYTIER0'
       WHEN ('EASYTIER')     SET &STORGRP = 'ESYTIER1'
       OTHERWISE             SET &STORGRP = &STORGRP
    END
ELSE ...
```

Storage Tiers
Invocation

- When DFSMShsm determines that a data set should be moved for a Class Transition, DFSMSdss is invoked to perform a Logical COPY with DELETE
  - DFSMSdss is the full data mover
    - Unlike migrate/recall and backup/recover where DSS is only the half data mover
  - DFSMSdss handles Copy Technique and Exit processing
    - After the movement, the data set retains all existing attributes and can be immediately accessed
- The catalog is updated as a part of the movement
  - *No* new DFSMShsm control data set records created for transitions
    - New FSR record type created for reporting purposes

      FSRTYPE = 24

      New DFSMS Report Generator sample report *provided*
    - DSR and VSR records are updated for DFSMShsm REPORT

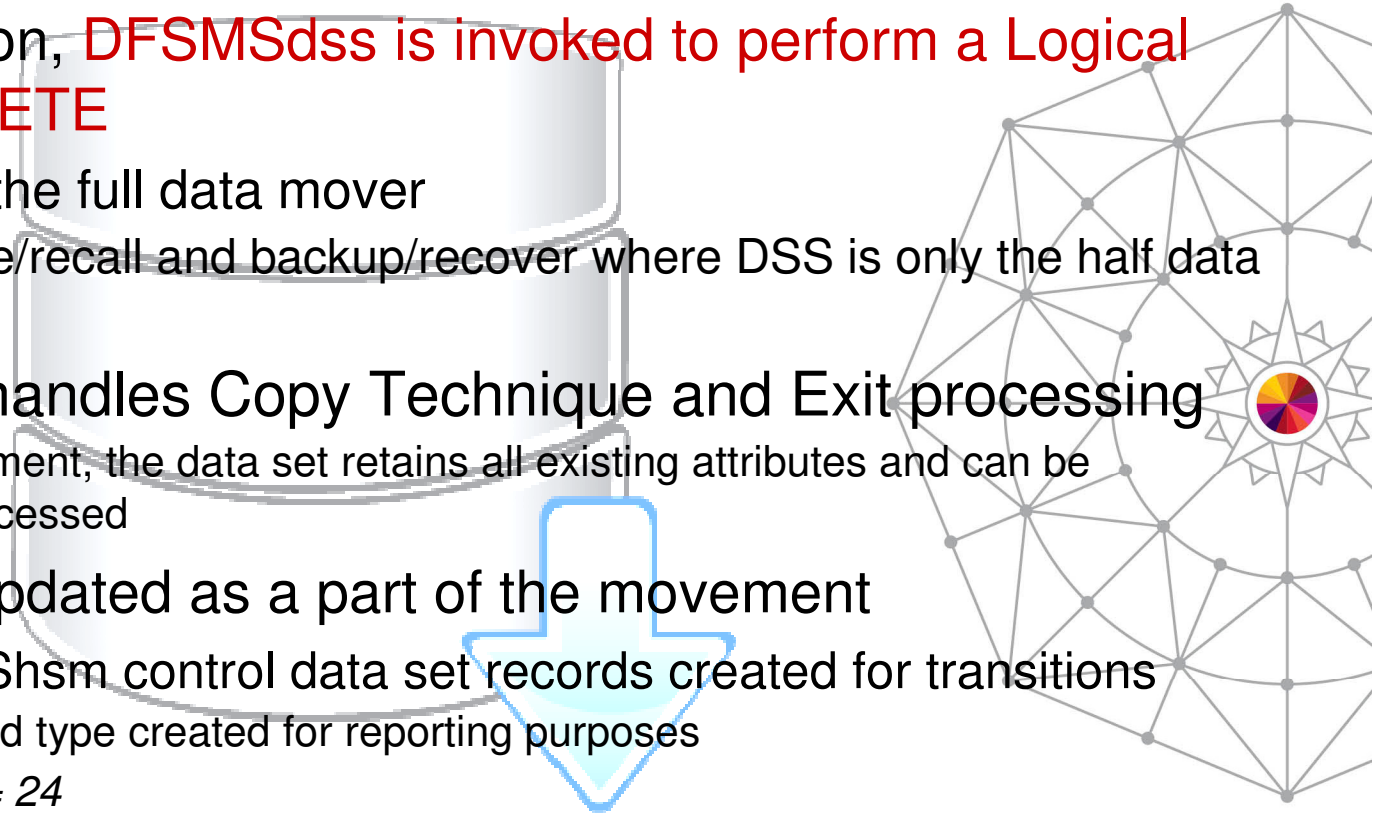

Storage Tiers
Invocation

- REPORT FUNCTION(TRANSITION) SYSOUT(A)
  - 'TRANSITION' is a new option

```
--DFSMSHSM STATISTICS REPORT ------- AT 08:05:07 ON 2011/09/13 FOR SYSTEM=2094
DAILY STATISTICS REPORT FOR 11/09/13
STARTUPS=000, SHUTDOWNS=000, ABENDS=000, WORK ELEMENTS PROCESSED=000005, BKUP VO
DATA SET MIGRATIONS BY VOLUME REQUEST= 0000000, DATA SET MIGRATIONS BY DATA SET
EXTENT REDUCTIONS= 0000000 RECALL MOUNTS AVOIDED= 00000 RECOVER MOUNTS AVOIDED=
DATA SET CLASS TRANSITION = 00000085 REQUESTED, 00000002 FAILED
                     NUMBER  ------READ--------  -----WRITTEN------  ------REQUES
   HSM FUNCTION     DATASETS TRK/BLK    BYTES    TRK/BLK    BYTES    SYSTEM USER
CLASS TRANSITION
 PRIMARY - PRIMARY 0000083  00053161 000632776K 00053161 000632776K 000085 00000
```

- New Report Generator Sample Report
  - ARCGS011 Statistics for Class Transitions

Storage Tiers
Invocation

- A new field was created in the NVR/VVR called 'Last Successful Class Transition Date' (LSCTD)
  - When a data set is successfully transitioned, the LSCTD is...
    - set to zero when the management class was changed
      -or-
    - set to the current date if the management class was not changed
  - LSCTD is used by DFSMShsm to know when a data set has already been successfully transitioned
  - DFSMShsm will attempt to transition a data set if it has met the transition criteria *AND* the LSCTD is zero
  - Exception for PERIODIC transitions. Data set will transition if the last transition was before the specified period

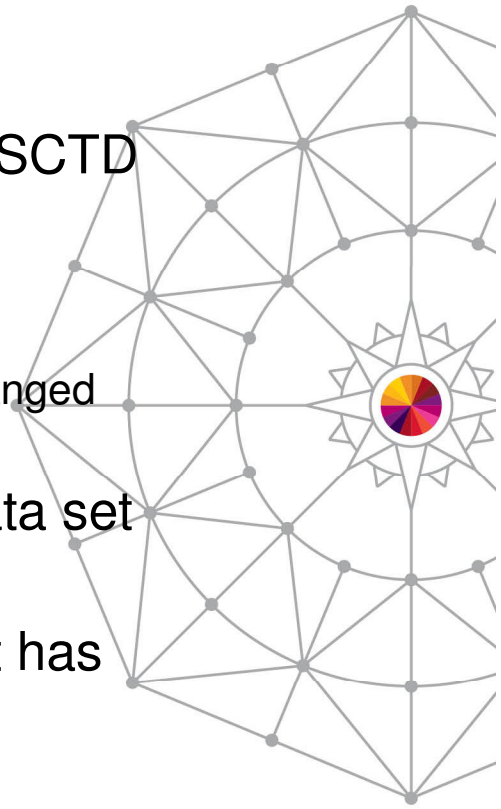

Storage Tiers
Invocation

- 'Last Successful Class Transition Date' (LSCTD) shown in LISTCAT output: *(when field is nonzero)*

```
NONVSAM ------- STORTIER.M01.S01.D01.N01.PSFB
    IN-CAT --- STRTRFVT.USERCAT
    HISTORY
      DATASET-OWNER-----(NULL)     CREATION--------2012.001
      RELEASE----------------2     EXPIRATION------0000.000
      ACCOUNT-INFO-----------------------------------(NULL)
    SMSDATA
      STORAGECLASS ---SCLASS22     MANAGEMENTCLASS-MCLASS01
      DATACLASS ------DCLASS01     LBACKUP ---2012.001.0701
      LAST TRANSITION-2012.013
    VOLUMES
      VOLSER-----------LSMS12     DEVTYPE------X'3010200F'
    ASSOCIATIONS--------(NULL)
    ATTRIBUTES
```

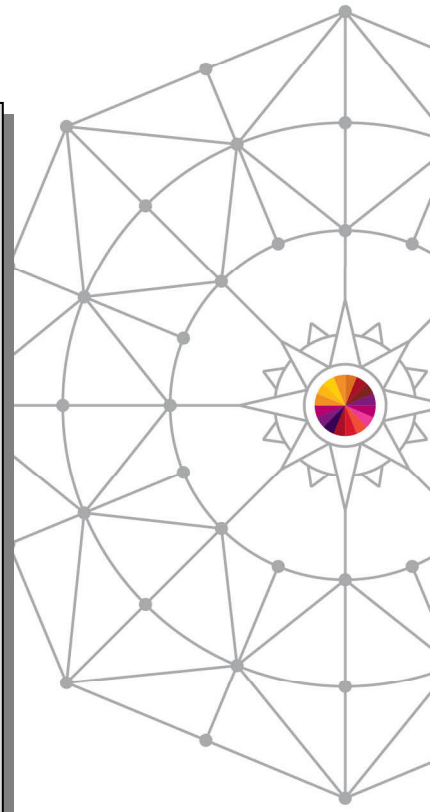

- Field is also available via DCOLLECT

Storage Tiers
Invocation

- Since Class Transitions are a part of the existing space management function...
  - The same tasking level that controls data set migrations controls data set transitions
    - If you expect more work during an existing space management window, then you may consider increasing your tasking level
  - The same HOLD, RELEASE, CANCEL, etc commands for space management control Class Transition processing also
  - ARC0734I (standard space management message) issued for class transitions
    - Insert updated to show 'ACTION=CLASS-TR'
    - Indicates From volume, To volume, RC, RSC, etc

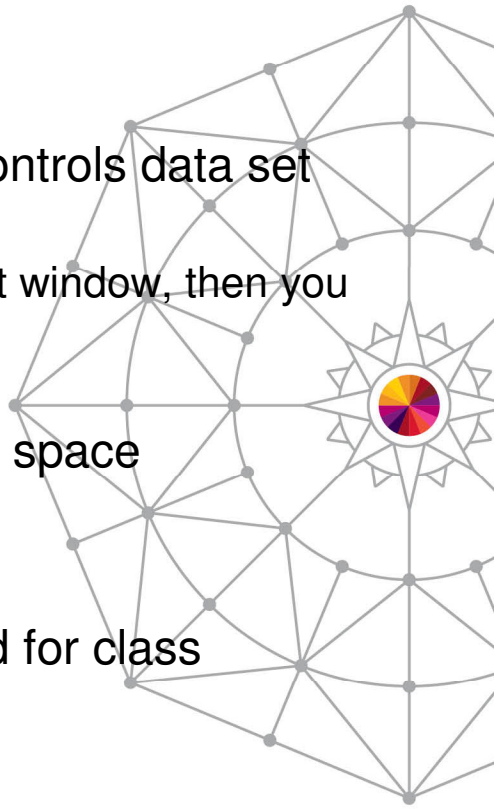

Storage Tiers
Invocation

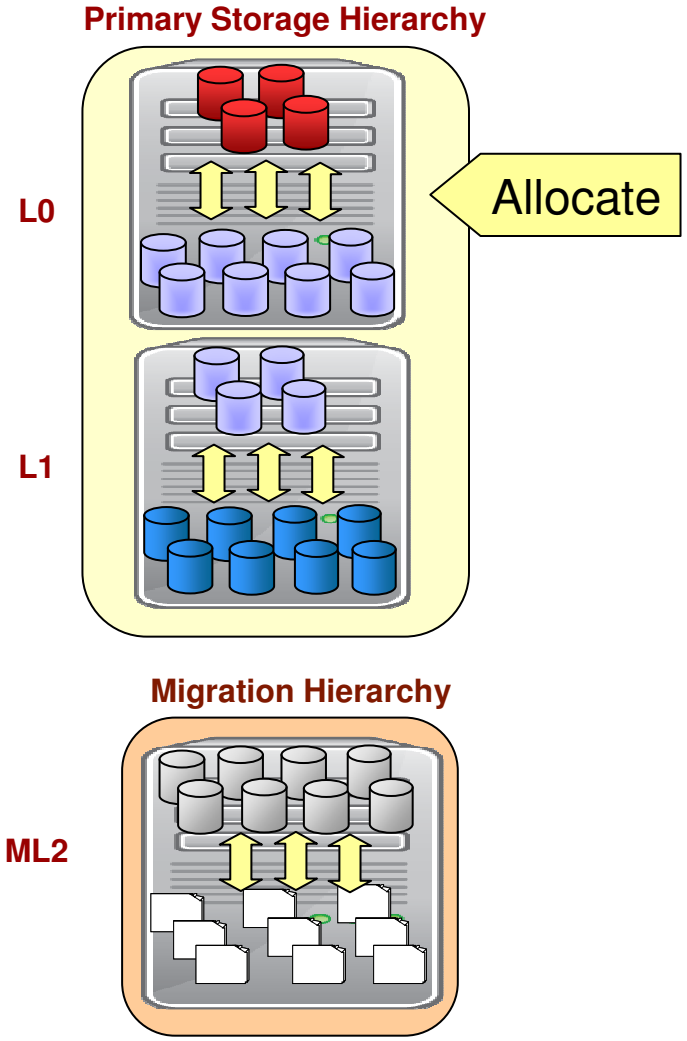

State 1 – Allocation

Expected access

Random reads and writes
with periods of high activity

Storage Group (SGLEV0)

Storage Tiered devices
High – SSD
Low - Enterprise

Management Class (ONLNSTMT)

No Migration
Data will be active
Even if not, no migration
Two Backup copies
Data is actively changing
Transition after 45 days
No Expiration

Storage Tiers Invocation

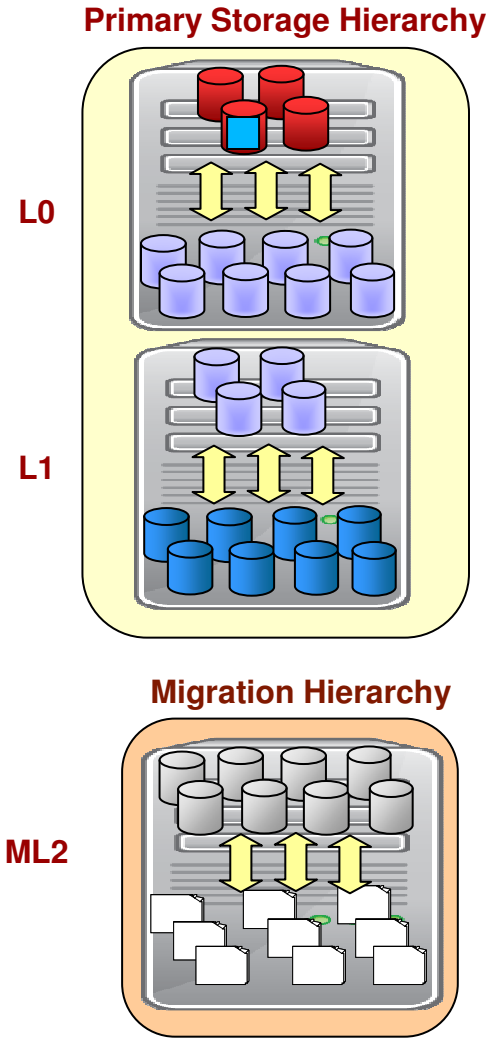

State 2 – First 45 Days

Storage Controller moves the data between SSD and Enterprise based on the data's heat map

DFSMShsm manages based on ONLNSTMT
Two Backup Copies

Storage Tiers Invocation

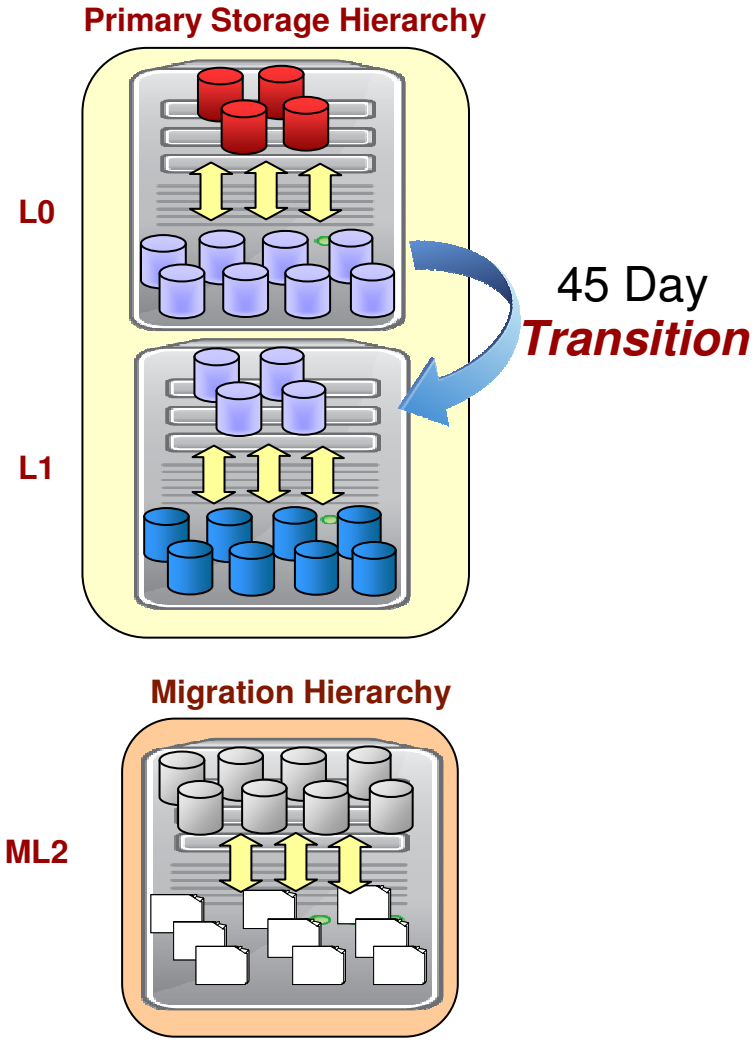

State 3 – 45 Day Transition

Expected access

After 45 days, the data is expected to be Historical, instead of Active.
No expected writes and mostly *sequential reads*.

Storage Group (SGLEV1)

Storage Tiered devices
- High – Enterprise
- Low - Nearline

Management Class (HISTSTMT)

Migration
- Migrate after 21 days of inactivity

Single Backup copy
- *Data is not changing*

No further transitions
Expire after 7 years

Storage Tiers Invocation

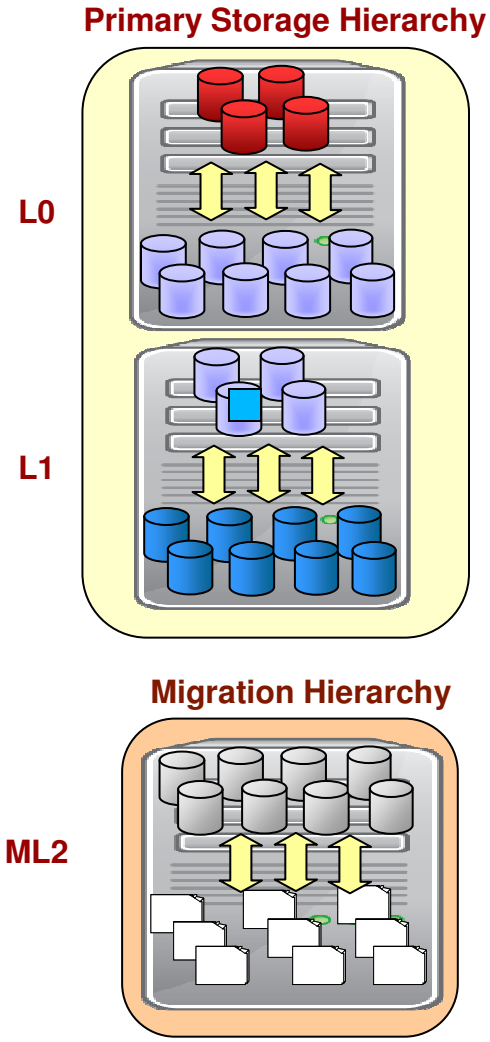

State 4 – The Golden Years

Storage Controller moves the data between Enterprise and Nearline based on the data's heat map

DFSMShsm manages based on HISTSTMT

- Single Backup Copy
- Migrate at 21 days inactive
  - *Storage Controller compresses the data and moves it from disk to tape*
- Expire after 7 years

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
- Storage Tiering
- Fast Replication
  - Consistency Groups
  - Recover Data Set Enhancements
  - Performance
  - Usability
  - UCB Refresh
- Miscellaneous

Fast Replication
Consistency Groups

- **Standard FlashCopy – 'Fuzzy' Backup**
- **DB2 ensures data consistency during recovery**

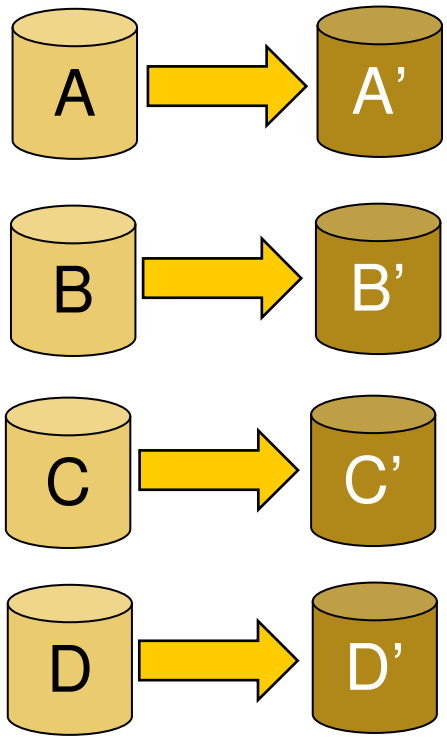

- **FC SrcA to TrgA'**
  - Writes held to SrcA during establish
  - Write activity on B, C, D
- **FC SrcB to TrgB'**
  - Writes held to SrcB during establish
  - Write activity on A, C, D
- **FC SrcC to TrgC'**
  - Writes held to SrcC during establish
  - Write activity on A, B, D
- **FC SrcD to TrgD'**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

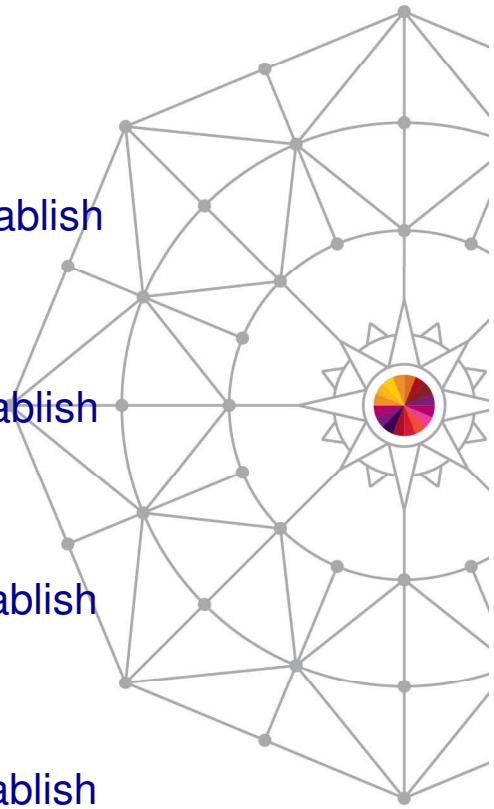

Fast Replication
Consistency Groups

- **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**

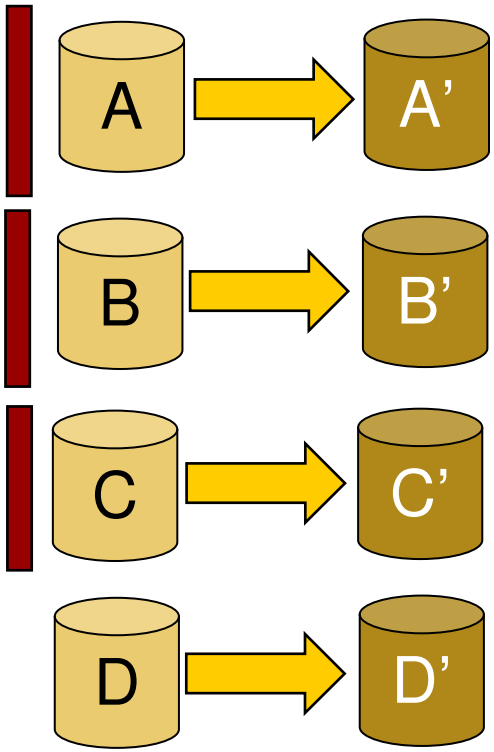

- **FC SrcA to TrgA' / CGFREEZE**
  - Writes held to SrcA by ELB
  - Non Dependent write activity on B, C, D
- **FC SrcB to TrgB'**
  - Writes held to SrcA, B by ELB
  - Non Dependent write activity on C, D
- **FC SrcC to TrgC'**
  - Writes held to SrcA, B, C by ELB
  - Non Dependent write activity on 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

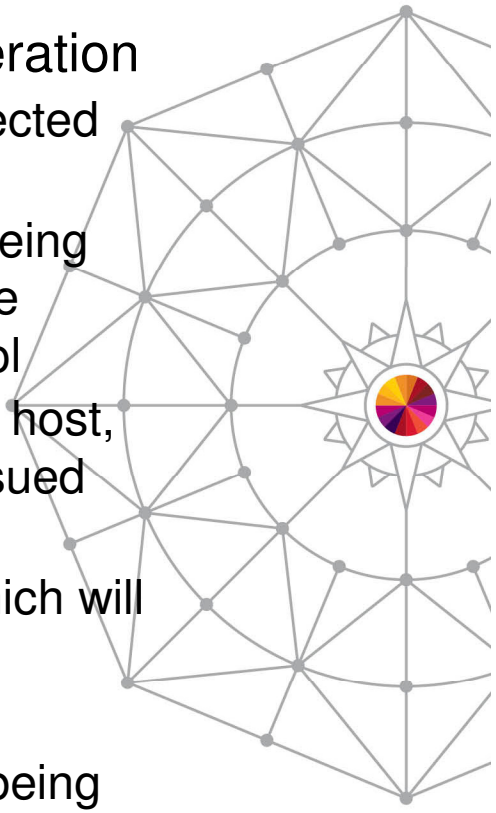

Fast Replication
Consistency Groups

```
Panel Utilities Scroll Help
------------------------------------------------------------------------------ DGTDCPP1 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.
```

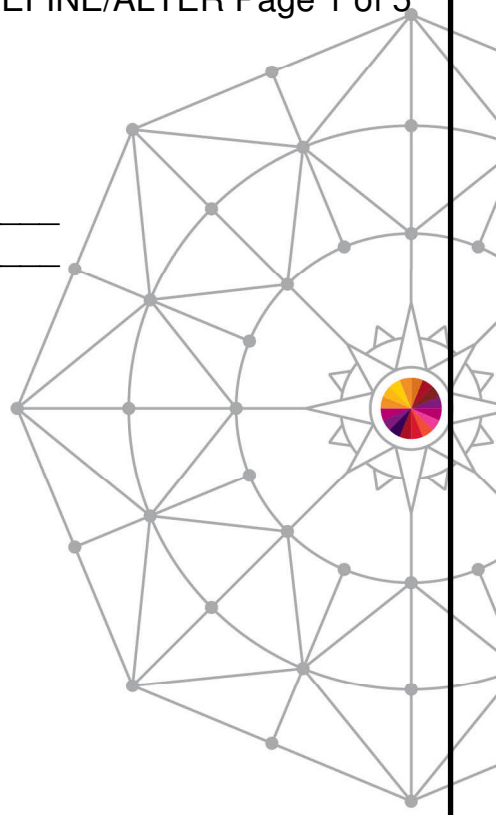

- **LIST COPYPOOL will have an indicator of the setting**

Fast Replication
Recover data set to any volume

- **Today**
  - FRRECOV DSNAME 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

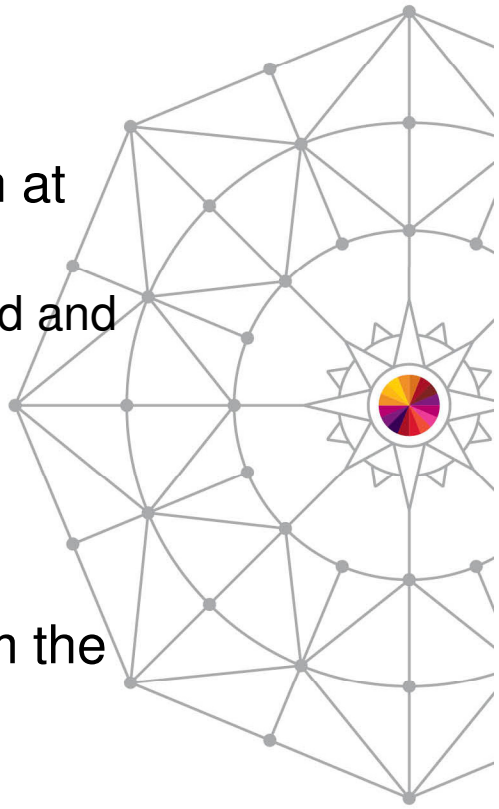

Fast Replication
NEWNAME

- **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

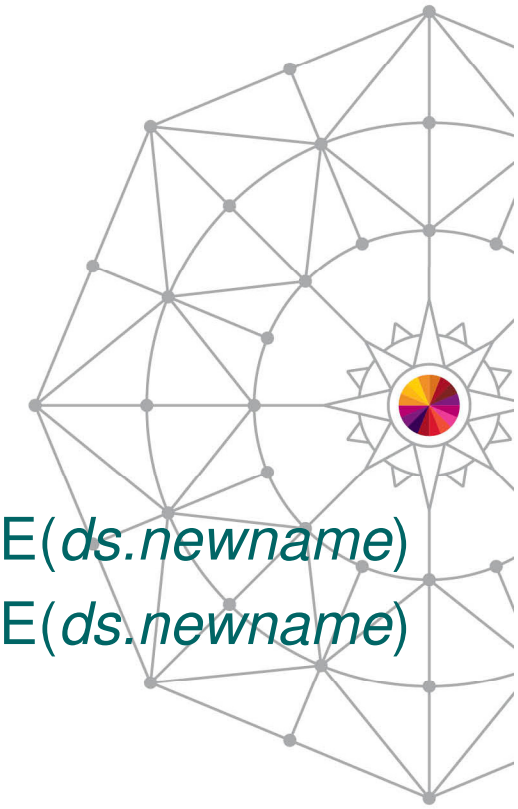

Fast Replication
Data Set Recovery Performance

- **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

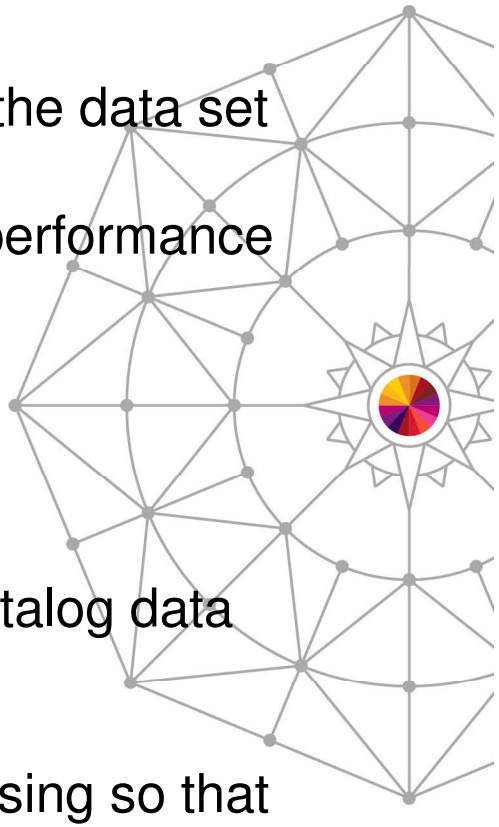

Fast Replication
D-Type APARS

OA39738

- For NOCOPY, volumes are always initialized after WITHDRAW, regardless of what the status of the volume it
- 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

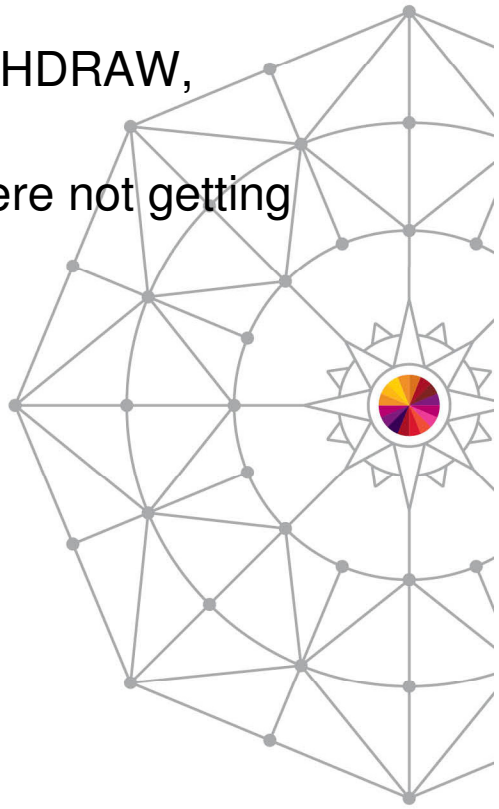

Agenda

- **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**

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

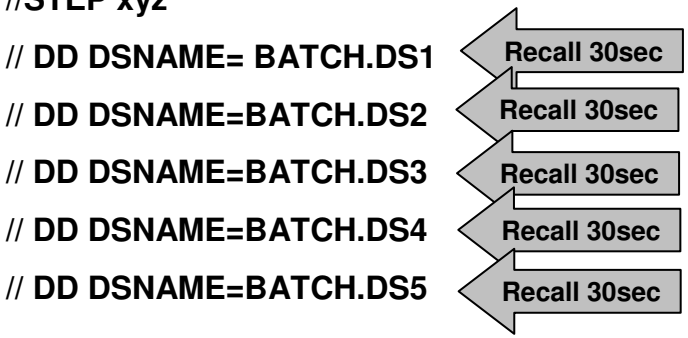

Elapsed Recall Time = 150 seconds

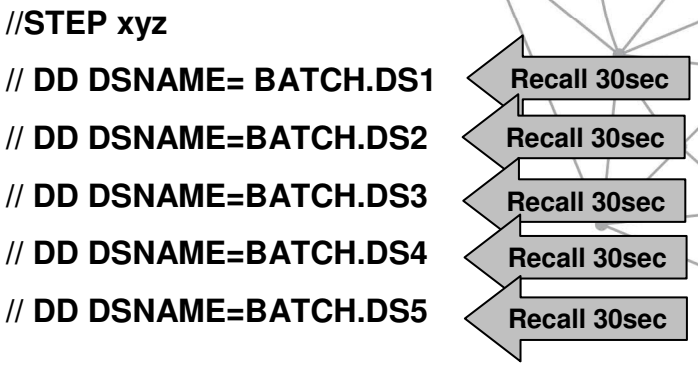

Elapsed Recall Time = 30 seconds

#SHAREorg

SHARE

What's New in DFSMShsm

Glenn Wilcock
IBM

March 11, 2014
Session Number 15076

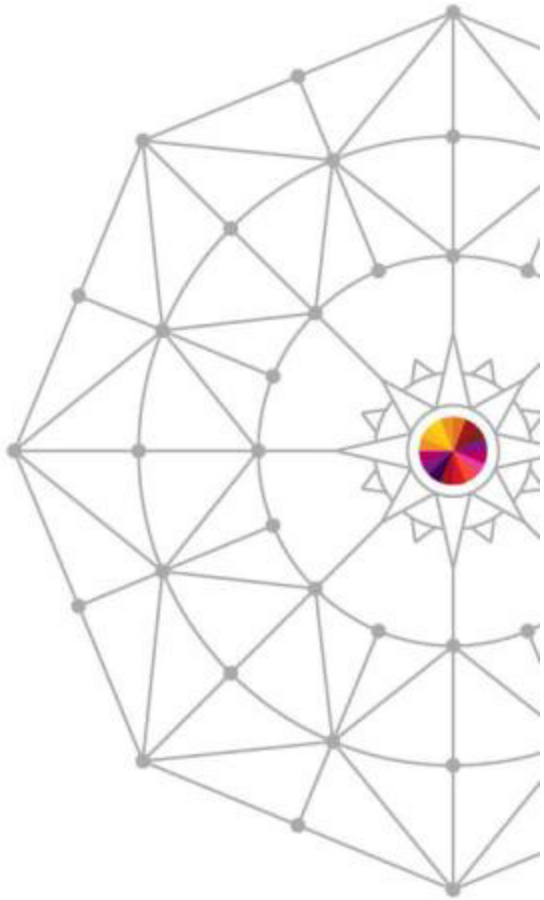

SHARE in Anaheim

Legal Disclaimer

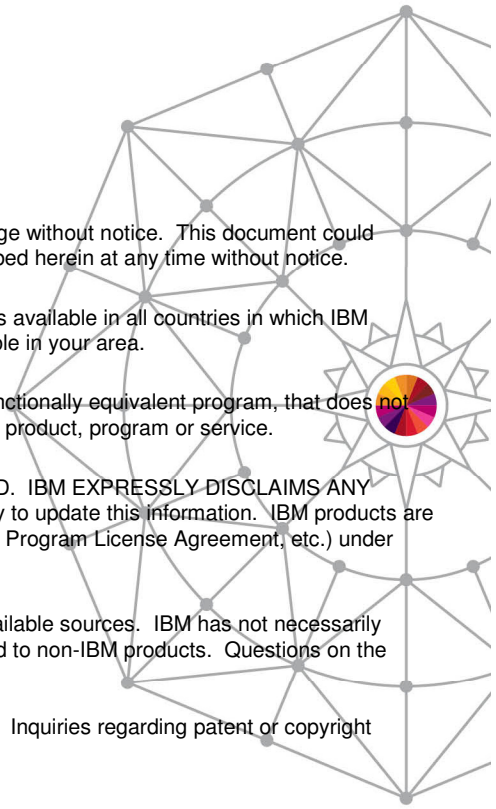

NOTICES AND DISCLAIMERS

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.

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.