



MVSS Project Opening and Keynote:

A New Frontier in the Evolution of Space Management...

Buckle In!

Russell Witt – CA
Glenn Wilcock - IBM

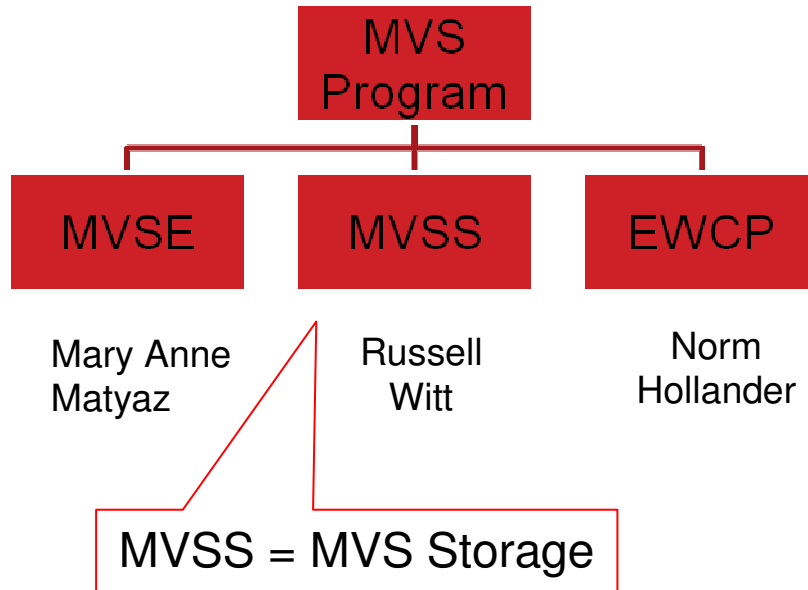
August 12, 2013
Session 14101



Our Purpose

The MVSS project focuses on the management of storage resources throughout the enterprise and the facilities necessary to provide these functions. In the z/OS environments, DFSMS and all of the DFSMS components and other vendor-developed (OEM) packages for providing these services are discussed. Among the many subjects covered are storage hardware supported by z/OS, access methods (including ICF/VSAM), ICF catalog topics and installation-developed tools and techniques. In the open systems arena we discuss storage-related aspects of USS, Tivoli Storage Manager on the z/OS platforms, Storage Area Networks, storage interfaces, and other open systems storage hardware and software issues. Topics that are addressed for all platforms are disaster recovery, data storage devices, data backup and archival systems and data integrity and recovery. The primary activities of the project are: quality IBM and user experience presentations at major SHAREs, requirements generation and prioritization and ad hoc task forces that influence IBM regarding future storage management directions by providing customer input.

MVS Program Organization & Leadership



MVS Program Leadership		
Ed Jaffe	Program-Manager	edjaffe@phoenixsoftware.com
Skip Robinson	Program Manager	JO.Skip.Robinson@sce.com
Dave Danner	Program Officer	ddanner9@yahoo.com
Cheryl Watson	Requirements Coordinator	
Bob Shannon	Program Officer	Bob.Shannon@rocketsoftware.com

Project Officers – MVSS



Name	Title	Email Address
Russell Witt – CA Technologies	Project Manager	Russell.witt@ca.com
David Astemborski – First Bank	Deputy Project Manager / Scheduling Coordinator	David.Astemborski@efirstbank.com
Jim Erdahl – U.S.Bank	Deputy Project Manager / Requirements Coordinator	James.Erdahl@usbank.com
William R. Smith – Hitachi Data Systems	Deputy Project Manager / Volunteer Coordinator	William.Smith@hds.com
Harold Durnford – RBC Financial Services	Deputy Project Manager / Newsletter	Harold.Dumford@rbc.com
Michael Friske – Fidelity	Project Officer	Michael.friske@fmr.com
Ed Petka – Fiserv Corp	Project Officer	Edward.petka@fiserv.com



IBM Representatives – MVSS



Name	Title	Email Address
Barbara McDonald	Lead Representative	bawhite@us.ibm.com
Stephen Branch	DFSMSdfp	branchs@us.ibm.com
Glenn Wilcock	DFSMSShsm	wilcock@us.ibm.com
Jeff Suarez	DFSMS SDM/DSS	jrsuarez@us.ibm.com
Horst Sinram	DFSMSrmm	sinram@de.ibm.com
DFSMS Service	Steve Huber or Neal Bohling	shuber@us.ibm.com bohling@us.ibm.com
DFSMS System Test	Arthur Bariska	bariska@us.ibm.com
Hardware Rep	Curtis Neal	curtisne@us.ibm.com



Join our Tribe

- SHARE is a Volunteer organization
- Volunteer as a CHAIR –
 - You get the best seat in the house
- Submit a Requirement (your chance to get IBM to do something your way)
 - Better then asking as an individual
 - Chance to work with IBM and other clients to better articulate your request
- Submit a topic for a presentation
 - Call for Presentations for Boston will open soon.

Items Of Interest:

- 11am Sessions duration may extend to 12:15pm
- MVSS: Free-For-All
 - Session 13781, Wednesday 6pm, Room 304
- Project Dinner @ <td>
 - After Free-For-All
 - Head Count
 - Location to meet <td>

A New Frontier in the Evolution of Space Management... Buckle In!



Agenda

- History 101
- Information Life Management 101
- Space Management 101
- Evolution of Space Management
 - General
 - Tape
 - Disk

History 101



United States Declaration of Independence, 1776 *Thomas Jefferson*

*“We hold these truths to be self-evident,
that all men are created equal,
that they are endowed by their Creator
with certain unalienable Rights, that among these are
Life, Liberty, and the Pursuit of Happiness.”*

Information Lifecycle Management 101



Storage Administrator's Declaration of Dependence, 1976*

*“We hold this truth to be self-evident,
that all data are not created equal!*

*That regardless of the will of the data's Creator,
the Storage Administrator shall determine all data's
Place of Residence, Availability, Security,
and End of Life, heretofore and forever after.”*

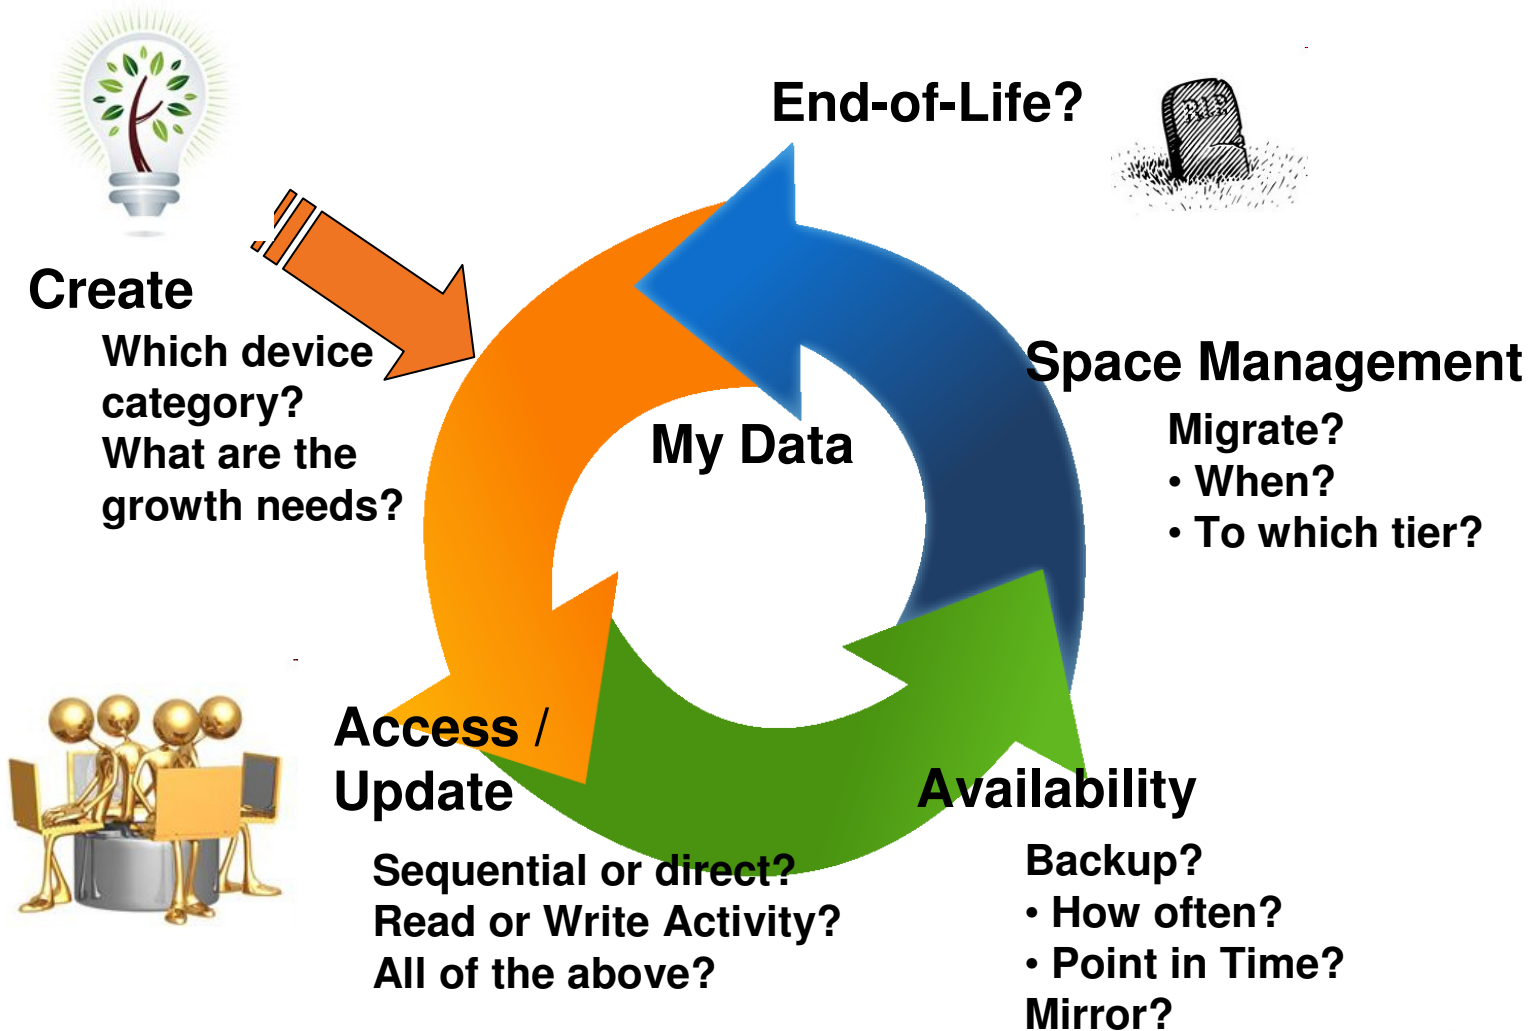
* Random date for affect

Information Lifecycle Management 101

- Data has various **Business Needs**
 - Business Critical
 - Performance Intensive
 - Time Sensitive
 - Production Data
 - Development Data
 - Test Data
 - Historical Data
- Various **Storage and Media Types**
 - Faster
 - Cheaper
 - Portable
 - Advanced Function



Information Lifecycle Management 101



Information Lifecycle Management 101



- **Information Lifecycle Management**

The process of managing information – from creation, through its useful life, to its eventual deletion – in a manner that aligns storage costs with a changing business value of information.

- DFSMS provides policy-based...

- Data Creation
- Backup / Recovery Management
- Expiration
- ★ Space Management



- *Today, we'll focus on the evolution of Space Management*

Brief History of DFSMSHsm

- **March 31, 2013 – DFSMSHsm’s 35th Birthday!**
 - *No, there will not be an anniversary reduction to 1978 pricing!*
 - *No, I do not plan on ever attending an HSM retirement party!*
- **Developed at IBM Research Center, beginning in 1975**
 - Named MASH – ‘MVS Automatic Storage Hierarchy’
 - MASH support (Multiple Address Space HSM) is a play on this acronym
- **Released in March 31, 1978 as ‘HSM’ (Hierarchical Storage Manger)**
 - Originally, product name was going to be ‘GPA’, for General Purpose Archiver
- **Renamed to DFHSM in 1980’s**
 - Data Facility Hierarchical Storage Manager
 - Some still refer to it as ‘DFHSM’
- **Became priced feature of DFSMS/MVS, DFSMSHsm, in early 90s**
 - Data Facility System Managed Storage hierarchical storage manager

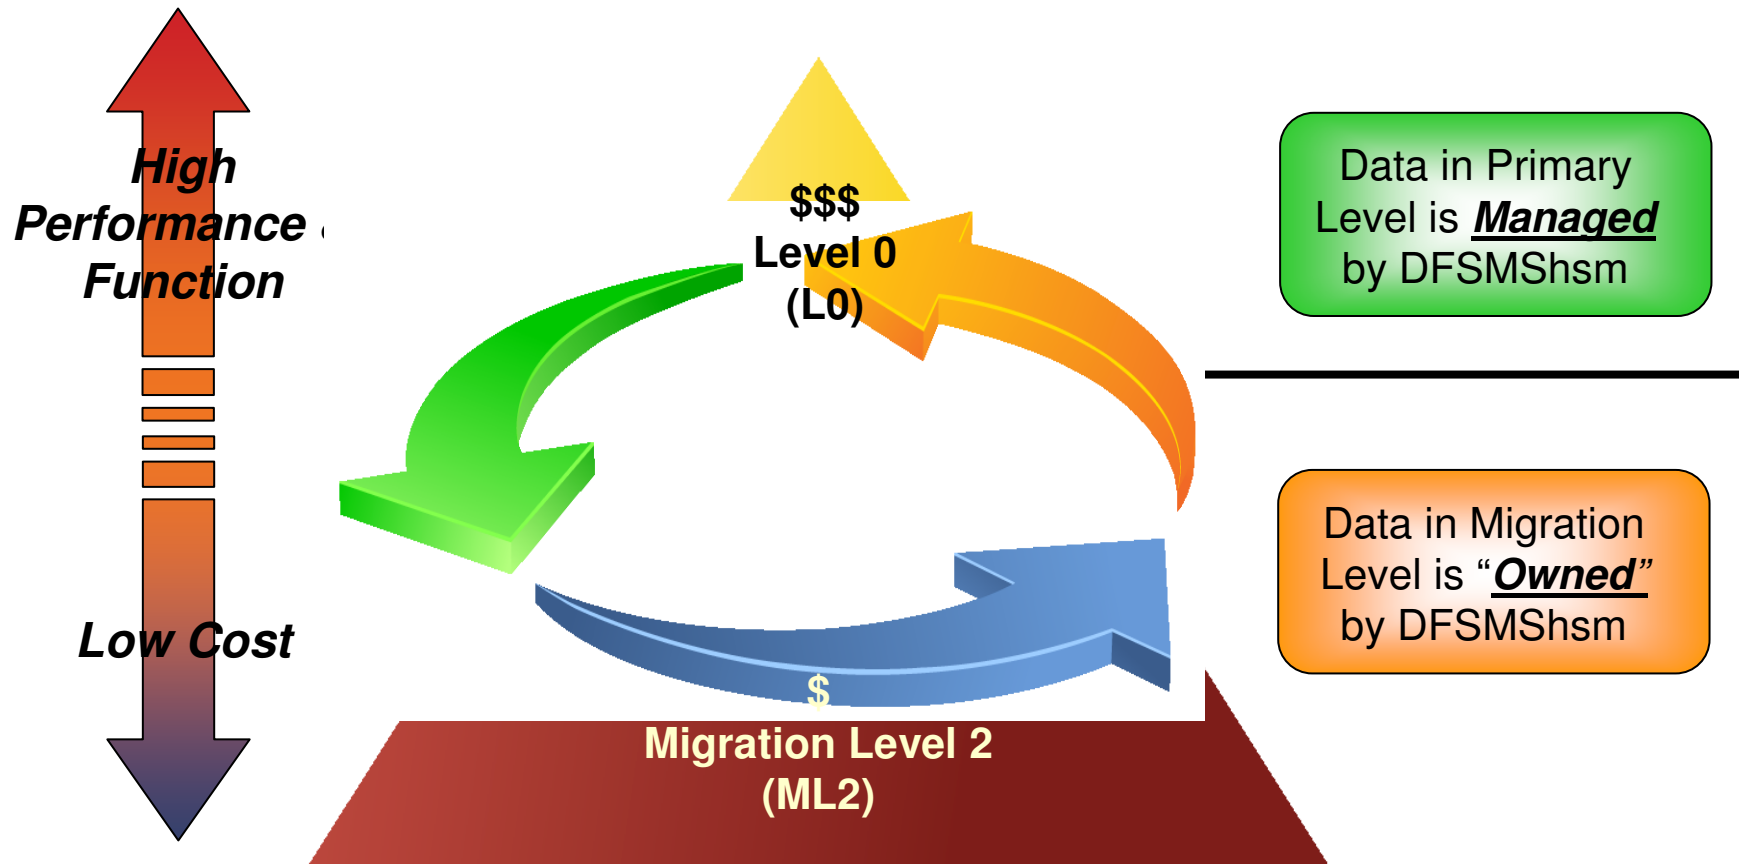


Space Management 101

- **Space Management** provides *automated, policy-based management* of your data in order to...
 - ★ Minimize the cost of storing data
 - ★ Ensure ample free space for data creation
- The major functions are...
 - Expiration
 - Migration / Recall
 - Extent Reduction

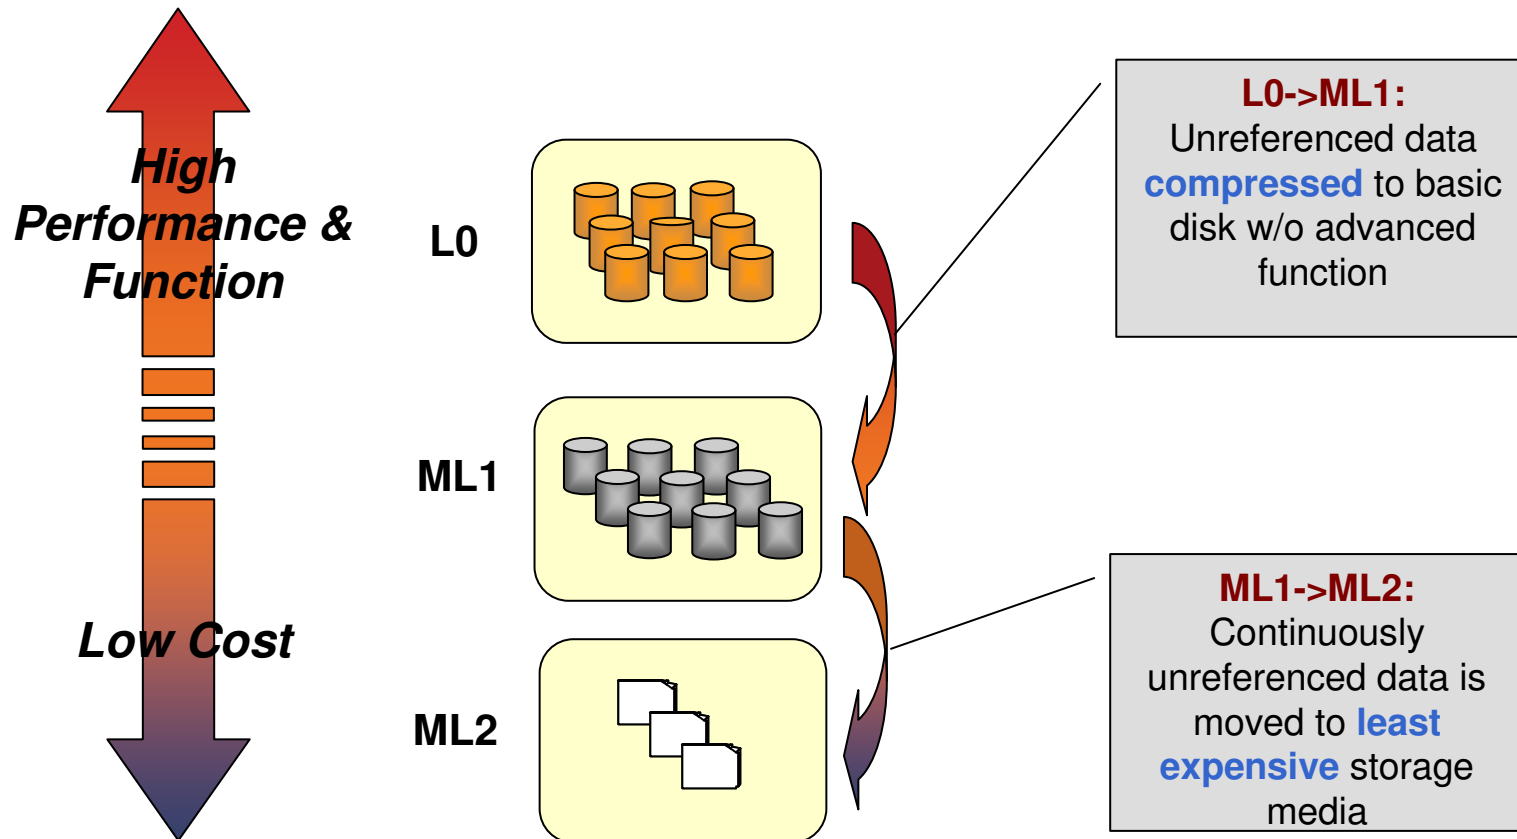
Space Management 101

- The *Classic* DFSMS Storage Hierarchy (35 years old)



Space Management 101

- The *Classic* DFSMS Storage Hierarchy (35 years old)



Space Management 101

- **Basic Advantages of Tape**
 - Media longevity
 - Tape 15-30 years, Disk 4-5 years
 - Portability
 - Lower Energy consumption
 - Total Cost of Ownership
- **Basic Advantages of Disk**
 - Advanced Functionality
 - Concurrent Access to data on same media
 - Immediate or Near-immediate access to data
 - i.e. – No dreaded tape mount

Space Management Evolution



- **Space Management Balancing Act**

- True, Space Management reduces the cost of storing data, but it comes at a price:
 - CPU overhead
 - *Compression*
 - *Run time*
 - Channel overhead of moving data
 - Delay in access to migrated data
 - Software expense



The Space Management environment continually evolves as we try to find the right Return on Investment and integrate new storage technologies

Space Management Evolution

General 201

- **Leave data on Level 0 longer**
 - The most basic change is to compare the cost of maintaining data on Level 0 to the cost of actually migrating the data
 - Small data sets
 - *Cost of L0 disk is getting so inexpensive, why not just keep these on L0?*
 - *Don't forget, SDSPs get up to 15:1 compaction!*
 - Keep data on L0 longer
 - *Extended duration of time that data must be inactive before it is eligible for migration*
 - *Eliminate thrashing*
 - *Evolve Migrate/Recall into more of an Archive function, where migrated data is not expected to be recalled*



Space Management Evolution

Tape 201


- **Virtual Tape**

- One of the fundamental weaknesses of tape has always been the delay in access to the data due to the tape mount
- Virtual Tape addresses this by integrating a disk cache that emulates tape
 - Academically speaking, tape mount delays could be eliminated if referenced data could always be staged into the disk cache at the appropriate time

Stand-alone Tape



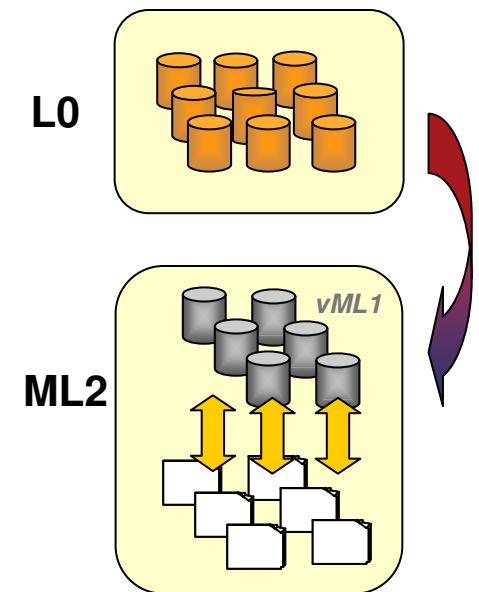
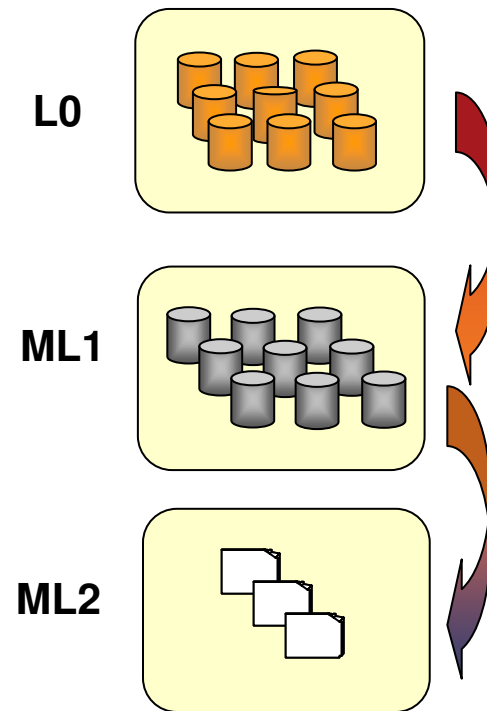
Virtual Tape



Space Management Evolution Tape 201

- **Significant Advantages of Virtual Tape**

- ★ Disk cache essentially replaces ML1 tier
- ★ Storage controller performs compaction instead of host!
- ★ Eliminates double data movement from L0->ML1 and then ML1->ML2!
- ★ Many logical drives
- Some ML1 is still needed for data that appropriately only goes to ML1 tier



Space Management Evolution

Tape 201

- **Disadvantages of Virtual Tape**

- HSM management of logical tapes does not ‘play well’ with the storage controller management of physical tapes
 - HSM RECYCLE function processes based on the % full of the logical tape, which does not take into account the % full of the physical tape
- HSM access to any data on a logical tape that resides on a physical tape requires the entire logical tape to be staged back to the disk cache before the data can be accessed
- *Prior to z/OS V2R1*, HSM will only span a data set across 40 tape volumes
 - Problematic for migration/backup of large data sets onto small logical tapes
 - ★ **V2R1 extends the limit to 254 volumes**

Space Management Evolution

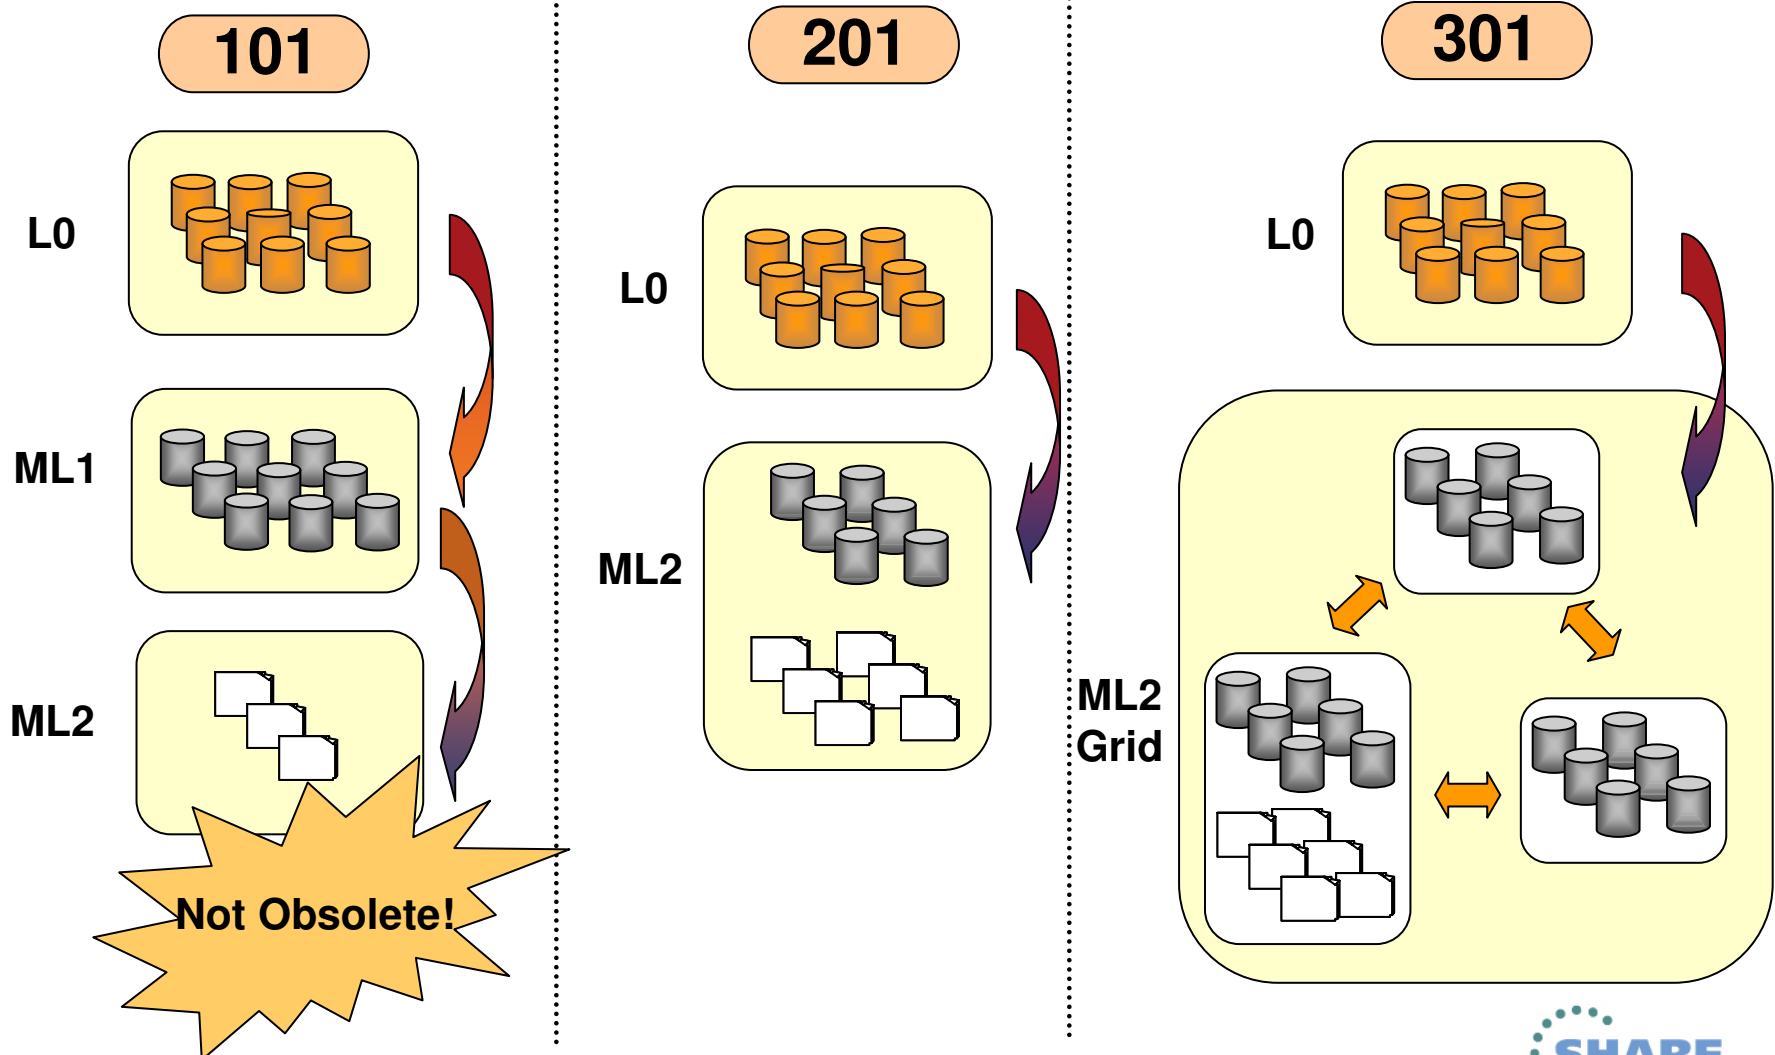
Tape 301



- **Disk Only Virtual Tape**

- *Just what the doctor ordered for virtual tape and HSM!*
- Disk only storage controller that emulates tape
 - ★ All of the advantages of Virtual Tape plus...
 - ★ Eliminates tape mounts / staging
 - ★ Eliminates RECYCLE issues
 - ★ Enables n copies of HSM tapes to be maintained
 - ★ *Enables reduction / elimination of HSM duplexing*
 - ! *Caution – Don't put all of your eggs in one basket. Hardware is not flawless. Make sure you have a backup copy of migrated data*
- ★ Some vendors enable real tape to be integrated into the grid.

Space Management Evolution Tape 301



Space Management Evolution

Tape 301

- **HSM Disk Only Virtual Tape Environment**

Actual numbers from Customer ABC who went from physical tape to DOVT

- ★ Reduced HSM's Total CPU Time by 19 hours, 45 minutes
 - ★ 13% Reduction in Total CPU Time used for HSM Migration to ML2, even though the number of Megabytes transferred increased by over 6TB or 680,000 data sets
 - ★ 64% Reduction in Average CPU Time used for HSM ML2 RECALLS
 - ★ 91% Reduction in Average ELAPSE Time for ML2 RECALL
 - ★ Reduction in duplicate recall requests (RC2 errors) from 2,431 to 264
- **Gotcha**
 - *Elapsed time for larger data sets can be degraded*

Space Management Evolution

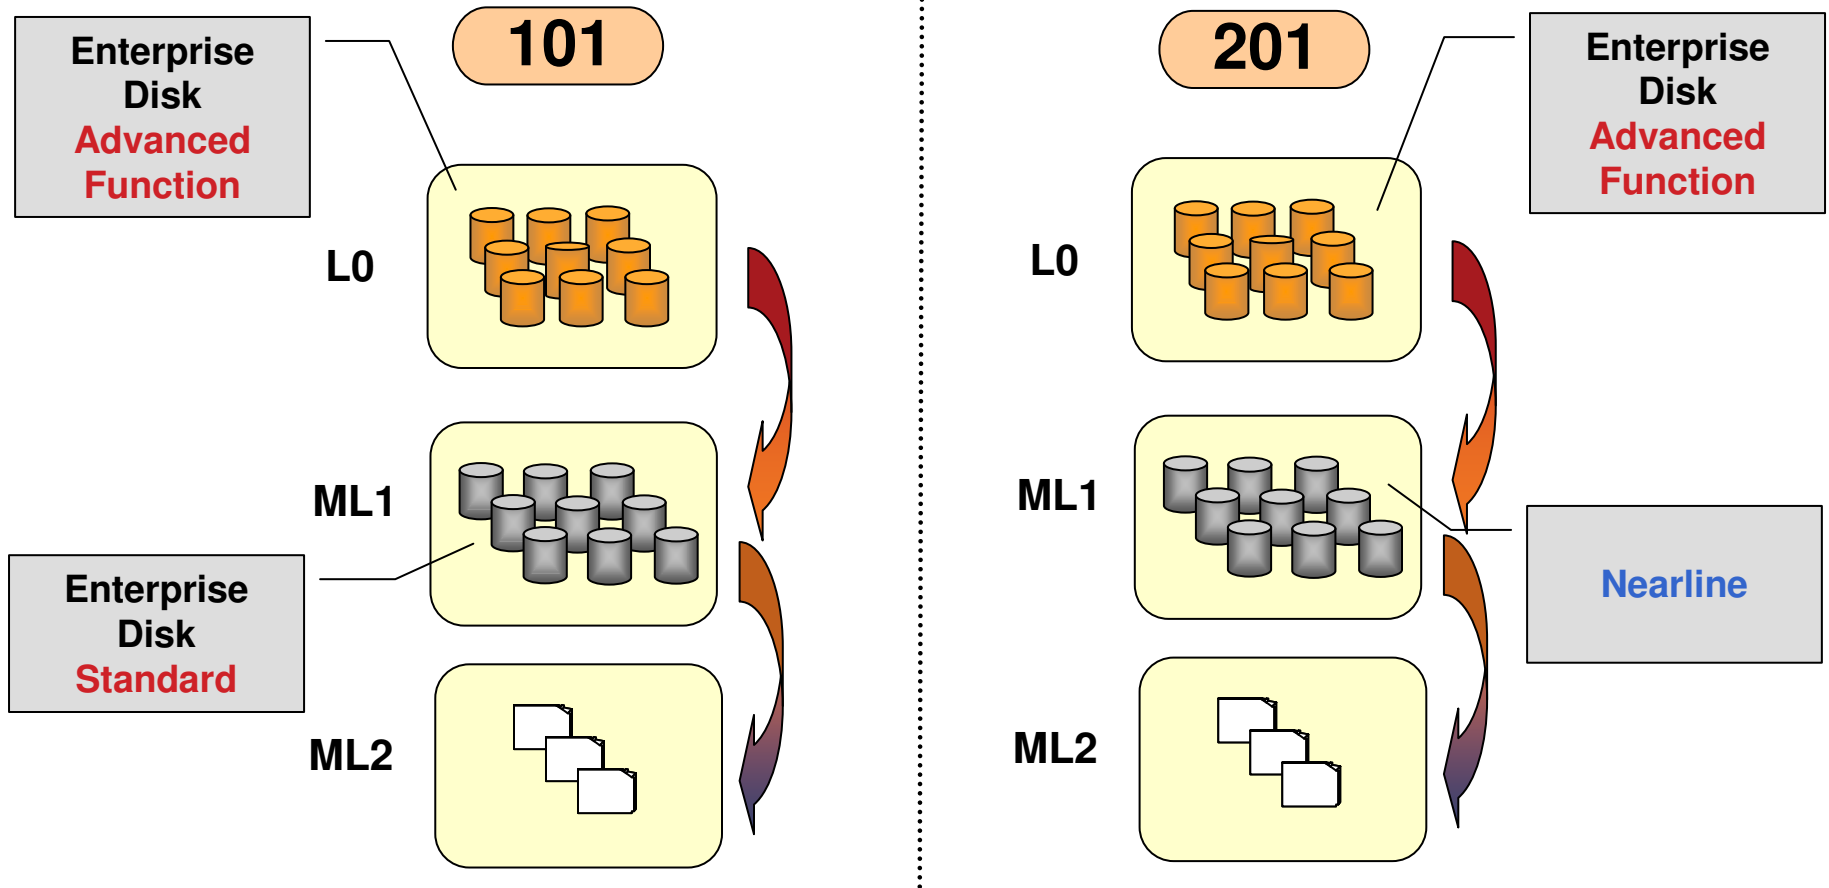
Tape 401

- **Evolving HSM Disk Only Virtual Tape Environment**
NOT a statement of direction
 - **Where do we go from here?**
 - Exploit the fact that the 'tape' data is really on disk
 - *Concurrent access to single tape*
 - *Software directed controller based data movement from volume to volume*
 - *Etc, etc...*



Space Management Evolution

Disk 201



Implement ML1 volumes using Nearline drives

Space Management Evolution

Disk 201

- **Controller based Storage Tiering**

- Disk controller *transparently* moves data from tier to tier based on the data's 'heat map'
 - 'Hot' data is moved to a higher tier
 - 'Cold' data is moved to a lower tier




- **Goal:** Significantly improve performance with minimal investment by dynamically placing 'hot' data in the highest tier

Space Management Evolution


Disk 201

Standard Logical Volume:



- All physical extents are of the same device type
- Data statically resides on physical extents

Tiered Logical Volume:



- Physical extents are different device types (tiers)
- Data dynamically moves between tiers (physical extents) based on the 'heat map'

Space Management Evolution

Disk 201

- *Quick overview of existing disk tiers...*
 - **Enterprise Disk:** Standard Fiber Channel spinning disk
 - Price point: **x**
 - **Solid State Devices (SSD):** Devices with significantly *faster direct read response times* and a significantly higher cost.
 - Price point: **5x**
 - **Nearline:** Devices that are *slower and less expensive*
 - Price point: **½x**




Space Management Evolution

Disk 201




- Storage controller uses heat map to transparently move data to the most appropriate tier...



Space Management Evolution Disk 201

- For example...

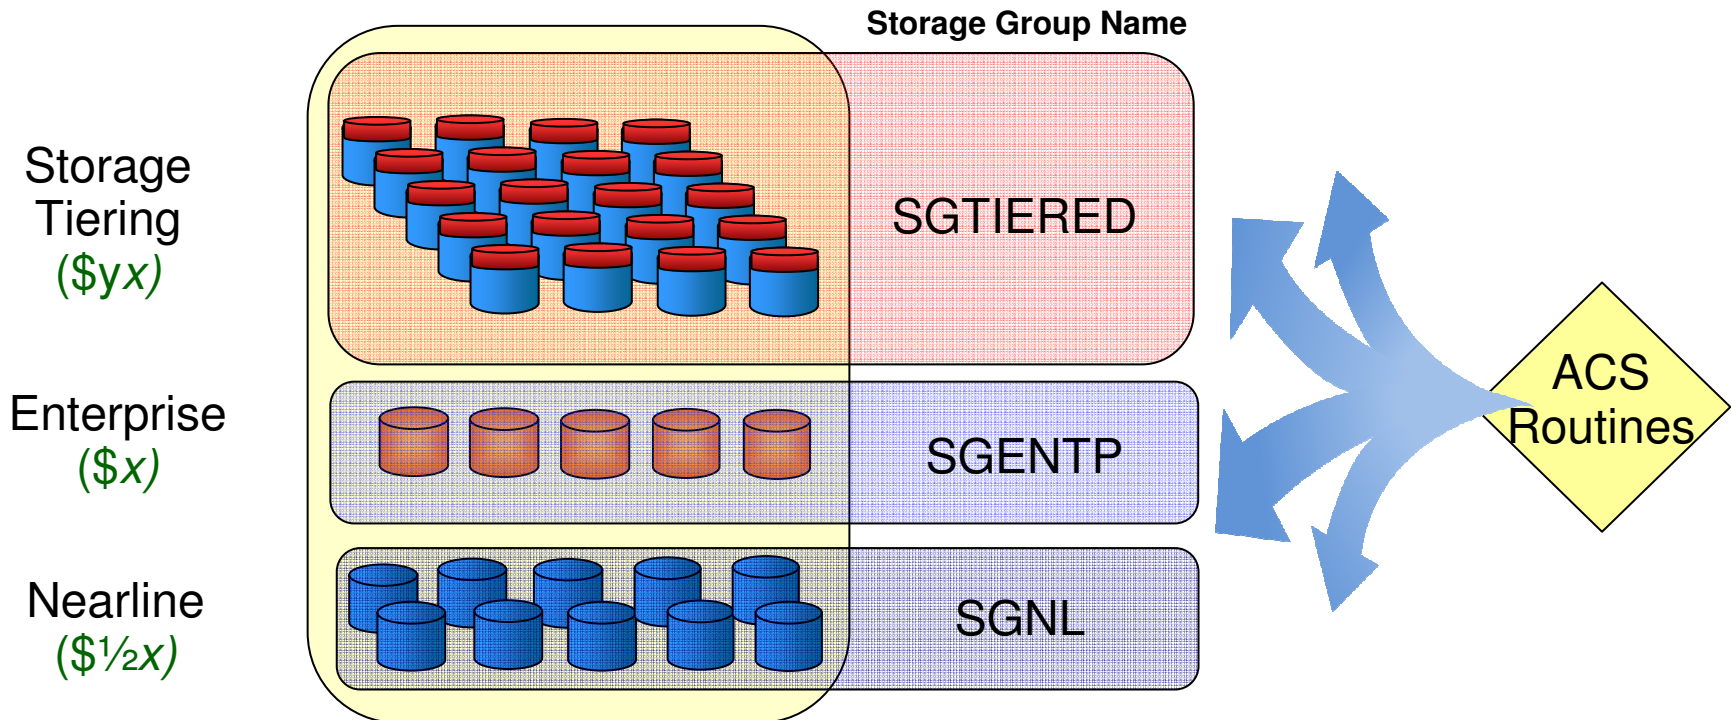


Storage performance council, IBM DS8700, 2010: http://www.storageperformance.org/results/benchmark_results_spc1#a00092

Space Management Evolution


Disk 201

- **How do SMS and Storage Tiering work together?**
 - Group your different classes of devices by Storage Group and use ACS Routines to direct data to the appropriate tier



Space Management Evolution

Disk 201



- **Not everything is rosy...**

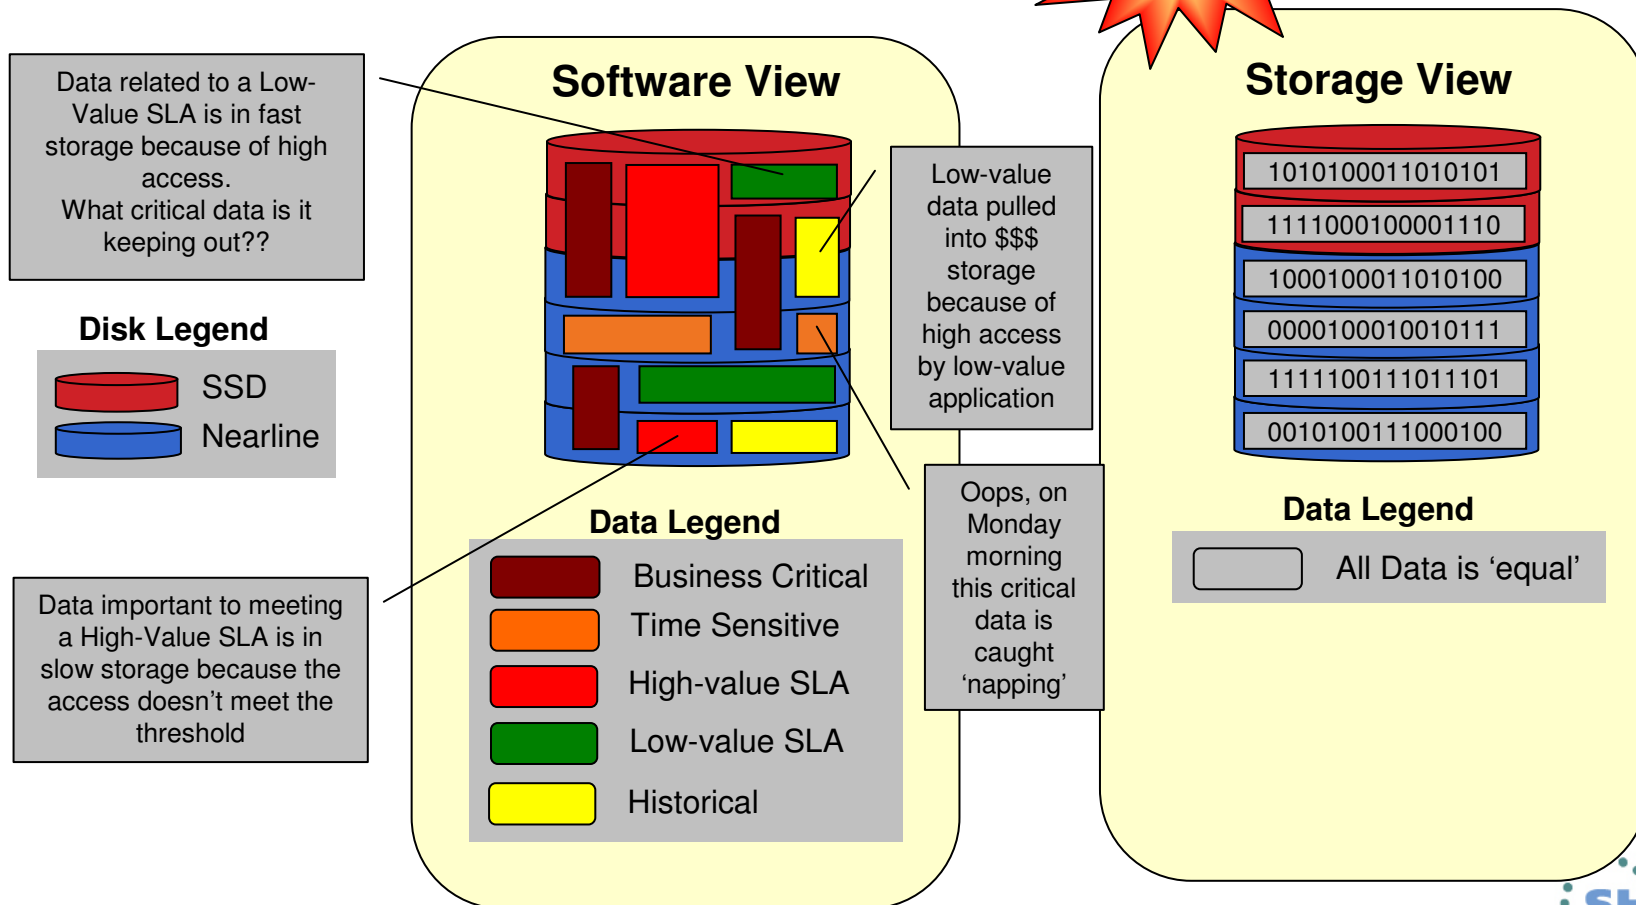
New technology that has room for growth

- Good technology for specific sets of data
 - Data that has relatively consistent data patterns that can be 'learned', such as database workloads
- Learning 'lost' when workload changes
 - For example, if weekend workload varies from weekday, then at 8 am on Monday morning, your data is on the wrong tier!
 - Reorgs, Defrags, Copies, etc
- Movement is done at a physical extent level...
 - *Not on data set boundaries*
- Disparate extent sizes across vendors

Space Management Evolution

Disk 201

- **When 'Smart' Tiering is not so smart...**
 - Where business value and heat maps collide...




Space Management Evolution

Disk 201



Data access patterns that don't match well with storage tiering...

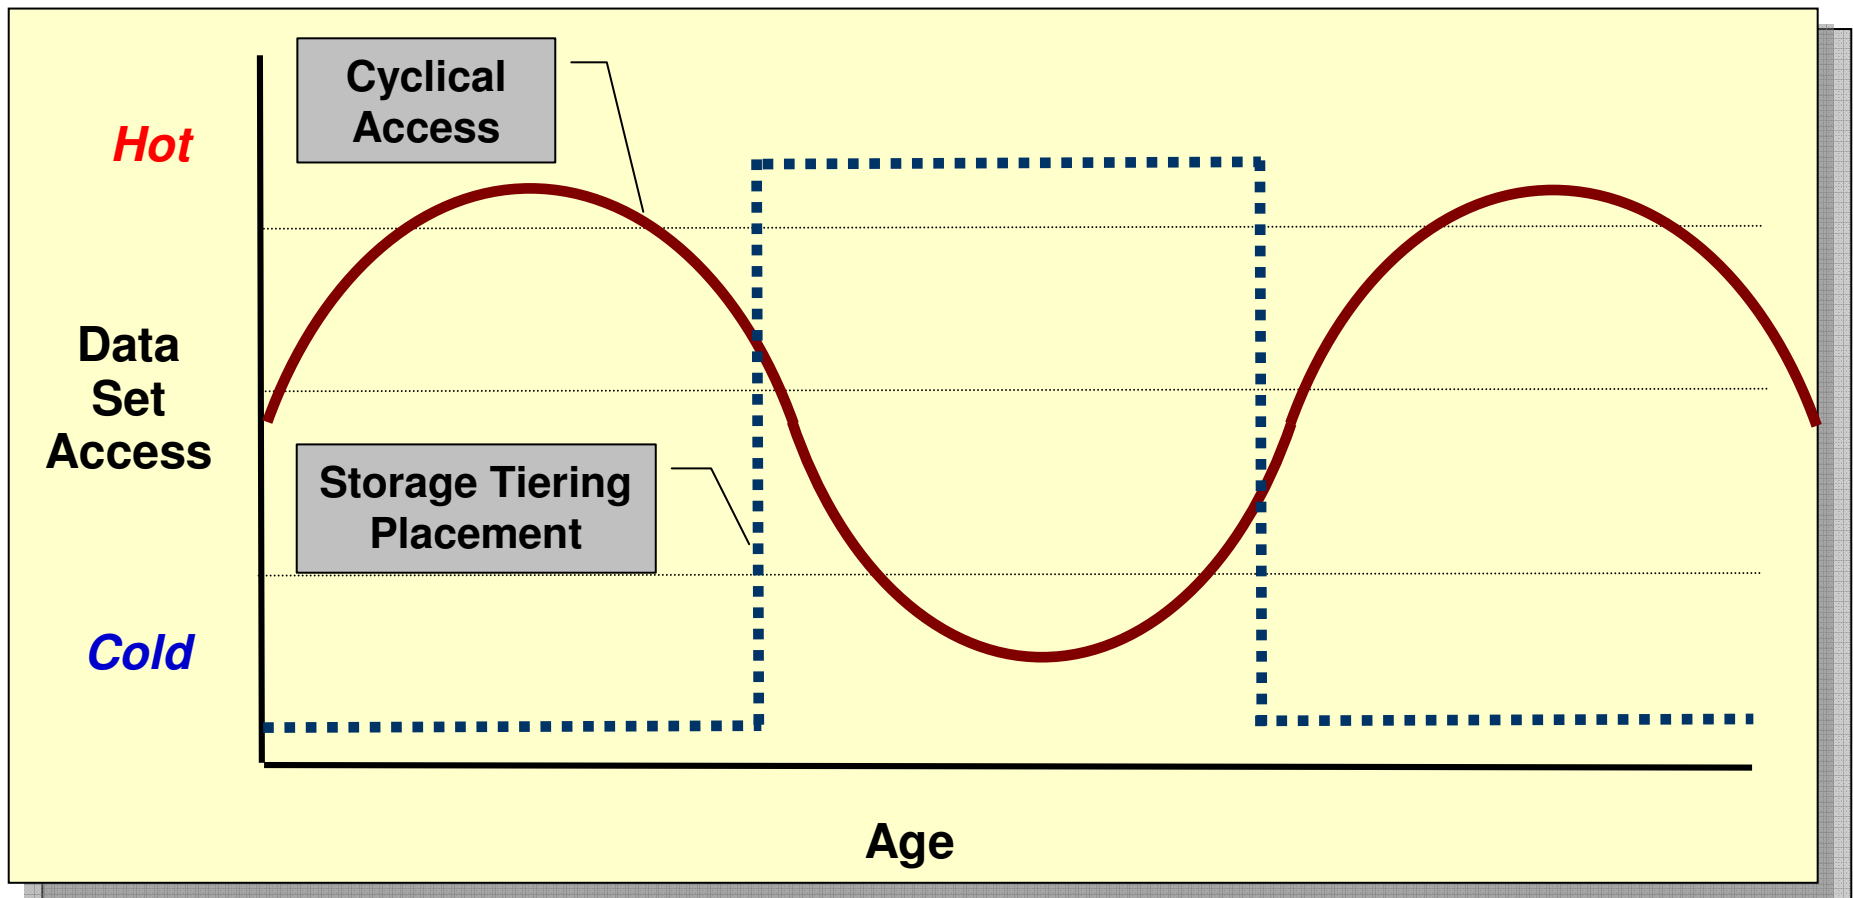


Space Management Evolution

Disk 201



Data access patterns that don't match well with storage tiering...



Space Management Evolution

Disk 301



“The reports of my death have been greatly exaggerated”

Mark Twain



*Smart storage tiering does not close the book on software based Space Management, but rather opens a new chapter in storage and software synergy to **take Space Management to a whole new level!***

Space Management Evolution

Disk 301



- *Just as Virtual Tape has addressed many key issues with Space Management, Smart Storage Tiering is just getting started...*

Issue	Virtual Tape	Smart Tiering
CPU - Compression		
CPU – Run time		+
Channel overhead		+
Delay in access		+
Software \$\$\$*		+

* Via decreased MIPS usage by space management software

Space Management Evolution

Disk 301

- **Evolving Smart Storage Tiering Environment**

Once again, NOT a statement of direction

- **Where do we go from here?**

- Direct synergy between hardware and Space Management software to drive automated, *policy-based*, transparent data movement within the storage controllers
 - *Data is at the right place at 8am Monday morning*
 - *Data is at the right place for High-Value SLA applications*
 - *Dynamic, on-demand, placement*
 - *Whenever an application runs*
 - *All related data*



DFSMS Storage Tiers z/OS V2R1

- **Information Lifecycle Management**

The process of managing information – from creation, through its useful life, to its eventual deletion – in a manner that aligns storage costs with a changing business value of information.

- **DFSMS Storage Tiers z/OS V2R1**

- ★ Better align storage costs with changing business value

- ★ Minimize the Total Cost of Ownership for System z Data by actively managing data on the lowest storage that meets the business needs of the data

- **Address two key shortcomings of today's functionality**

- ✓ No policy-based automation for moving data within the Primary Storage Hierarchy (Level 0)

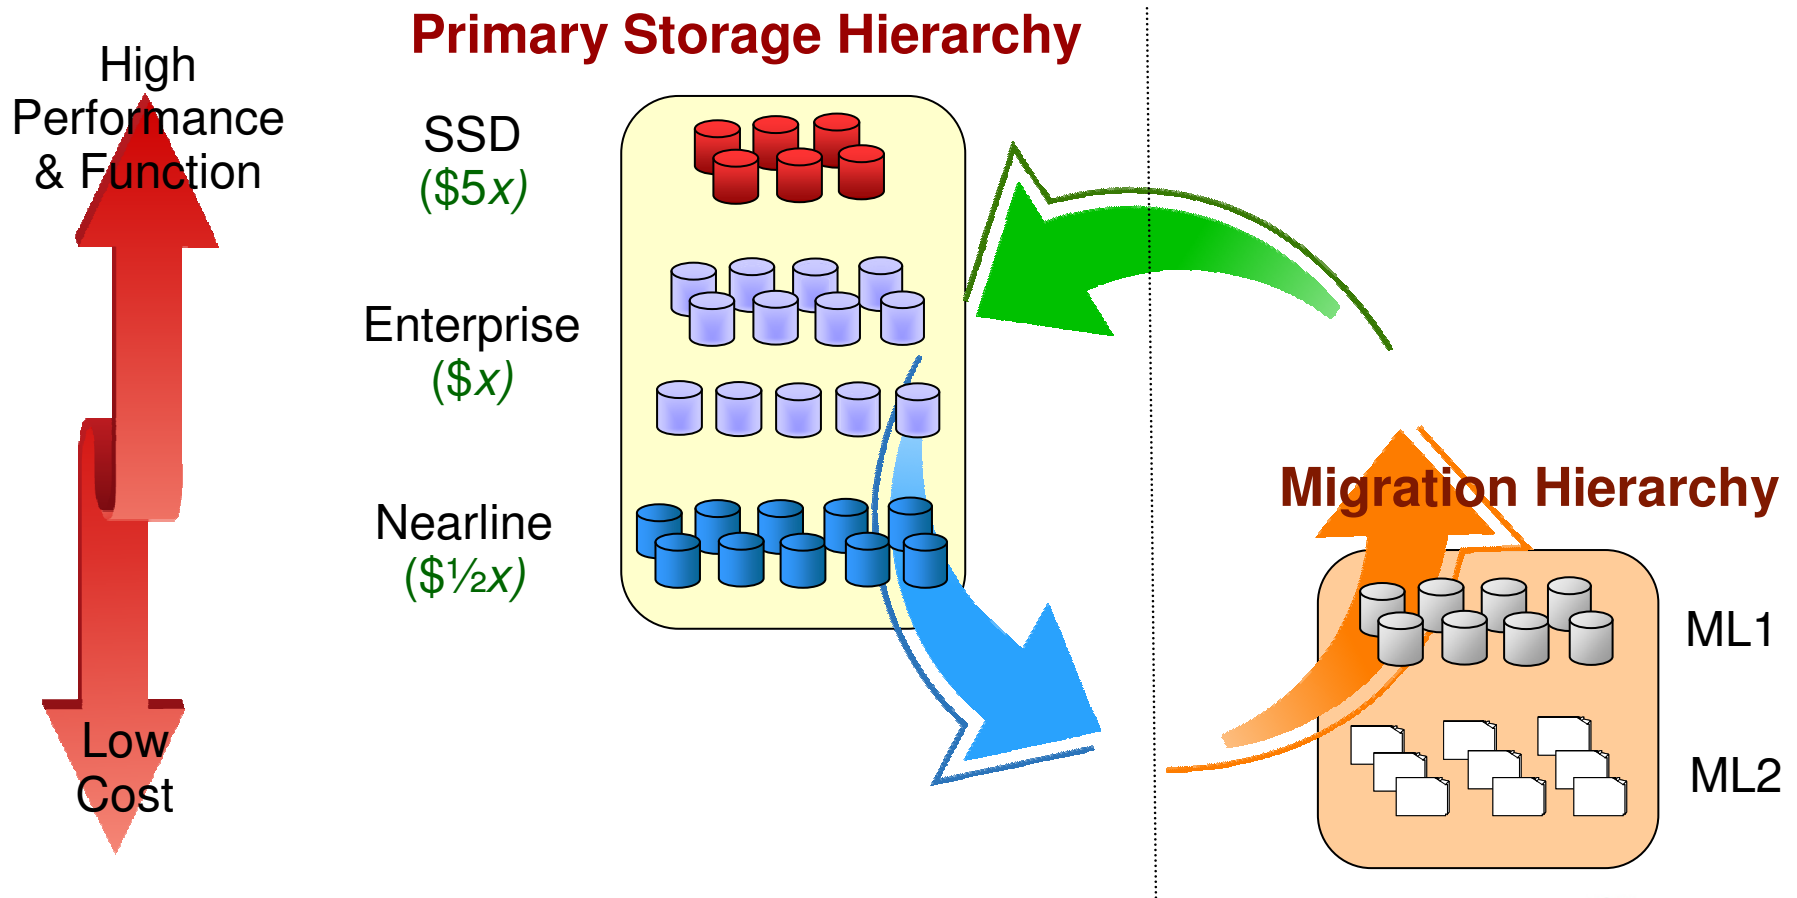
- ✓ No policy-based management of Active (open) data

DFSMS Storage Tiers z/OS V2R1

- **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 existing Class Transition policies and updated Management Class policies
- ✓ Enhanced support for DB2, CICS and zFS data
 - Open data sets are temporarily closed to enable movement

DFSMS Storage Tiers z/OS V2R1

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

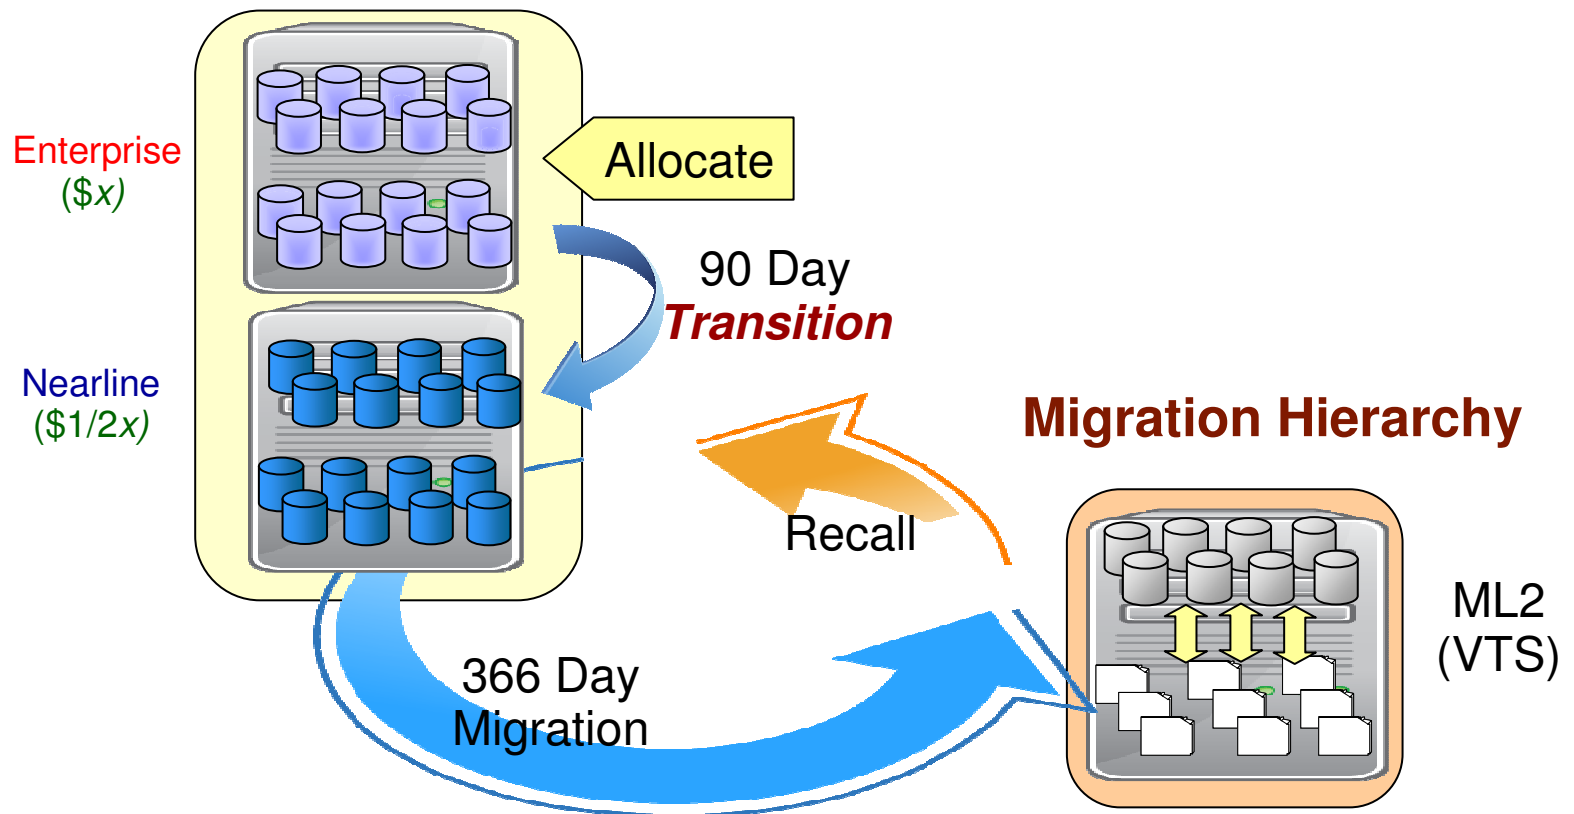


DFSMS Storage Tiers z/OS V2R1



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

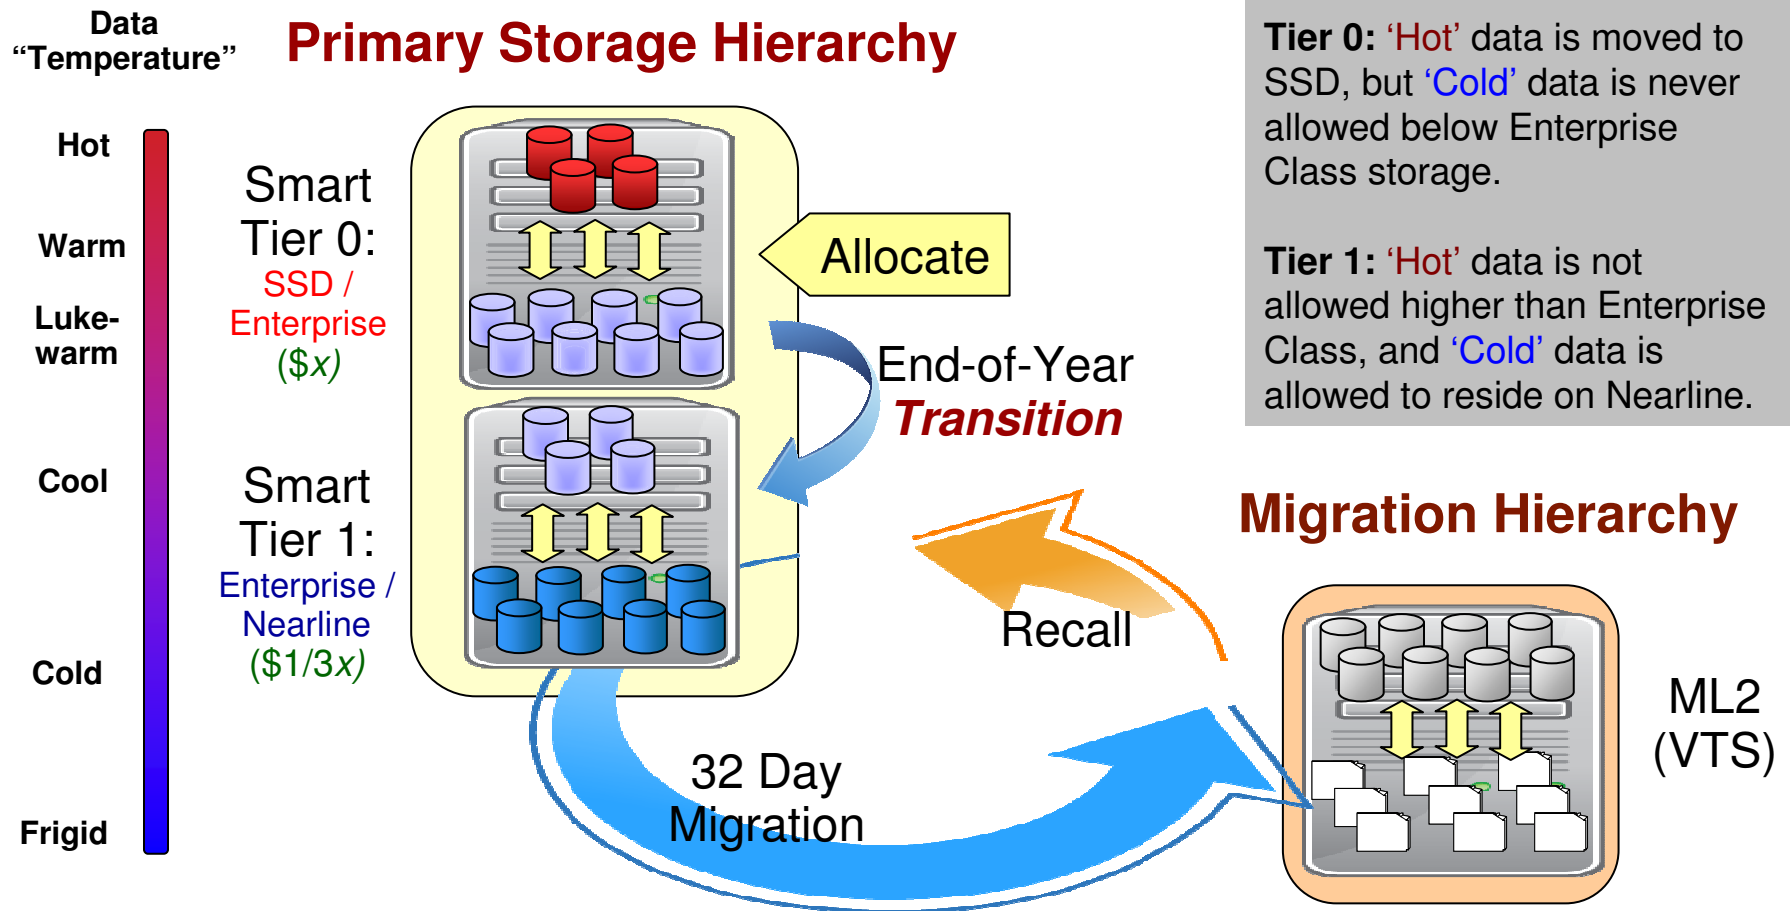
Primary Storage Hierarchy



DFSMS Storage Tiers z/OS 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 Nearline.


DFSMS Storage Tiers z/OS V2R1

So, *if* a vendor introduces Tape as a Tier, I'll be able to get rid of my Migration Hierarchy, correct?

Not so Fast! Since the tape will look like a disk, before the data can be accessed, it must first be cached back to a real disk. That delay will be manifested in the I/O response time. How many of your applications can accept a 2.5 minute, or worse, I/O delay? For the most part, tape as a tier will need to be closely integrated with application behavior and storage software support.

(Remember, the delay on a Recall is on the OPEN not a read or write I/O)

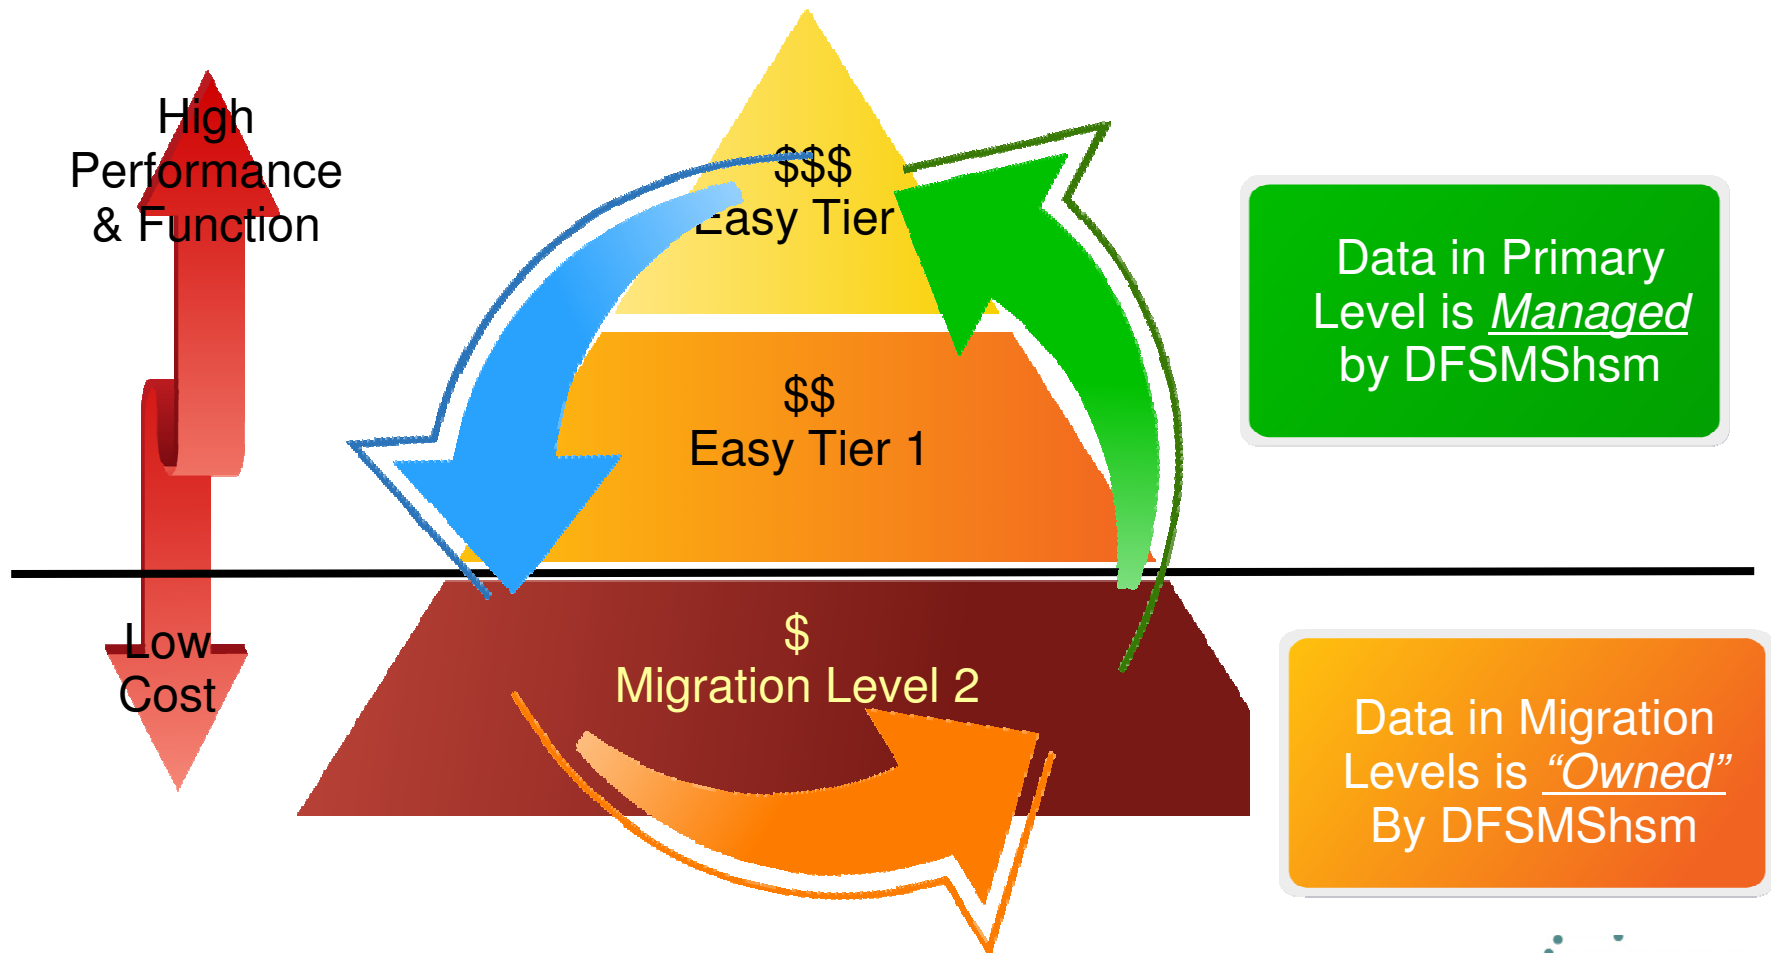
Primary Storage Hierarchy



DFSMS Storage Tiers z/OS V2R1



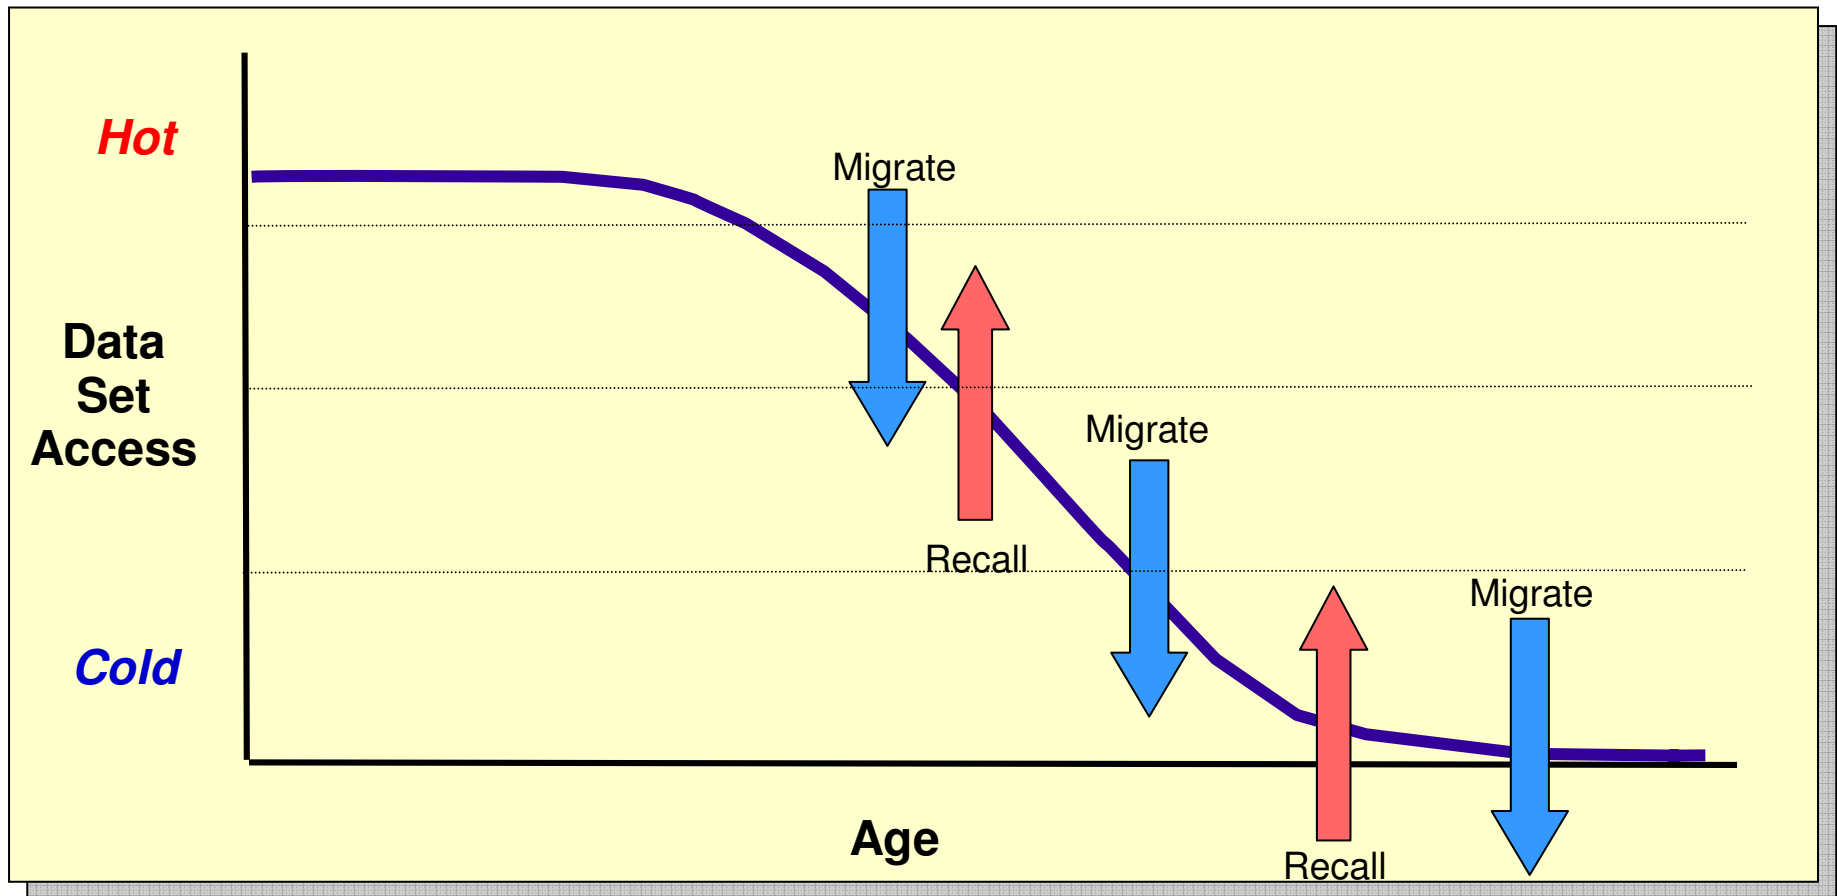
The *enhanced* DFSMS storage hierarchy enables “the bar” to be in the most appropriate location for each customer environment.



DFSMS Storage Tiers z/OS V2R1



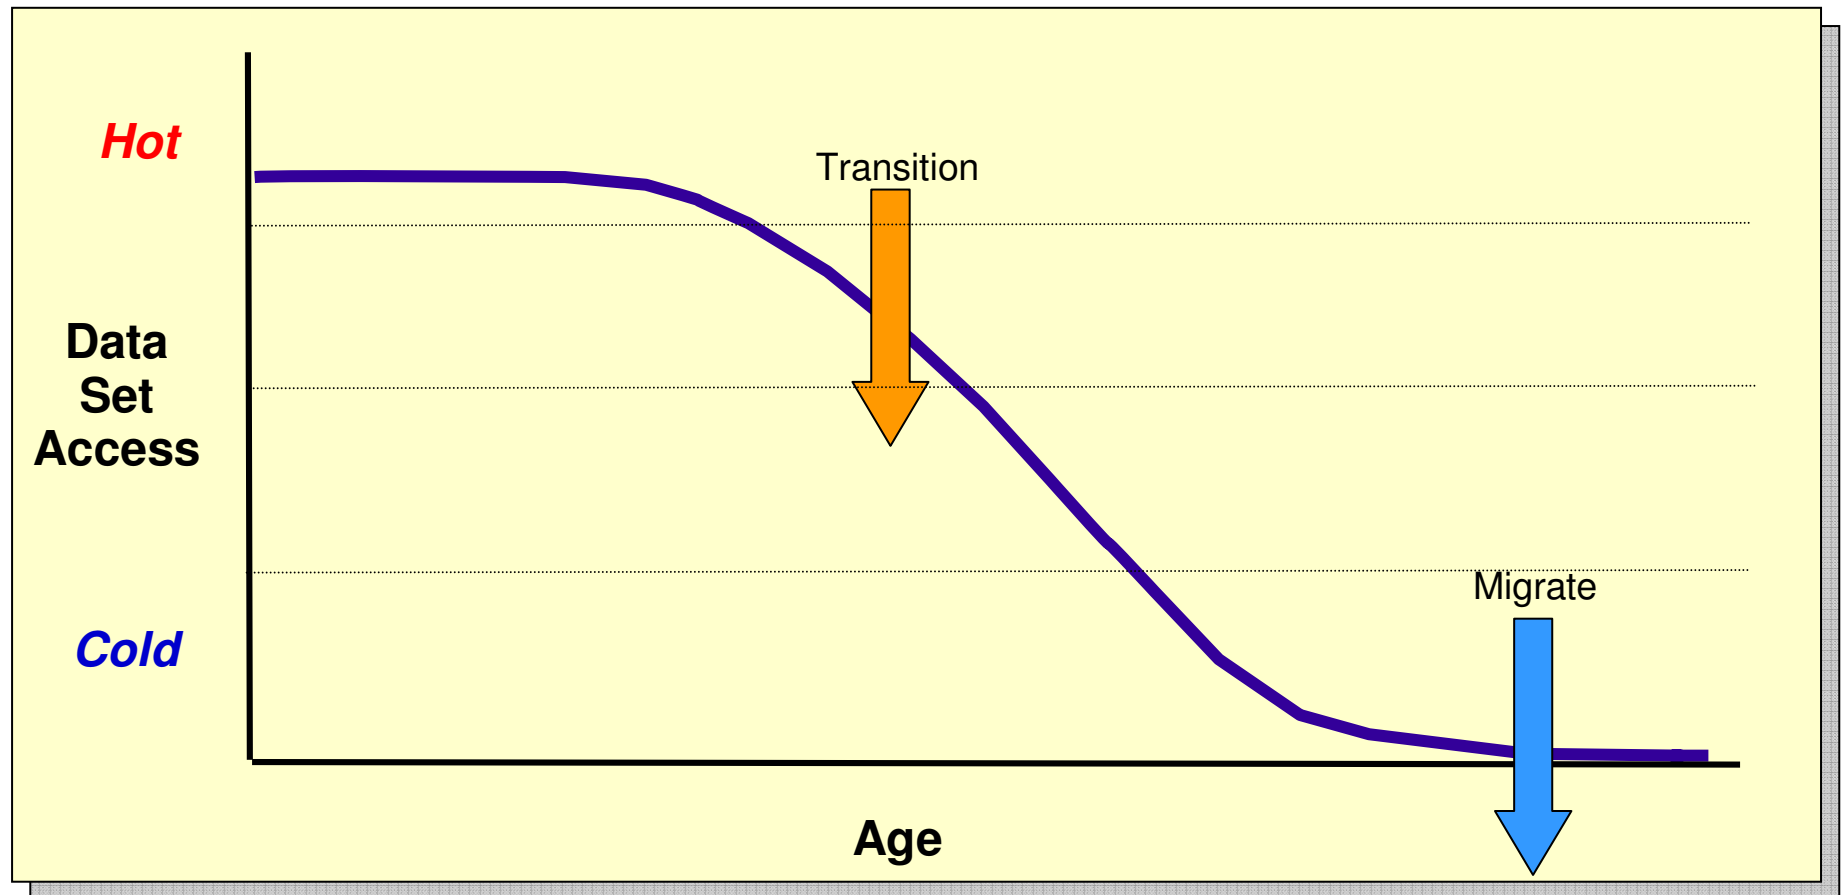
As a data set ages and goes through cycles of activity and inactivity, it can go through many migration and recall iterations.



DFSMS Storage Tiers z/OS V2R1



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

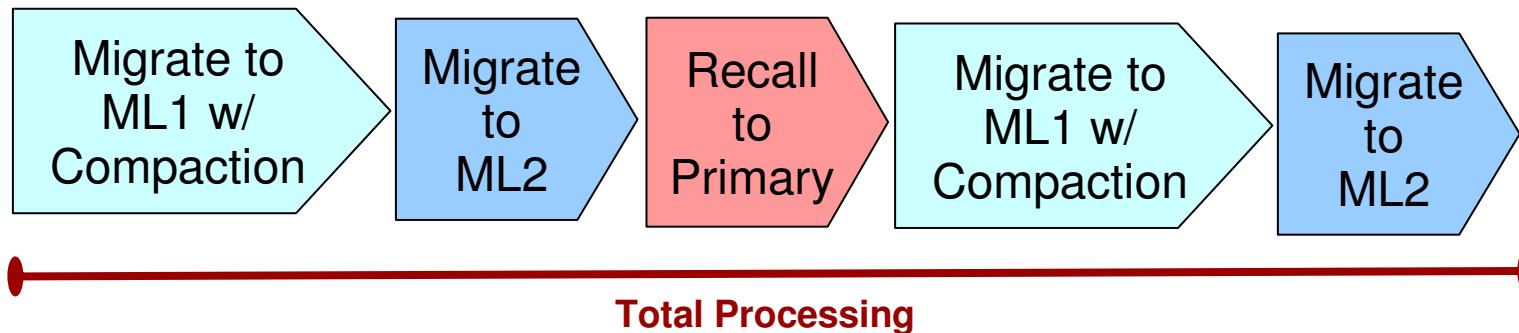


DFSMS Storage Tiers z/OS V2R1




Contrast the Space Management processing required when ...

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



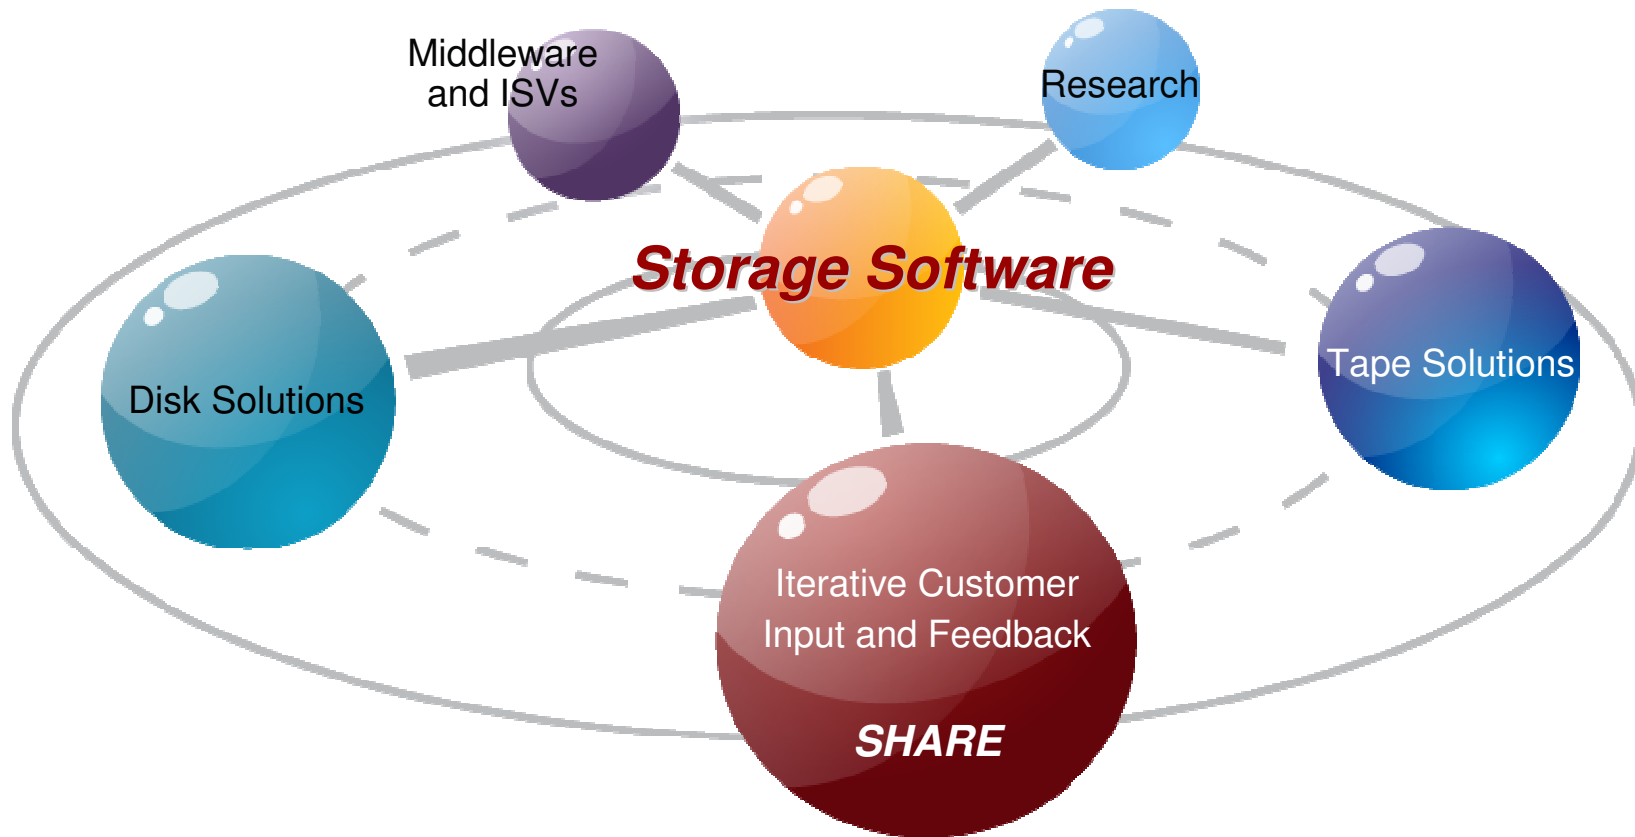
- A data set is transitioned using **FlashCopy** and then migrated directly to ML2



DFSMS Storage Tiers z/OS V2R1



Storage Software is uniquely positioned to integrate all of the pieces and provide automated, policy-based solutions...



Complete your sessions evaluation online at SHARE.org/BostonEval



Evolution of Space Management



Want to find out more??

Get all of the juicy details 1:30 pm Wednesday afternoon...

Session 14135, Hynes Room 304

*A New Frontier in the Evolution
Of Space Management... Blast Off!!*



This Session:
Session 14101
QR Code ->



System z Social Media Channels

- Top Facebook pages related to System z:

- [IBM System z](#)
- [IBM Academic Initiative System z](#)
- [IBM Master the Mainframe Contest](#)
- [IBM Destination z](#)
- [Millennial Mainframer](#)
- [IBM Smarter Computing](#)

- Top LinkedIn groups related to System z:

- [System z Advocates](#)
- [SAP on System z](#)
- [IBM Mainframe- Unofficial Group](#)
- [IBM System z Events](#)
- [Mainframe Experts Network](#)
- [System z Linux](#)
- [Enterprise Systems](#)
- [Mainframe Security Gurus](#)

- Twitter profiles related to System z:

- [IBM System z](#)
- [IBM System z Events](#)
- [IBM DB2 on System z](#)
- [Millennial Mainframer](#)
- [Destination z](#)
- [IBM Smarter Computing](#)

- YouTube accounts related to System z:

- [IBM System z](#)
- [Destination z](#)
- [IBM Smarter Computing](#)

- Top System z blogs to check out:

- [Mainframe Insights](#)
- [Smarter Computing](#)
- [Millennial Mainframer](#)
- [Mainframe & Hybrid Computing](#)
- [The Mainframe Blog](#)
- [Mainframe Watch Belgium](#)
- [Mainframe Update](#)
- [Enterprise Systems Media Blog](#)
- [Dancing Dinosaur](#)
- [DB2 for z/OS](#)
- [IBM Destination z](#)
- [DB2utor](#)



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.