#SHAREorg

SHARE
Technology · Connections · Results

Hitachi Dynamic Tiering

Storage Tiering for the Mainframe

William Smith
Hitachi Data Systems

February 5, 2013
Session Number 12680

SHARE in San Francisco 2013

AGENDA

- Virtual Storage Platform – Architecture Designed for Tiering

Hitachi Dynamic Provisioning (HDP)

Hitachi Dynamic Tiering (HDT)

Summary

Virtual Storage Platform - 3D Scaling Architecture

- 5th generation multi-dimentional internal crossbar switch
- Native 6G SAS backend
- 192 x FC, 96 x FICON / FCoE
- Independent Scale up, scale out, scale deep growth path
- 1TB Global Cache
- 3rd generation imbedded virtualization (Open Systems and Mainframe)
- Granular non-disruptive upgrade paths
- 100% non-disruptive maintenance (7 x 24 x forever)

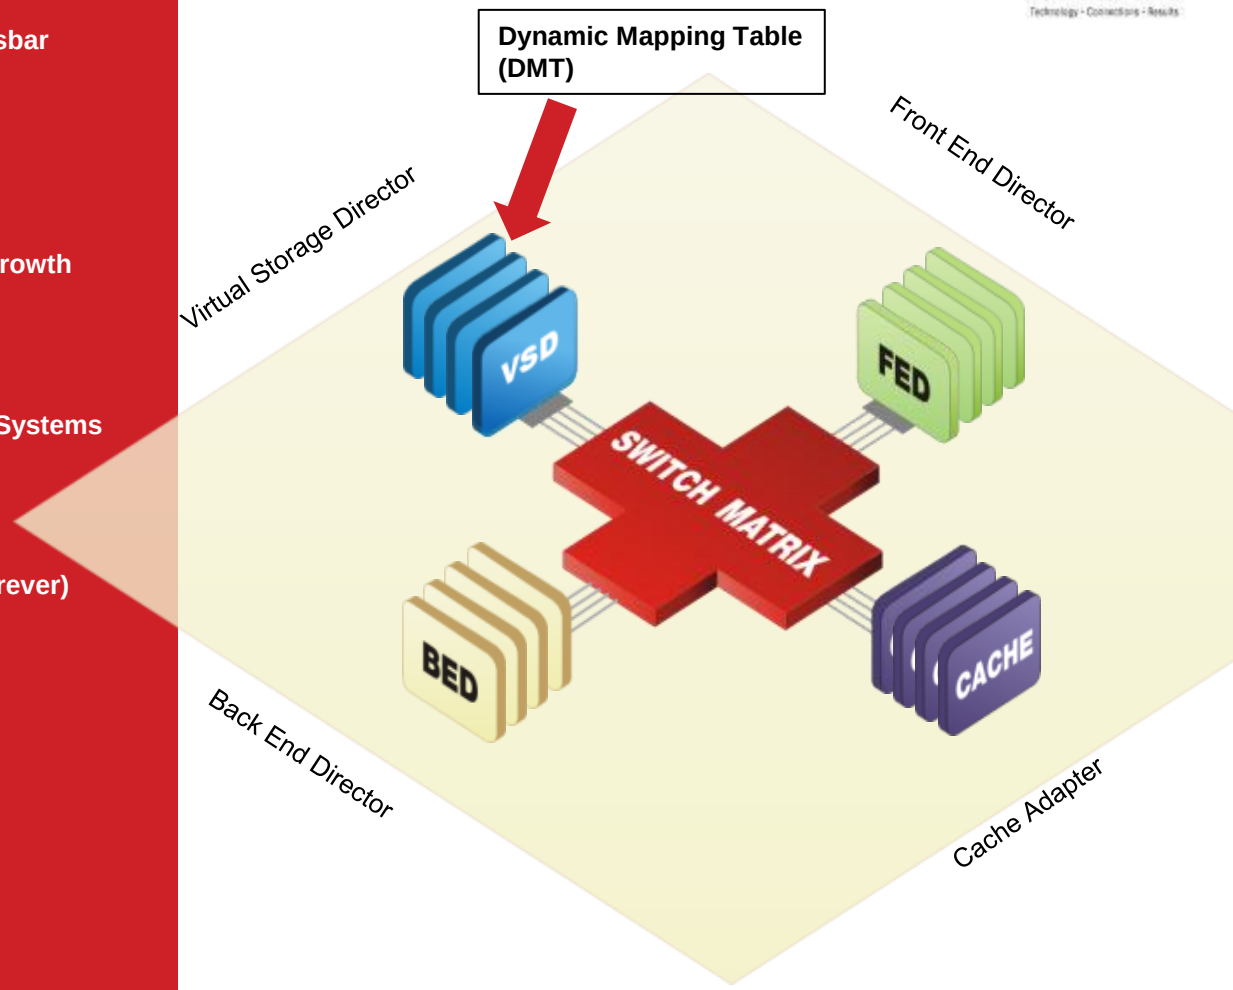

Drive Types and Capacities

- Current Drive Technology:
  - 2.5” 10K 300GB SAS
  - 2.5” 10K 600GB SAS
  - 2.5” 10K 900GB SAS
  - 2.5” 15K 146GB SAS
  - 2.5” Flash SLC 200GB 3G and 6G SAS
  - 3.5” Flash SLC 400GB 3G and 6G SAS
  - 3.5” 2TB Nearline SAS / 7200 RPM

- New Drive Technology (last quarter):
  - 2.5” 15K 300GB SAS
  - 2.5” Flash MLC 200GB 6G SAS
  - 2.5” Flash MLC 400GB 6G SAS
  - 3.5” Flash MLC 400GB 6G SAS
  - 3.5” 3TB Nearline SAS / 7200 RPM (later this month)
  - New High Density Flash Drives

New Hitachi Flash Technology

- The Flash Memory Drive (FMD) is Hitachi built specifically for Hitachi arrays.
- Each FMD is as follows:
  - Raw capacity option of 1.6TB
  - eMLC based
  - Each FMD uses an Hitachi designed and built System-On-Chip (SOC) ASIC
    - Multi-Core ASIC performs flash memory management, data transfer, RAS etc.
    - High bandwidth PCIe Gen2 data paths to eMLC memory
    - Much improved performance over industry SSD

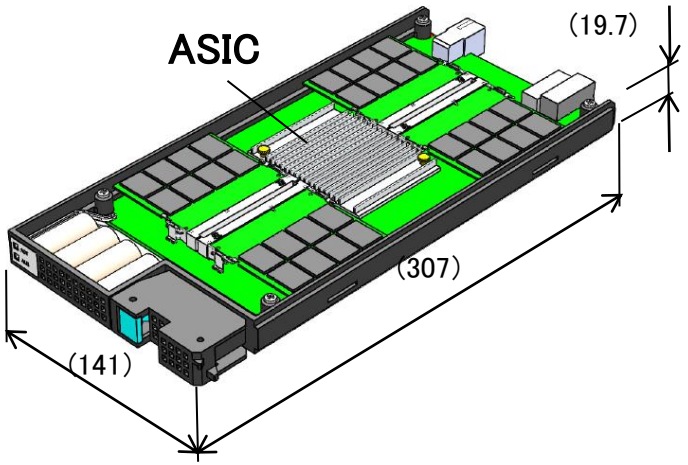

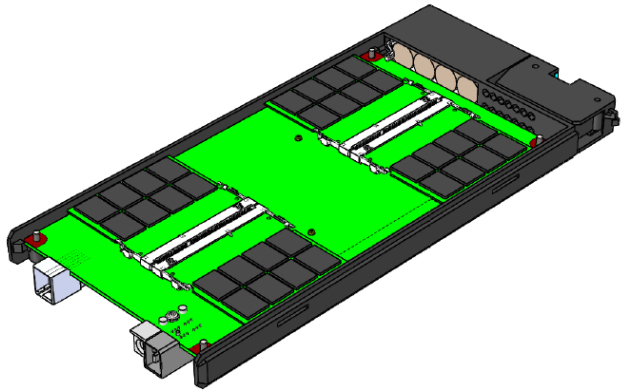

New Hitachi Flash Technology

- VSP will offer granularity of 4 Flash Memory Drives (FMD) minimum, in 4 separate FMU (Flash Memory Units)
- VSP can support a total of 192 FMD's
  - 12 Flash FMD's per FMU
  - 4 FMU per Flash Box (FBX - similar to current Disk Chassis)
  - 4 FBX per VSP (2 per Control Chassis)
  - Total raw capacity: **307TB** (1.6TB FM)
- Target performance is >1 Million Random Read IOPS per FMU
- Available for all VSP's and non-disruptive install on live VSP's
- Game changer in flash technology
  - Cost effectiveness per TB improved considerably
  - Expect market adoption to be very high

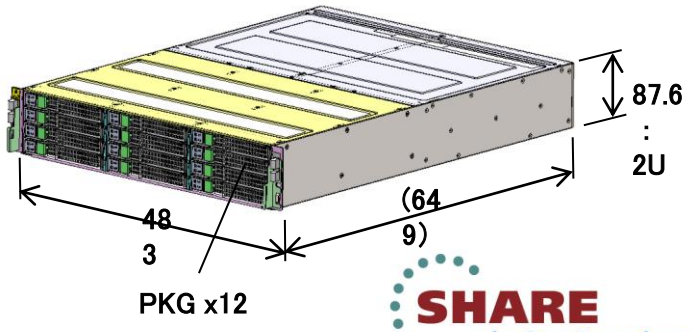

AGENDA

Virtual Storage Platform – Architecture Designed for Tiering

- Hitachi Dynamic Provisioning (HDP)

Hitachi Dynamic Tiering (HDT)

Summary

Dynamic Provisioning for Mainframe

- **Dynamic Provisioning**
  - **OPEN Systems – 2nd Generation**
  - **Mainframe – August 2011**
  - **Internal and external (virtualized) storage**
  - **Managed in pages (42MB for OPEN, 38MB for Mainframe)**

- **Cost savings through:**
  - **Improved storage utilization**
  - **Improved performance through ultra-wide striping**
  - **Non-disruptive addition of storage into the pool**
  - **Integration with SMS**

Hitachi Dynamic Provisioning (HDP) for the Mainframe

- So, just what is a page?
  - For the VSP, there are two page sizes
    - OPEN – 42MB
    - Mainframe – 38MB
  - They are different to fit the best use of storage space for the underlying OS structures they support
  - For Mainframe, 38MB fits the 3390 architecture with minimal wasted space
- Each storage supplier has a different implementation
  - Some are in MB
  - Some are sub-LUN
  - Some are partial volumes
  - Some are Chunks / Chunklets
- Each is designed to best fit their architecture

Hitachi Dynamic Provisioning for Mainframe

- Optimize Storage performance by spreading the I/O across all available Disk Drives
- Optimize Storage Capacity by only allocating capacity which is actually used
- Reduce OPEX by automatic thin moves and copies of fat volumes
- Achieve Flash Copy Space Efficiency for Target Volumes
- Leverage Dynamic Volume Extension (DVE) to seamlessly grow physical and logical capacity
- Leverage Extended Address Volumes (EAV)

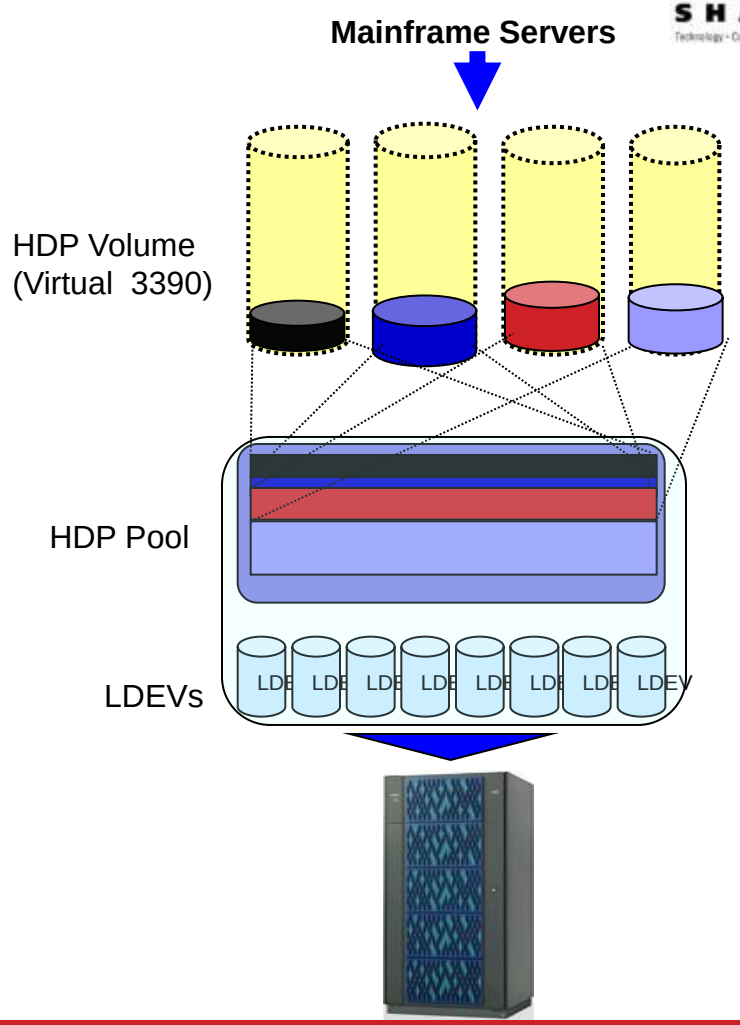

Simplifies Capacity Expansion; no IOGEN, no documentation updates, no changes to GDPS Parms, XRC Sessions, Flashcopy configs etc.

HDP Mainframe — Basic Concepts

- **Key Functionalities**
  - Thin Provisioning
    - Pool Rebalance
    - Pool Volume Reduction
    - Space Reclaim
  - Support Mainframe Features
    - PAV, HyperPAV
    - z/HPF
  - Support Replication Products
    - TrueCopy and HUR
    - ShadowImage and FlashCopy V2
  - Support New Features
    - EAV and DVE
    - FlashCopy SE
    - Hitachi Dynamic Tiering

- **With HDP for Mainframe**
  - Optimize storage performance by spreading the I/O across all available physical disk drives

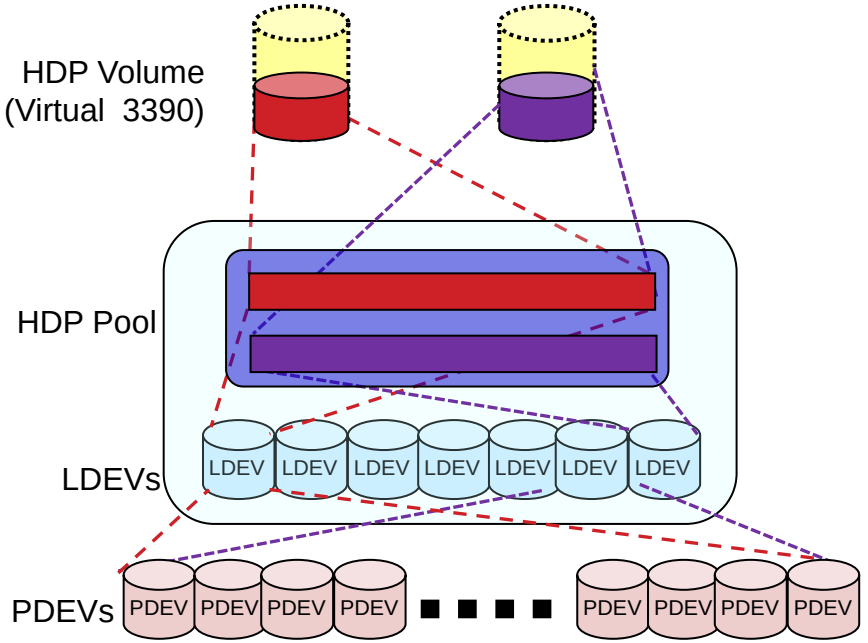

  - Spreading data across more physical volumes reduces the I/Os per volume and improves performance

AGENDA

Virtual Storage Platform – Architecture Designed for Tiering

Hitachi Dynamic Provisioning (HDP)

- Hitachi Dynamic Tiering (HDT)

Summary

Hitachi Dynamic Tiering for Mainframe

- Dynamic Tiering
  - Open Systems available since September 2010
  - Mainframe Available since January 2012
  - Automated performance management
  - Cost effective use of storage tiers
  - Internal and external (virtualized) storage
  - Managed in pages (42MB for OPEN, 38MB for Mainframe)
- Cost savings through:
  - Improved storage utilization
  - Reduced storage requirements for Tier 1
  - Less time spent on performance management
  - Re-deploy personnel
  - Integration with SMS

Virtual Storage Platform: Page Level Tiering

- Different tiers of storage are in one pool of pages – up to 128 pools
- Automatically detects and assigns Tiers based on media type (includes the new HAFs SSD).
- Dynamically:
  - add or remove Tiers
  - expand or shrink Tiers
  - expand 3390A volumes
- Dynamic Tiering Policies at volume level:
  - Default - All
  - Level 1 - Tier 1 only
  - Level 2 - Tier 1 and Tier 2 only
  - Level 3 - Tier 2 only
  - Level 4 - Tier 2 and Tier 3 only
  - Level 5 - Tier 3 only
  - Tier 3 External can have 3 tiers
  - 26 new policies added (% tier)

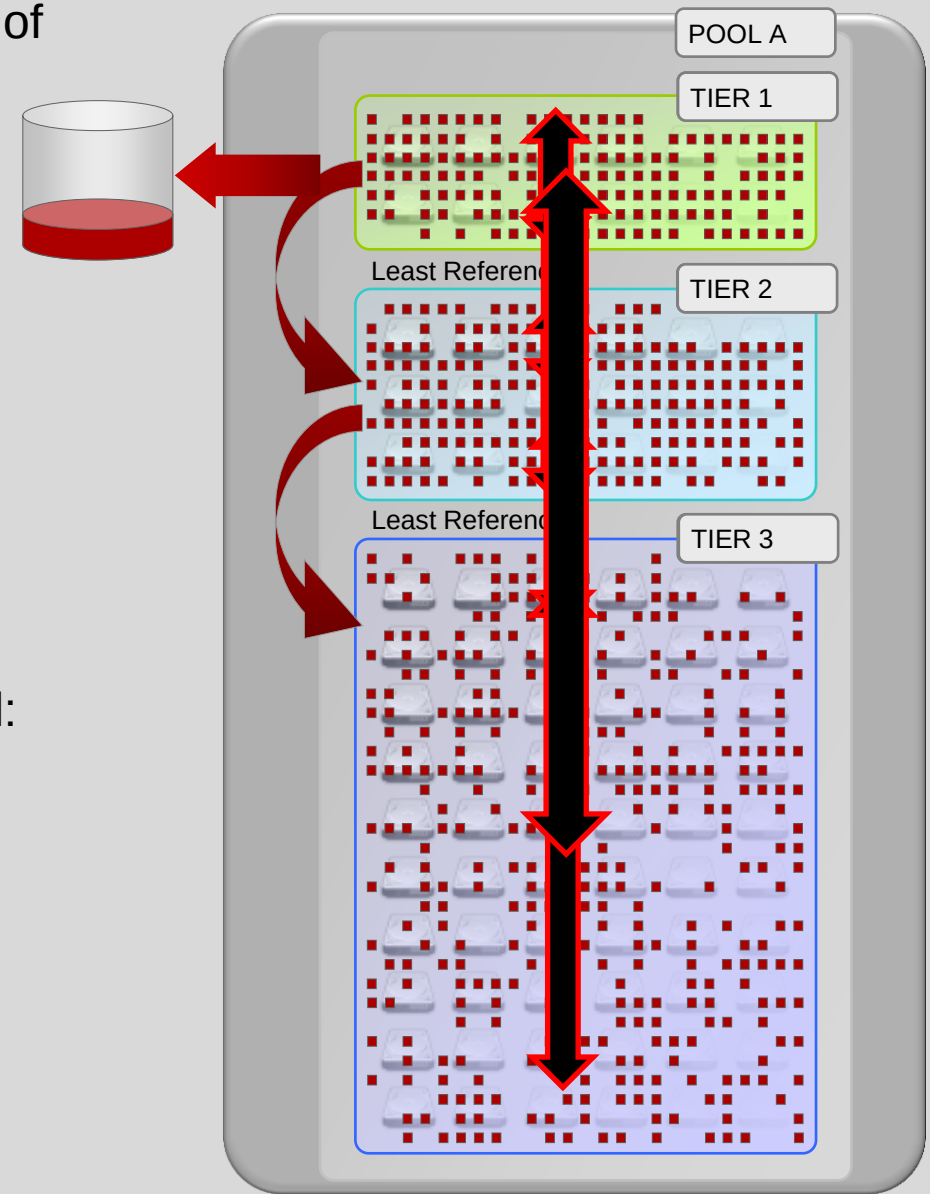

Virtual Storage Platform: Page Level Tiering

- Cycles
  - Customer defines strategy, it is executed automatically
  - 24 hour
    - defined part of 24 hour
    - ½, 1, 2, 4 or 8 hourly
    - All aligned to midnight
  - Weekly or longer
- Use Hitachi Tier Storage Manager (HTSM)
  - Set tiering policy schedules to match business cycles and requirements
  - Report on data mobility (promotions and demotions)
  - Dynamically change or adjust cycle, policies, tiers etc…

Media tier groupings supported by Hitachi VSP	Order of tiers
2.5" SSD (200GB) SLC or eMLC 2.5" SSD (400GB) SLC or eMLC	1
2.5" SAS15Krpm (146GB) 2.5" SAS15Krpm (300GB)	2
2.5" SAS10Krpm (300GB) 2.5" SAS10Krpm (600GB) 2.5" SAS10Krpm (900GB)	3
2.5" SAS 7.2Krpm (2TB, 3TB)	4
3.5" SATA (2TB)	5
External	Lowest
External (June '12)	Any Tier

Hitachi Dynamic Tiering

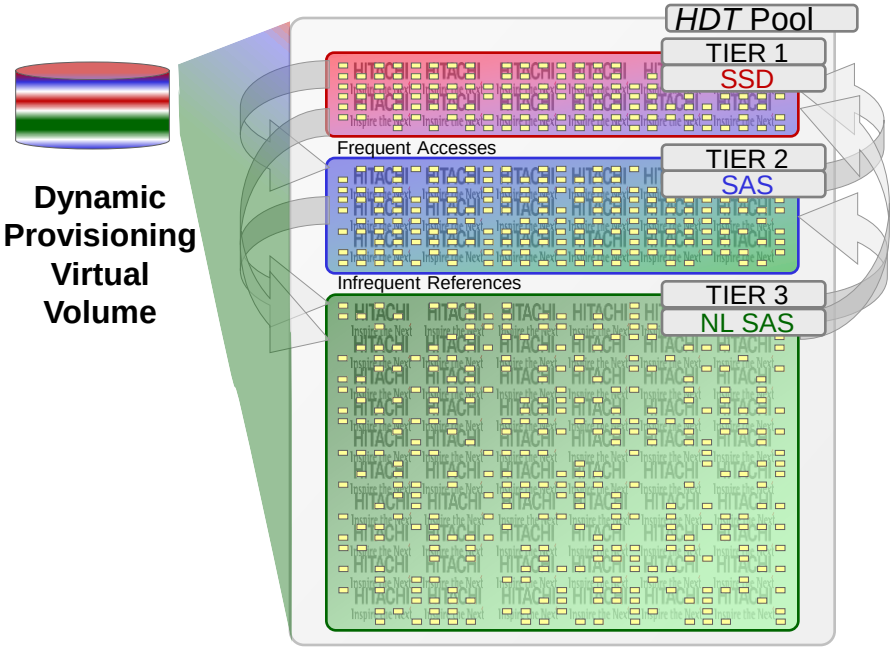

- What determines if a page moves up or down?
- When does the relocation happen?

Page Relocation

- At the end of a monitor cycle the counters are recalculated
  - Either IOPH (Period) or weighted average (Continuous)
- Page counters with similar IOPH values are grouped together
- IOPH groupings are ordered from highest to lowest
- Tier capacity is overlaid upon the IOPH groupings to decide on values for Tier Ranges
  - Tier Range is the ‘break point’ in IOPH between tiers
- Relocation processes page by page looking for pages on the ‘wrong’ side of a Tier Range value
  - i.e. high IOPH in a lower tier
- You can see the IOPH groupings and Tier Range values in HDT reporting “Pool Tier Properties”

Pool Tier Properties

View Tier Properties

View Tier Properties

HITACHI

View Tier Properties

Tier1	Drive Type/RPM	SSD/-
	Capacity(Used/Total)	2.45 TB / 3.13 TB [78 %]
	Performance Utilization	5 %
Tier2	Drive Type/RPM	SAS/15k
	Capacity(Used/Total)	34.96 TB / 43.92 TB [79 %]
	Performance Utilization	62 %
Tier3	Drive Type/RPM	SATA/7.2k
	Capacity(Used/Total)	10.42 TB / 42.99 TB [24 %]
	Performance Utilization	100 %

Performance Graph (SPC-1(0))

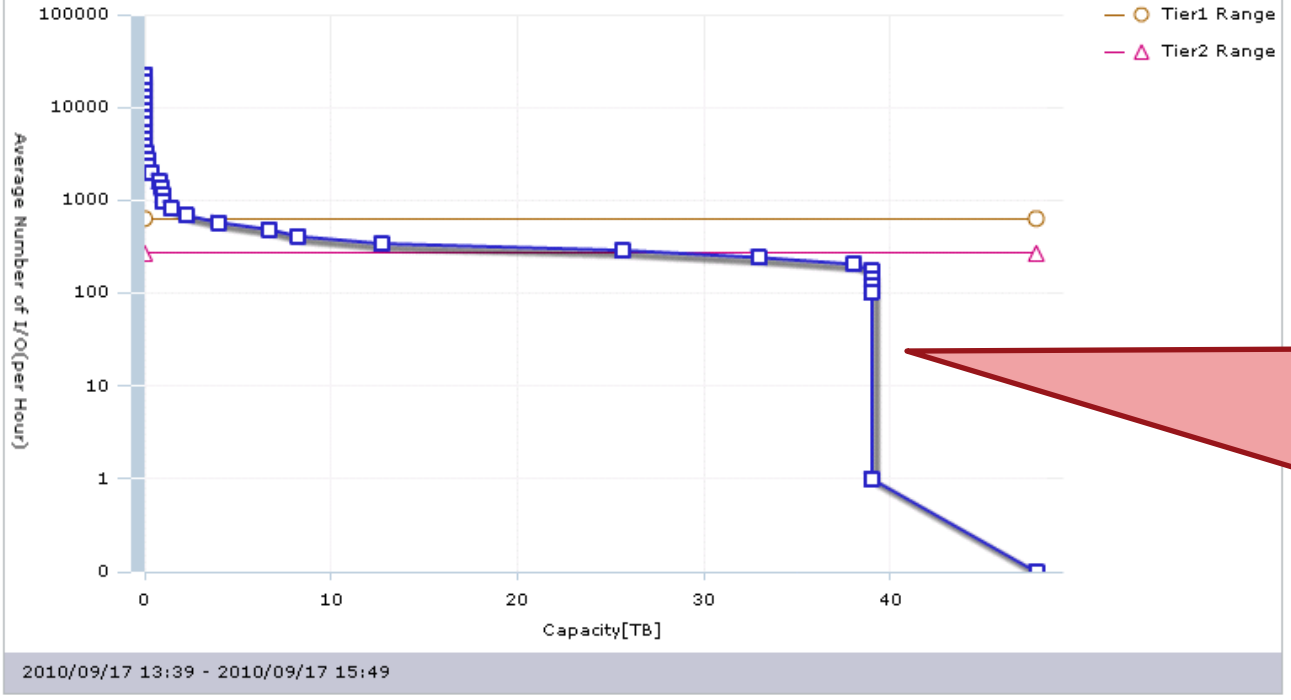

2010/09/17 13:39 - 2010/09/17 15:49

Close Help

What is being used **now** in the pool in terms of capacity **and** performance

The I/O distribution across all pages in the pool.

Combined with tier range this is what HDT is using to decide “where the pages should go to”

MF-HDT Cycle

- Cycle Time set at the MF-HDT Pool Level
- Manual Mode
  - User can start and stop performance monitoring using any interval up to 7 days
- Automatic Mode
  - Continuous Monitoring followed by Relocation cycles
  - Monitor Interval from 30min, 1h, 2h, 4h, 8h or 24h(default)

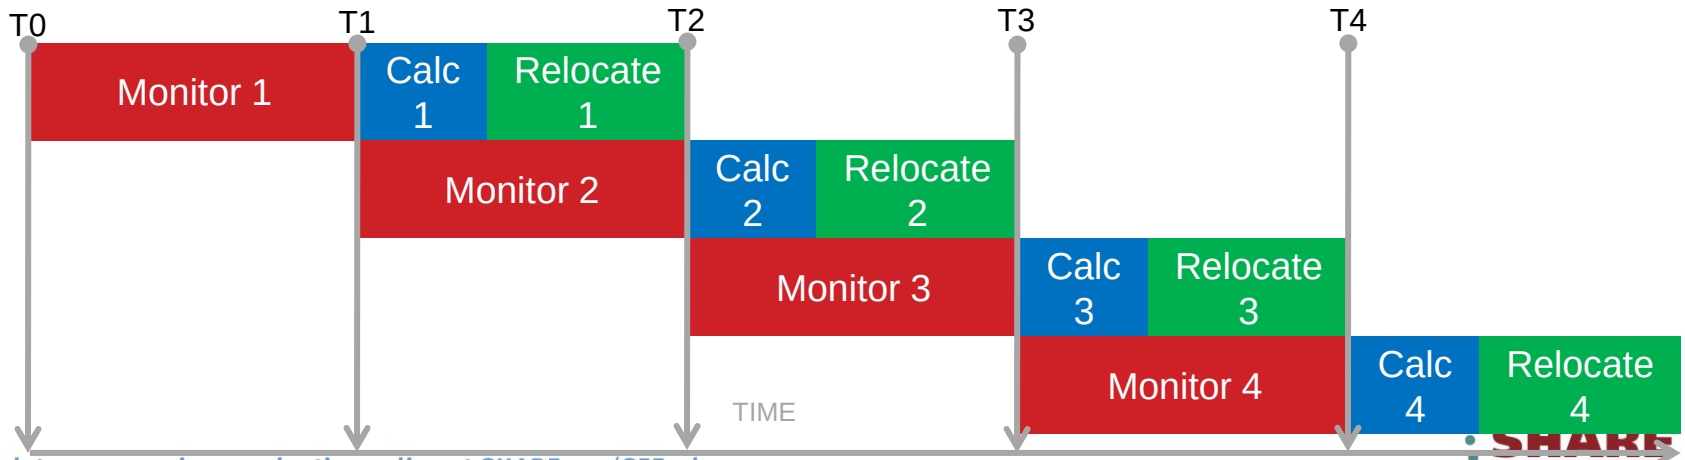

MF-HDT Monitoring Modes

- Monitoring Modes set at the MF-HDP Pool Level
- Period
  - The value used in the Calculation cycle is the Actual IO Load on DP-Volume Page from Previous Monitoring Cycle
- Continuous
  - The value used in the Calculation cycle is the Weighted Average of Multiple Previous Monitoring Cycles for DP-Volume Page
  - Reduces Page Thrashing
  - May Slow Migration to upper Tiers

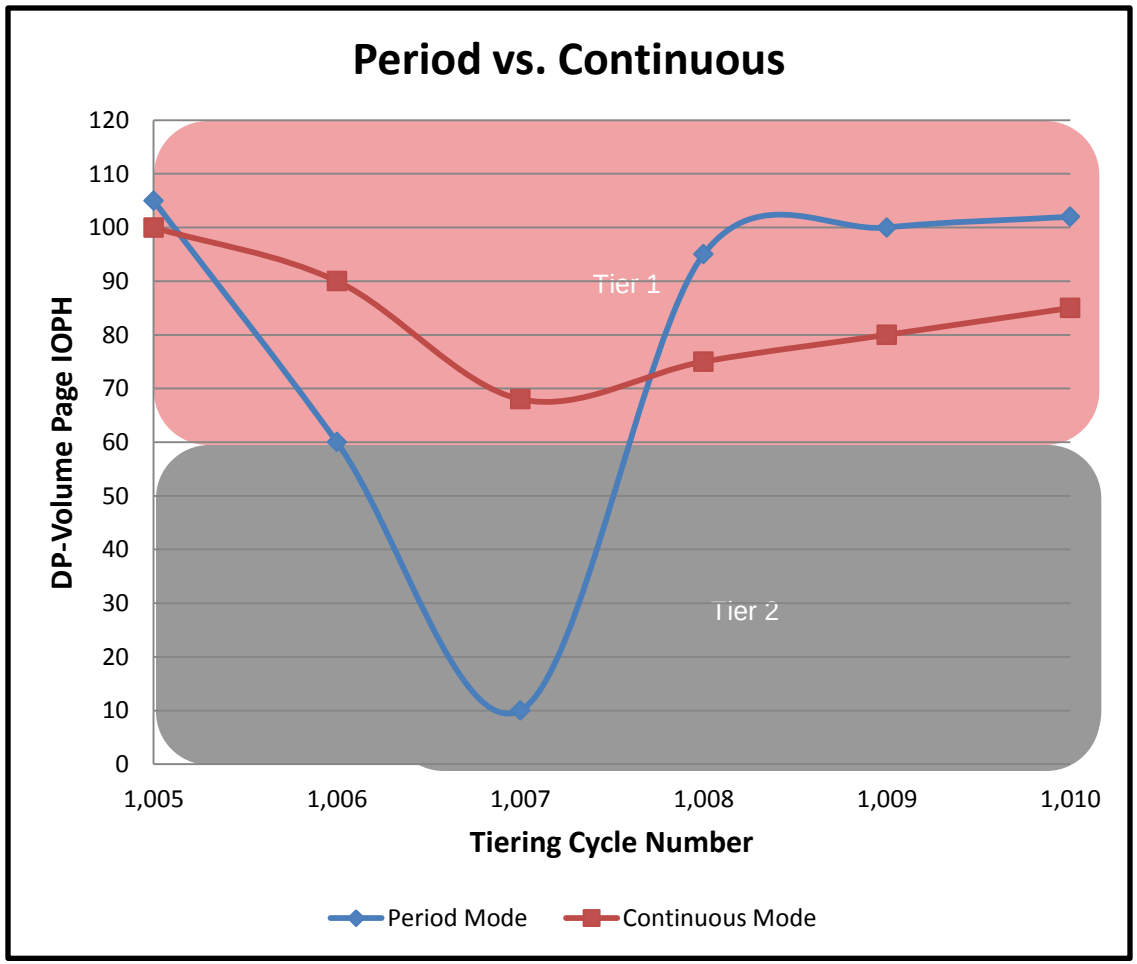

DFSMS Storage Groups and MF HDP and HDT

- Tiered Storage can be implemented with DFSMS Storage Groups
- DFSMS Attributes such as Data Class and Storage Class are used to direct allocation to the right Tier

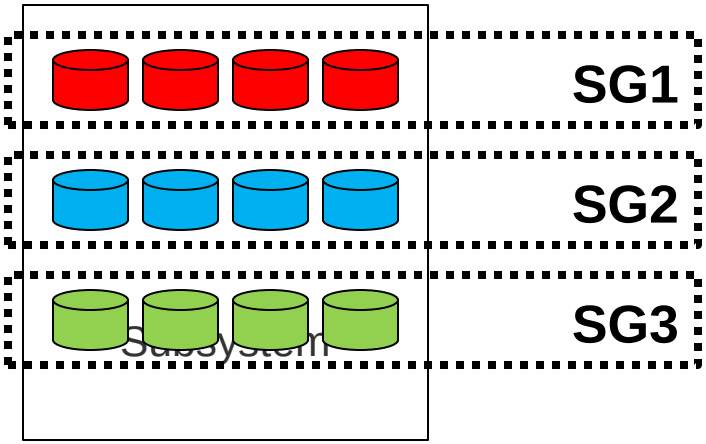

- With MF HDT may combine all SG's into one and use HDT to move data to the right Tier automatically based on Workload performance

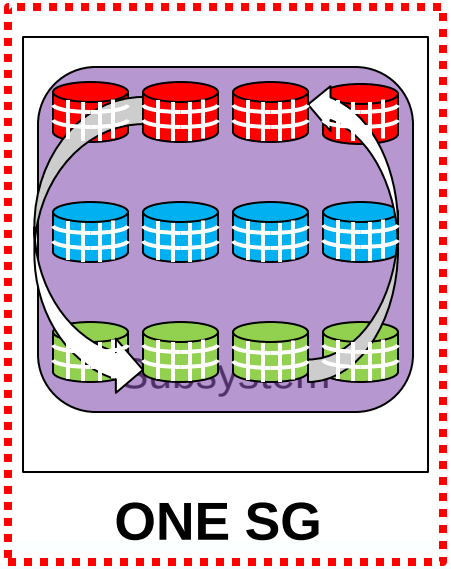

HDP Volumes are managed at page-level

DFSMS Storage Groups and MF HDT

- HDT can be applied to selected Storage Groups only
- DFSMS Horizontal Storage Groups and HDT Pools with Dynamic Page based Tiering can be aligned

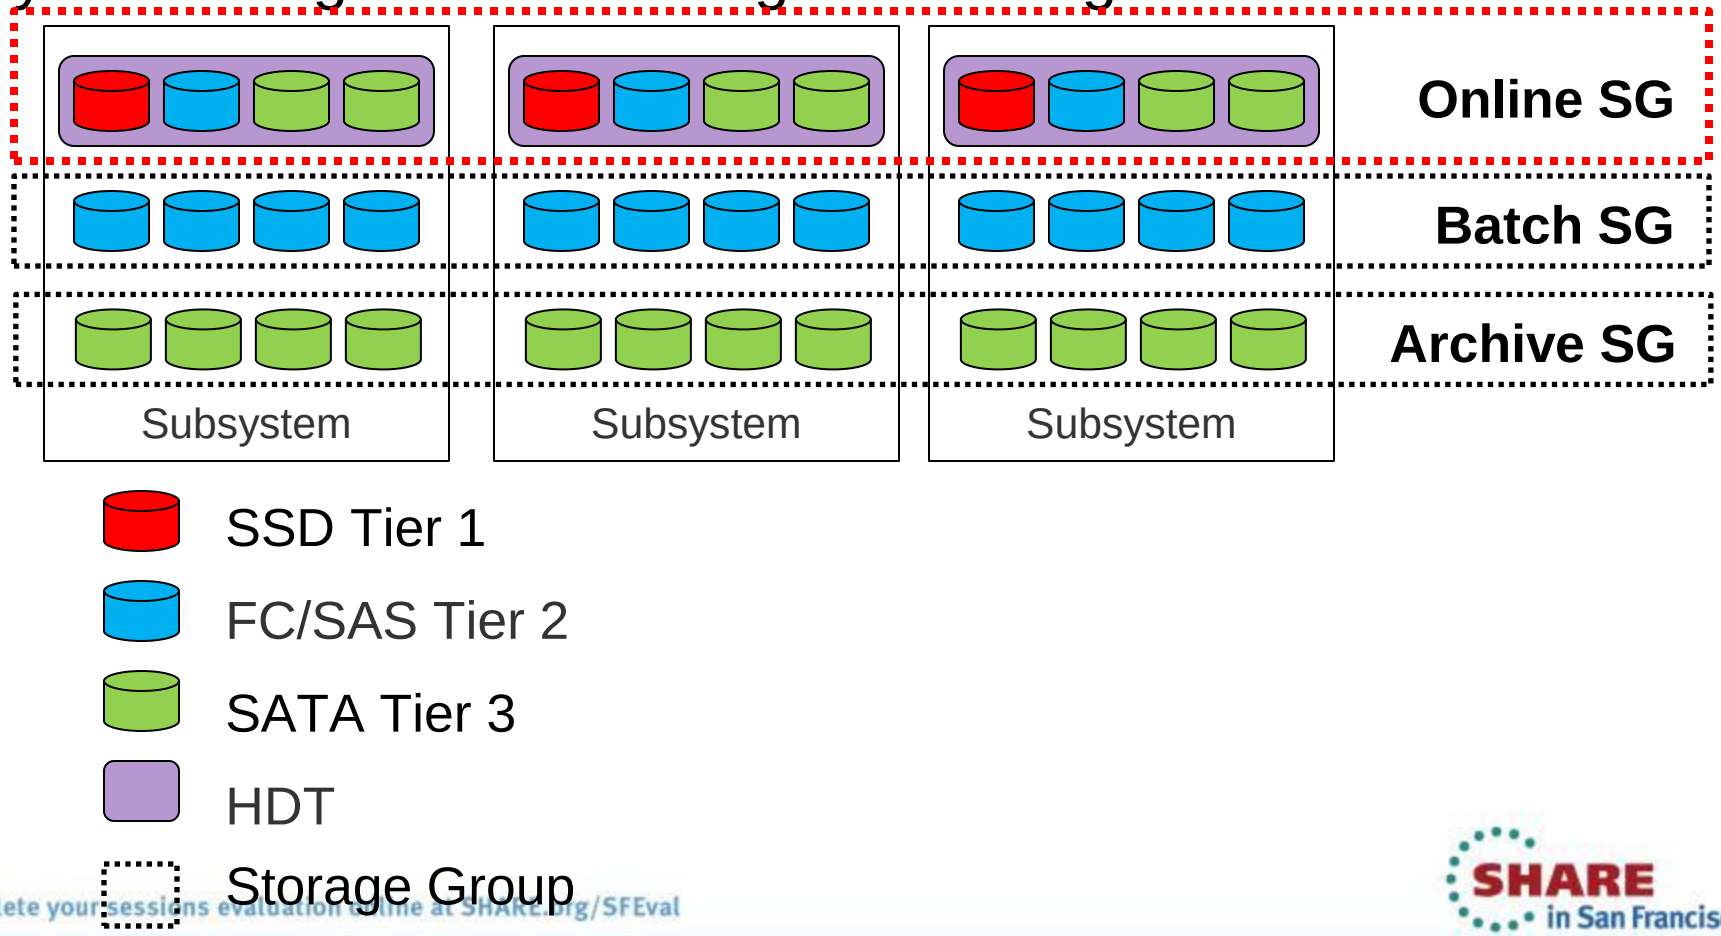

DFSMS Storage Groups and MF HDT

- HDT can be applied to selected Storage Subsystems only
- DFSMS Vertical Storage Groups and HDP Pools with Dynamic Page based Tiering can be aligned

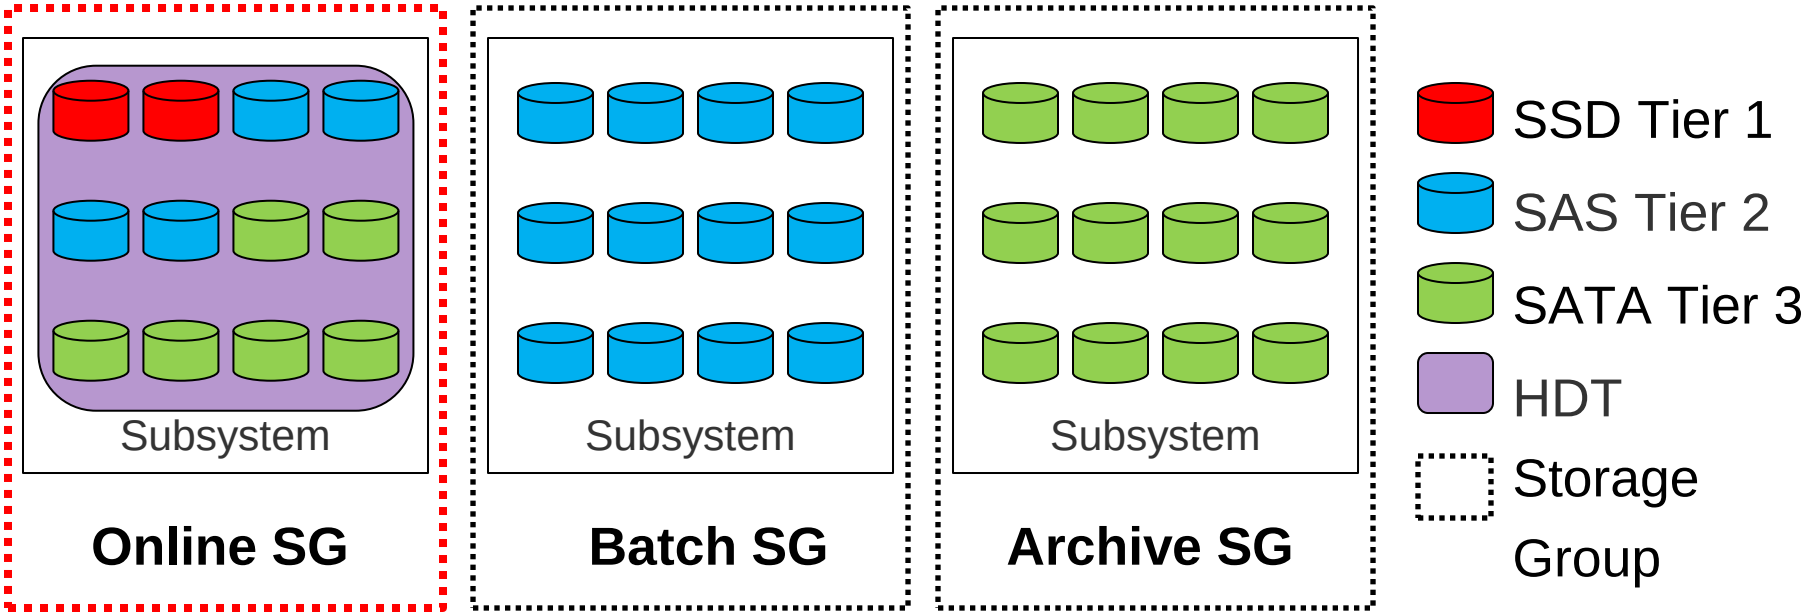

AGENDA

Virtual Storage Platform – Architecture Designed for Tiering

Hitachi Dynamic Provisioning (HDP)

Hitachi Dynamic Tiering (HDT)

- Summary

Tiered Infrastructure example

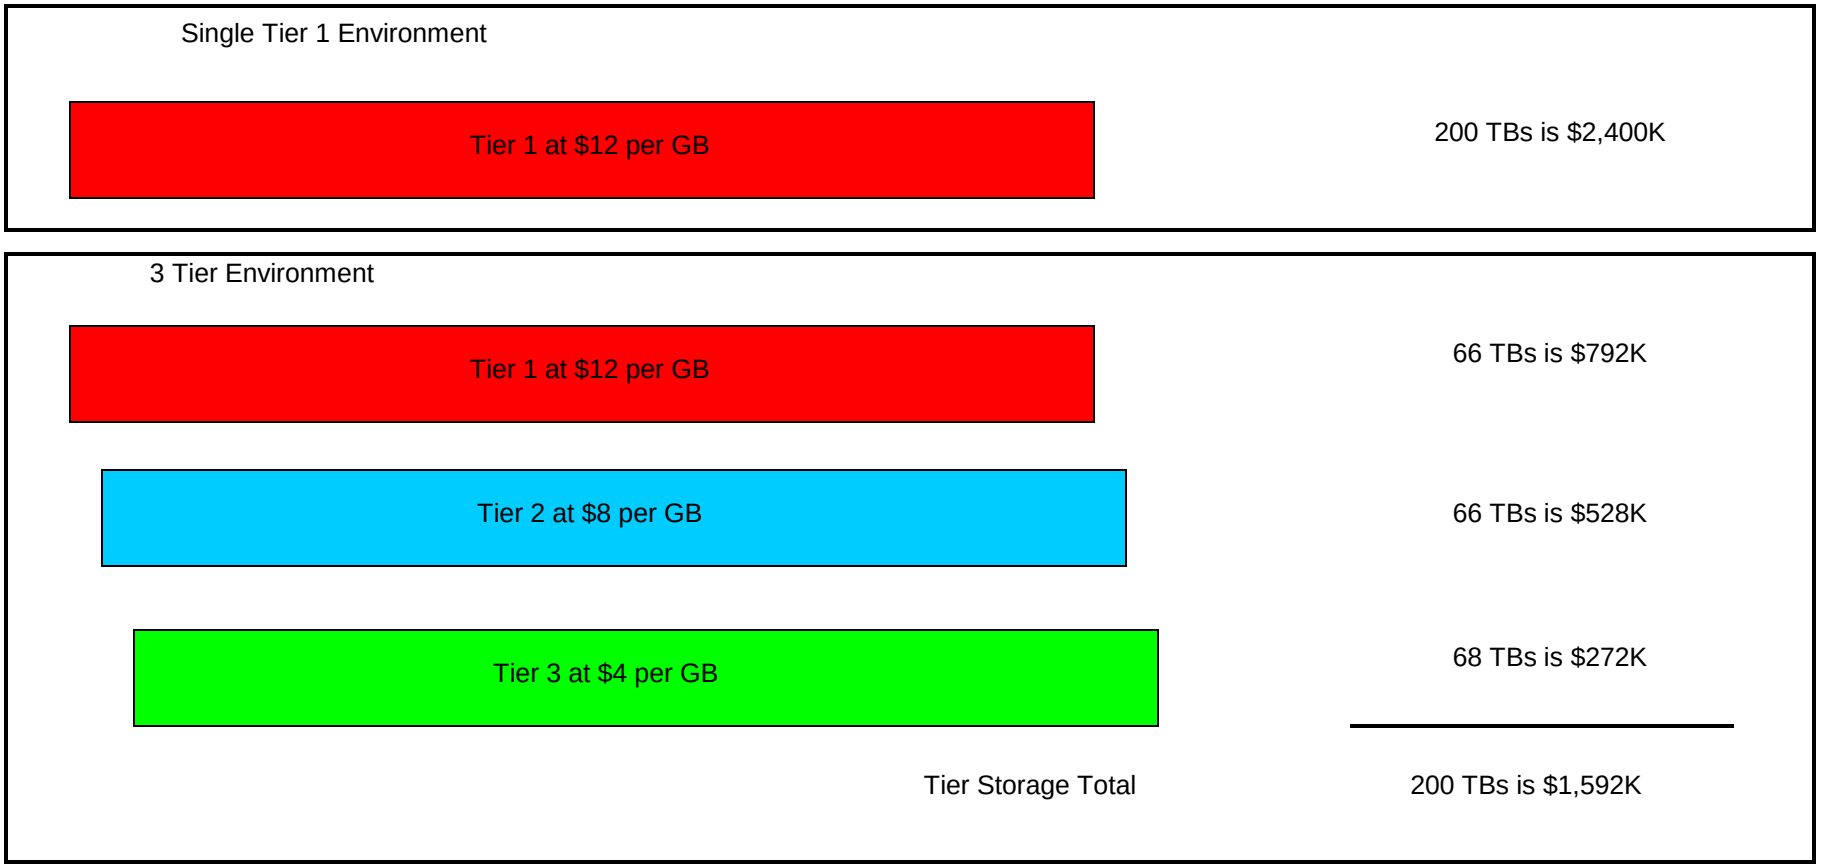

Tiering allows "Buy the same capacity of storage for less"

Summary

Let Tiering Learn your work Load

- Two different Page Sizes
  - 38MB Page Size for Mainframe
  - 42MB Page Size for OPEN
- Up to 128 unique pools
- All features and functions the same for OPEN and Mainframe
- Tiering Policies to better control data placement and movement
  - Default - All
  - Level 1 - Tier 1 only
  - Level 2 - Tier 1 and Tier 2 only
  - Level 3 - Tier 2 only
  - Level 4 - Tier 2 and Tier 3 only
  - Level 5 - Tier 3 only
  - 26 new policies recently added.
- Full support of Mainframe standard features including DFSMS

Thank You.