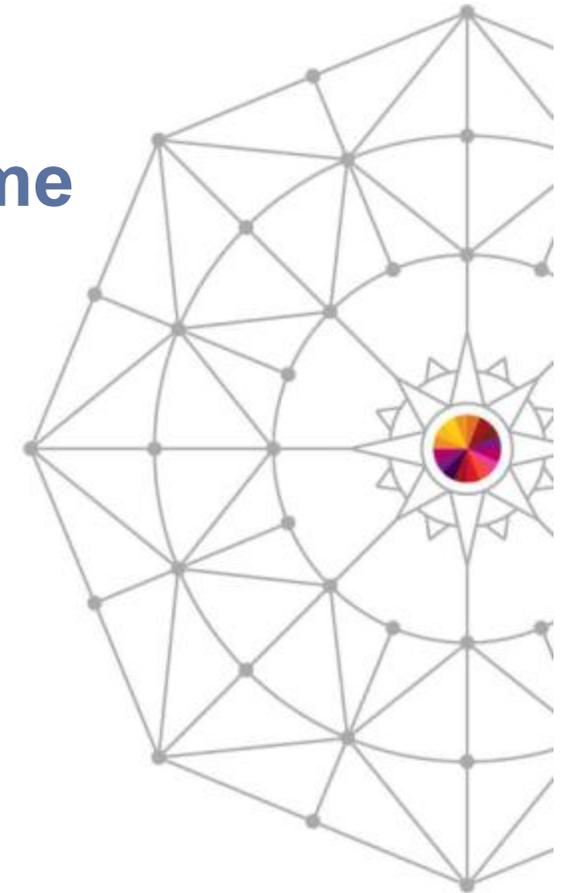




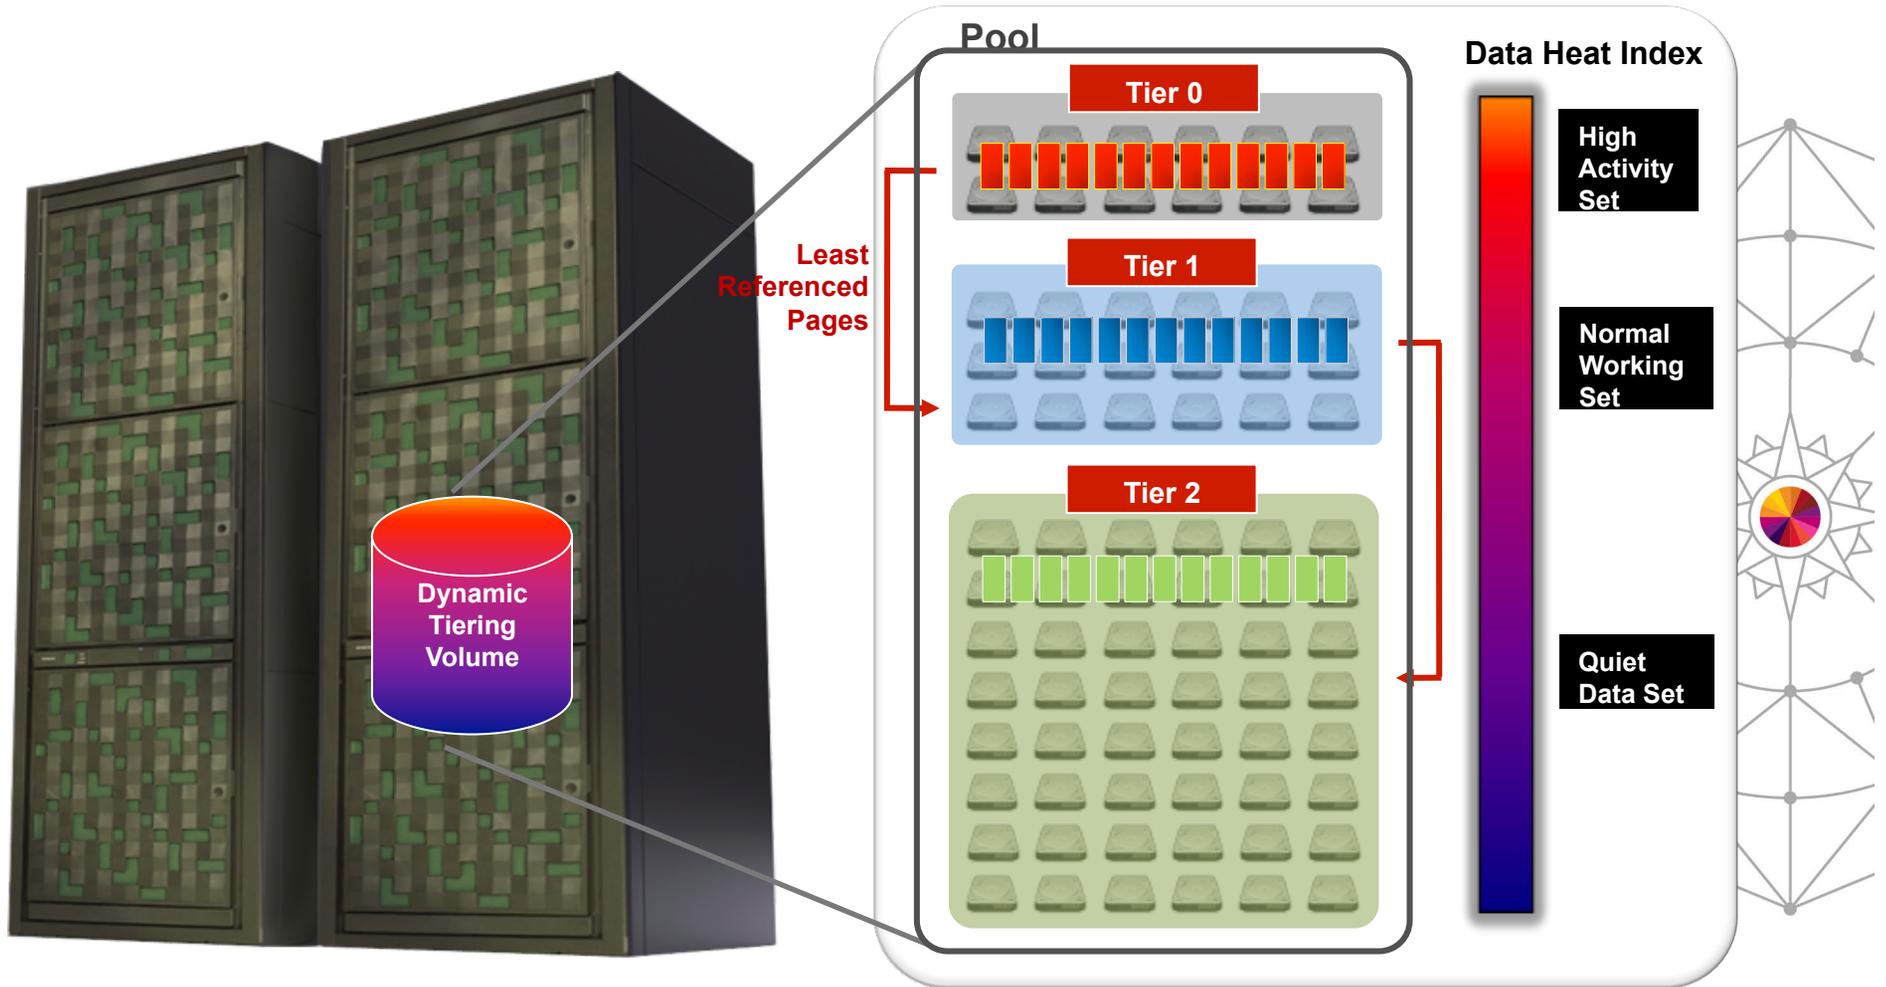
Hitachi Dynamic Tiering for mainframe *and* Hitachi Tiered Storage Manager for Mainframe

Ros Schulman
HDS (ros.schulman@hds.com)

March 11th, 2014
Session Number (14811)



Hitachi Dynamic Tiering (HDT)

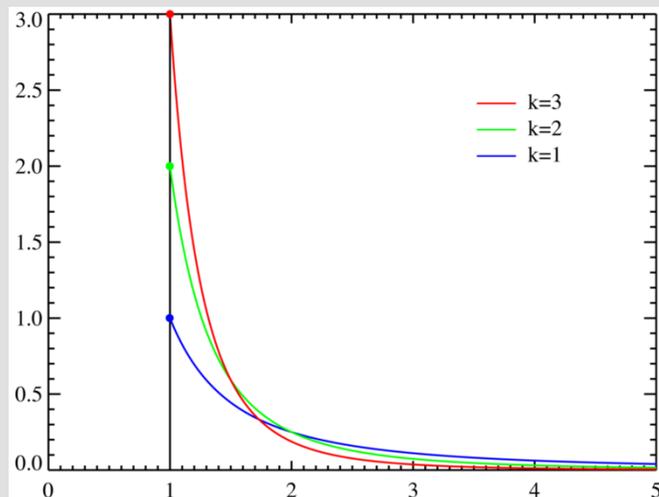


Complete your session evaluations online at www.SHARE.org/Anaheim-Eval

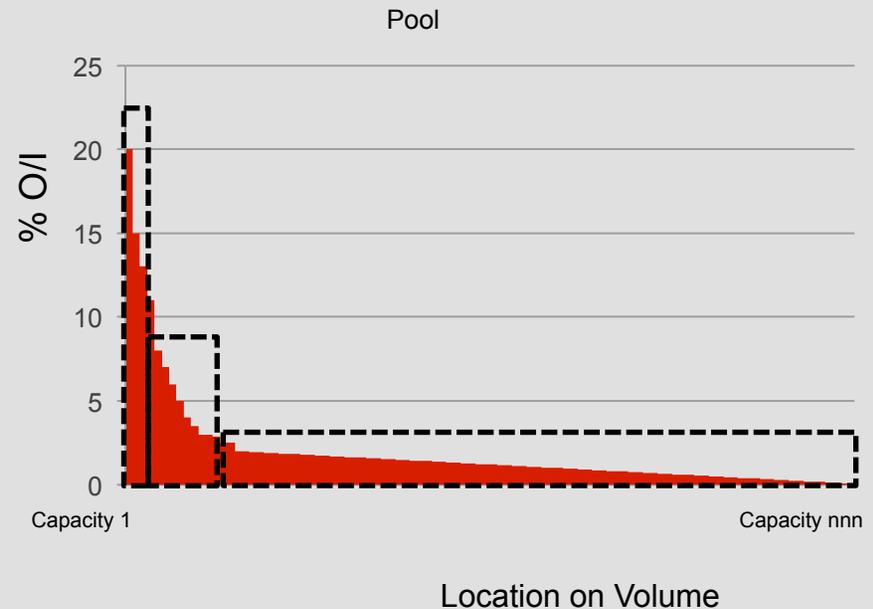


Improved Performance at Reduced Cost: Data Locality and Pareto Distributions

Classic Pareto Distributions
(Also Known as the 80/20 Rule)



Actual Volume Workload



Why does it work so well?

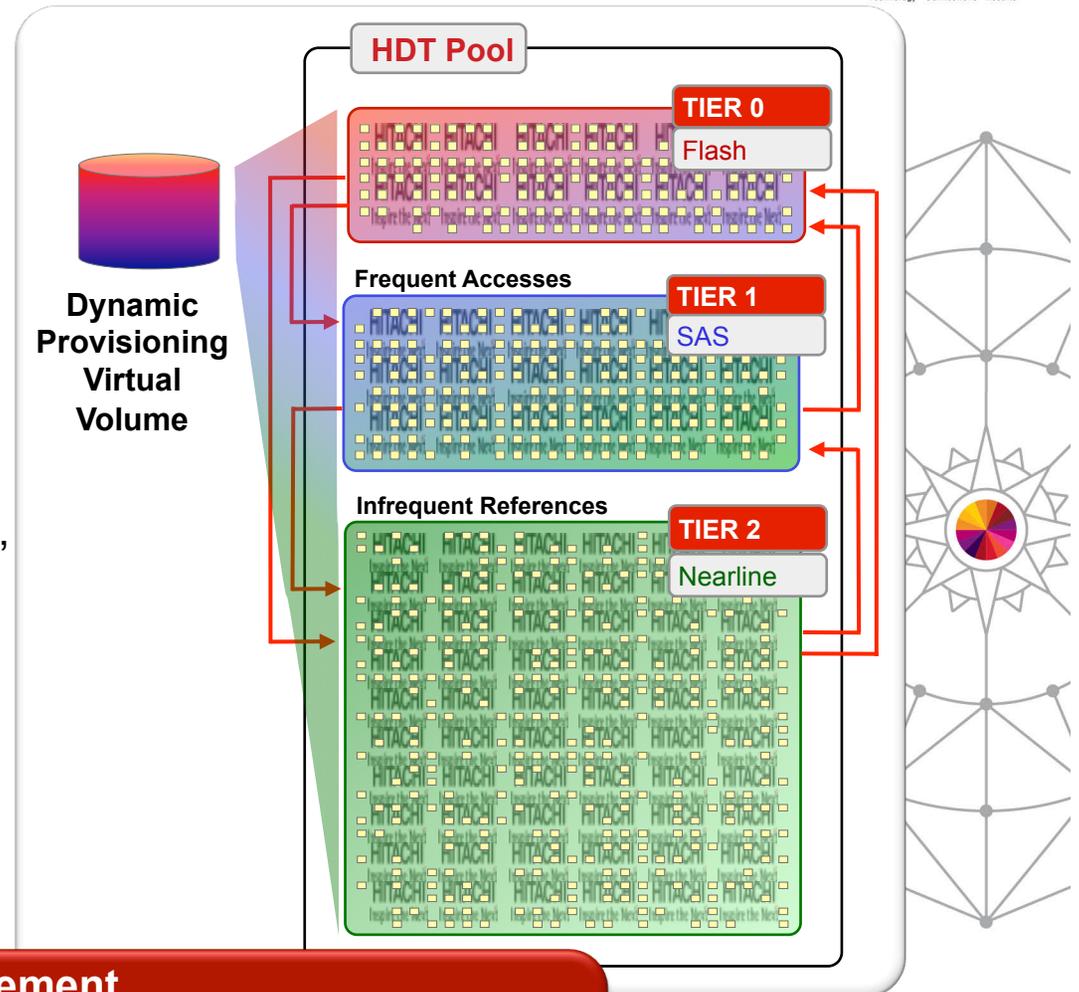
- Skew – At any time, only a small address range is active
- Persistence – When an address range is accessed it tends to remain so for a while

Hitachi Dynamic Tiering

Automated Optimized Tiered Storage Management



- **Before: Tiered storage and provisioning**
 - Labor intensive
 - Data classification before tiering
 - Complicated management of multiple storage tiers
- **Now: Dynamic tiering and provisioning**
 - Controller-based automation
 - Single, self-managed, self-healing, efficient pool of data
 - All the benefits of tiered storage
 - All the benefits of dynamic provisioning
 - No need for data classification



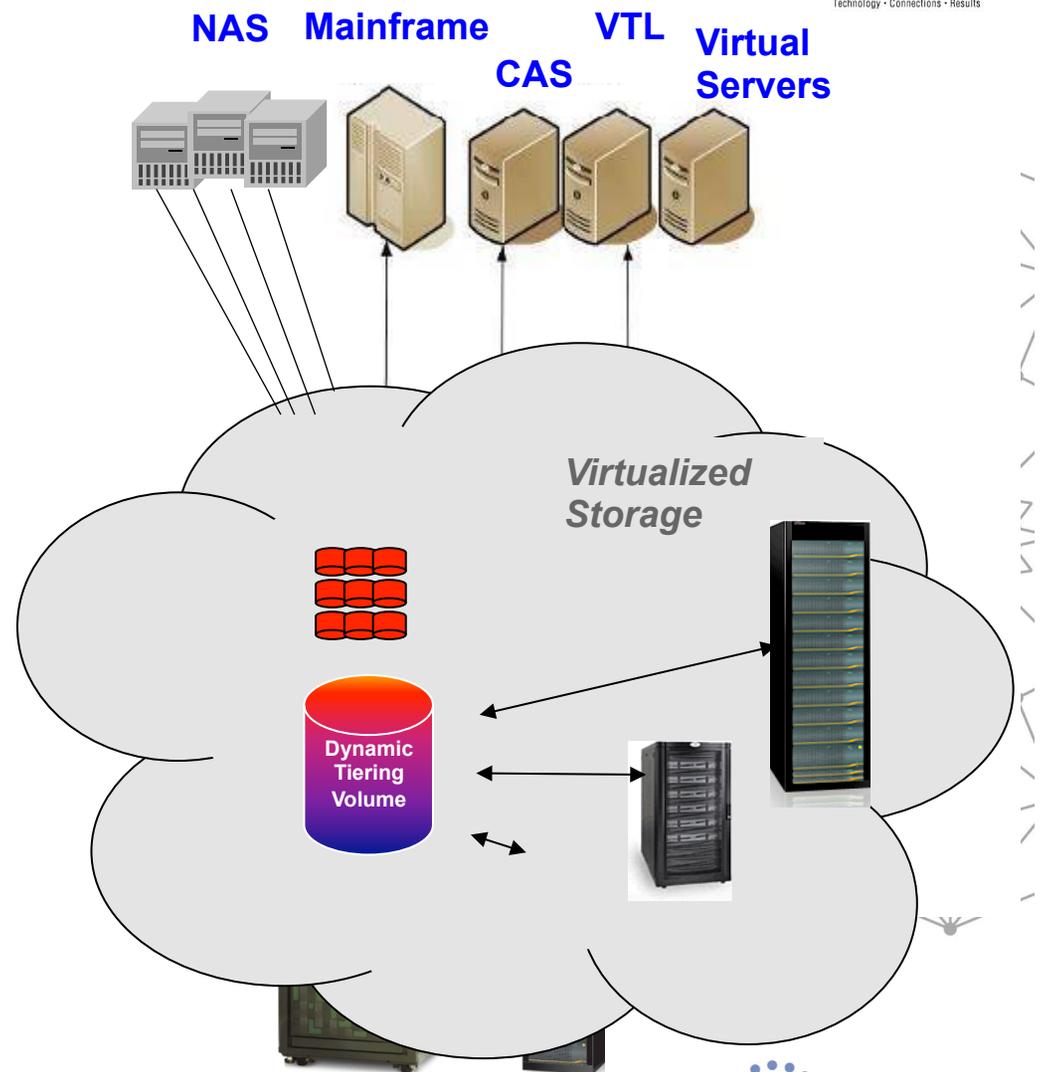
- **Simplifies operations and data management**
- **Reduces opex, capex, and TCO**
- **Data Moves in 38MB Pages**
- **Datasets can span multiple Tiers**



Hitachi Dynamic Tiering Supports Virtualized Storage



- With HDT, Hitachi Virtual Storage Platform (VSP) provides automated tiered storage management and performance acceleration for multiple tiers of heterogeneous external storage
 - As an SSD accelerator
 - As a storage virtualization controller
 - Heterogeneous storage business continuity and disaster recovery solution
 - Leveraging the breadth of offerings on VSP

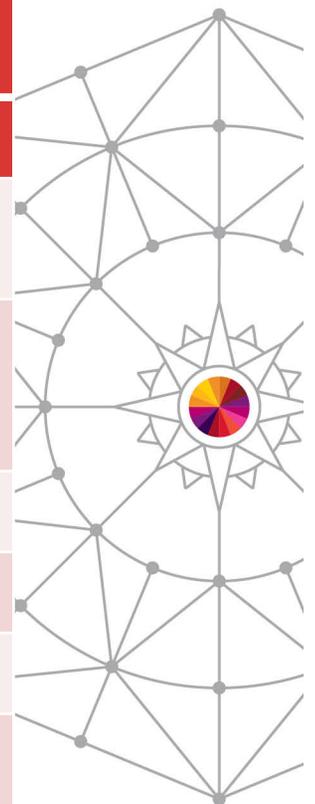


Relationship between number of tiers and Tiering policy



- Configurable DP Volume Tiering Policies

| Tiering | | Number of Tiers in Pool | | Note |
|---------|------------|-------------------------|--|---|
| Policy | 2 Tier | 3 Tiers | | |
| All | All tiers | All tiers | | Default value; data is stored in all tiers. |
| Level 1 | Tier 1 | Tier 1 | | Data is always stored in the highest-speed tier |
| Level 2 | All tiers* | Tier 1 and Tier 2 | | |
| Level 3 | All tiers* | Tier 2 | | |
| Level 4 | All tiers* | Tier 2 and Tier 3 | | |
| Level 5 | Tier 2 | Tier 3 | | Data is always stored in the lowest-speed tier |



* Data is stored in all tiers as in the case of All specified in the tiering policy



New Page Assignment policy



CONTROL WHICH TIER NEW PAGES ARE ASSIGNED TO

| New page allocation tier level | Description |
|--------------------------------|--|
| High | The new page is assigned from the higher tier of tiers set in the tiering policy |
| Middle | The new page is assigned from the middle tier of tiers set in the tiering policy |
| Low | The new page is assigned from the lower tier of tiers set in the tiering policy |



Complete your session evaluations online at www.SHARE.org/Anaheim-Eval

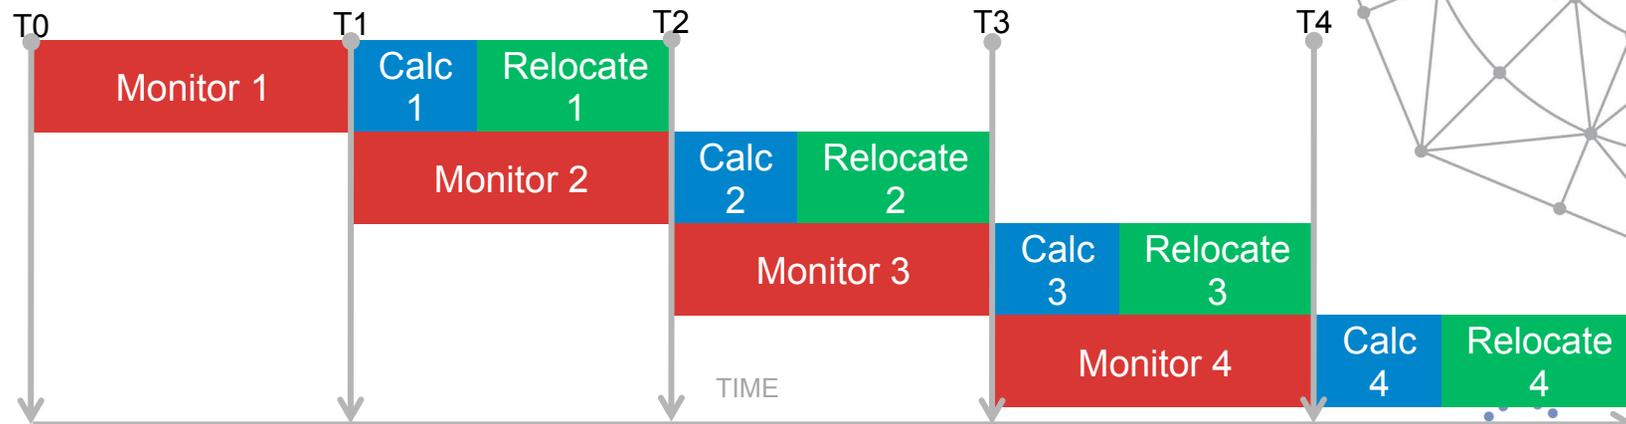


HDT Cycles



HOW TIERING LEARNS YOUR WORKLOAD

- Cycle time set at the HDP pool level
- Manual mode
 - User can start and stop performance monitoring using any interval up to 7 days
 - Scripting can set complex schedules to custom fit to priority work periods
- Automatic mode
 - Customer defines strategy; it is then executed automatically
 - Continuous monitoring followed by relocation cycles
 - Monitor interval from 30 minutes to 1, 2, 4, 8 or 24 hours (default)



Complete your session evaluations online at www.SHARE.org/Anaheim-Eval

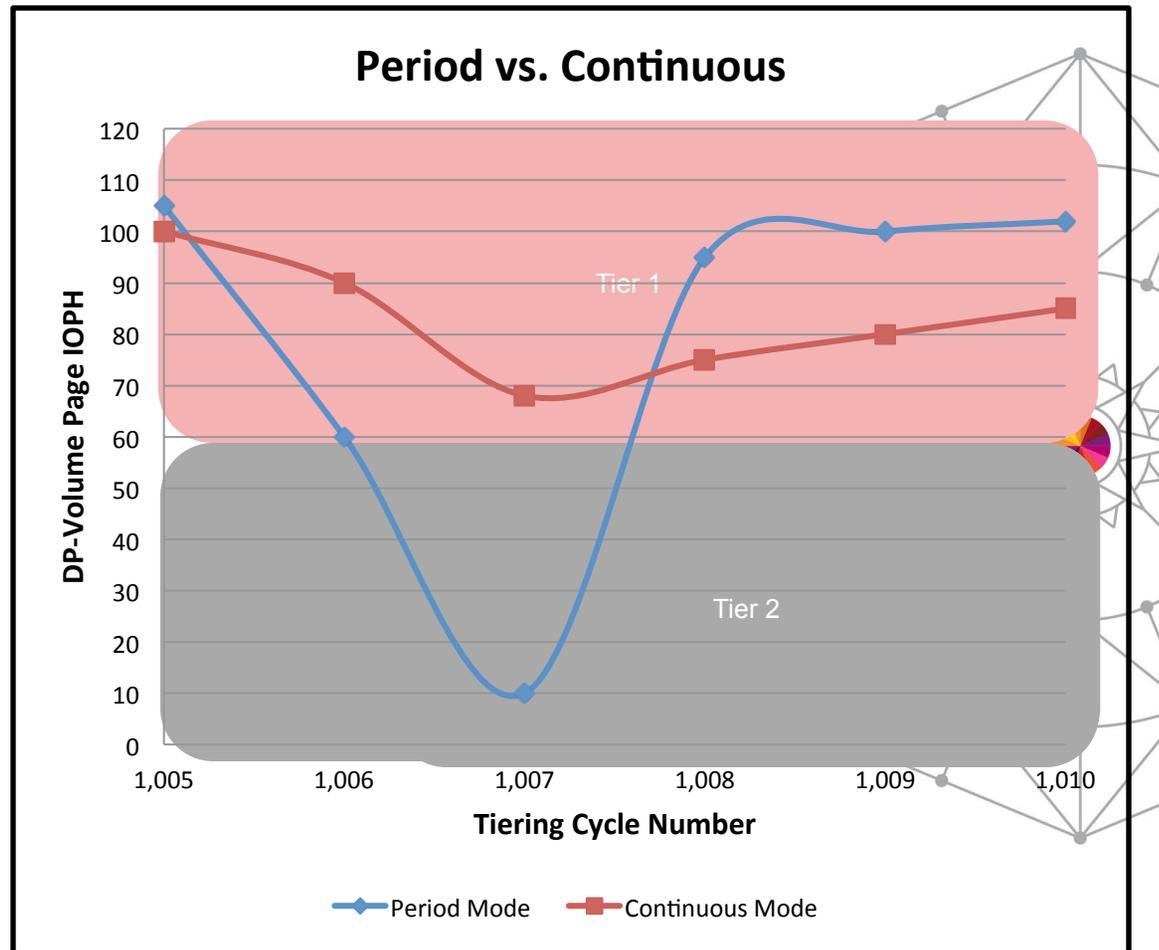


HDT Monitoring Modes



CONTROL HOW TIERING LEARNS YOUR WORKLOAD

- Monitoring modes set at HDP/HDT pool level
- Period mode
 - The value used in the calculation cycle is the actual I/O load on DP volume page from previous monitoring cycle
- Continuous mode
 - The value used in the calculation cycle is the weighted average of multiple previous monitoring cycles for DP volume page
 - Reduces page trashing
 - May slow migration to upper tiers



Hitachi Dynamic Tiering for Mainframe



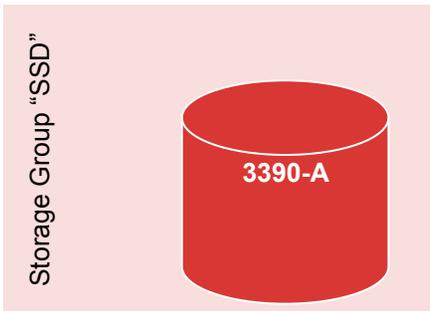
Hitachi Dynamic Tiering for Mainframe and DFSMS



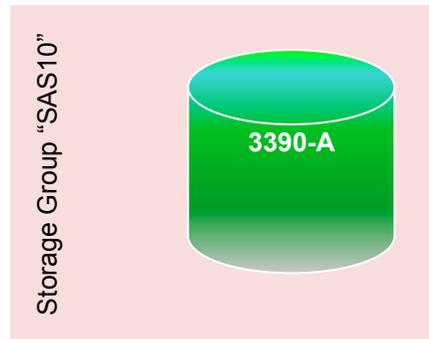
Complete your session evaluations online at www.SHARE.org/Anaheim-Eval



LIMITS OF SMS Storage Groups and ACS Routines

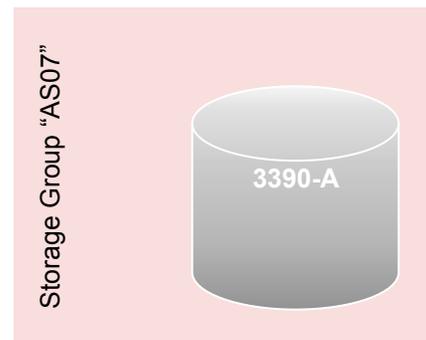


- 3390 volumes are "fixed" to a single tier
- To transition a volume's data to another tier requires ACS work, then migration and recall

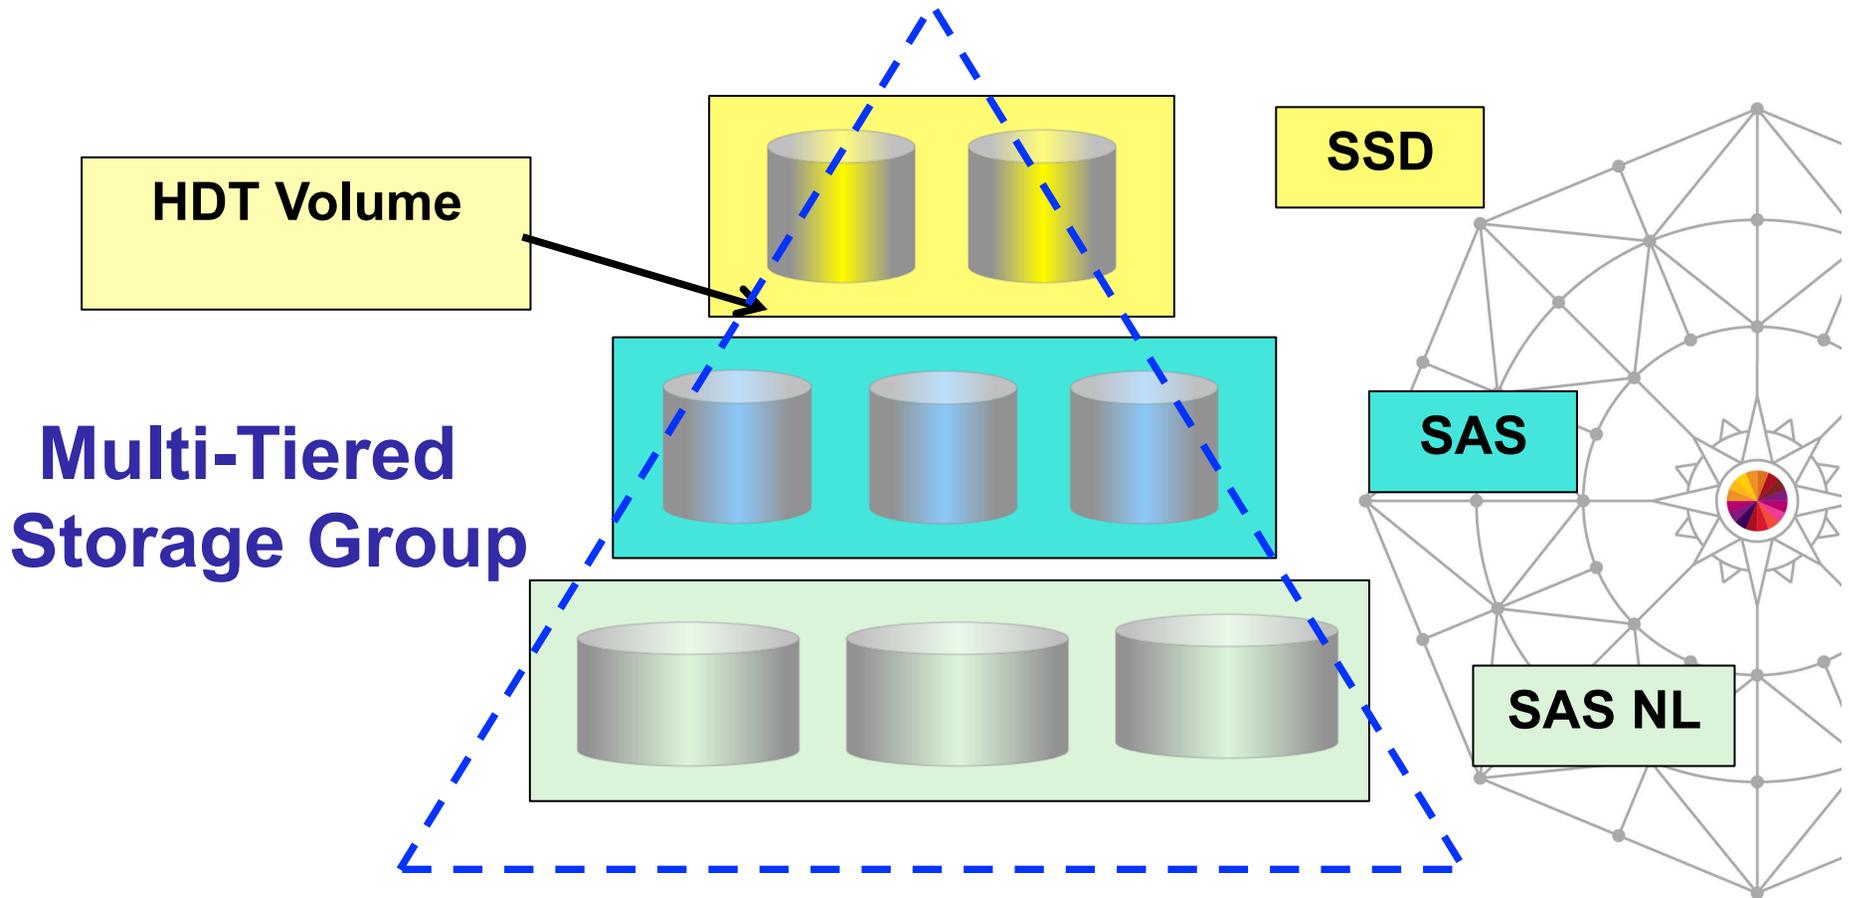


- Stale datasets are treated the same as active ones until HSM migration

- Performance problems need intervention to migrate to "higher" storage group
- Host-based volume movement has high overhead cost

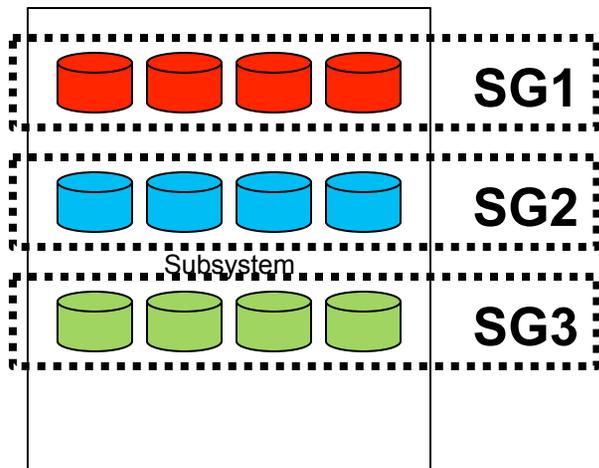


New! HDT Volume-provisioned storage group



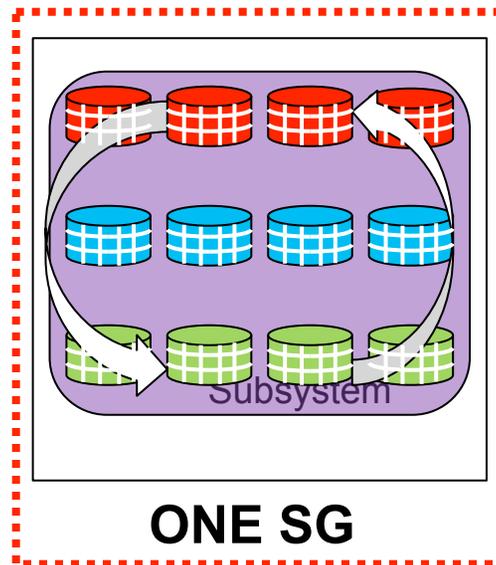
DFSMS

Storage Groups (SG) and HDT for Mainframe storage



- SSD Tier 1
- SAS Tier 2
- NL SAS Tier 3
- HDT Pool
- Storage Group

- Now with HDT for Mainframe storage, all storage tiers may be combined into a single storage group
- HDT automatically moves data within a storage group to the correct tier based on Workload performance



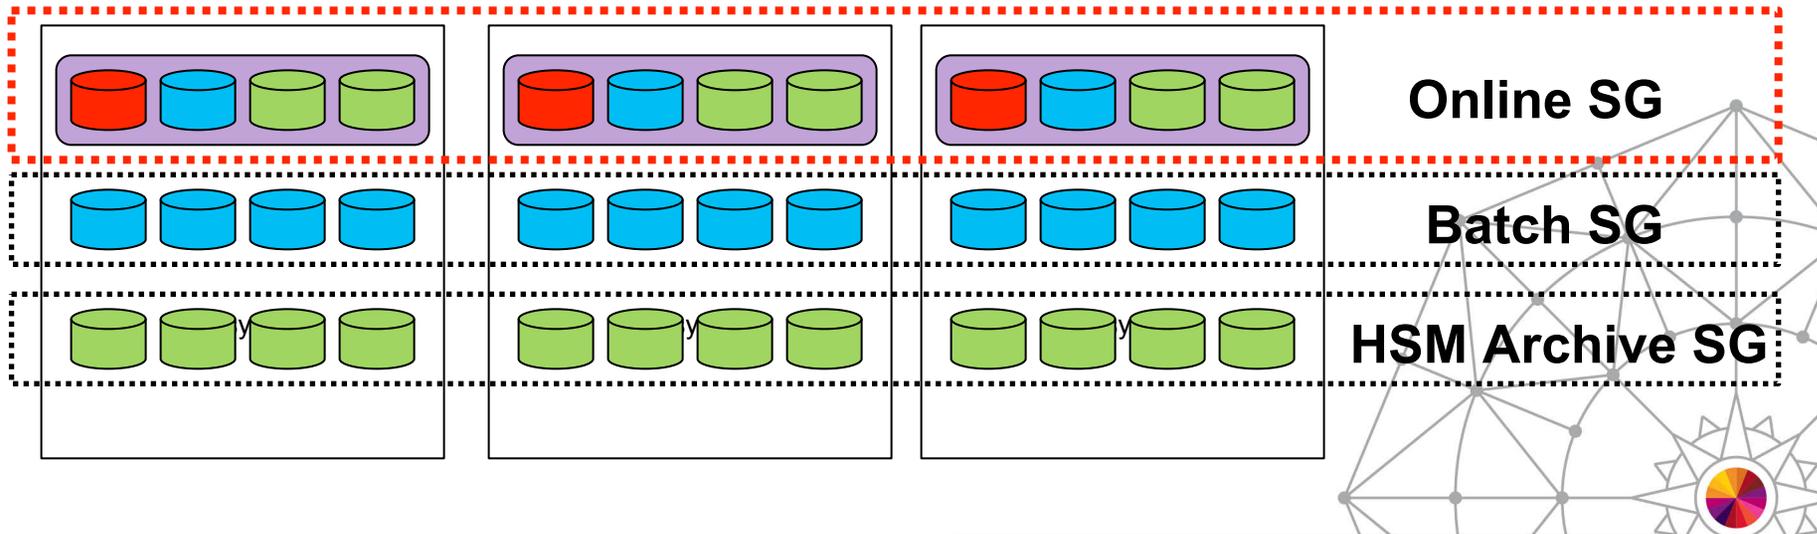
HDT Volumes are optimally managed dynamically at the page level

HDT Operational Impact



| Working with Storage | Before HDT | With HDT |
|--|--|---|
| Add physical capacity | Add 3390-X volumes into storage groups | Add capacity into pool |
| Balance use over new capacity | Manually use HSM migration/recall | No actions are needed |
| Direct specific applications to specific storage resources | Code ACS routines, follow-up with HSM migrations and recalls | Set 3390-A to an HDT policy – use same ACS routines but no HSM needed |
| Address performance problems by moving datasets or volumes | Code ACS routines and use HSM migration/recall | HDT relocation has likely prevented the issue; otherwise use HDT policy |
| Maintain SMS storage groups and ACS routines | Manual effort to keep updated with rules describing exceptions | Fewer exceptions since HDT keeps tiers properly populated |
| Demote data to lower tiers | HSM moves <i>datasets</i> to an ML “tier” that hasn’t been <i>opened</i> for a while | HDT automatically moves <i>pages</i> that haven’t been <i>used</i> |

DFSMS Storage Groups and HDT for Mainframe Storage



-  FLASH Tier 0
-  SAS Tier 1
-  NL SAS Tier 2
-  HDT Pool
-  Storage Group

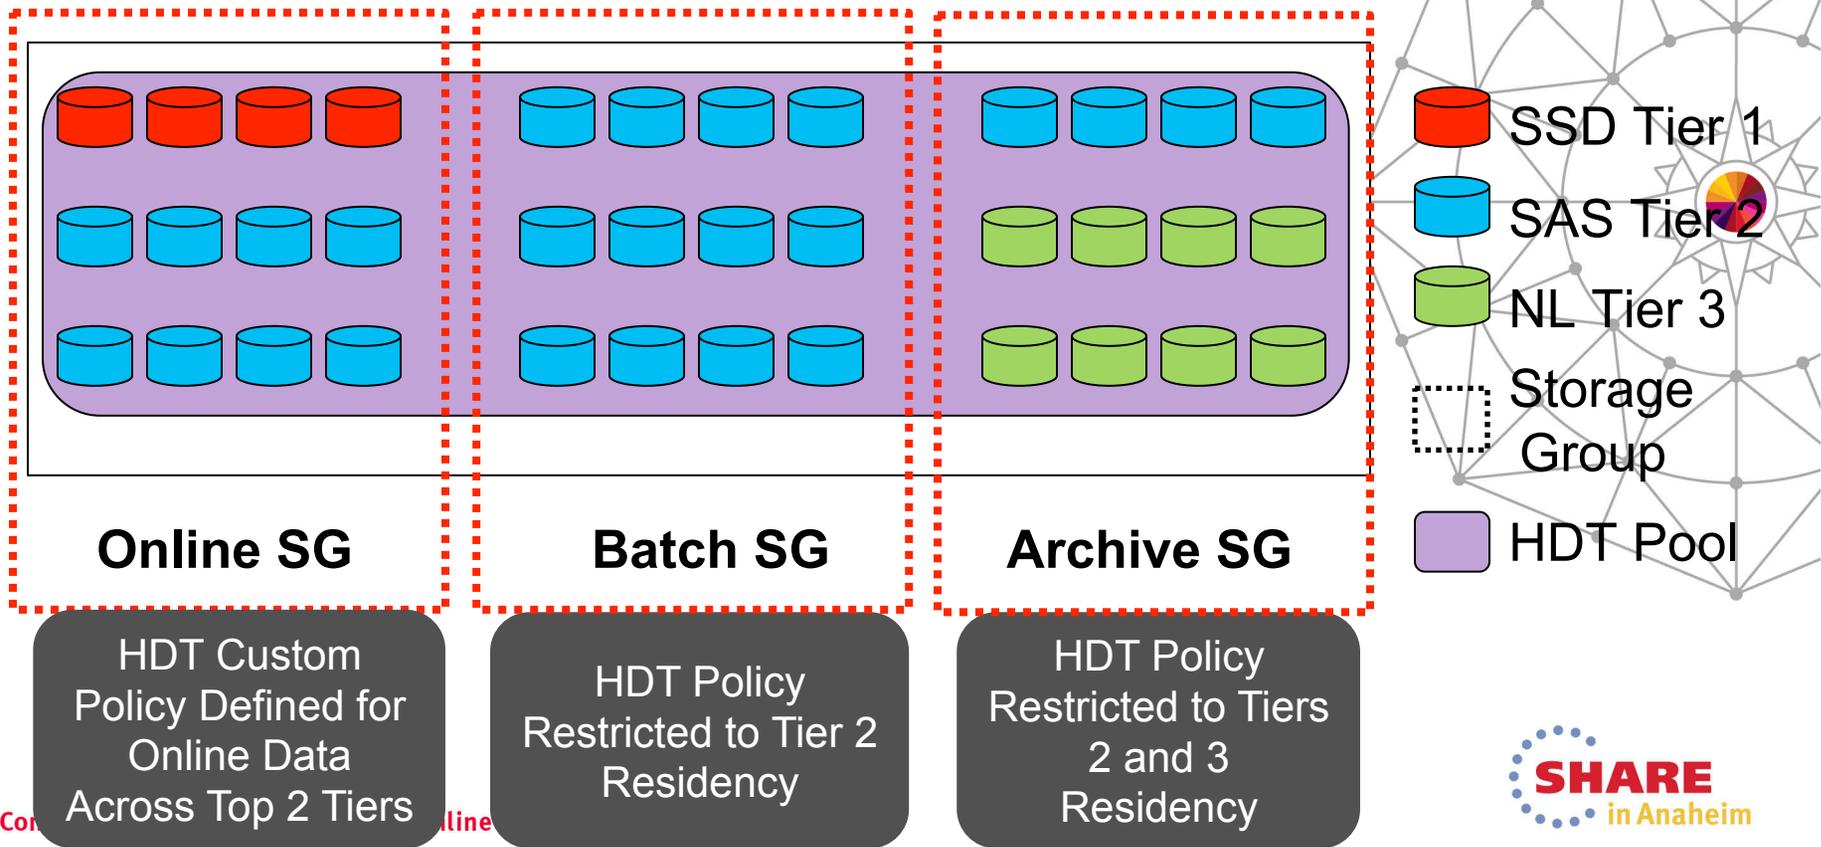
- Simplifies integration
- HDT can be applied to selected Storage Groups only
- DFSMS Horizontal Storage Groups and Dynamic Page based Tiering volumes can be aligned

DFSMS

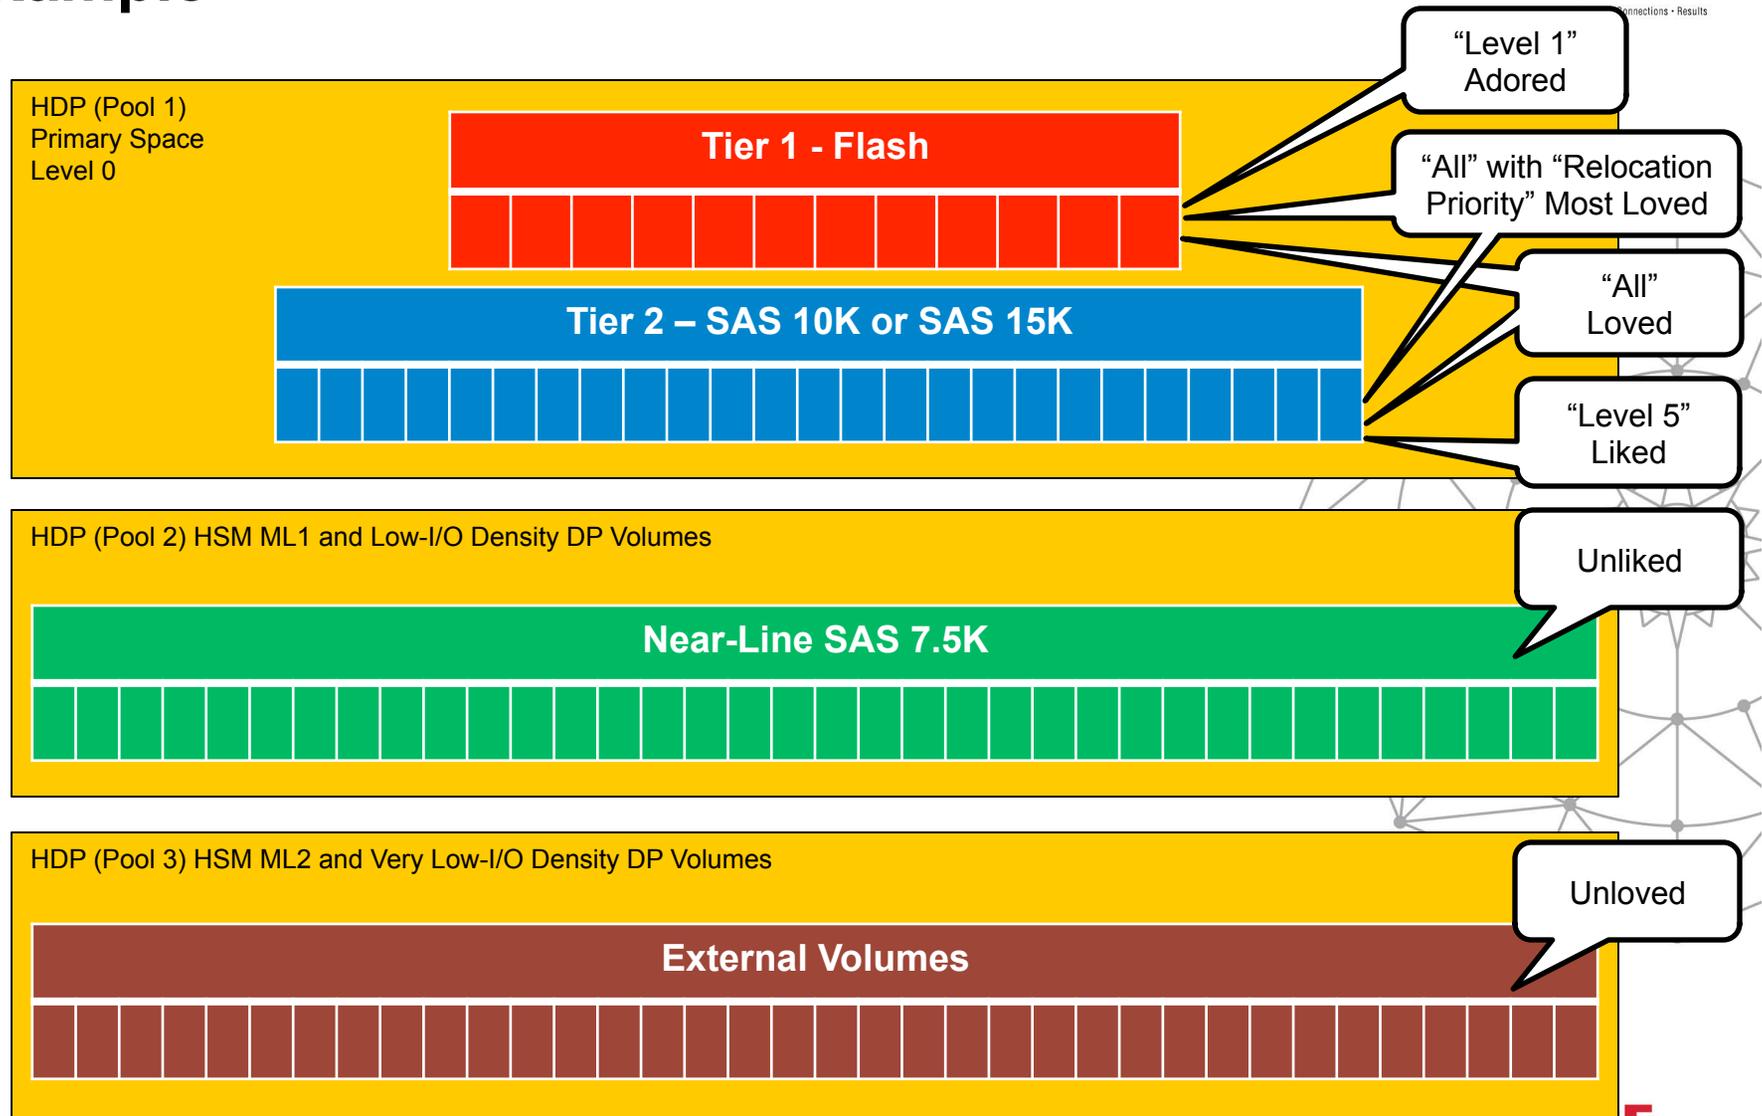
Storage Groups and HDT for mainframe storage



- With HDT for Mainframe storage policies, individual policies can be defined for volumes mapped to different storage groups
- Policies are supported based on tier ranges, analysis/migration periods, initial tier page assignments and relocation priority



HDT and DFSMS/HSM integration Example



Hitachi Tiered Storage Manager for Mainframe



Hitachi Dynamic Tiering for Mainframe Management



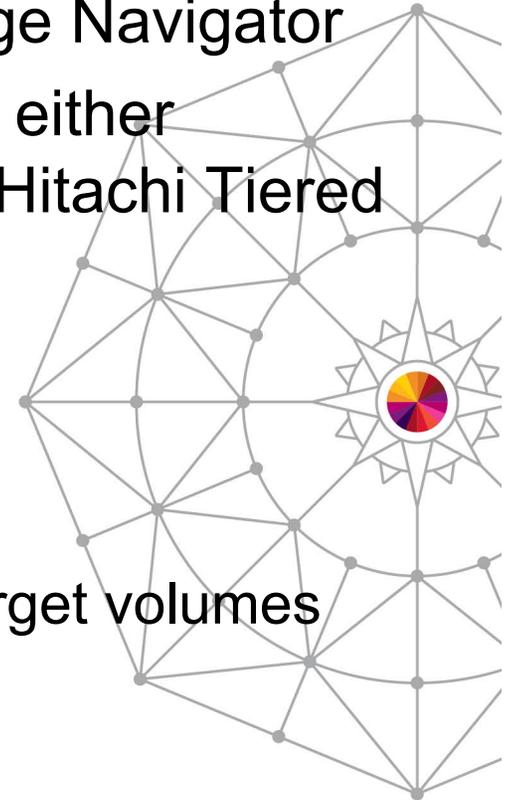
Complete your session evaluations online at www.SHARE.org/Anaheim-Eval



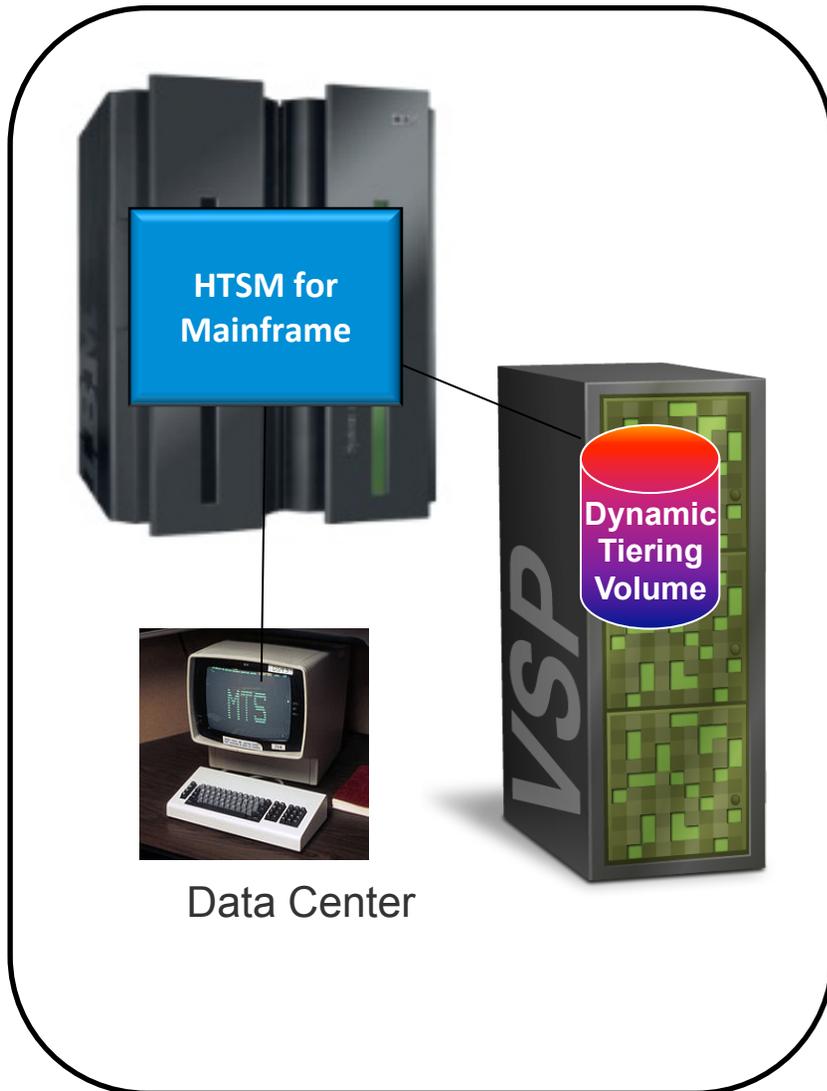
Options for Managing Dynamic Tiering for Mainframe



- Setup and maintenance operations are done via VSP Service Processor (SVP) with Hitachi Storage Navigator
- Operational management can be done from either Storage Navigator or native from z/OS with Hitachi Tiered Storage Manager for Mainframe (HTSM)
- HTSM features
 - Storage SLA management via HDT policies
 - Native mainframe operations
 - Group operations make it easy to manage target volumes
 - Linkage/integration with DFSMS
 - Customization via REXX scripting
- **ISPF interface – April 2014 NEW**



Hitachi Tiered Storage Manager for Mainframe – Z/OS HDT management



Host-based software that provides:

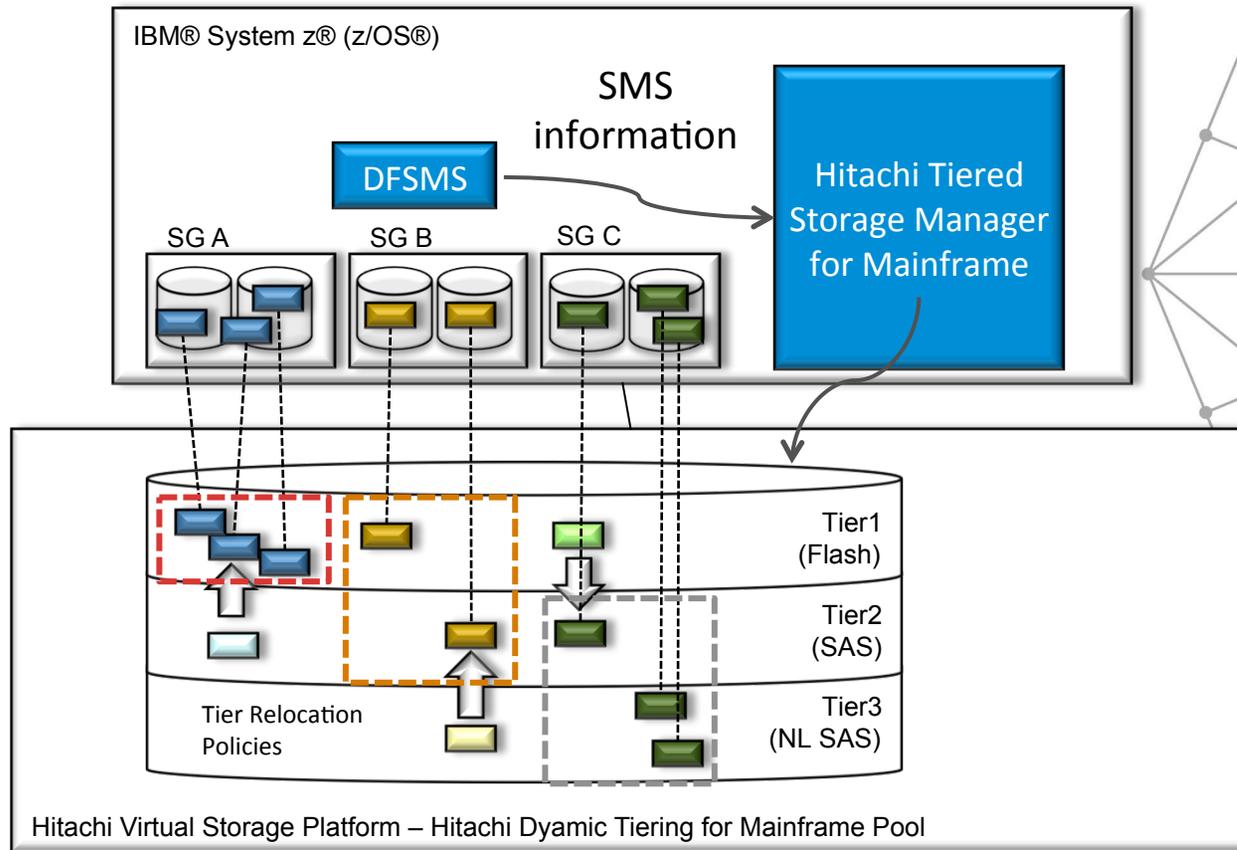
- Centralized and unified mainframe management of Hitachi Dynamic Tiering
 - Automation
 - Integration with DFSMS and storage groups
- Online storage service level controls
 - Increase application performance
 - Improves problem avoidance
- Single, consistent interface
 - Command based, script driven
- Auto-discovery eliminates errors
 - Accelerates deployment
- Enables reporting and automatic notifications



What is Tiered storage Manager for Mainframe?



- Hitachi Tiered Storage Manager for Mainframe
 - Manages operational aspects of HDT, including tiering policies
 - Integrates with z/OS SMS (storage groups)

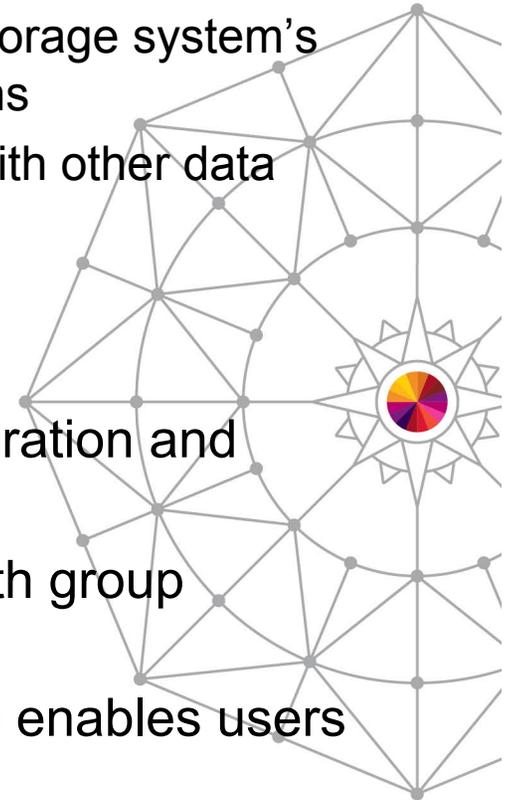


Why HTSM for mainframe?



HTSM provides important benefits

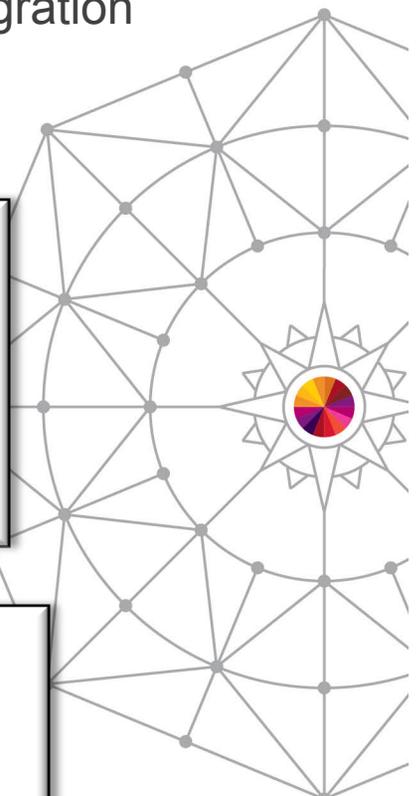
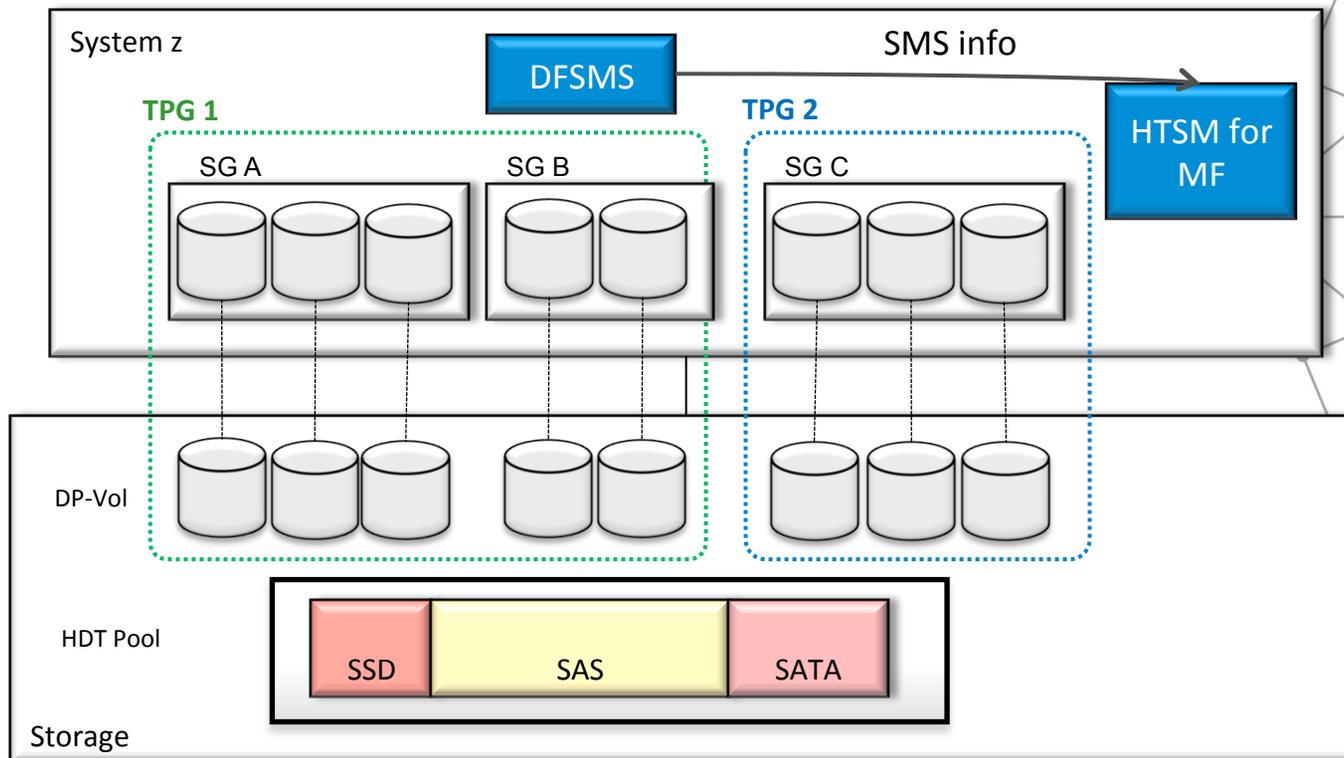
- Native management from z/OS
 - Control for data location from host's point of view, not storage system's point of view – improves control and simplifies operations
 - Ability to utilize HDT performance and relocation data with other data (such as SMF records)
 - Reduce dependency on open server-based operations
- Control of storage service levels using HDT policies
- Linkage with z/OS SMS (storage group) speeds integration and reduces opex
- Simplify management in large-scale environments with group operations
- Flexible command-line interface (CLI) (TSO/E REXX) enables users to get the most out of Dynamic Tiering



Linkage with z/OS SMS Feature



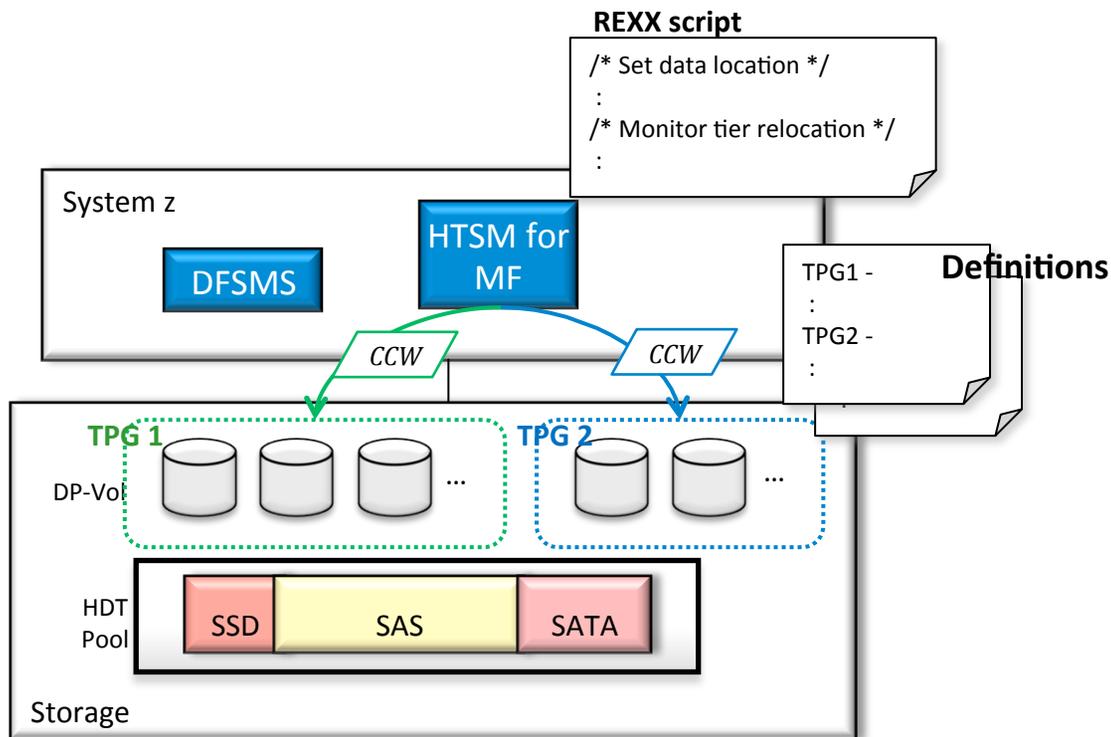
HTSM works with SMS construct (storage group). Each tiering policy group can have 1 or more storage groups. This feature gives users the capability to manage the HDT environment from a z/OS point of view. SMS integration makes it easy to add Dynamic Tiering to existing SMS operations



Customization via REXX Scripting



- HTSM provides a CLI (TSO/E REXX) that enables users to manage their environments. The CLI also has the ability to cooperate with OS services (such as TSO/E service) and other products.
- For improved manageability HTSM for MF has separated operation procedures (REXX script) and configuration file-based objects (such as group definition)



HTSM Mainframe Reporting



- TPG_QUERY_STATISTICS Sample Output

```

Tiering Policy ID: PROD
Date: 10 Jul 2013
Time: 00:07:23

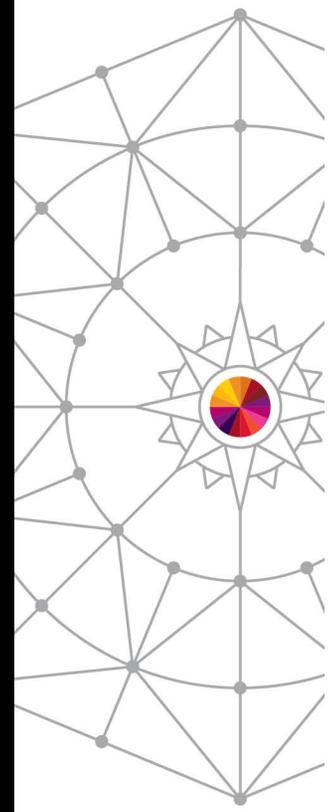
***** Query TPG Tier Metrics *****

=====
TPG Total or          Tier1      Tier2      Tier3      Total
SN:PoolID or         Used      Used      Used      Used
*StorGrp* or         Pages    Pages    Pages    Pages
Volser      or      / UsedGB  / UsedGB  / UsedGB  / UsedGB
Volser Prefix / Used%  / Used%  / Used%  / Used%
=====
TPG Total              340         0         0         340
                    12.9GB      0GB       0GB       12.9GB
                    100%        0%        0%        100%
-----
SN53004:81             340         0         0         340
                    12.9GB      0GB       0GB       12.9GB
                    100%        0%        0%        100%

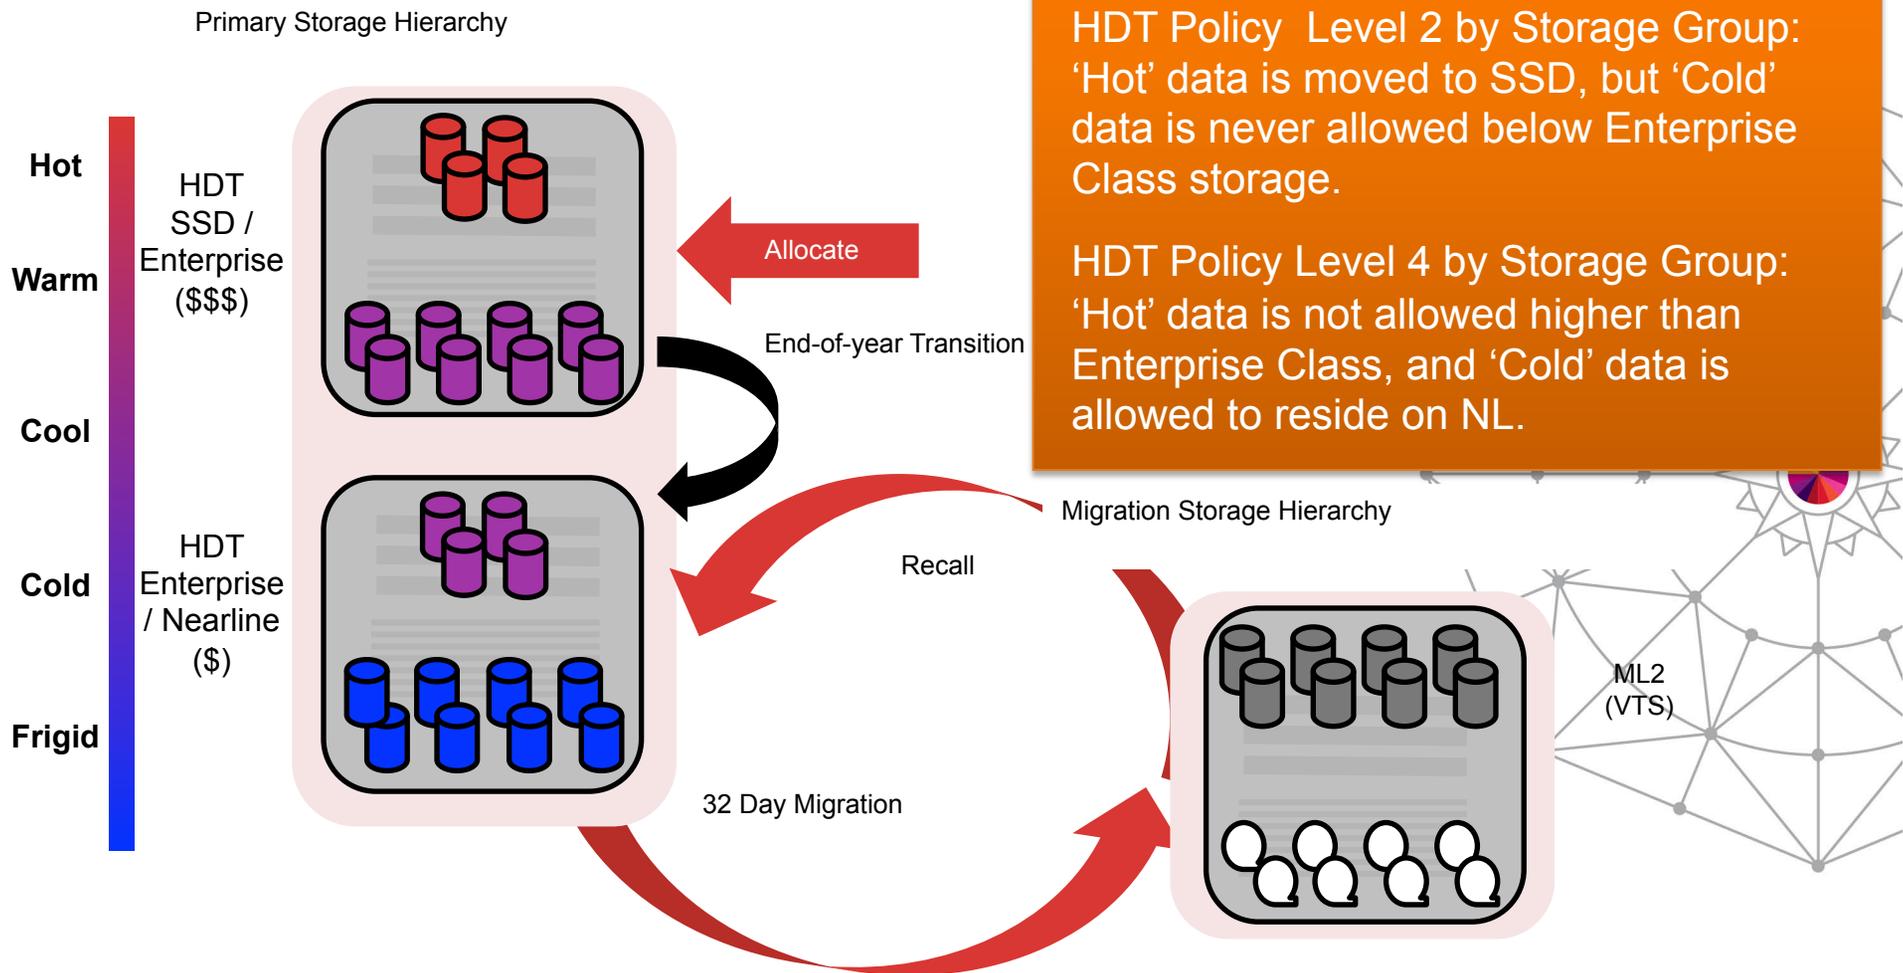
    Used% of Pool      25.4%        0%        0%        8.67%

    Pool Pages         1340        1240        1340        3920
    Pool GB            50.9GB     47.1GB     50.9GB     149GB
    Tier% of Pool      34.2%      31.6%      34.2%      100%
-----
*ALPHA*                340         0         0         340
                    12.9GB      0GB       0GB       12.9GB
                    100%        0%        0%        100%
-----
GSE*                   340         0         0         340
                    12.9GB      0GB       0GB       12.9GB
                    100%        0%        0%        100%
-----

10 Jul 2013 00:07:23 *** Action TPG_QUERY_TIERS Successful
    
```



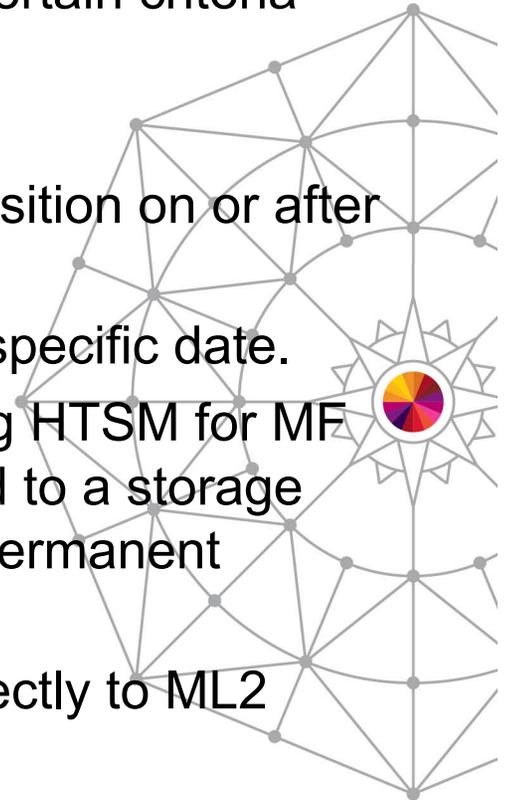
Hitachi Dynamic Tiering and HTSM for MF with DFSMShsm 2.1



DFSMSHsm 2.1 AND HTSM FOR MF



- Datasets may be eligible for class transitions when certain criteria are met specified by the management class
 - Time since Creation
 - Time since Last Use: Data set is eligible for a transition on or after this time
 - Periodic: Data set is eligible for a transition on a specific date.
- Data sets could be allocated to a storage group using HTSM for MF that uses Tier 1 and Tier 2 and then later transitioned to a storage group that uses less expensive class of storage for permanent retention
 - After a certain period of time, can be migrated directly to ML2



Hitachi Dynamic Tiering for Mainframe



Hitachi Dynamic Tiering for Mainframe Use Case Slides



Complete your session evaluations online at www.SHARE.org/Anaheim-Eval



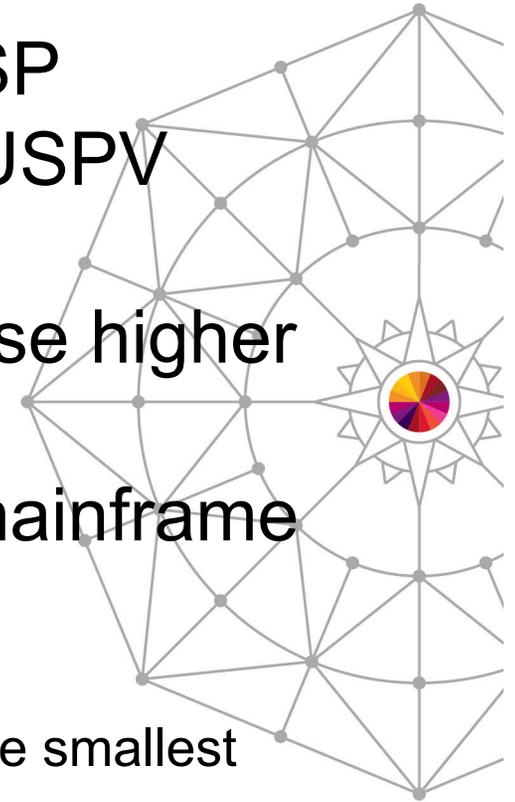
VSP – USPV Performance Comparison



Study Objectives

- Establish baseline performance of VSP
- Determine performance level versus USPV environment
- Prove the capability of being able to use higher density drives
- Prove the feasibility of using HDT in mainframe environment

Are we maximizing the greatest number of IOs with the smallest amount of SSDs?



ENVIRONMENTAL COMPARISON



USPV – Total Cap 231TB

| SN | (kg) | (lb) | (kW) | BTU/hr | kVA | SQ FT | Frames |
|----|-------|-------|-------|--------|-------|-------|--------|
| | 3,227 | 7,099 | 23.93 | 81,713 | 24.93 | 33.64 | 5 |
| | 3,267 | 7,187 | 23.39 | 79,899 | 24.36 | 33.64 | 5 |



VSP – Total Cap 682TB

| (kg) | (lb) | (kW) | BTU/hr | kVA | SQ FT | Frames |
|-------|-------|-------|--------|-------|-------|--------|
| 2,893 | 6,371 | 17.80 | 60,666 | 18.70 | 35.63 | 5 |



SAVINGS

| | (kg) | (lb) | (kW) | BTU/hr | kVA | SQ FT | Frames |
|-----------|--------|--------|--------|--------|--------|--------|--------|
| VSP | 55.45% | 55.40% | 62.38% | 62.46% | 62.06% | 47.04% | 50.00% |
| % Savings | | | | | | | |

With 3X+ Capacity

Complete your session evaluations online at www.SHARE.org/Anaheim-Eval

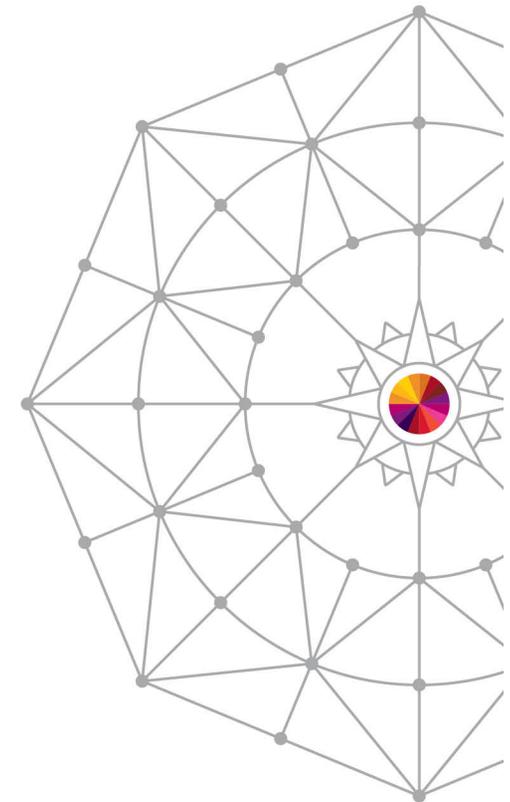
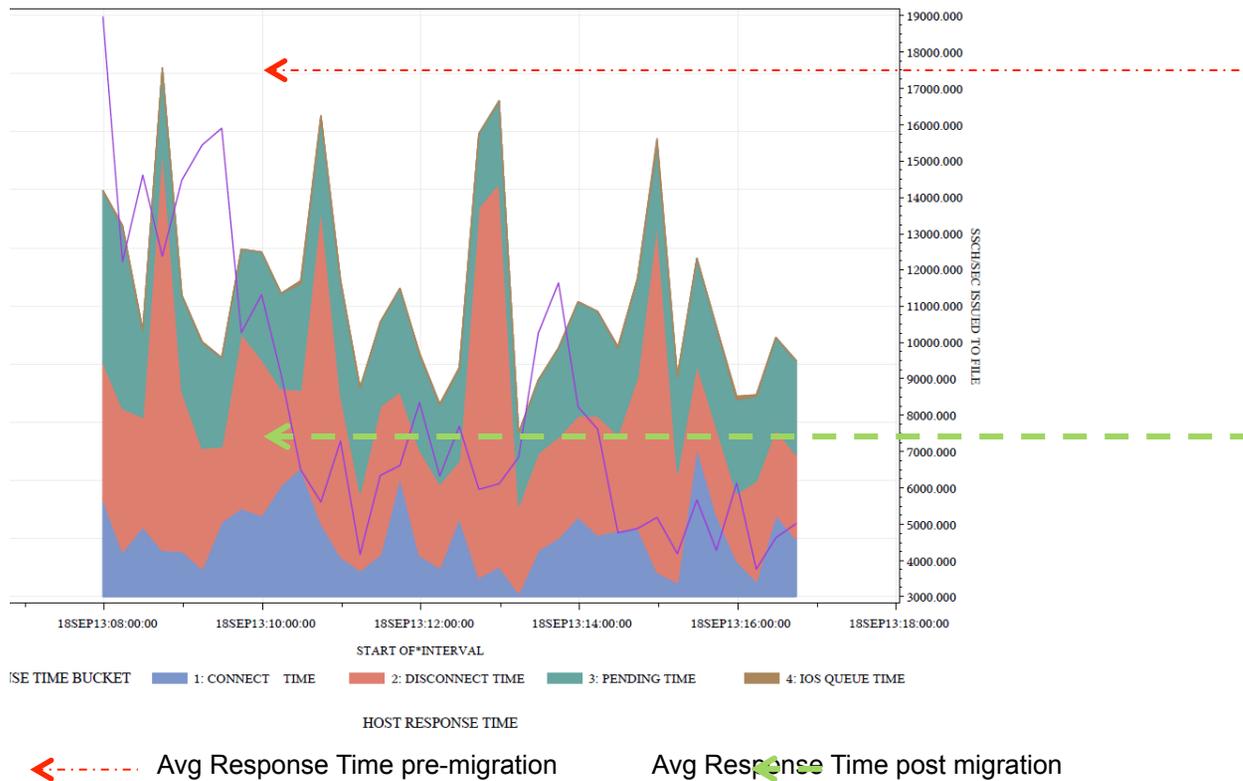


Frontend Response – RMF Type 74



Response Time vs. SSCH Rate
HTC 93621
9501

September 25, 2013 1
VSP 1



Close to 2X Faster Response Time over USP V
Up to 50% more IO

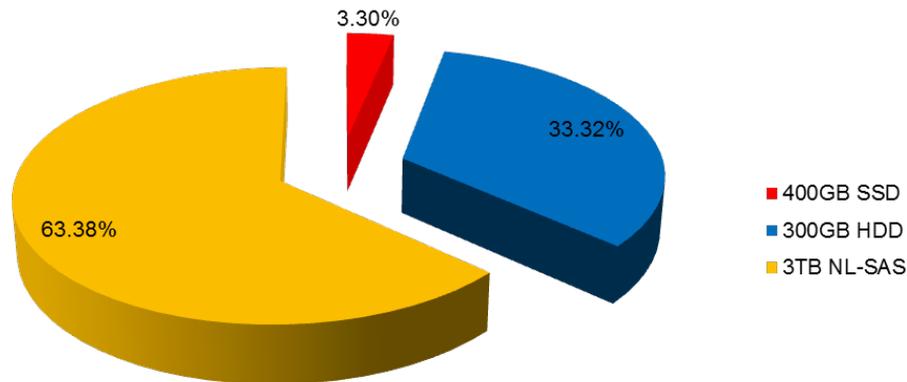
Complete your session evaluations online at www.SHARE.org/Anaheim-Eval



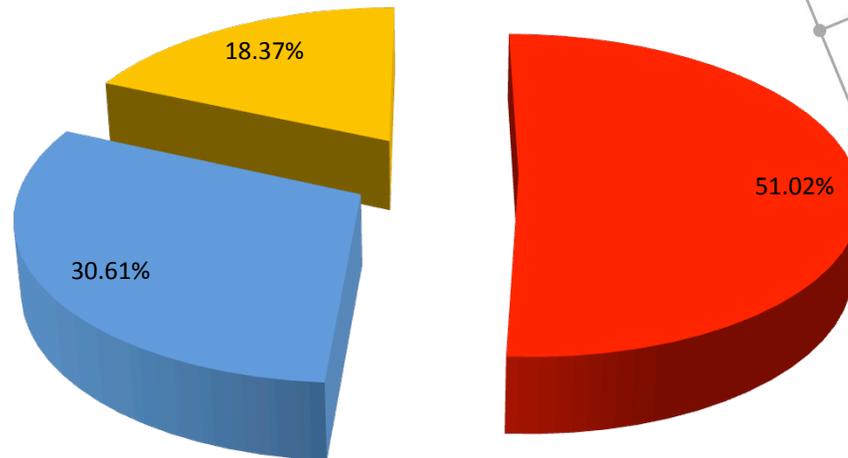
Backend IO Distribution vs Physical Capacity



Physical Capacity Distribution



IO Rate Distribution



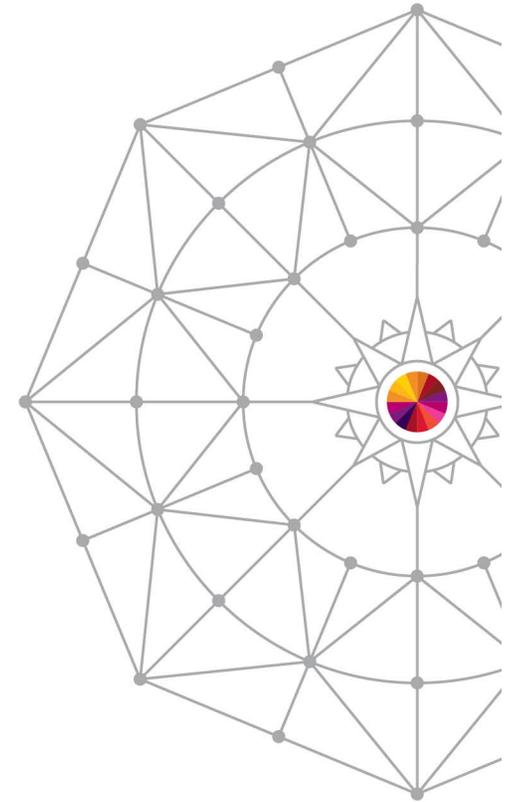
Complete your session evaluations online at www.SHARE.org/Anaheim-Eval



HDT and HTSM for MF



Summary



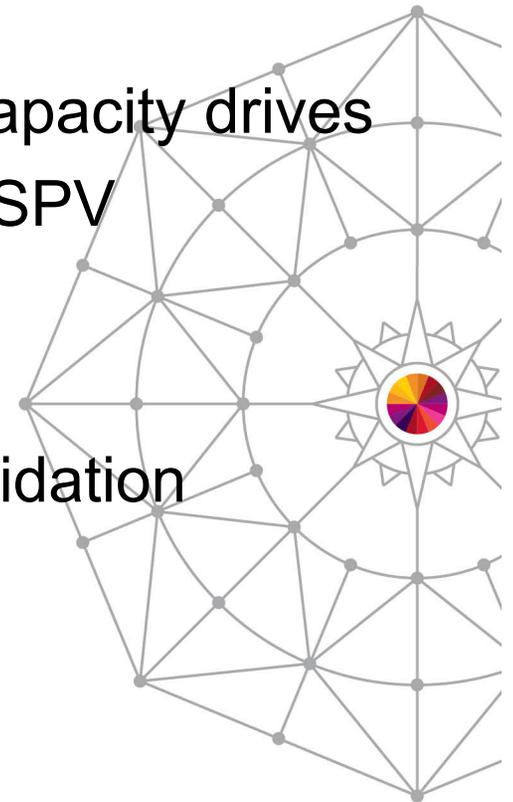
Complete your session evaluations online at www.SHARE.org/Anaheim-Eval



VSP HDT Benefit Summary

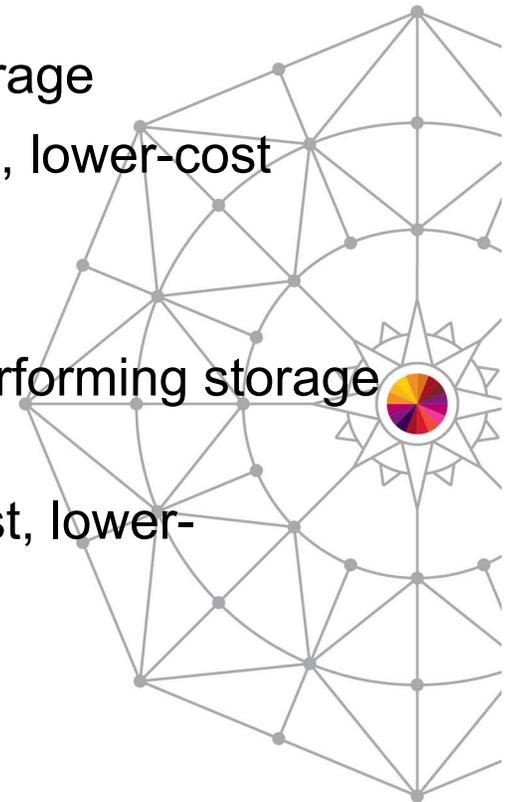


- Excellent performance benefit with a small amount of SSD
- Proven ability to take advantage of larger capacity drives
- 2X improved response time compared to USPV
- Increased workload of up to 55% more IO
- 50%+ Environmental Cost Savings
- Significant software savings through consolidation
- Ease of performance management

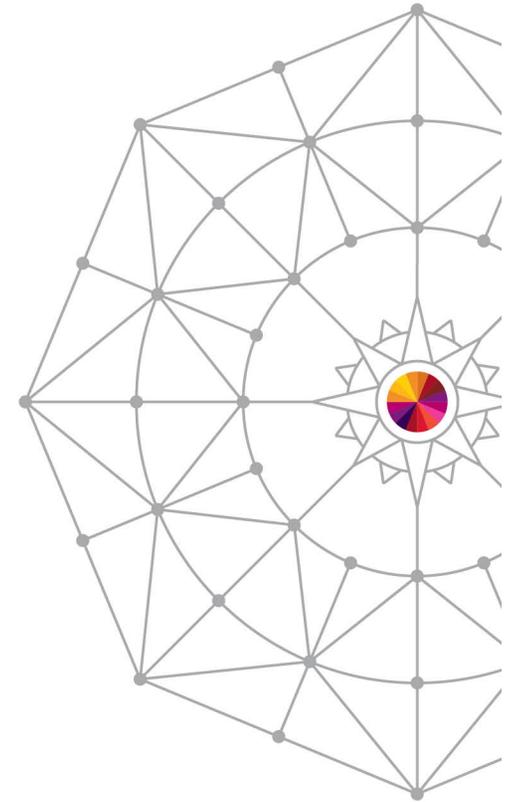


HDT for Mainframe and HTSM for Mainframe Storage Benefits

- Enables automation and more efficient use of storage
 - Data that is highly used is in high-performance storage
 - Lower priority jobs can be kept in lower-performing, lower-cost storage
- Improves ability to manage SLAs
 - High-priority jobs using more expensive, higher-performing storage can be given it and charged accordingly
 - Jobs with less-demanding SLAs can use lower-cost, lower-performance storage
- Ability to manage via DFSMS storage groups
- Reduced overhead of migrations to ML1



Questions and discussion



Thank you

