# Hitachi Dynamic Tiering for Mainframe
# Hitachi Tiered Storage Manager for Mainframe

SHARE
Technology · Connections · Results

*Session 13979*

## Hitachi Dynamic Tiering, Managing Your Mainframe Storage Easily and Effectively

Ros Schulman – Data Protection Product Line Manager
John Harker - Senior Product Marketing Manager

# Hitachi Dynamic Tiering (HDT)

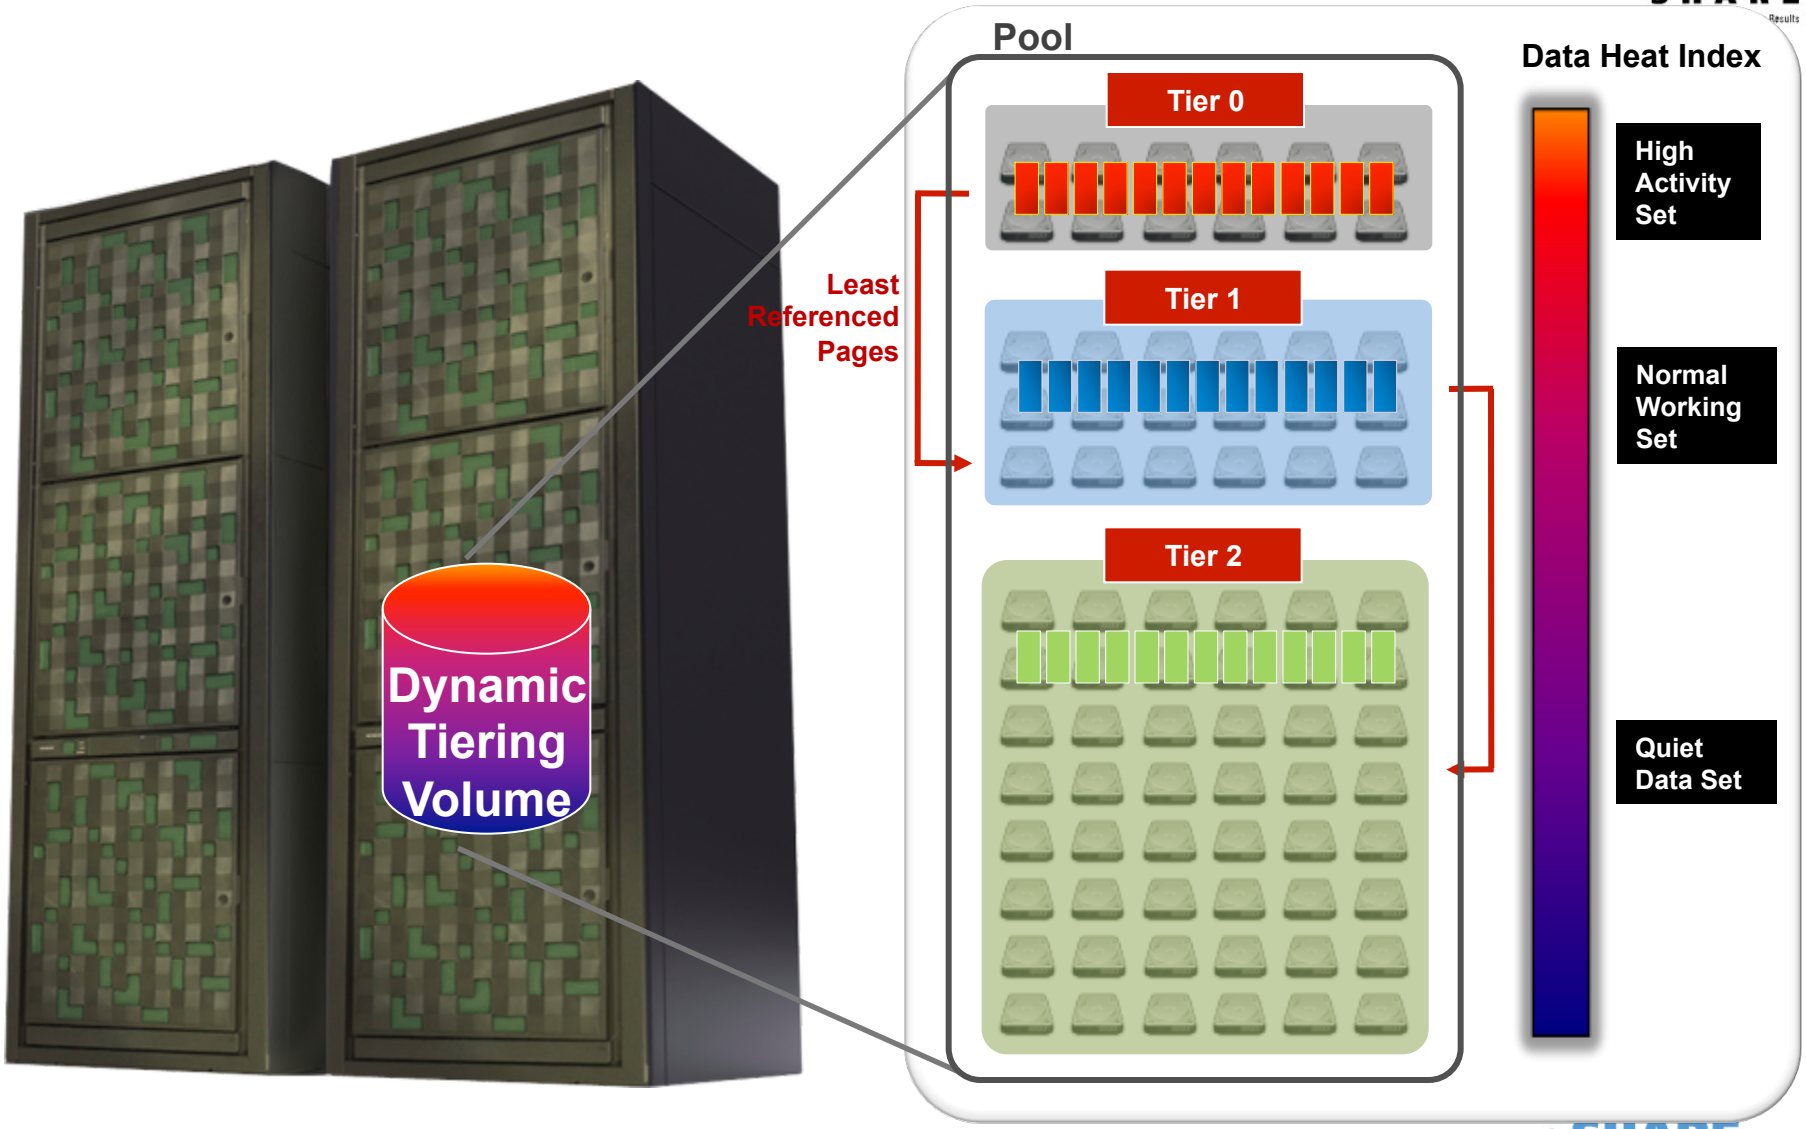

# Improved Performance at Reduced Cost: Data Locality and Pareto Distributions

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

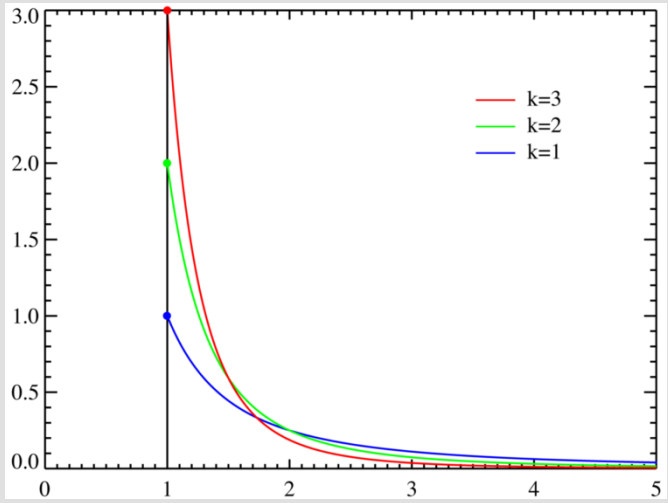


## Actual Volume Workload

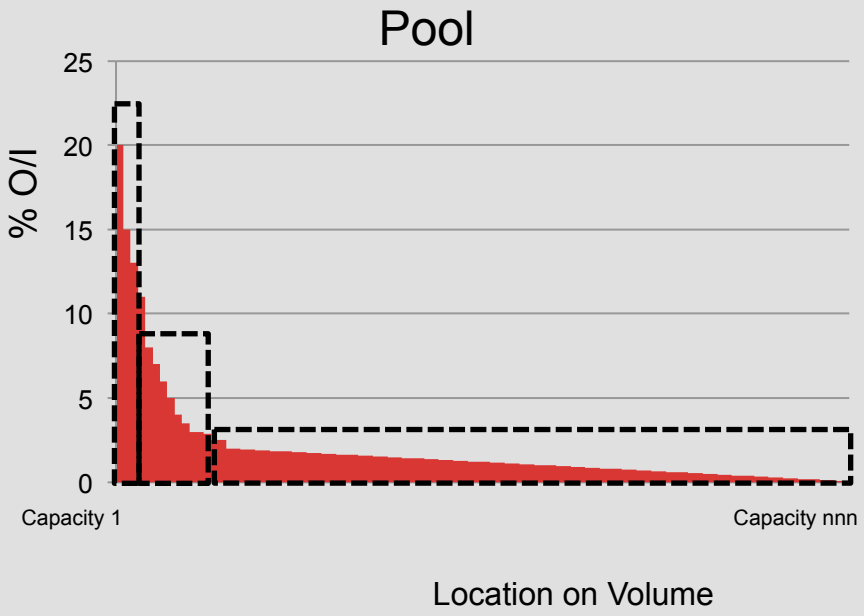


## Why does it work so well?

- 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 PROVISIONING FOR MAINFRAME STORAGE Architecture

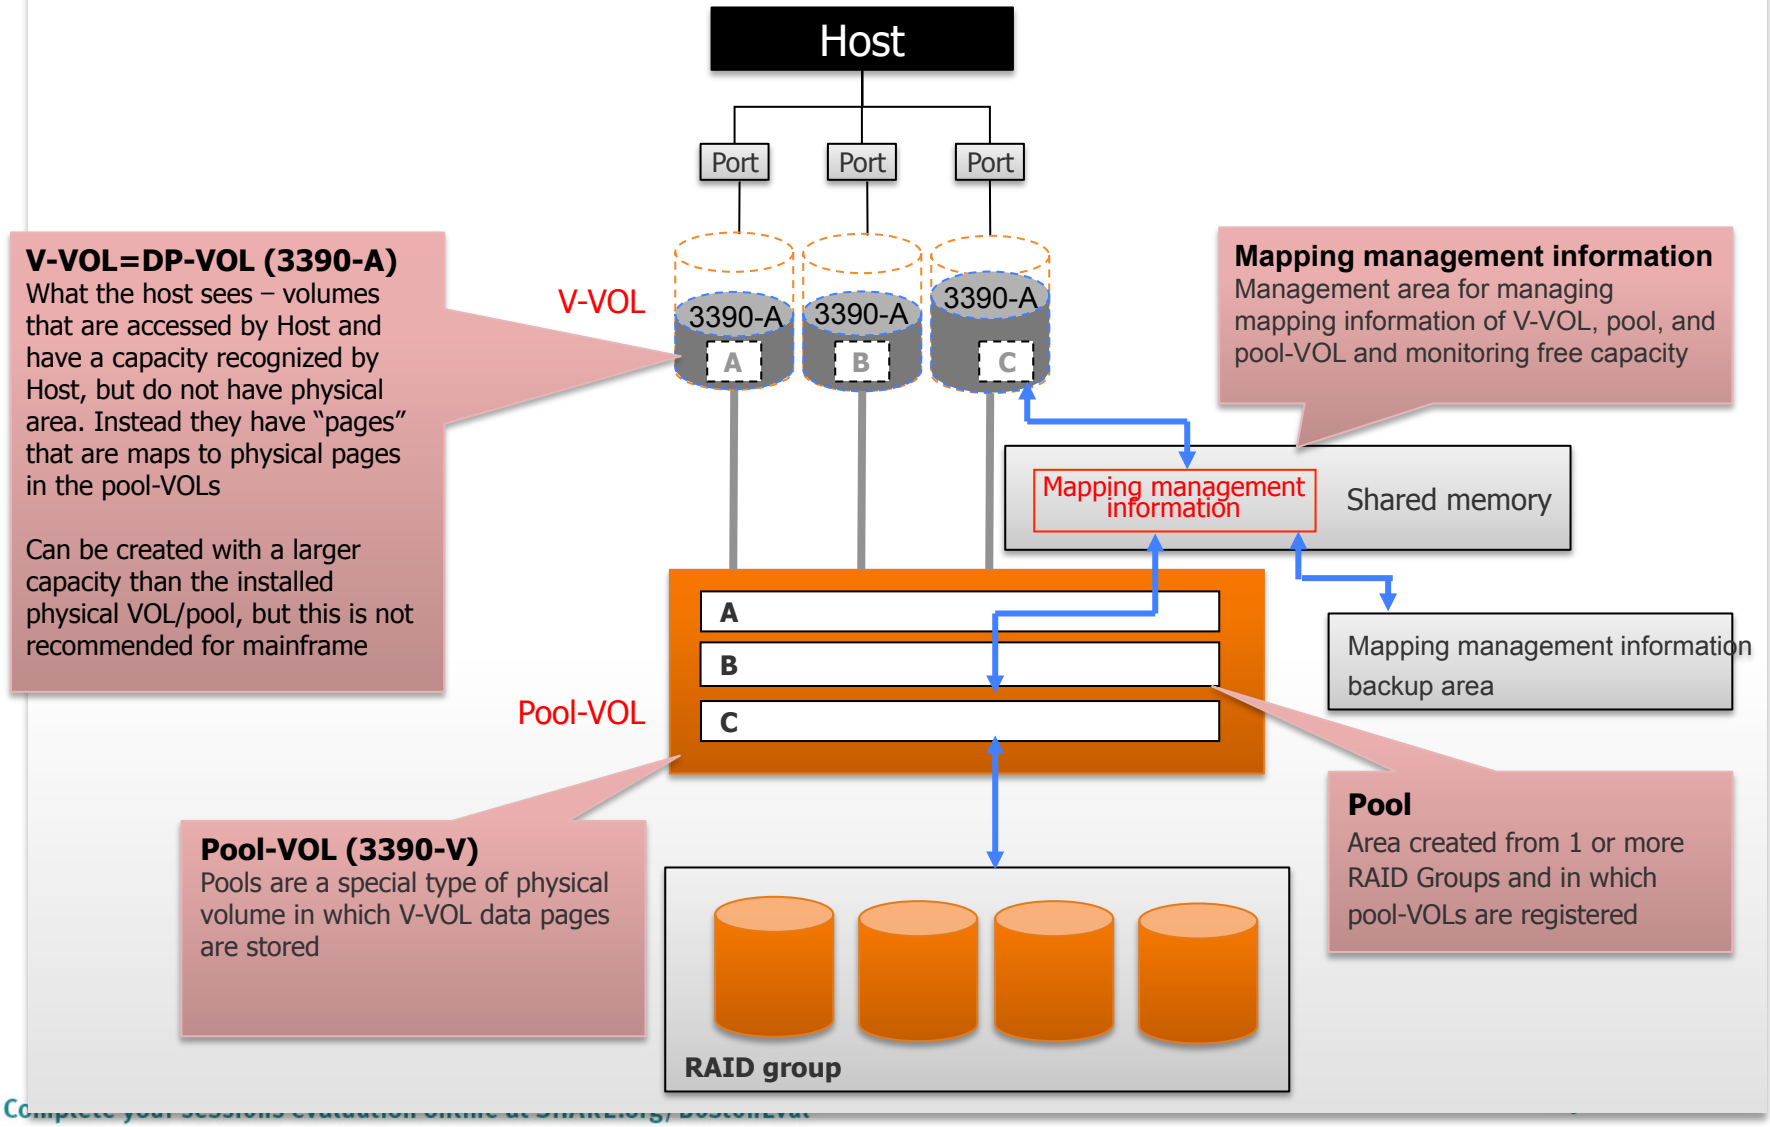


**V-VOL=DP-VOL (3390-A)**
What the host sees – volumes that are accessed by Host and have a capacity recognized by Host, but do not have physical area. Instead they have "pages" that are maps to physical pages in the pool-VOLs

Can be created with a larger capacity than the installed physical VOL/pool, but this is not recommended for mainframe

**Mapping management information**
Management area for managing mapping information of V-VOL, pool, and pool-VOL and monitoring free capacity

**Pool-VOL (3390-V)**
Pools are a special type of physical volume in which V-VOL data pages are stored

**Pool**
Area created from 1 or more RAID Groups and in which pool-VOLs are registered

# 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 dynamic provisioning
  - Optimized use of Flash storage
  - No need for data classification

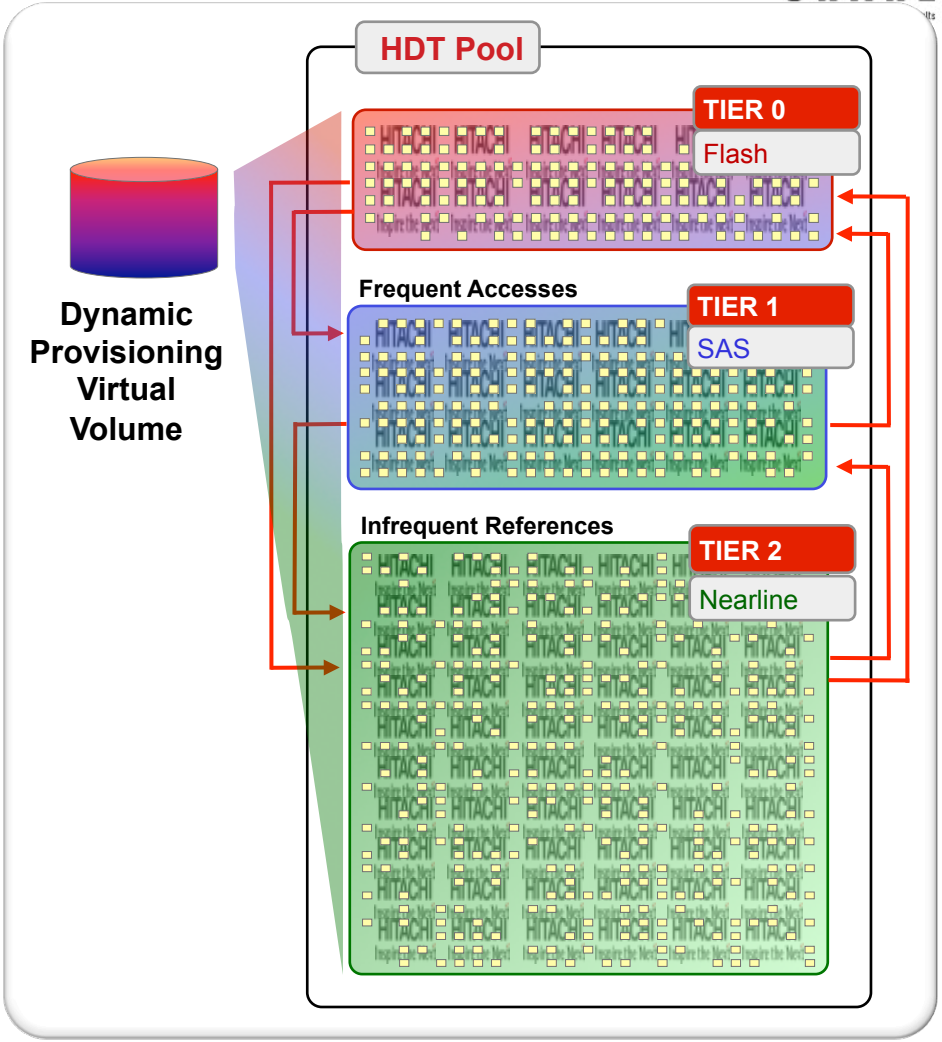


- **Simplifies operations and data management**
- **Reduces opex, capex, and TCO**

# 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

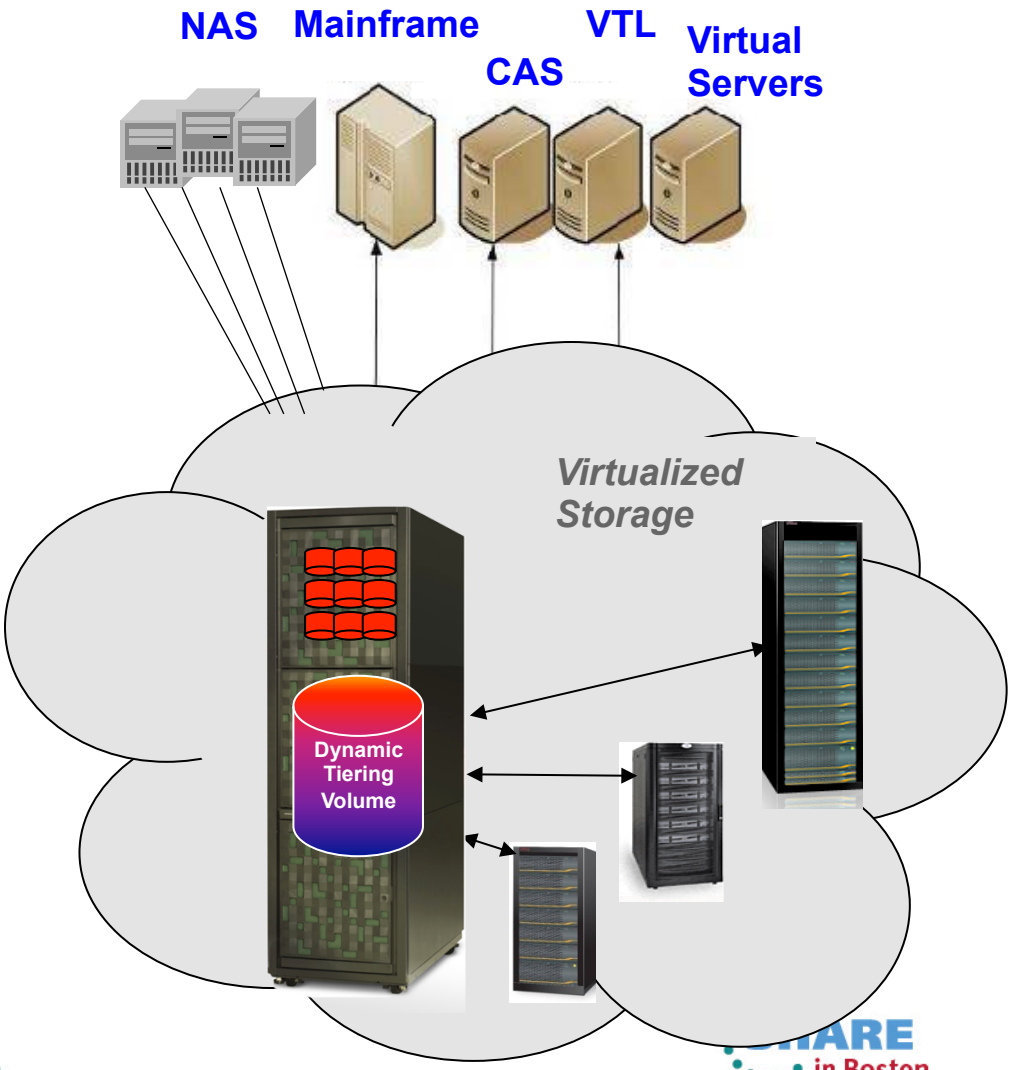

# Hitachi Dynamic Tiering for Mainframe and DFSMS

# LIMITS OF SMS Storage Groups and ACS Routines

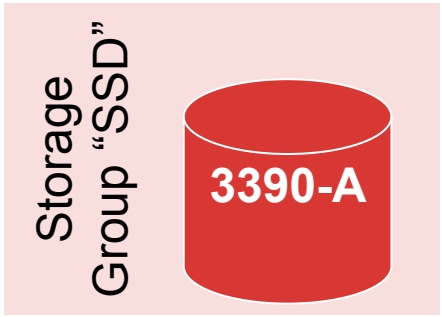


- 3390 volumes are "fixed" to a single tier
- To transition a volume's data to another tier requires ACS work, then migration and recall

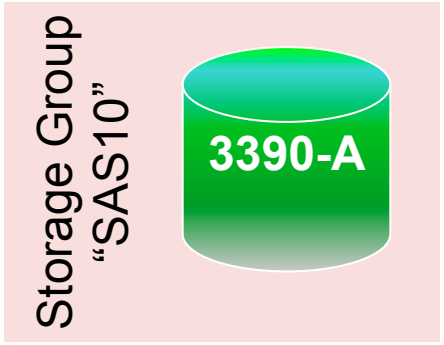


- Stale datasets are treated the same as active ones until HSM migration

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

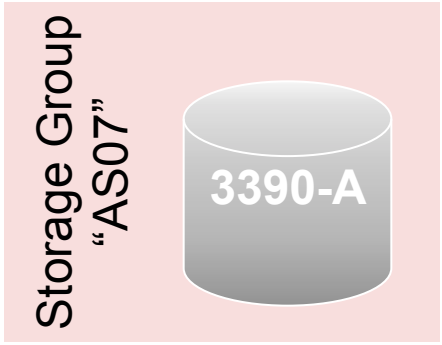

# DFSMS
## Storage Groups (SG) and HDT for Mainframe storage

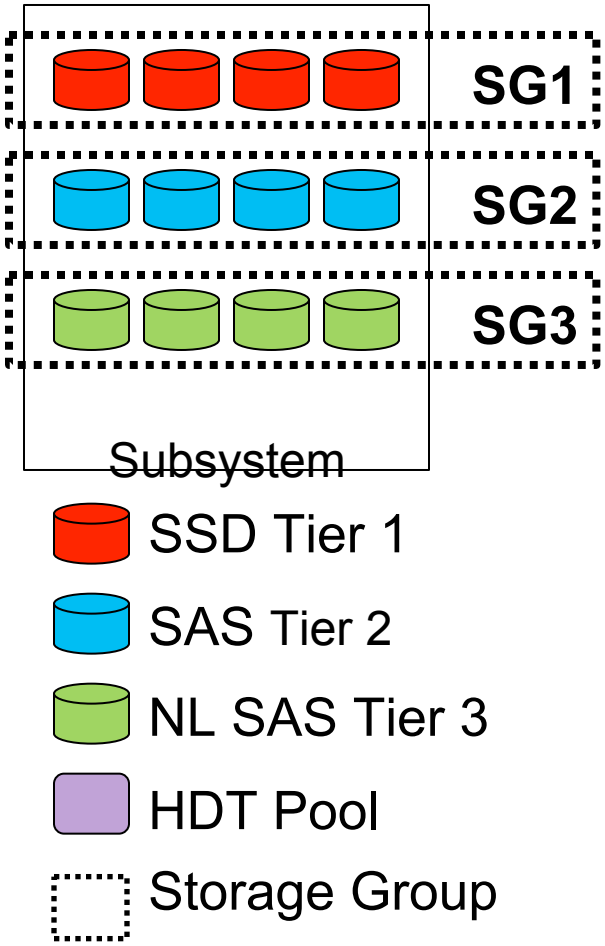


- Now with HDT for Mainframe storage, storage tiers may be combined into a single storage group
- HDT automatically moves data within a storage group to the correct tier based on Workload performance

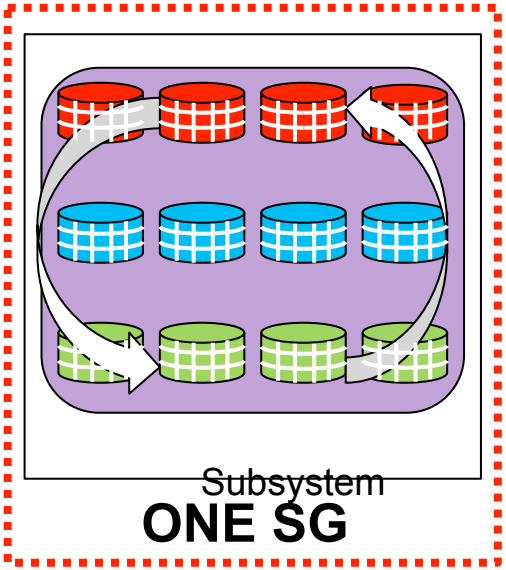


HDT Volumes are optimally managed dynamically at the page level

# DFSMS Storage Groups and HDT for Mainframe Storage

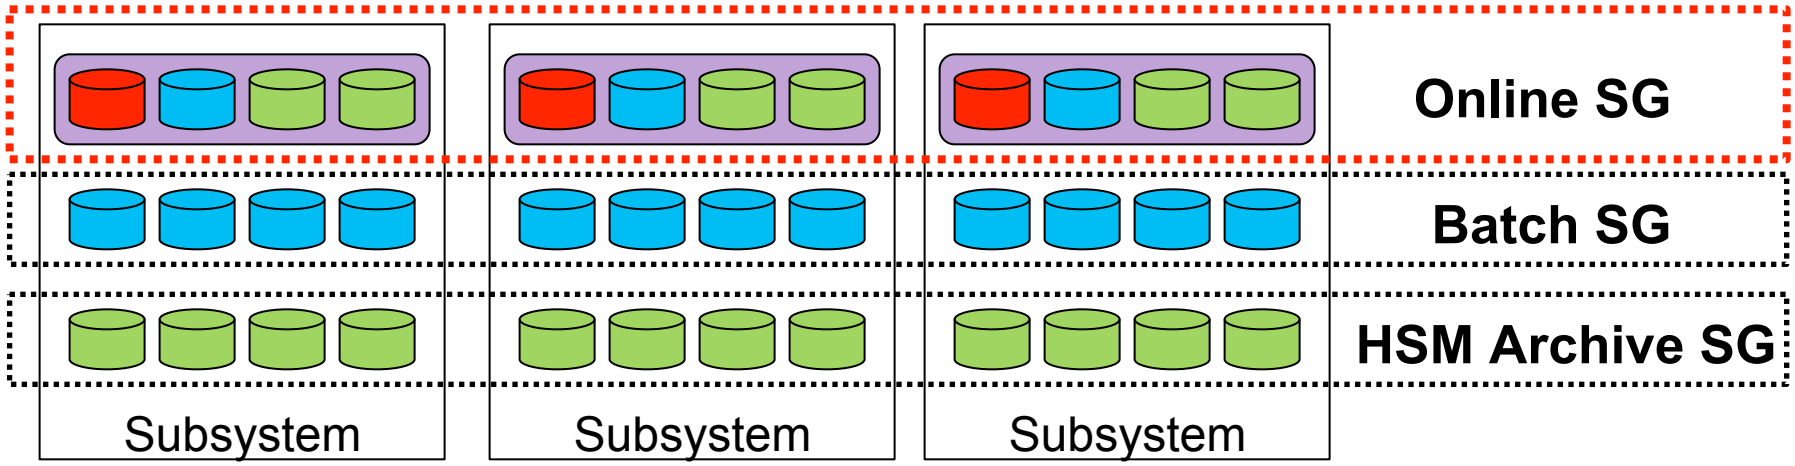


FLASH Tier 0

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

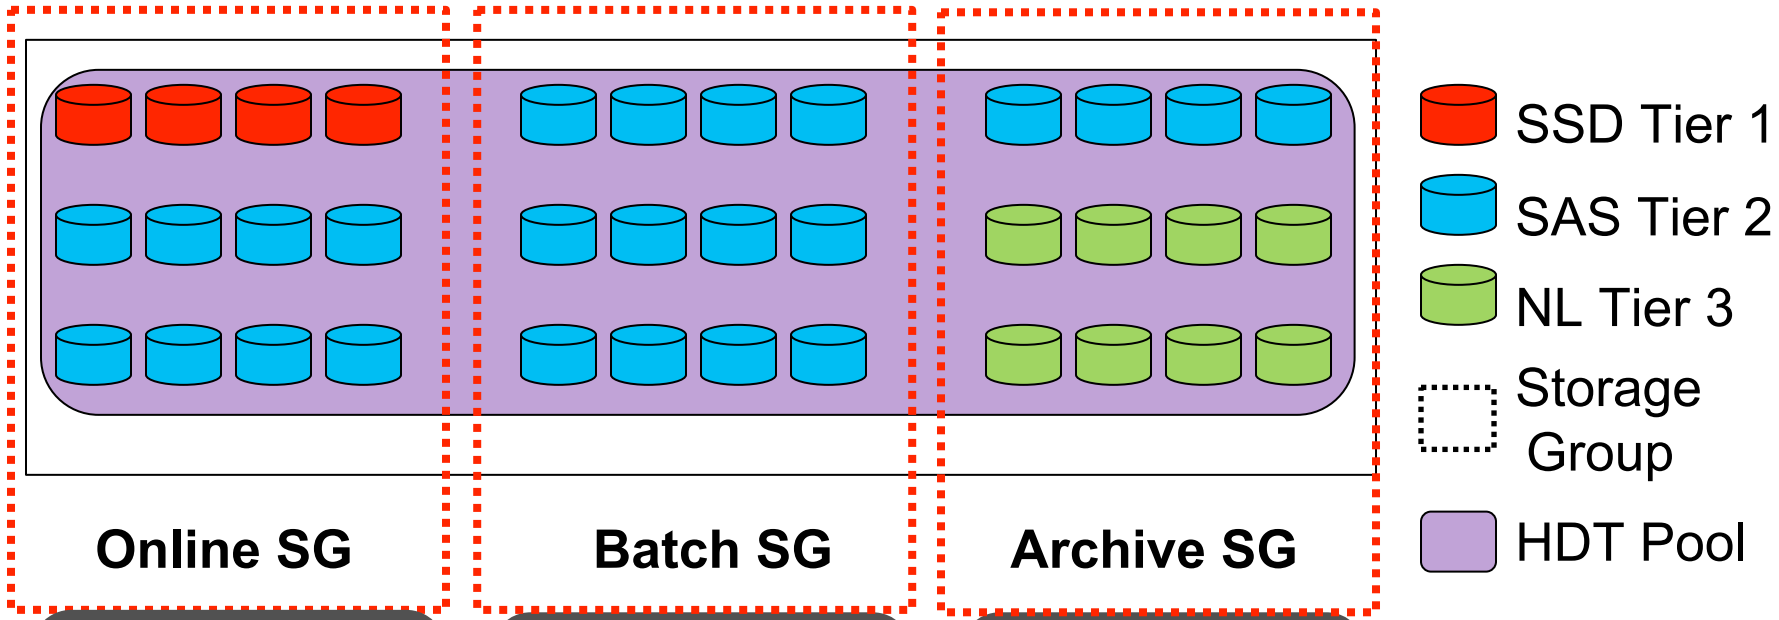


HDT Custom Policy Defined for Online Data Across Top 2 Tiers

HDT Policy Restricted to Tier 2 Residency

HDT Policy Restricted to Tiers 2 and 3 Residency

# HDT Operational Impact #1

| Operation | DFSMS/HSM | Dynamic Tiering |
|---|---|---|
| What does it take to move between tiers? | HSM or manual; high MIPs overhead | Automated; offloads mainframe overhead |
| How responsive is a tier adjustment? | Slow; entire volumes are moved | Short as 30 minute cycles; only pages are moved |
| How to proactively avoid problems? | Static ACS routines, manual change & monitoring required | Automated; fine-grained and self-optimizing |
| How well does wide striping work for performance? | No wide striping | Improved performance, automatic dynamic optimization of all available spindles |
| Effort to effectively manage tiers with SMS and IBM® z/OS®? | Manual calculations and programming | Automatic |

# HDT Operational Impact #2


| 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 | Constant challenge to keep updated with rules describing exceptions | Fewer exceptions since HDT keeps tiers properly populated |
| Demote data to lower tiers | HSM moves *datasets* to an ML "tier" that hasn't been *opened* for a while | HDT automatically moves *pages* that haven't been *used* |

# Hitachi Dynamic Tiering Management

## Hitachi Tiered Storage Manager for Mainframe

# Management of Hitachi Dynamic Tiering

- Automatically detects and assigns Tiers based on media type and speed
- Control relocation of data based on either most recent IO patterns or longer term averaging
- Control relocation and analysis periods (from 30 min to 24 hours), relocation priority, tier for new allocations and the range of tiers a volume can use
- User also can dynamically:
  - Add or remove Tiers
  - Expand or shrink Tiers
  - Replace media and RAID types
  - Expand LUNs
  - Move LUNs between pools
  - Add or remove any media or RAID type

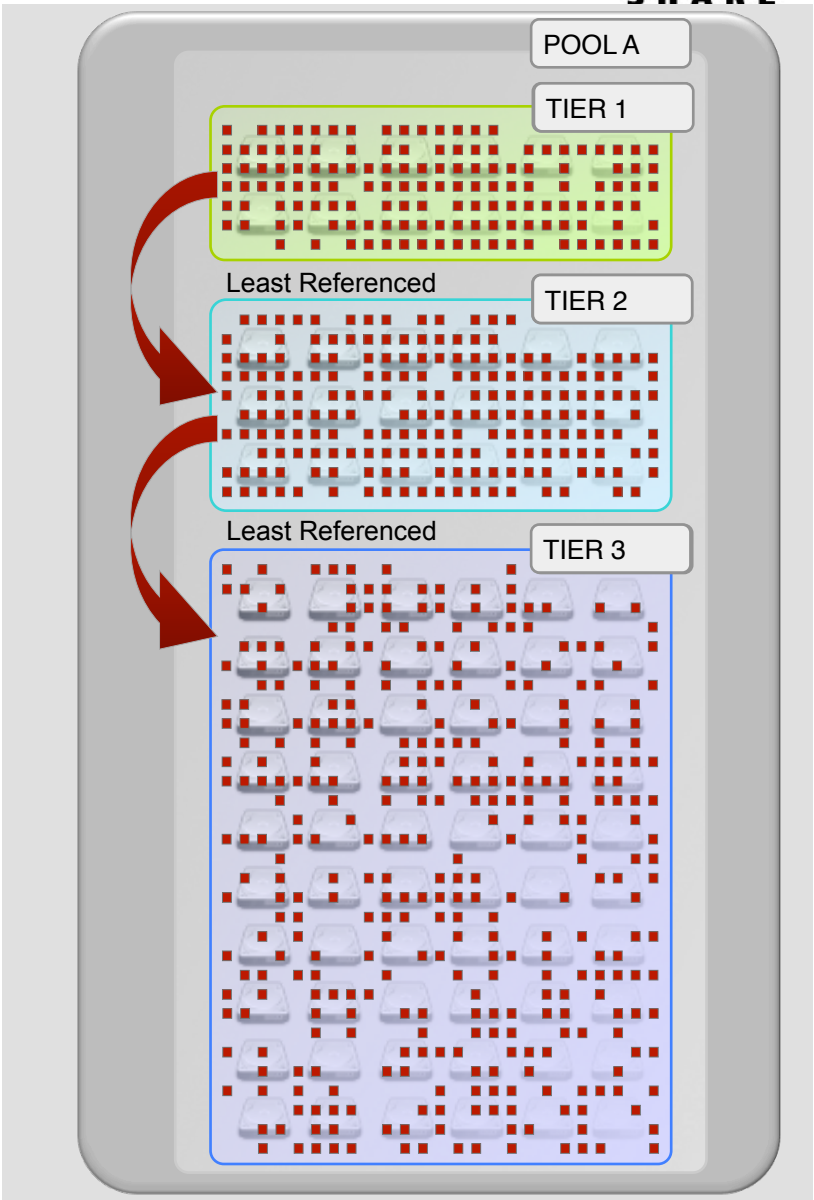

# Hitachi Tiered Storage Manager for Mainframe
## z/OS-based Management of Dynamic Tiering

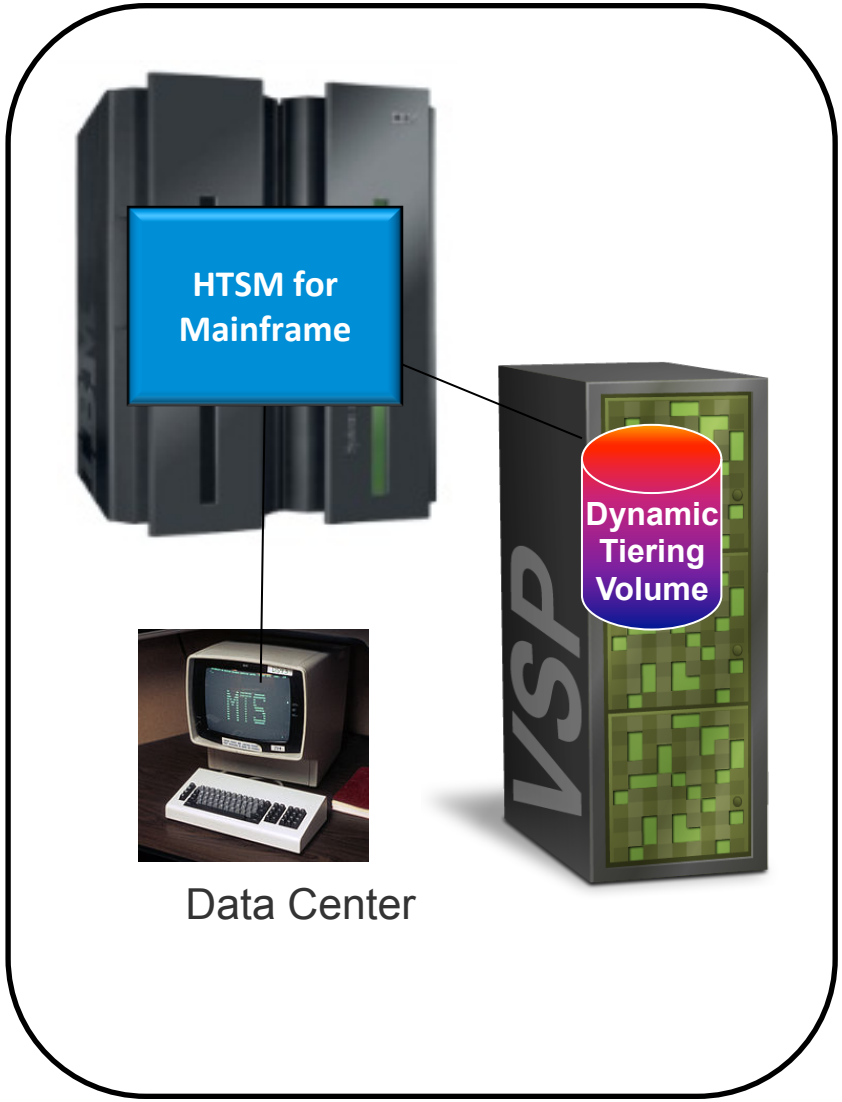


## Native z/OS host-based software provides:

- Online storage service level controls
  - Increase application performance
  - Improves problem avoidance
- Centralized and unified mainframe management of Hitachi Dynamic Tiering
  - Automation
  - Integration with DFSMS and storage groups
- Enables reporting and automatic notifications

# Hitachi Tiered Storage Manager for Mainframe
## z/OS-based Dynamic Tiering Management

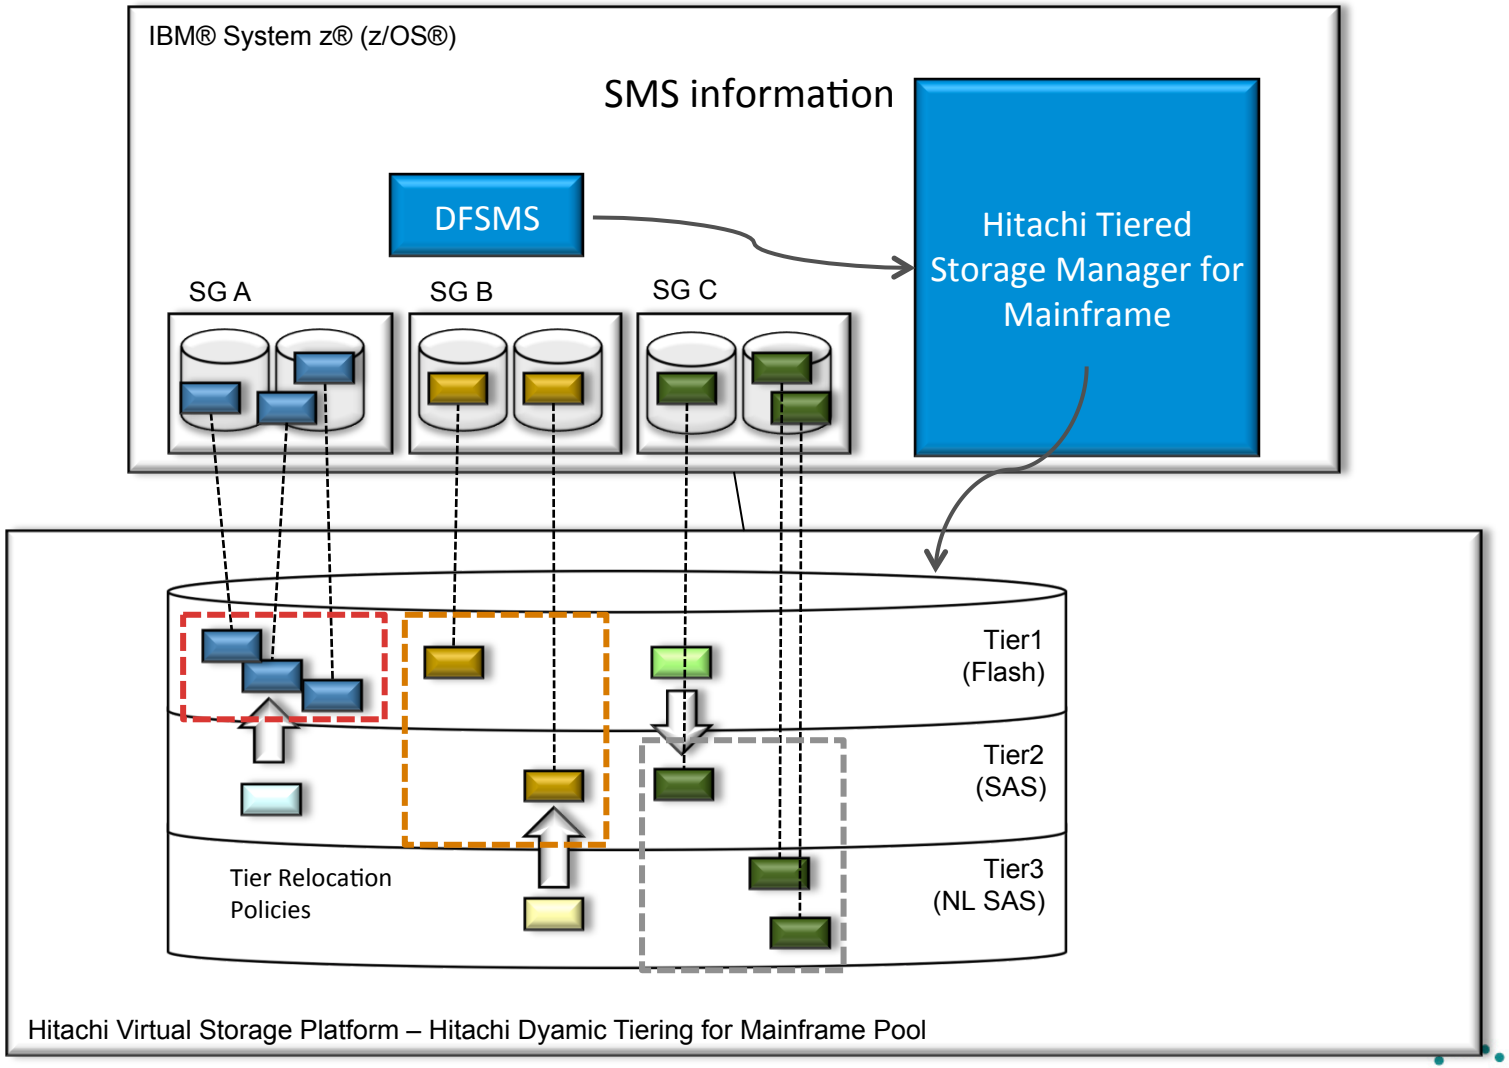

# Why HTSM for mainframe?

## HTSM simplifies and improves operations

- 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)
  - Reduced dependency on open server-based operations
- Control of storage service levels using Dynamic Tiering 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

# Easy-To-Manage Target Volumes

HTSM defines target volumes as a group. A single operation for the group applies to all volumes belonging to the group (called Tiering Policy Group [TPG]). Once the group is defined, subsequent management can be done via a group operation

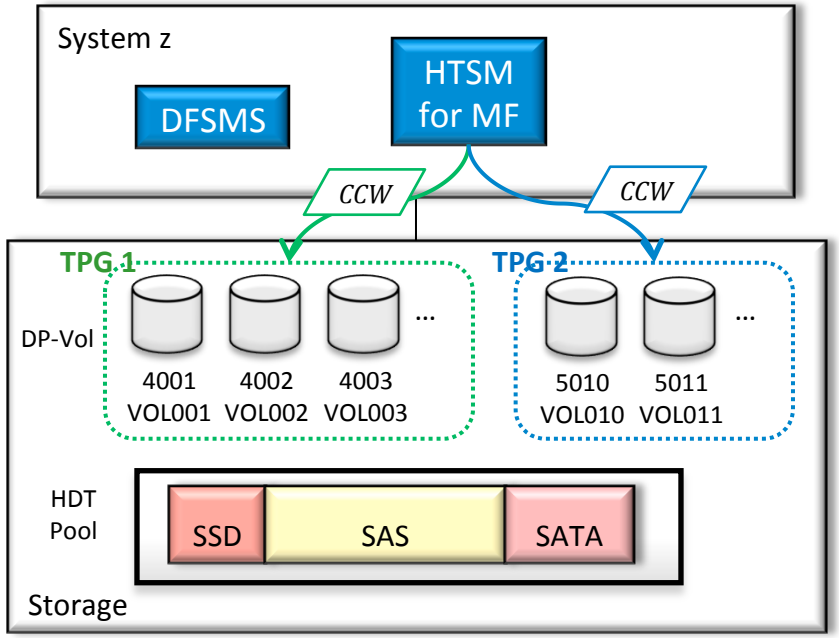

# Linkage with z/OS DFSMS

HTSM works with SMS storage groups and ACS routines. Each Tiering Policy Group can have 1 or more SMS Storage Groups. This feature gives users the capability to manage the HDT environment from a SMS point of view and makes it easy to add Dynamic Tiering to existing operations

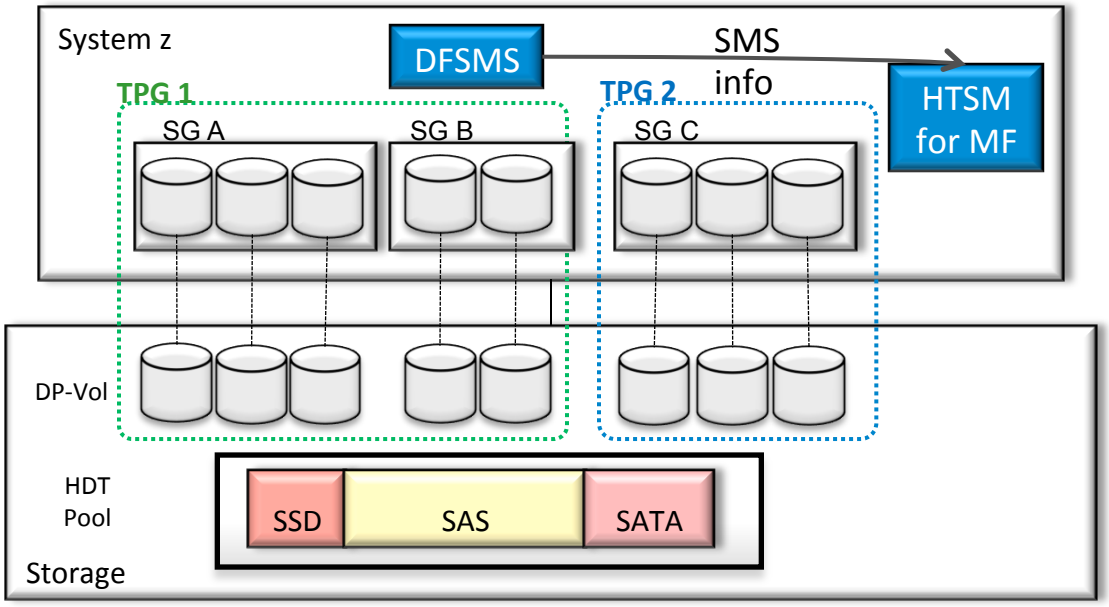

# DFSMS and HDT and HTSM for Mainframe

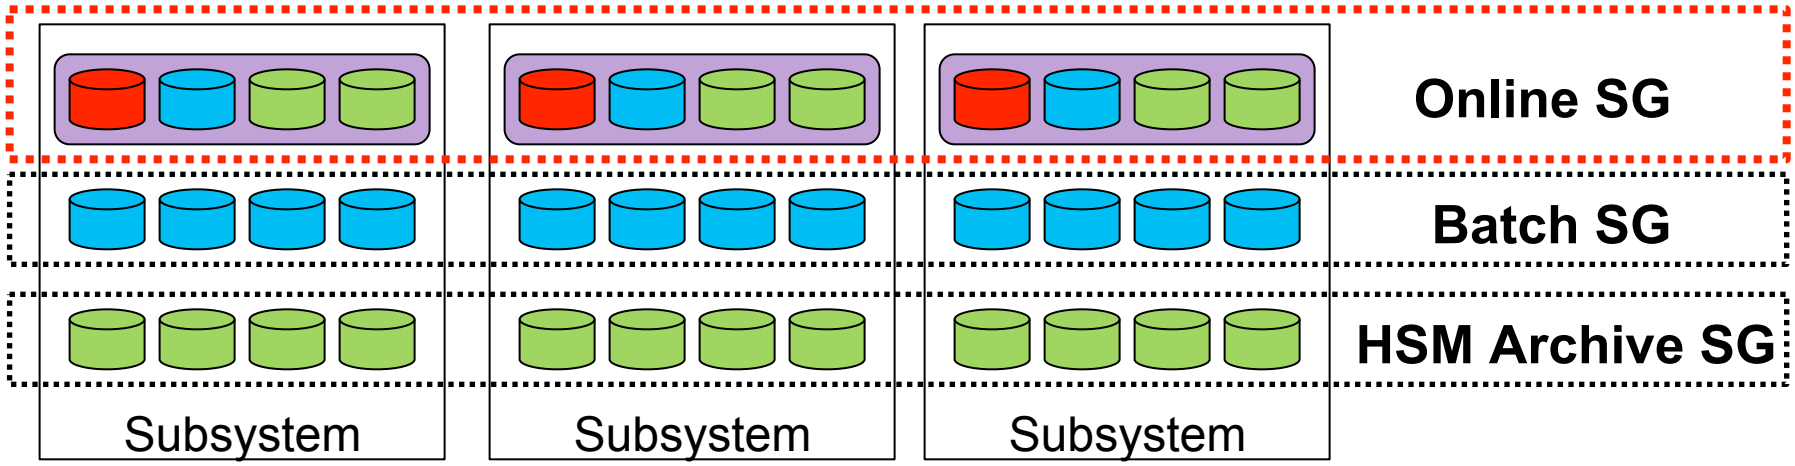


SSD Tier 1

SAS Tier 2

NL SAS Tier 3

HDT Pool

Storage Group

- HDT can be applied to selected Storage Groups only
- HTSM for MF can be used to apply the Tiering policies by Storage Group
- DFSMS Horizontal Storage Groups and Dynamic Page based Tiering volumes can be aligned

# Outline of HTSM configuration files

- The input of HTSM for MF is CSV file. The output are TPG config and disk config.

Sample of tiering policy information CSV file

```
#comment
#type,SMS Storage Group,TieringPolicyName,TieringPolicy
STORGRP,SG1,TPG1,2
DEVN,1101,TPG1,1 ,HIGH
STORGRP,SG2,TPG1,5,LOW
```

Dakota for MF

Sample of Tiering Policy Group definition file

```
<?xml version="1.0" encoding="ebcdic-cp-us"?>
<APIInfo Level="7.5.0">
<TieringPolicyGroup ID="TPG1" />
<TieringPolicy Level="1" Entry="High">
<DiskDevice SerialNum="FGHK2" CUNum="01"
SSID="2340" CCA="11" POOLID="01" />
</TieringPolicy>
<TieringPolicy Level="2" Entry="Middle">
<DiskDevice SerialNum="FGHK2" CUNum="01"
SSID="2340" CCA="01" POOLID="01" />
<DiskDevice SerialNum="FGHK2" CUNum="01"
SSID="2340" CCA="02" POOLID="01" />
</TieringPolicy>
<TieringPolicy Level="5" Entry="Low">
<DiskDevice SerialNum="FGHK2" CUNum="02"
SSID="2340" CCA="01" POOLID="02" />
<DiskDevice SerialNum="FGHK2" CUNum="02"
SSID="2340" CCA="02" POOLID="02" />
</TieringPolicy >
</TieringPolicyGroup>
</APIInfo>
```

Sample of Disk configuration definition file

```
<?xml version="1.0" encoding="ebcdic-cp-us"?>
<APIInfo Level="7.5.0">
<HostDiscoveredArray SerialNum="FGHK2"
Model="VSP"
Microcode="700000FF" IFType="3131"
PhysicalSerialNum="14002">
<POOL ID="01" NAME="POOL1" Type="DT" />
<POOL ID="02" NAME="POOL2" Type="DT" />
<HostAddressedDisk Devn="1101" >
<DiskDevice SerialNum="FGHK2" CUNum="01"
SSID="2340" CCA="12" POOLID="01" />
</HostAddressedDisk>
<HostAddressedDisk Devn="1201" >
<DiskDevice SerialNum="FGHK2" CUNum="01"
SSID="2340" CCA="01" POOLID="01" />
</HostAddressedDisk>
<HostAddressedDisk Devn="1202" >
<DiskDevice SerialNum="FGHK2" CUNum="01"
SSID="2340" CCA="02" POOLID="01" />
</HostAddressedDisk>
<HostAddressedDisk Devn="2501" >
<DiskDevice SerialNum="FGHK2" CUNum="02"
SSID="2340" CCA="01" POOLID="02" />
</HostAddressedDisk>
<HostAddressedDisk Devn="2502" >
<DiskDevice SerialNum="FGHK2" CUNum="02"
SSID="2340" CCA="02" POOLID="02" />
</HostAddressedDisk>
</HostDiscoveredArray>
</APIInfo>
```

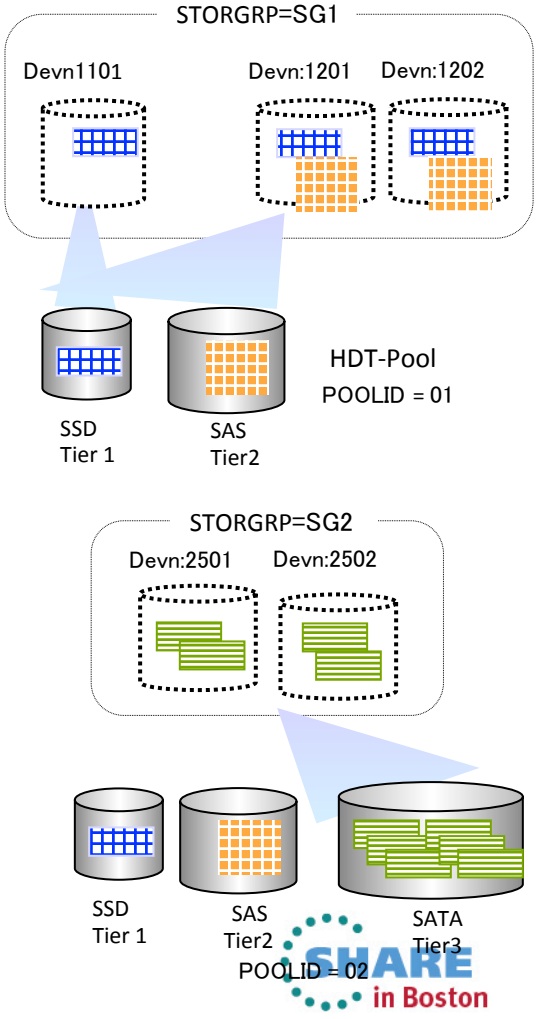

# HTSM Customization via REXX Scripting

HTSM provides a CLI (TSO/E REXX) that enables users to tailor the operation to 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 objects (such as group definition)

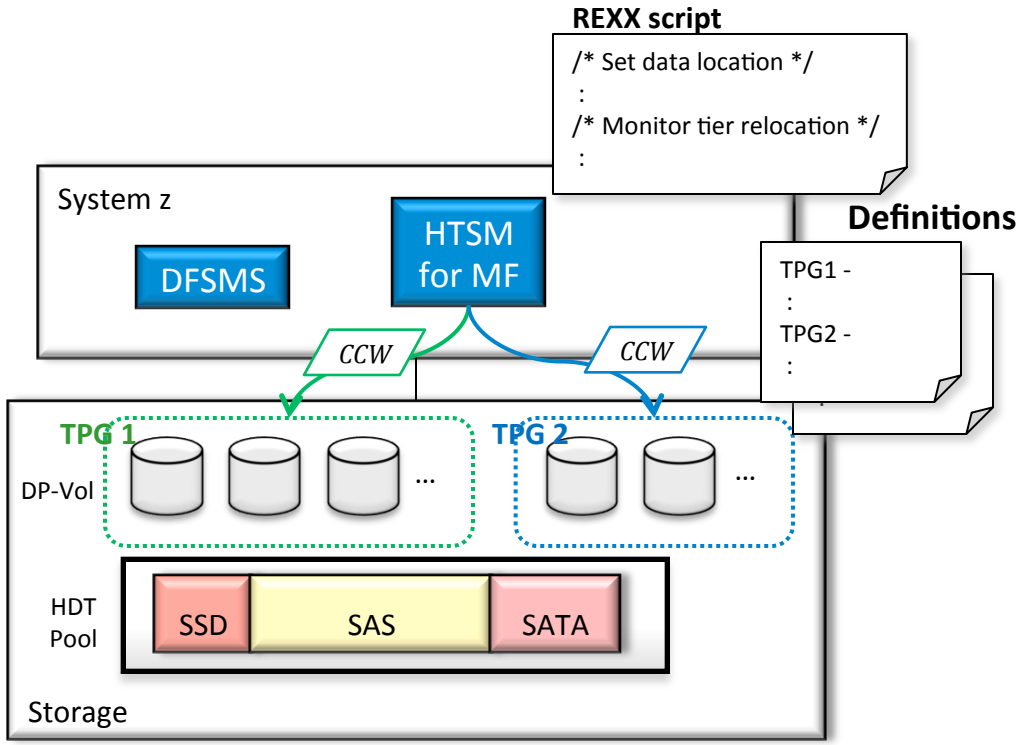

# HTSM for MF Reporting - Query Relocation

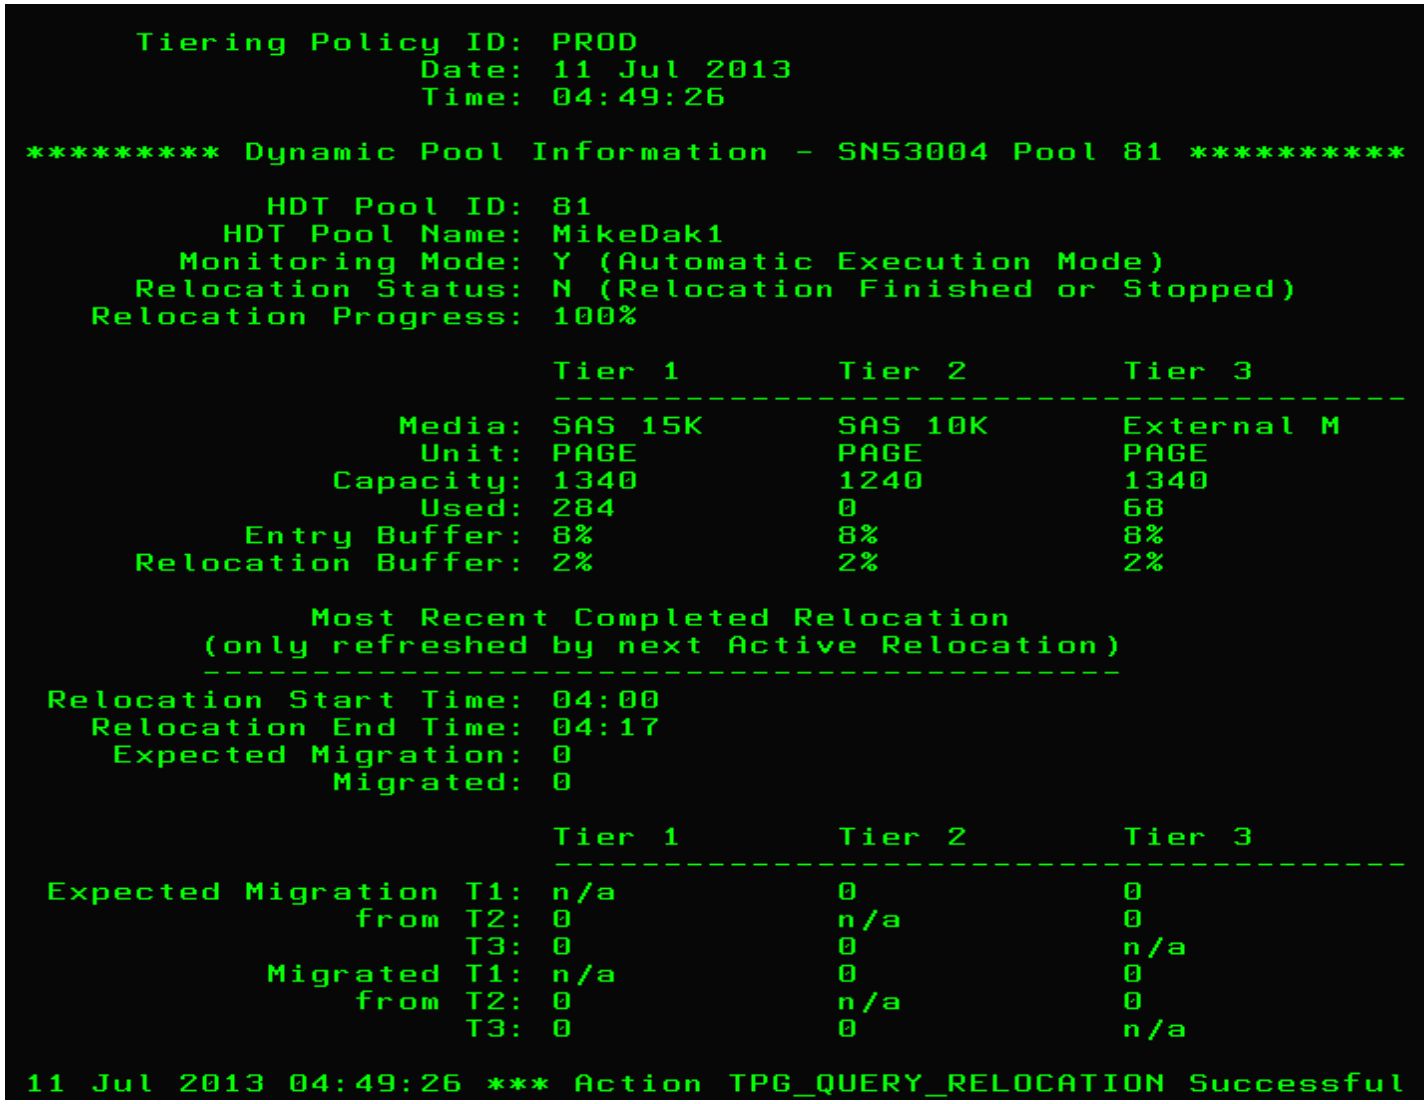

```
     Tiering Policy ID: PROD
                  Date: 11 Jul 2013
                  Time: 04:49:26

********** Dynamic Pool Information - SN53004 Pool 81 **********

           HDT Pool ID: 81
         HDT Pool Name: MikeDak1
       Monitoring Mode: Y (Automatic Execution Mode)
     Relocation Status: N (Relocation Finished or Stopped)
   Relocation Progress: 100%

                        Tier 1        Tier 2        Tier 3
                        ------------------------------------
                 Media: SAS 15K       SAS 10K       External M
                  Unit: PAGE          PAGE          PAGE
              Capacity: 1340          1240          1340
                  Used: 284           0             68
          Entry Buffer: 8%            8%            8%
     Relocation Buffer: 2%            2%            2%

            Most Recent Completed Relocation
        (only refreshed by next Active Relocation)
        ------------------------------------------
 Relocation Start Time: 04:00
   Relocation End Time: 04:17
    Expected Migration: 0
              Migrated: 0

                        Tier 1        Tier 2        Tier 3
                        ------------------------------------
 Expected Migration T1: n/a           0             0
                from T2: 0           n/a           0
                      T3: 0          0             n/a
           Migrated T1: n/a          0             0
                from T2: 0           n/a           0
                      T3: 0          0             n/a

11 Jul 2013 04:49:26 *** Action TPG_QUERY_RELOCATION Successful
```

# HTSM for MF Reporting - Query Policy

```
     Tiering Policy ID: PROD
                  Date: 11 Jul 2013
                  Time: 03:10:54

************************ Query TPG ***************************

Note: Any '*' Character below indicates a Tiering Policy
      Mismatch between the TPG Definition and the Active
      HDT-MF Settings on the DKC(s).

====================== Defined in TPG ======================
             Storage  Serial Pool Tier  Entry  Tier
Devn Volser  Group    Number ID   Level Tier   Relocation
============================================================
E941 GSE941 ALPHA    53004  81   2     HIGH   Y (Enabled)
E942 GSE942 ALPHA    53004  81   2     HIGH   Y (Enabled)
E943 GSE943 ALPHA    53004  81   2     HIGH   Y (Enabled)
E944 GSE944 ALPHA    53004  81   2     HIGH   Y (Enabled)
E945 GSE945 ALPHA    53004  81   2     HIGH   Y (Enabled)

==================== Active on DKC(s) ======================
             Storage  Serial Pool Tier  Entry  Tier
Devn Volser  Group    Number ID   Level Tier   Relocation
============================================================
E941 GSE941 ALPHA    53004  81 * 0     MIDDLE Y (Enabled)
E942 GSE942 ALPHA    53004  81 * 5     LOW    Y (Enabled)
E943 GSE943 ALPHA    53004  81 * 0     MIDDLE Y (Enabled)
E944 GSE944 ALPHA    53004  81 * 0     MIDDLE Y (Enabled)
E945 GSE945 ALPHA    53004  81 * 0     MIDDLE Y (Enabled)

11 Jul 2013 03:10:54 *** Action TPG_QUERY_POLICY Successful
```

# HTSM for MF Reporting - Query TPG

```
HTSM001I
HTSM001I       Tiering Policy ID: APACMC
HTSM001I                    Date: 27 Jul 2013
HTSM001I                    Time: 04:02:36
HTSM001I
HTSM001I ********************* Query TPG Tier Metrics ***********************
HTSM001I
HTSM001I ============================================================
HTSM001I TPG Total or          Tier1      Tier2      Tier3      Total
HTSM001I SN:PoolID or           Used       Used       Used       Used
HTSM001I *StorGrp* or          Pages      Pages      Pages      Pages
HTSM001I Volser    or      / UsedGB   / UsedGB   / UsedGB   / UsedGB
HTSM001I Volser Prefix     / Used%    / Used%    / Used%    / Used%
HTSM001I ============================================================
HTSM001I TPG Total               352          0          0        352
HTSM001I                      13.4GB       0GB        0GB     13.4GB
HTSM001I                        100%         0%         0%       100%
HTSM001I ------------------------------------------------------------
HTSM001I SN53004:81              352          0          0        352
HTSM001I                      13.4GB       0GB        0GB     13.4GB
HTSM001I                        100%         0%         0%       100%
HTSM001I
HTSM001I   Used% of Pool      26.3%         0%         0%      8.98%
HTSM001I
HTSM001I       Pool Pages      1340       1240       1340       3920
HTSM001I          Pool GB    50.9GB     47.1GB     50.9GB      149GB
HTSM001I   Tier% of Pool      34.2%      31.6%      34.2%       100%
HTSM001I ------------------------------------------------------------
HTSM001I *ALPHA*                 352          0          0        352
HTSM001I                      13.4GB       0GB        0GB     13.4GB
HTSM001I                        100%         0%         0%       100%
HTSM001I ------------------------------------------------------------
HTSM001I GSE94*                  352          0          0        352
HTSM001I                      13.4GB       0GB        0GB     13.4GB
HTSM001I                        100%         0%         0%       100%
HTSM001I ------------------------------------------------------------
HTSM001I
HTSM001I Action TPG_QUERY_TIERS Successful
```

# HDS HTSM for Mainframe Scripting Solution

From SDSF using z/OS SYSREXX

```
  Display  Filter  View  Print  Options  Search  Help
 ------------------------------------------------------------------------------
 SDSF SYSLOG  24996.101 GS03 GS03 06/11/2013 0W          93,450  COLUMNS 02- 81
 COMMAND INPUT ===> /                                          SCROLL ===> CSR
N 4                                                                          D
N C                          System Command Extension                        CT
NC0                                                                          $H
SC    Type or complete typing a system command, then press Enter.
X K                                                                          TP
X K   ===> rexx htsmcmdc tpg_query_status tpg(alphatpg) prefix(varend.htsm
X K   ===> cmd)
X K
X K   Place the cursor on a command and press Enter to retrieve it.
X K                                                            More:      +
X K   =>  D A,L                                                             TP
X K   =>  DS QD,SSID=ALL
X K   =>  rexx htsmcmdc tpg_query_relocation prefix(varend.htsmcmd) tpg(a
X K   =>  rexx htsmcmdc tpg_activate tpg(alphatpg) prefix(varend.htsmcmd)
X K   =>  rexx htsmcmdc tpg_deactivate tpg(alphatpg) prefix(varend.htsmcm
X K   =>  rexx htsmcmdc tpg_query_status tpg(alphatpg) prefix(varend.htsm
X K   =>  rexx htsmcmdc tpg_query_status tpg(alphatpg)
X K
X K   _  Wait 1 second to display responses (specify with SET DELAY)
X K   _  Do not save commands for the next SDSF session
X K                                                                          TP
X K   F1=Help  F5=FullScr  F7=Backward F8=Forward  F11=ClearLst  F12=Cancel
X K
X KTC311I 21:46:37 KTQRYTPG RC=0
X KTC310I 21:46:41 KTLOAD STEM(KTSTEM.) MSG(KTMSG.) PREFIX(VAREND.HTSMCMD) TPG(A
X KTC311I 21:46:41 KTLOAD RC=0
X KTC310I 21:46:41 KTQRYTPG STEM(KTSTEM.) MSG(KTMSG.)
X KTC311I 21:46:41 KTQRYTPG RC=0
X KTC310I 21:46:41 KTLOAD STEM(KTSTEM.) MSG(KTMSG.) PREFIX(VAREND.HTSMCMD) TPG(A
X KTC311I 21:46:41 KTLOAD RC=0
X KTC310I 21:46:41 KTDCTTPG STEM(KTSTEM.) MSG(KTMSG.)
X KTC311I 21:46:41 KTDCTTPG RC=0
X KTC310I 21:46:41 KTQRYTPG STEM(KTSTEM.) MSG(KTMSG.)
X KTC311I 21:46:41 KTQRYTPG RC=0
X KTC310I 21:48:10 KTLOAD PREFIX(VAREND.HTSMCMD) STEM(tpg.TPG00001.) TPG(ALPHATP
X KTC311I 21:48:10 KTLOAD RC=0
X KTC310I 21:48:10 KTQRYTPG STEM(tpg.TPG00001.) MSG(msg.)
X KTC311I 21:48:10 KTQRYTPG RC=0
****************************** BOTTOM OF DATA ******************************
4A                  1 Sess-1    172.17.51.30                TCP00019        10/12
```

# HDT for Mainframe and HTSM for Mainframe Storage - Summary

- Enables automation and more efficient use of tiered storage – self optimizes
- Improves ability to manage SLAs
- Improves performance

# Hitachi Dynamic Tiering for Mainframe Performance Examples

## HDT Performance Example 1

- This example shows results from HDT testing
  - Demonstrates how HDT learns your workload
- Scenario: Customer reluctant to upgrade from 300GB to 600GB HDD
- Same capacity of HDD (not Including SSD)
  - (128) 300GB SAS
  - (64) 600GB SAS + (8) 400GB SSD
- IMPORTANT NOTE: SSD drives are added to the pool after all data sets are created

***Disclaimer:*** *PAIO workload tests used only demonstrate Dynamic Tiering behavior and concept. These tests are not intended for benchmarking purposes and results may vary depending on the workload used and the systems in the environment.*

# Basic Configuration

| Config. Name | RAID Type | LCU | DP-VOL per Pool | PAIO Data-set | Base/ Alias | Dev. Num. | Desc. |
|---|---|---|---|---|---|---|---|
| HDT3HF | RAID-6(6 D+2P) | 00 - 03 | 256 | 1024 | 64/192 | 70xx | (128) 300GB SAS HDP pool |
| HDT6HF | RAID-6(6 D+2P) | 08 – 0B | 256 | 1024 | 64/192 | 72xx | (64) 600GB SAS HDP pool |
| HDT6HF Run 1 through Run 4 | RAID-6(6 D+2P) | 08 – 0B | 256 | 1024 | 64/192 | 72xx | HDT pool (8) 400GB SSD (64) 600GB SAS |

# 300Gb and 600GB HDP BaselineS have been run (NO SSD Drives)

HDP (64)
600GB HDD

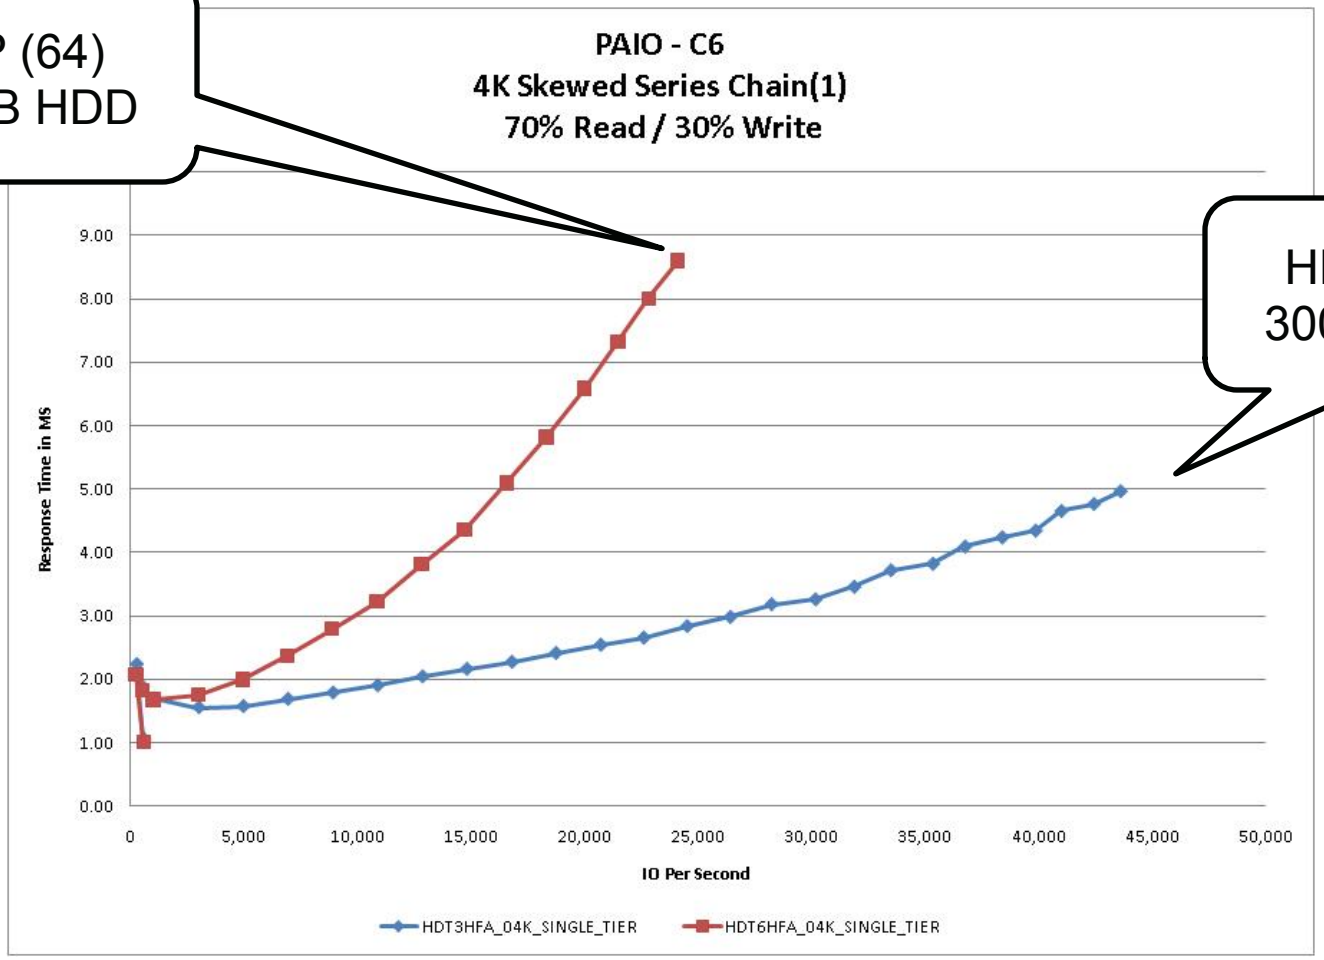


HDP (128)
300GB HDD

# First Run: 600GB Tier 2 + SSD Tier 1 – 0 minutes – No Learning

LET TIERING LEARN YOUR WORKLOAD

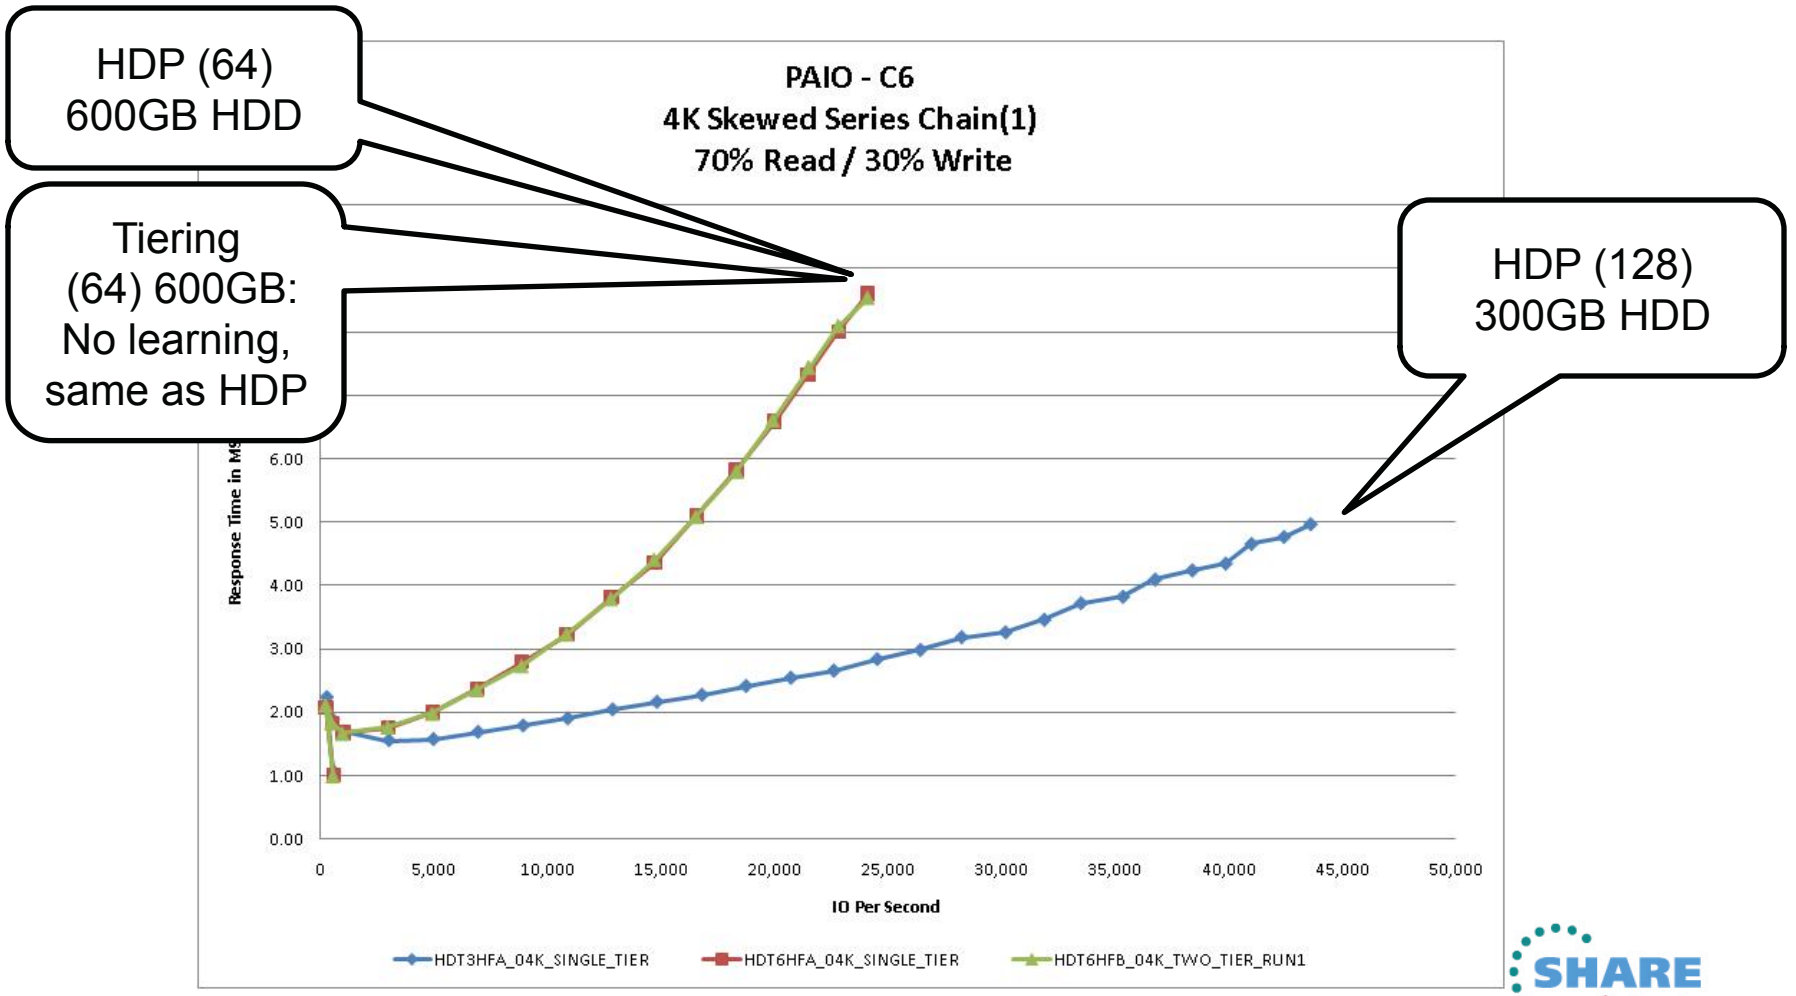

# Second Run: 600GB Tier 2 + SSD Tier 1 – 30 Minutes of Rest after Run 1

LET TIERING LEARN YOUR WORKLOAD

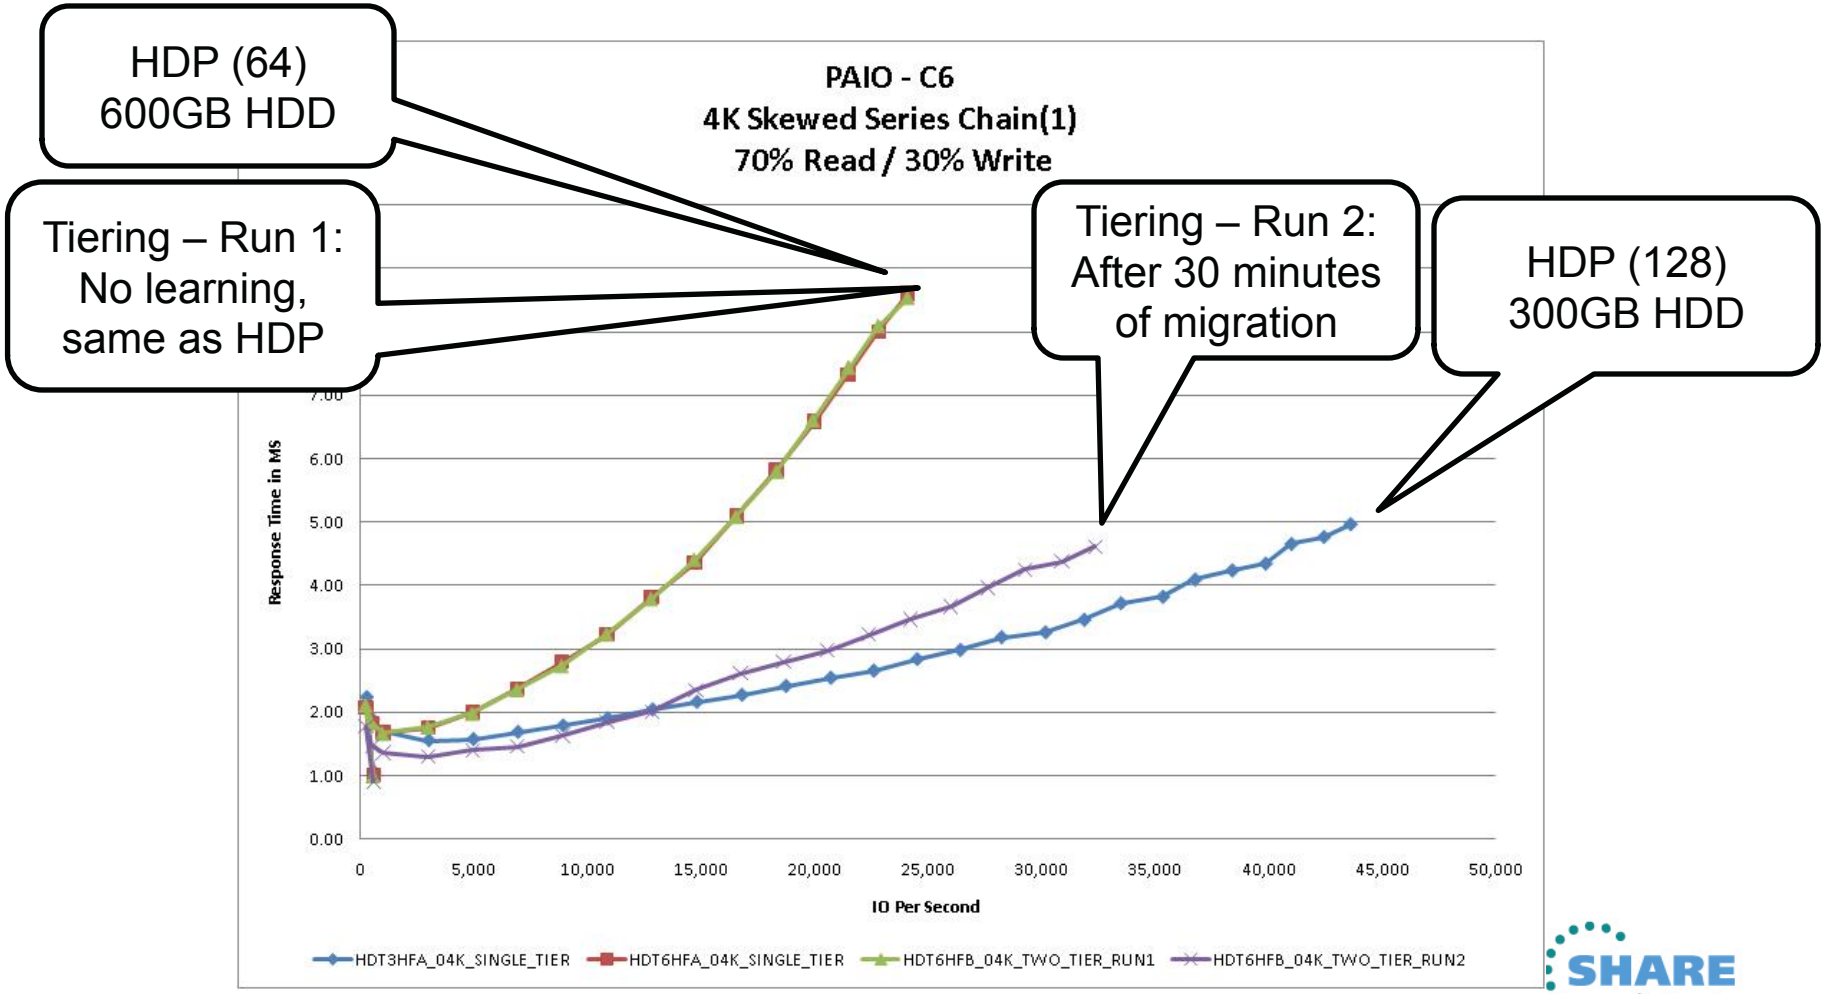

# Fourth Run: 600GB Tier 2 + SSD Tier 1 – 30 Minutes of Rest after Run 3

LET TIERING LEARN YOUR WORKLOAD

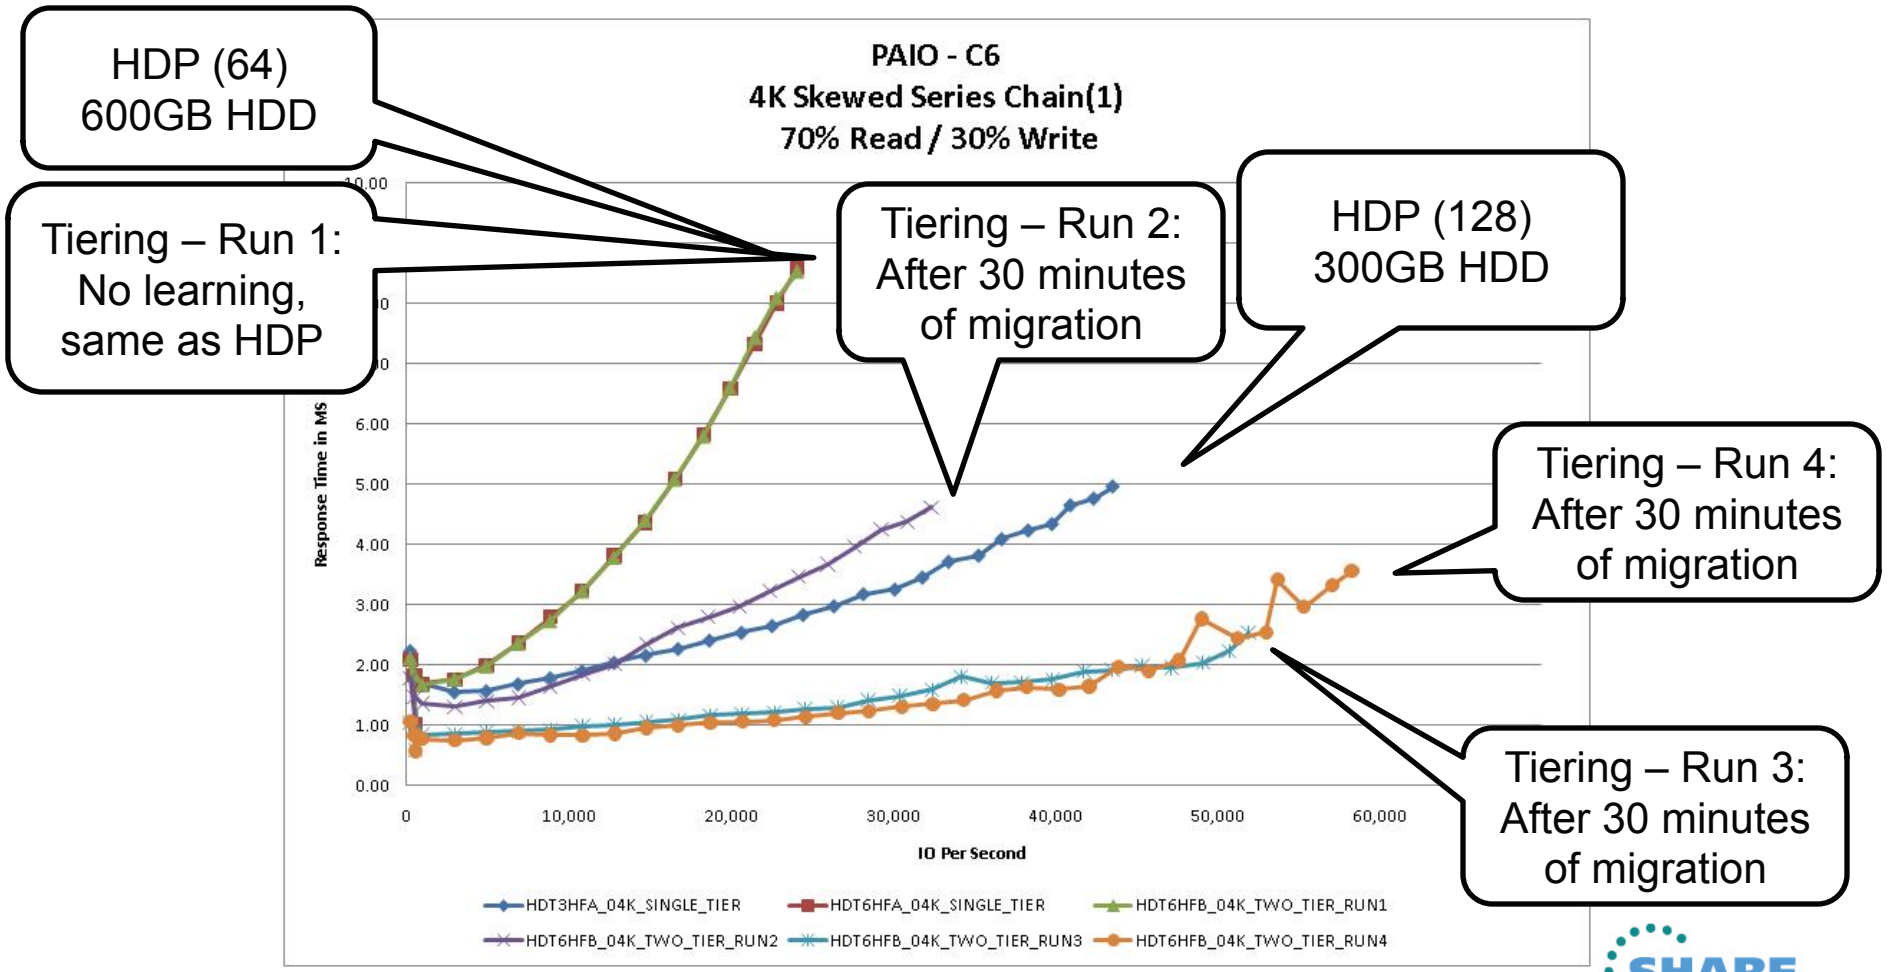

# HDT Performance Example 2

- This example is another way to show how HDT learns your workload
- Same 600GB tier as previous experiment except at a steady state of 24K IOPS
  - (64) 600GB SAS drives + (8) 400GB SSD
- IMPORTANT NOTE: SSDs are added to the pool after all data sets are created

# 24K IOP Steady State

## LET TIERING LEARN YOUR WORKLOAD

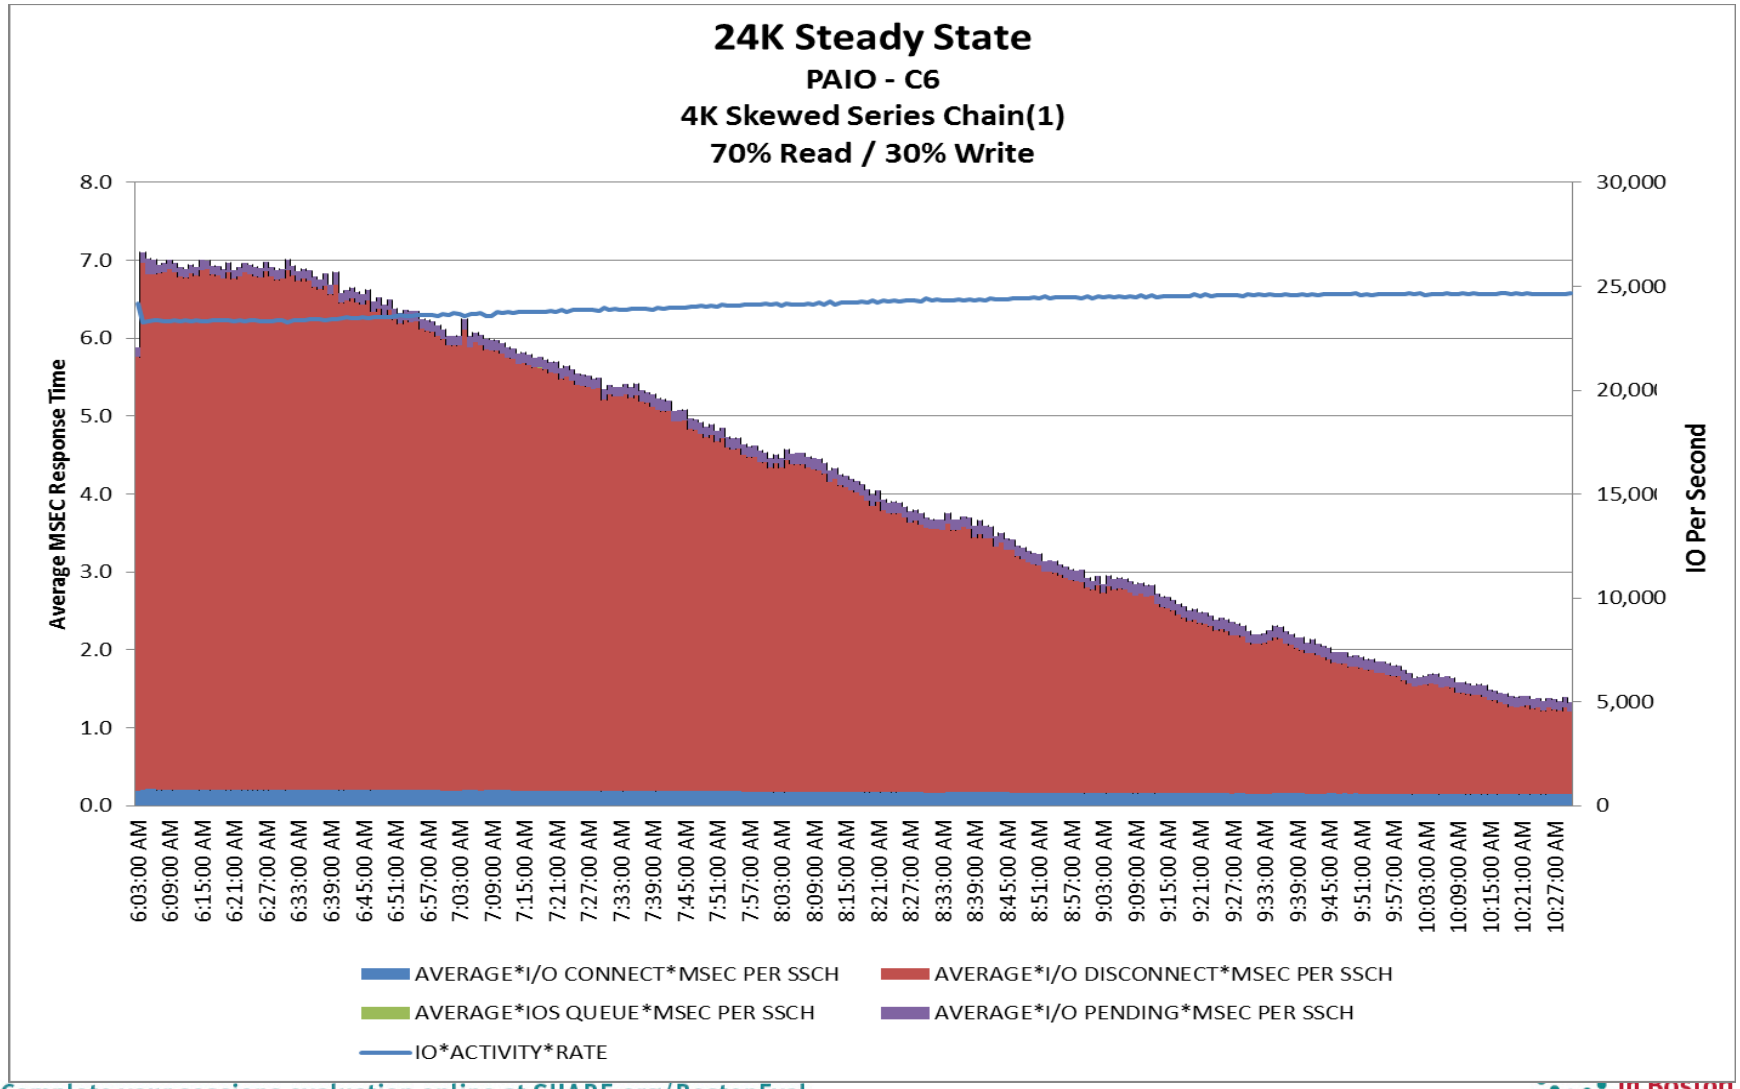

# Observations

- After HDT "learned" the access patterns, the throughput doubled on a smaller storage footprint
- 90% of the active dataset area was migrated to Tier 1
  - 10% of data did not meet the Tier 1 IO/hour criteria
- HDT can resolve HDD sibling pend contention
  - Migrate to a 2-tier HDT configuration
  - The VSP will learn and migrate pages that will benefit from SSD performance in HDT structure
  - A single SSD parity group can improve performance in a short time

# Questions and Discussion

*Contact the speakers:*

**Ros.Schulman at hds.com**

**John.Harker at hds.com**

# Thank you

*Contact the speakers:*

**Ros.Schulman@hds.com**

**John.Harker@hds.com**