Tiering Solutions from Hitachi

Ros Schulman
HDS (Ros.Schulman@hds.com)
@roselinda_s

August 5th, 2014
Session Number (16058)

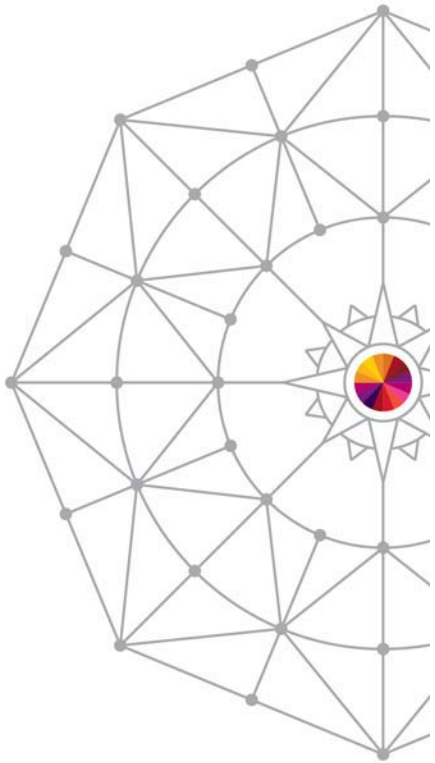

#SHAREorg

SHARE is an independent volunteer-run information technology association that provides education, professional networking and industry influence.

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

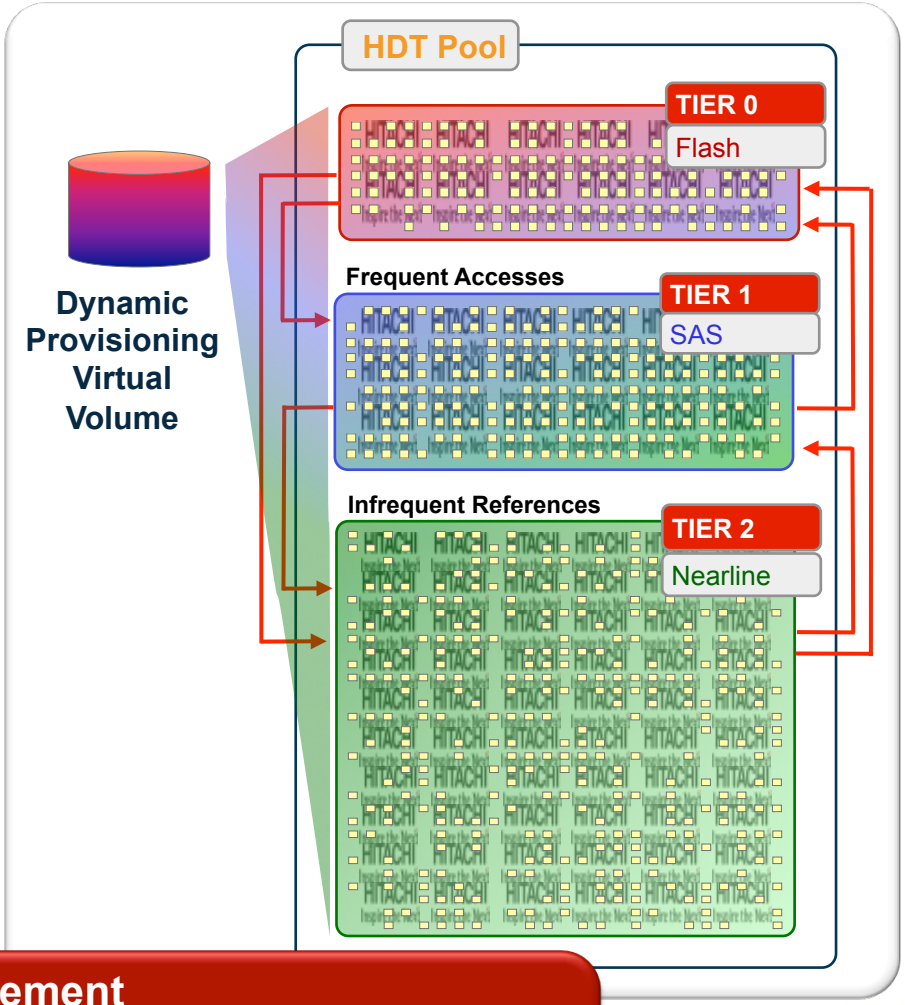

- **Simplifies operations and data management**
- **Reduces opex, capex, and TCO**
- **Data Moves in 38MB Pages**
- **Datasets can span multiple Tiers**

Improved Performance at Reduced Cost: Data Locality and Pareto Distributions

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

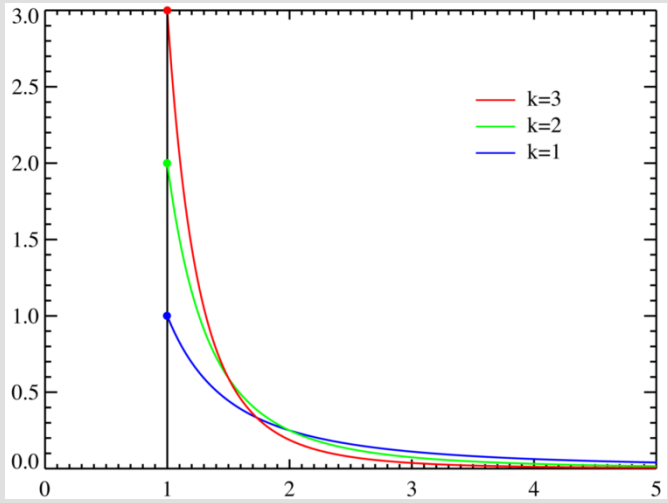

Actual Volume Workload

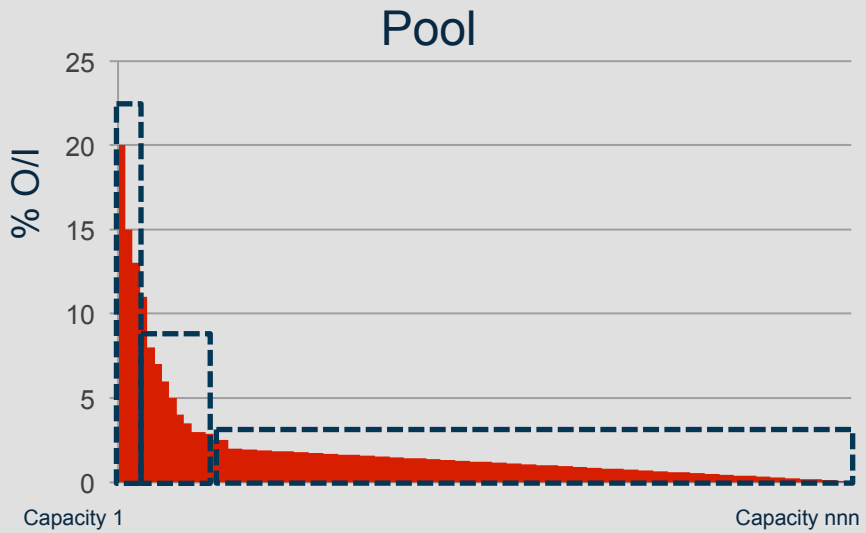

Why does it work so well?

- 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

Backend IO Distribution vs Physical Capacity

Physical Capacity Distribution

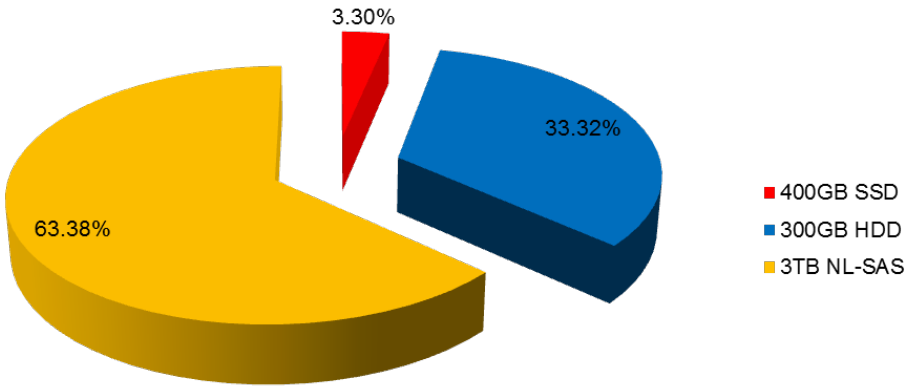

IO Rate Distribution

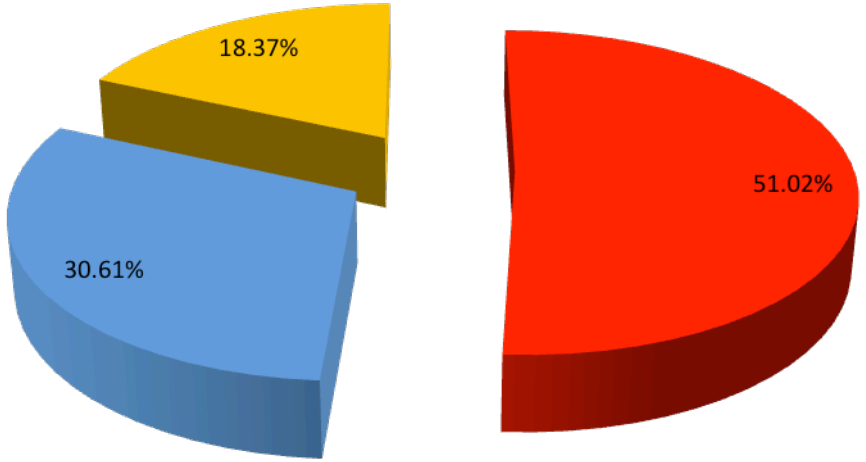

ENVIRONMENTAL COMPARISON

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 % Savings	55.45%	55.40%	62.38%	62.46%	62.06%	47.04%	50.00%

With 3X+ Capacity

LIMITS OF SMS Storage Groups and ACS Routines

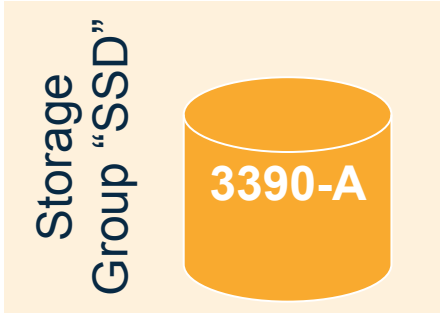

- 3390 volumes are "fixed" to a single tier
- To transition a volume's data to another tier requires ACS work, then migration and recall

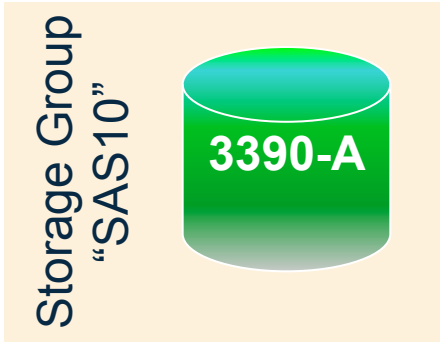

- Stale datasets are treated the same as active ones until HSM migration

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

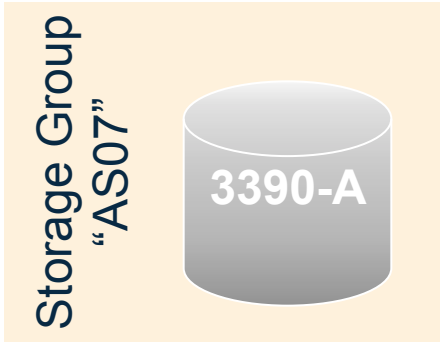

New! HDT Volume-provisioned storage group

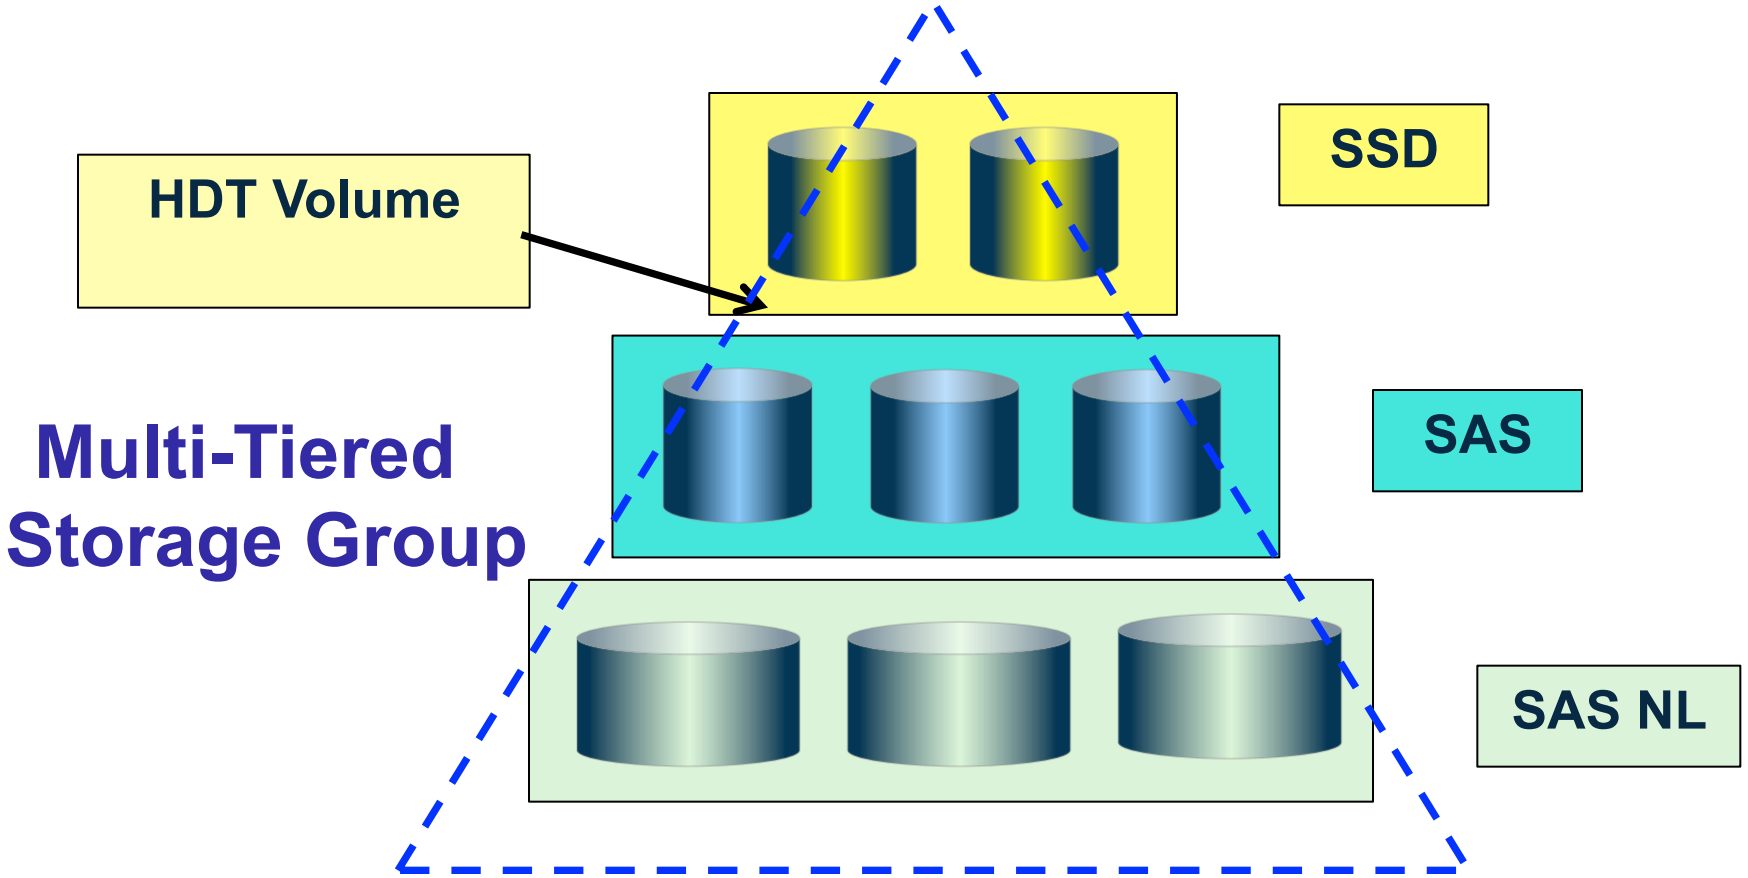

DFSMS

Storage Groups 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

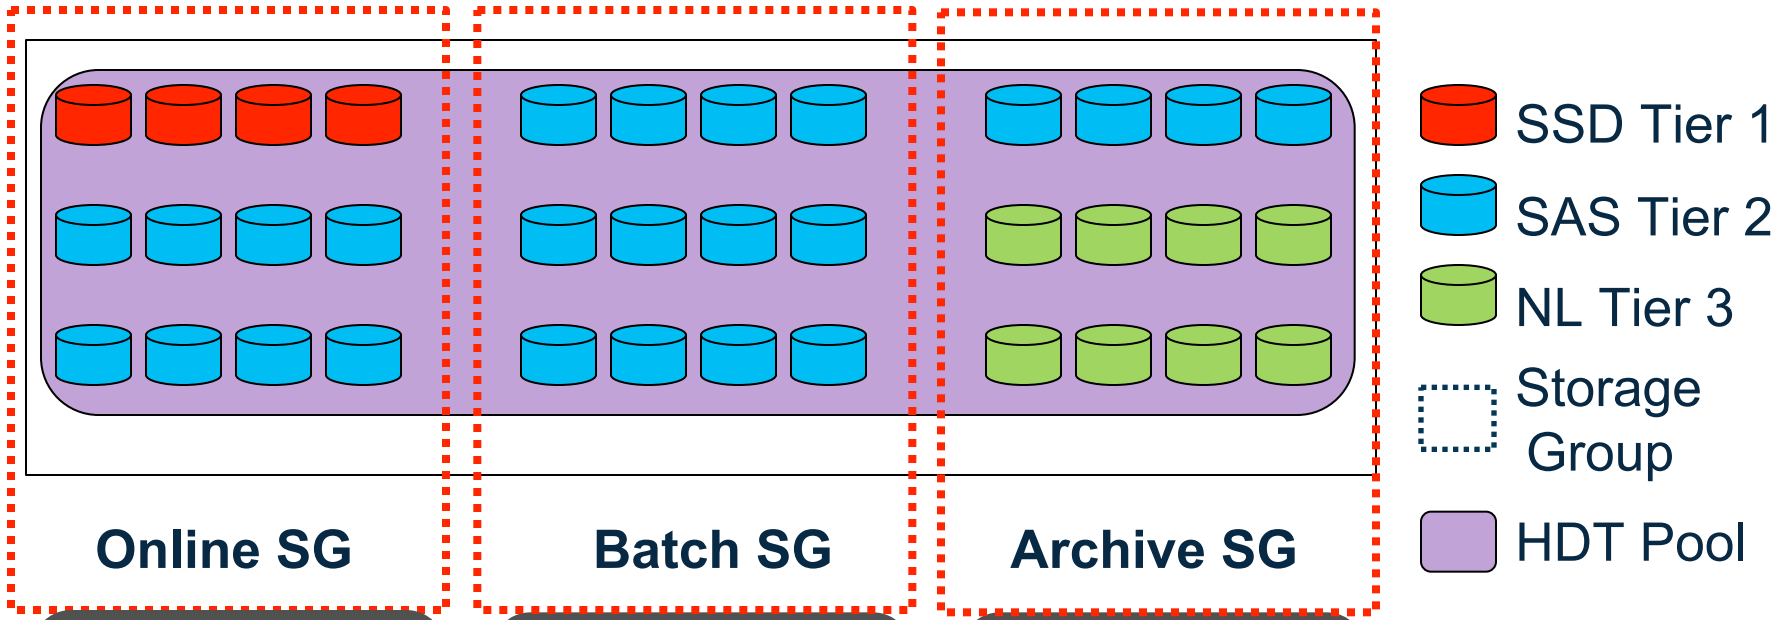

HDT Custom Policy Defined for Online Data Across Top 2 Tiers

HDT Policy Restricted to Tier 2 Residency

HDT Policy Restricted to Tiers 2 and 3 Residency

HDT and DFSMS/HSM integration Example

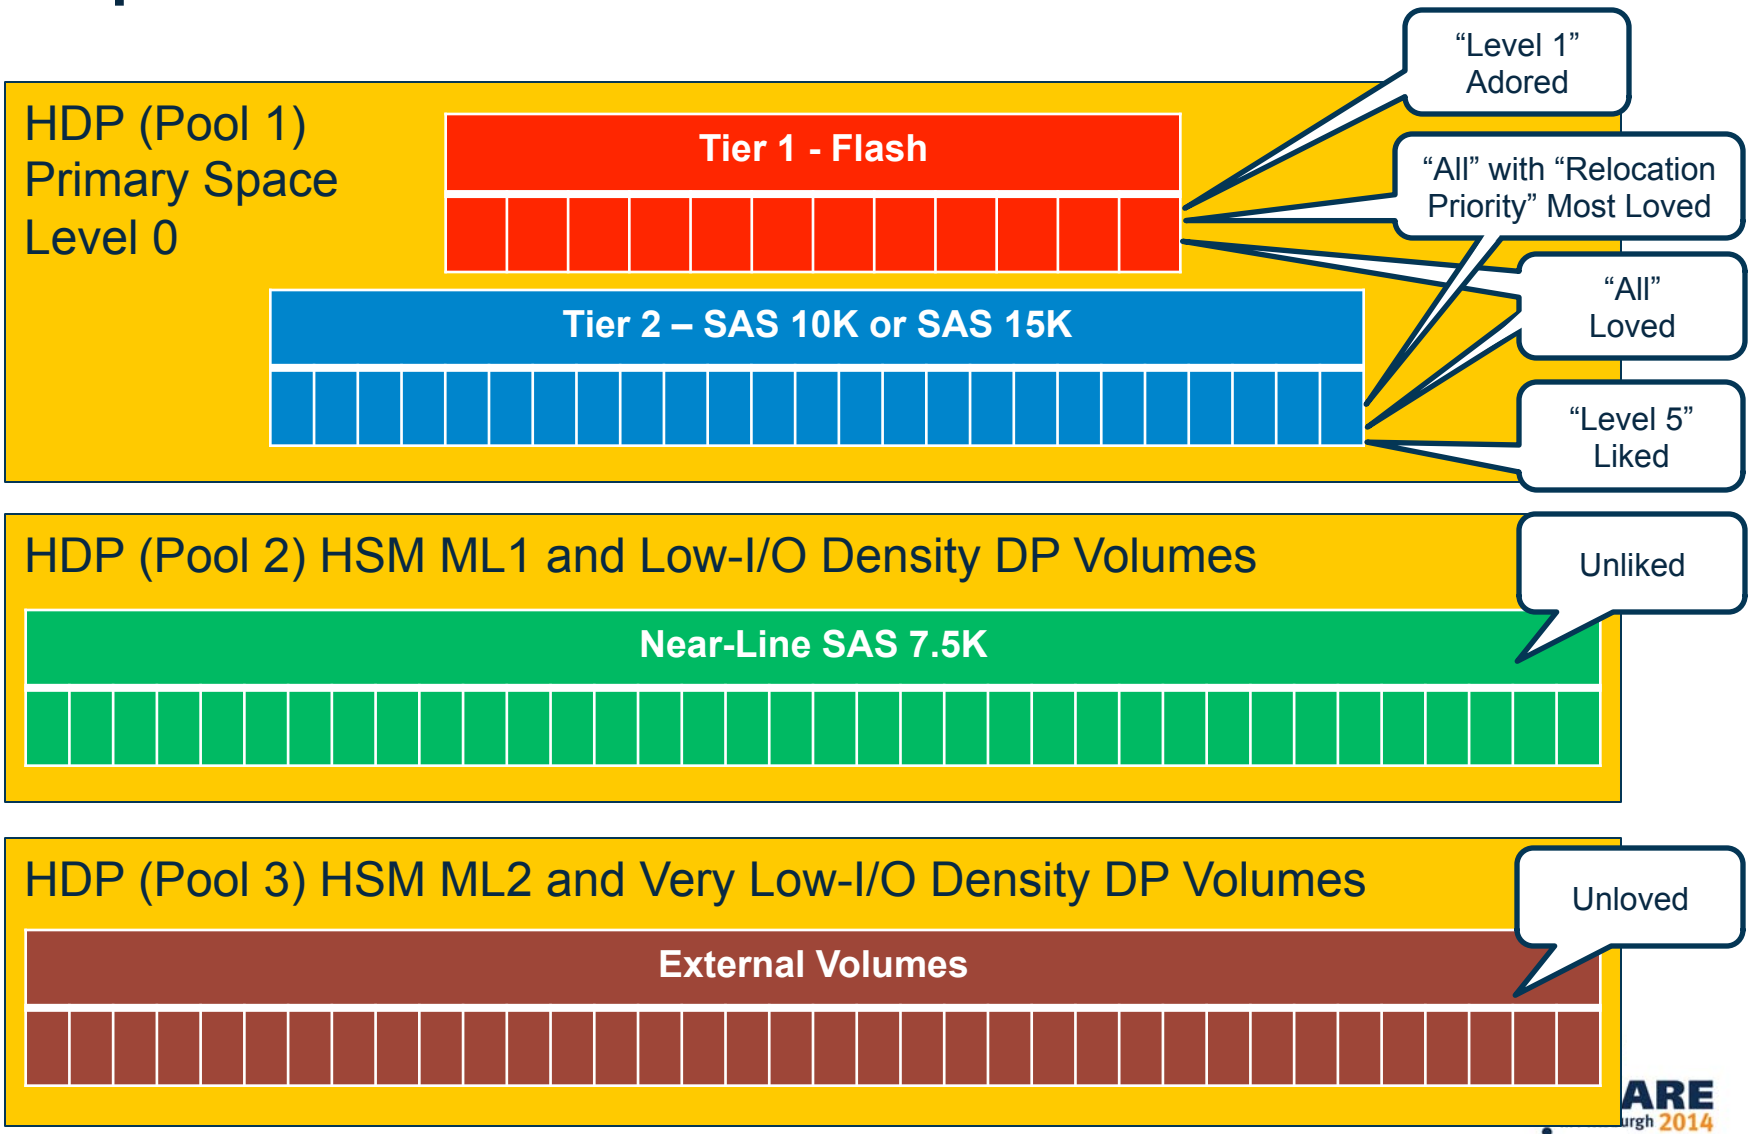

Hitachi Dynamic Tiering and HTSM for MF with DFSMShsm 2.1

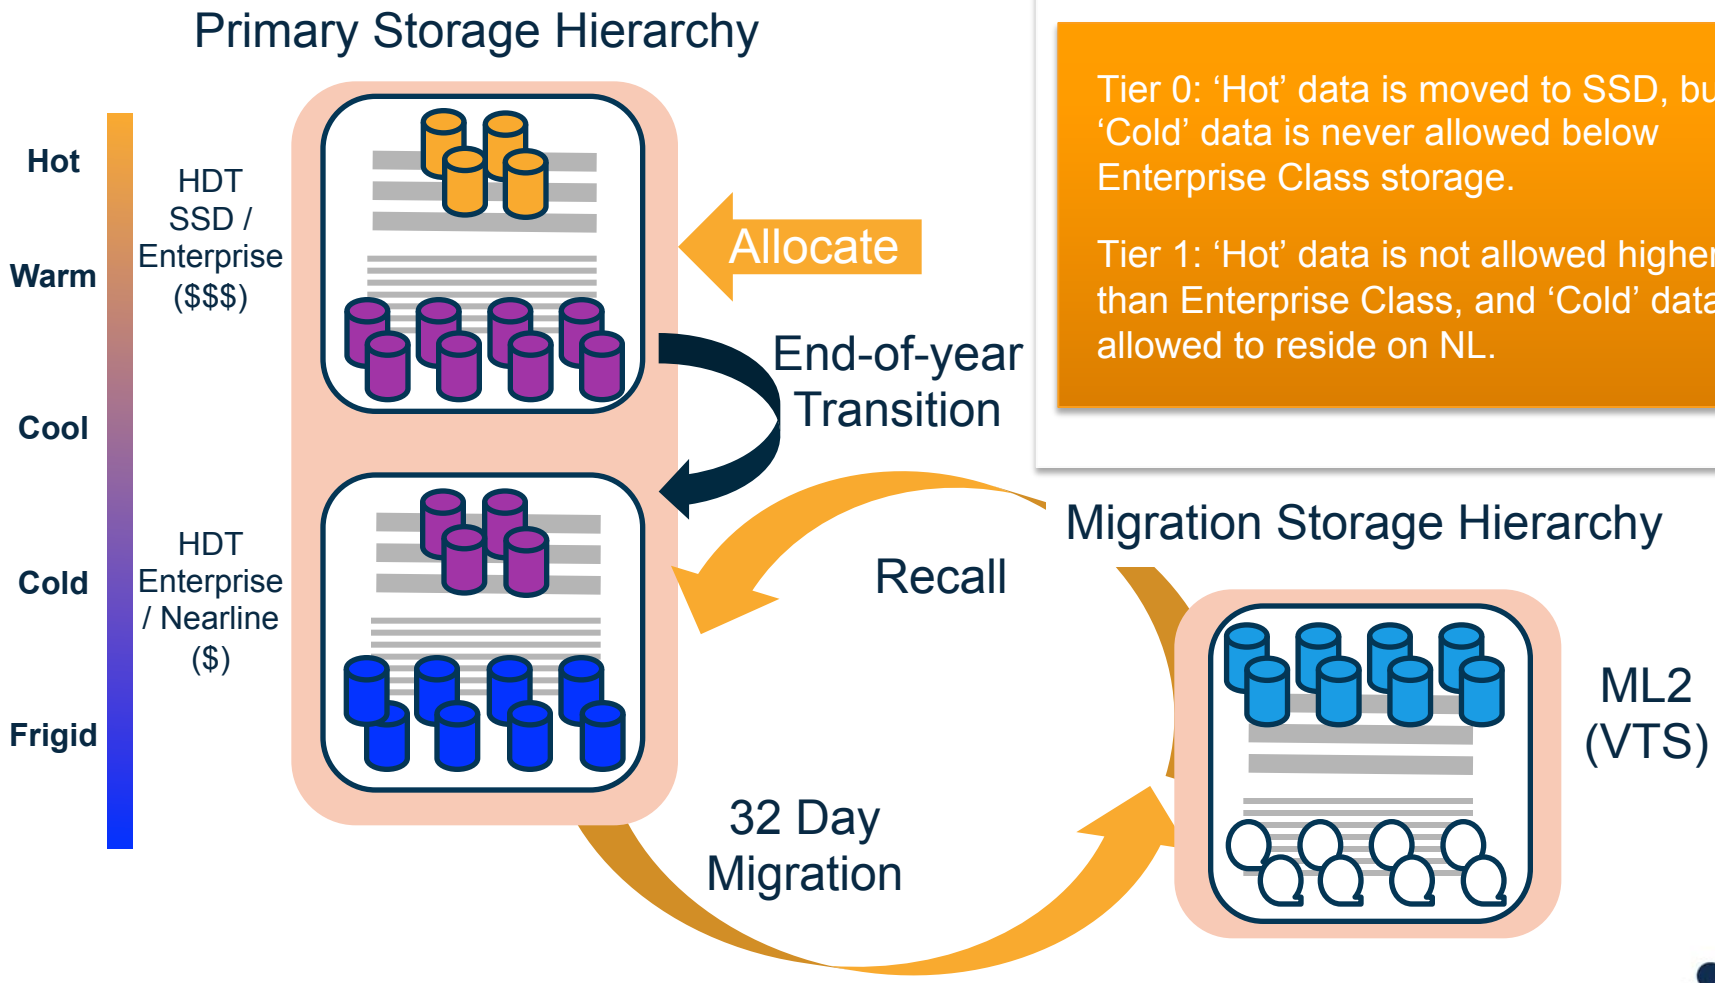

Tier 0: 'Hot' data is moved to SSD, but 'Cold' data is never allowed below Enterprise Class storage.

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

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
  - ISPF interface
- Auto-discovery eliminates errors
  - Accelerates deployment
- Enables reporting and automatic notifications

ISPF ease of use with Point and Shoot

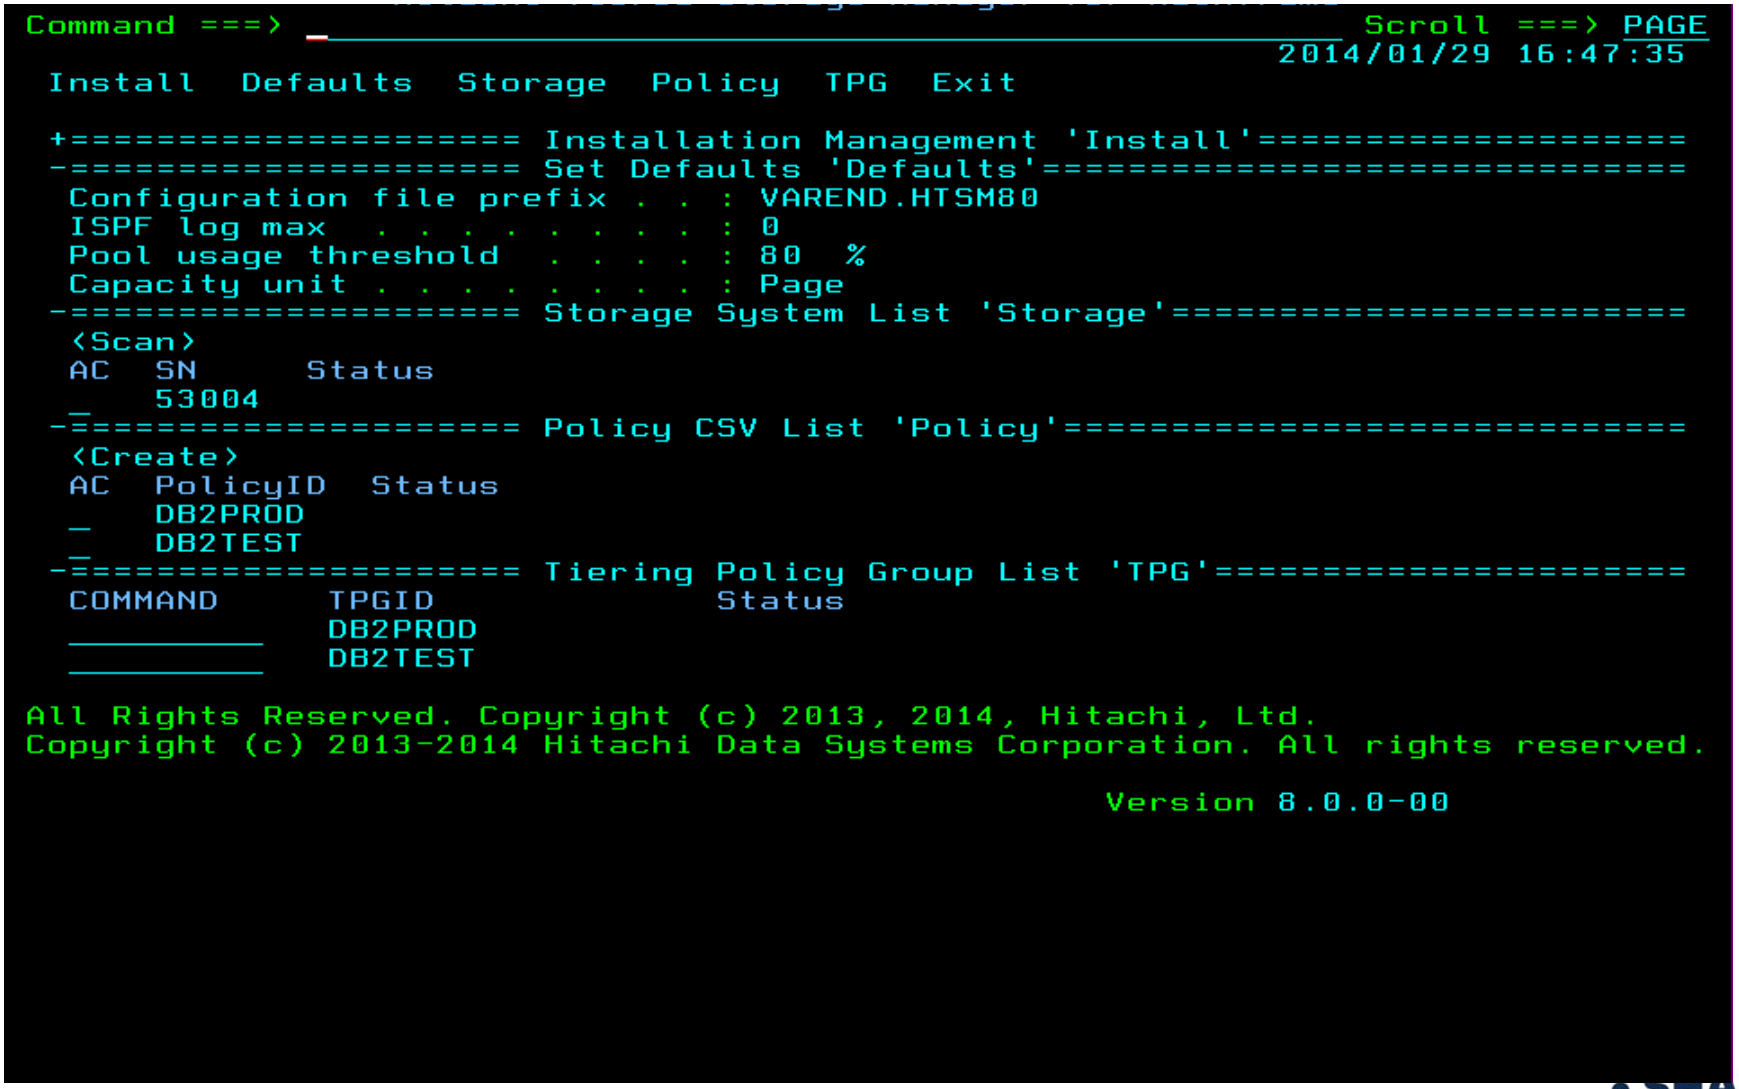

- **TPG_QUERY_STATISTICS** Sample Output

HTSM Mainframe Reporting

```
     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
```

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

Thank You

Ros Schulman
HDS (Ros.Schulman@hds.com)
@roselinda_s

August 5th, 2014
Session Number (16058)

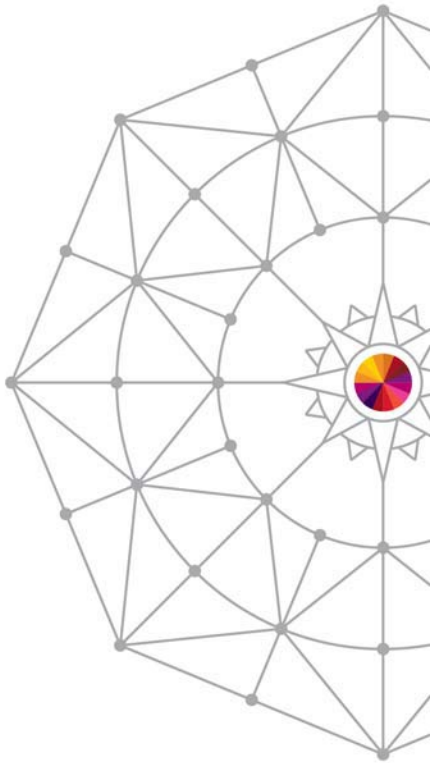

Insert Custom Session QR if Desired.

#SHAREorg

SHARE is an independent volunteer-run information technology association that provides education, professional networking and industry influence.