Easy Tier for z/OS Deep Dive

Nick Clayton

Solution Architect for DS8000

IBM Systems Division

#SHAREorg

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

SHARE in Seattle 2015

Agenda

Actions | Filter

- Tiering Functionality for z/OS
- Easy Tier review
- New Easy Tier functionality
  - Easy Tier Policy and Controls
  - Easy Tier z/OS integration
- Tiered storage configurations for z/OS
- Easy Tier reporting
- Summary

Tiering Functionality for z/OS

- DFSMS
  - Storage groups provide ability to assign a dataset to a group of volumes
  - Policy based criteria control allocation and management
  - Now available for distributed storage with Spectrum Scale functionality
- DFHSM
  - DFHSM provides ability to migrate and recall data from offline storage
  - Migrated data is not accessible by user until recalled to primary storage
  - Storage Tiers functionality provides transition between multiple online tiers
- Disk Subsystem based tiering
  - Volume based tiering assigns volumes to distinct classes of storage
  - SubLUN based tiering (Easy Tier) allows a volume to reside on multiple tiers

Easy Tier automated tiering

- Optimisation of backend storage resources based on historical performance data
- SubLUN granularity using native DS8000 extents for any volume type
- Flexible configurations with any combination of drives of any size and speed
- Easy Tier Application provides APIs for policy and proactive actions
- Easy Tier HeatMap transfer enables workload history to be transferred for replication scenarios (DR, migration etc)

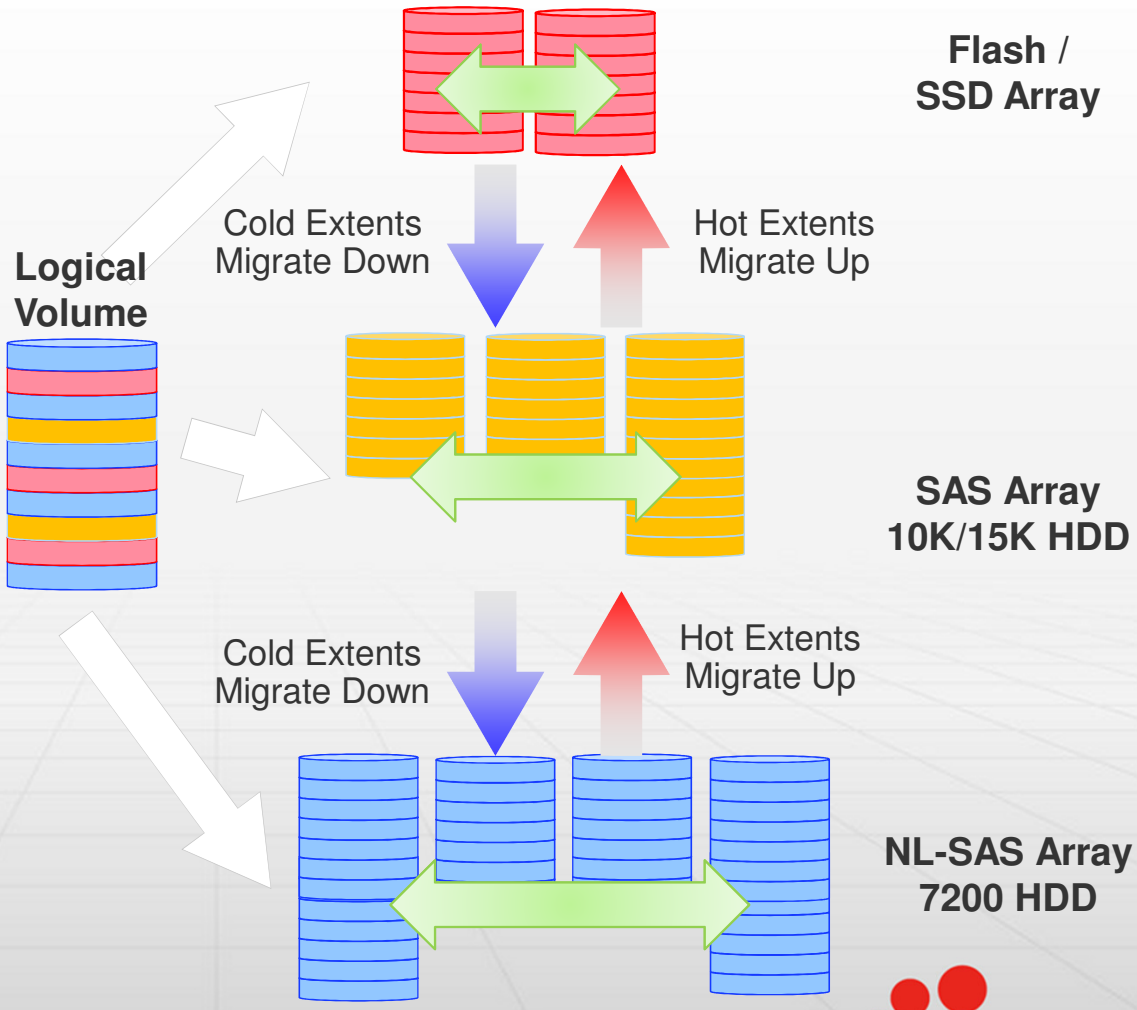

Easy Tier Processing Cycle

Performance data collected ever 5 minutes and incorporated into history

Data collected is for backend activity not IO from hosts

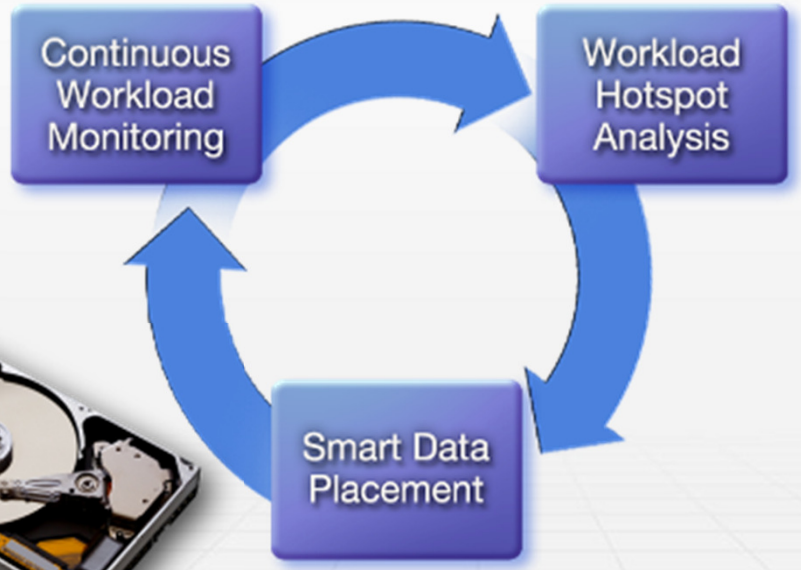

Workload analysis performed at various intervals – 5 minutes, ~6 hours and ~24 hours

Extent “heat” is categorised based on small and large IO activity

Movement of extents scheduled

Cost benefit analysis to ensure excessive movement is avoided

Easy Tier Terminology - Heat and Skew

- Hot data
  - Hot data is simply extents that have more I/O workload, relatively speaking, when compared to other extents within the pool and tier
- Cold data
  - Cold data either has low (or no) I/O workload. Cold data would not benefit from a higher tier and thus is not promoted and is considered for demotion
- Warm data
  - Warm data is the rest of the workload that is not considered hot or cold. Warm data could be promoted – but that would depend on the workload level and available resources
  - Hot and Warm data will reside on Flash/SSD to maximize capacity
- Skew
  - Highly skewed workload has a small number of hot extents
  - Low skewed workload has a more even distribution of workload to extents

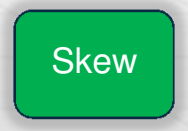

Workload skew drives Easy Tier benefits

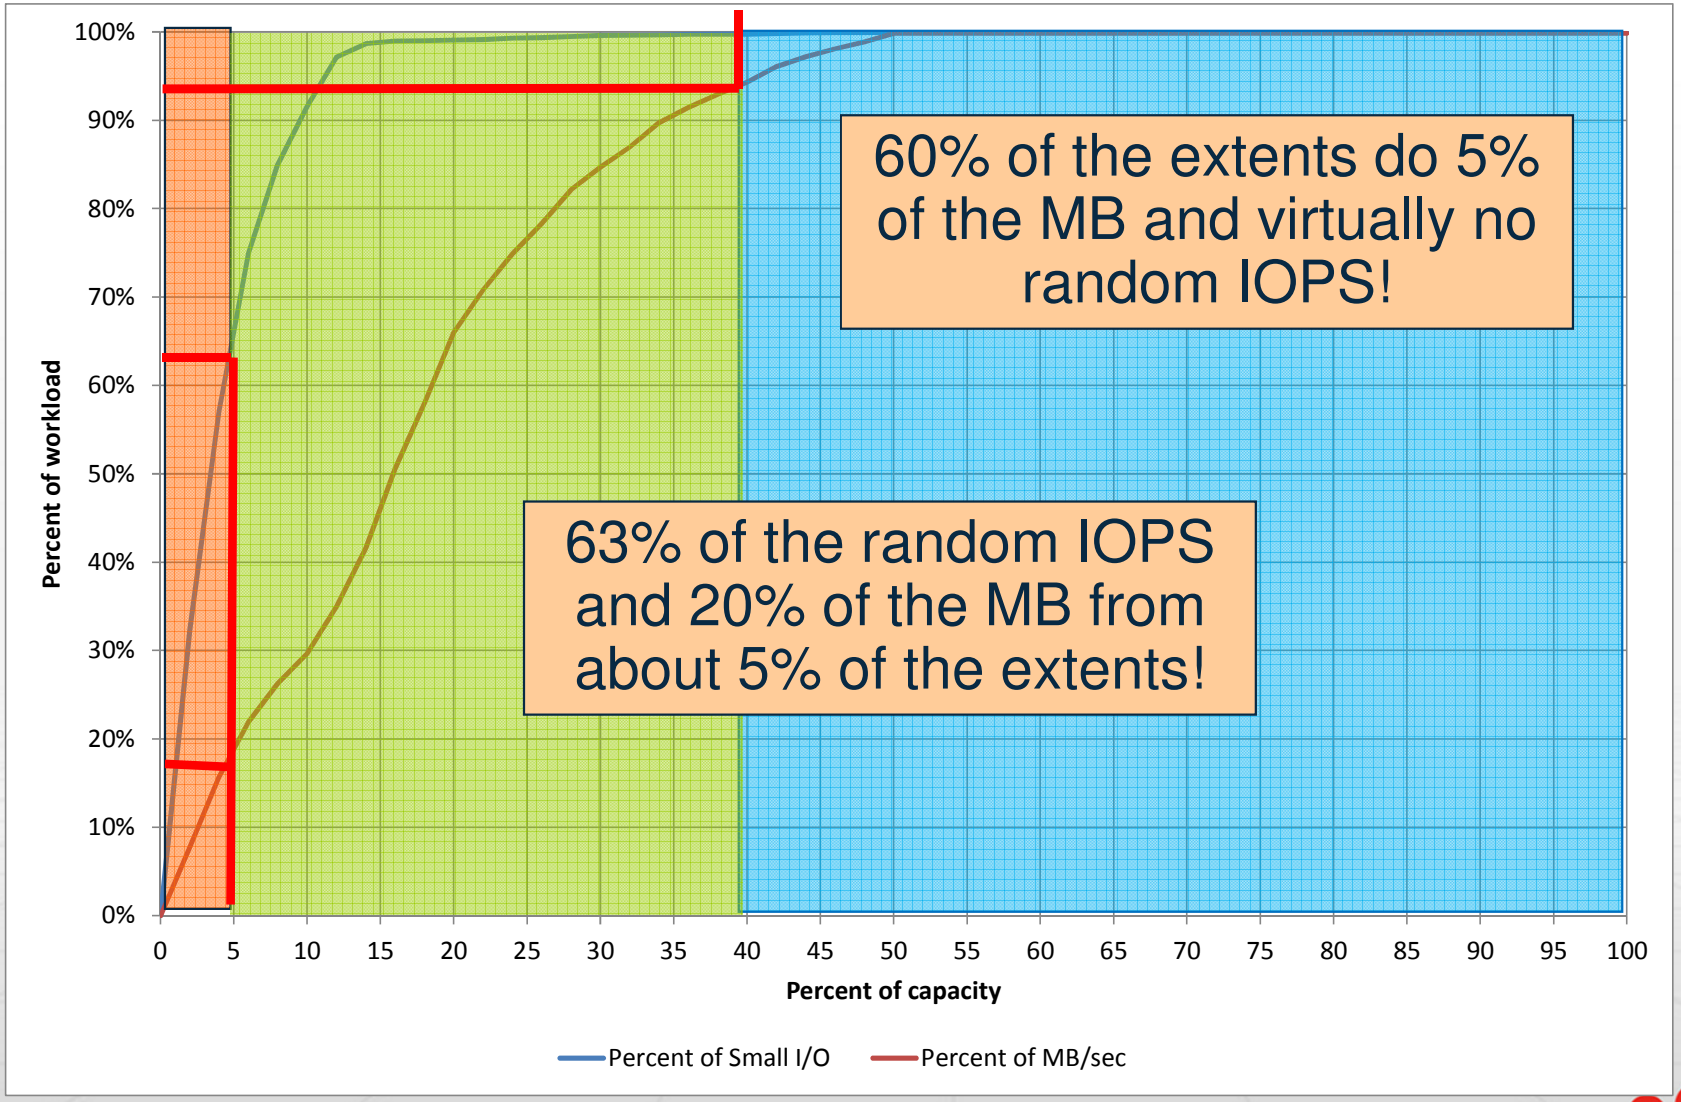

z/OS data from banking environment

Easy Tier Policies

- New Exclude Nearline tier assignment policy
- Prevents the extents of a volume from being demoted to Nearline arrays
- If data is already on Nearline it will be promoted to Enterprise drives
- Three common use cases for Easy Tier Application policies
  - Default – optimise use of all tiers
  - Exclude Nearline – avoid potential low performance
  - Assign Flash – high performance guaranteed
- Also possible to assign to Enterprise or assign to Nearline but less common use cases

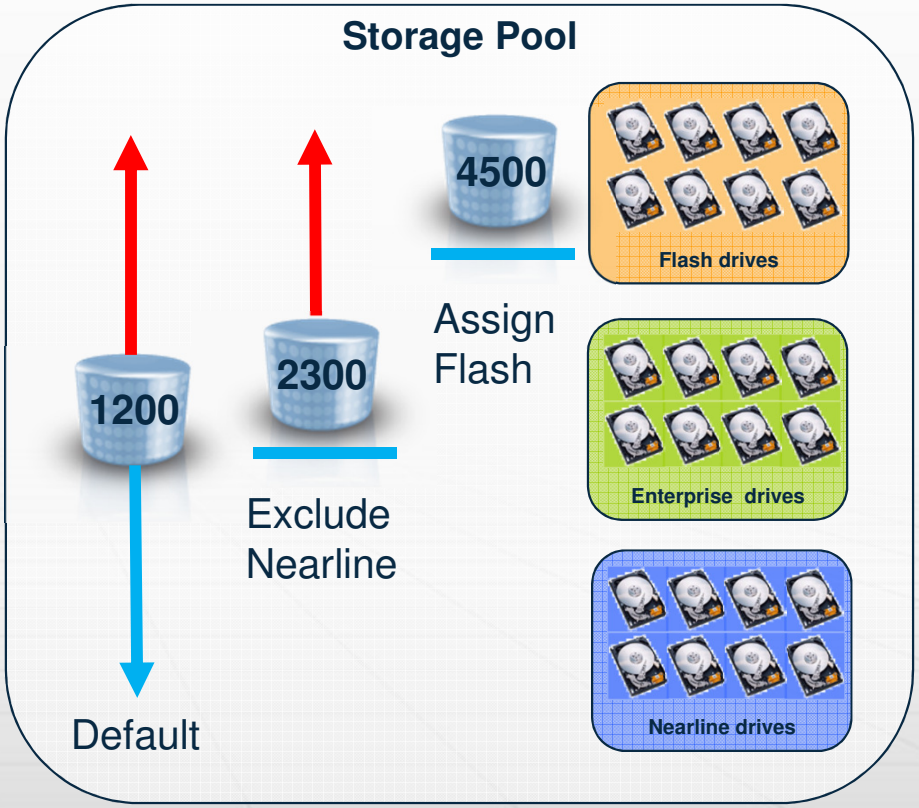

Easy Tier Controls

- In the majority of environments Easy Tier is able to successfully use the history of workload performance to predict the future requirements
  - There are however cases where this is not true
- Easy Tier Controls provide mechanisms for proactively and reactively modifying Easy Tier behaviour to handle these situations
- Controls include
  - Pause and Resume Easy Tier learning for volume or pool
  - Reset Easy Tier learning for volume or pool
  - Pause and Resume Easy Tier migration for a pool

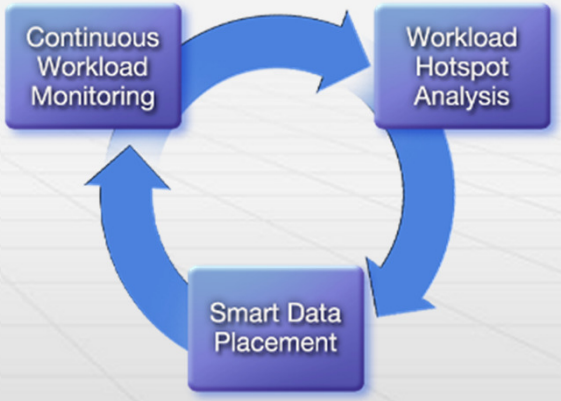

Easy Tier Application
Integration with DFSMS and DB2

- Easy Tier currently optimises data placement and tiering based on workload history and this does not always reflect the future performance requirements of the data
- Easy Tier provides interfaces to enable software such as DFSMS and DB2 to provide hints when data has been created, moved or deleted
- This will avoid performance degradation following maintenance activities such as database reorganisation
- DB2 integration PTFs not yet available

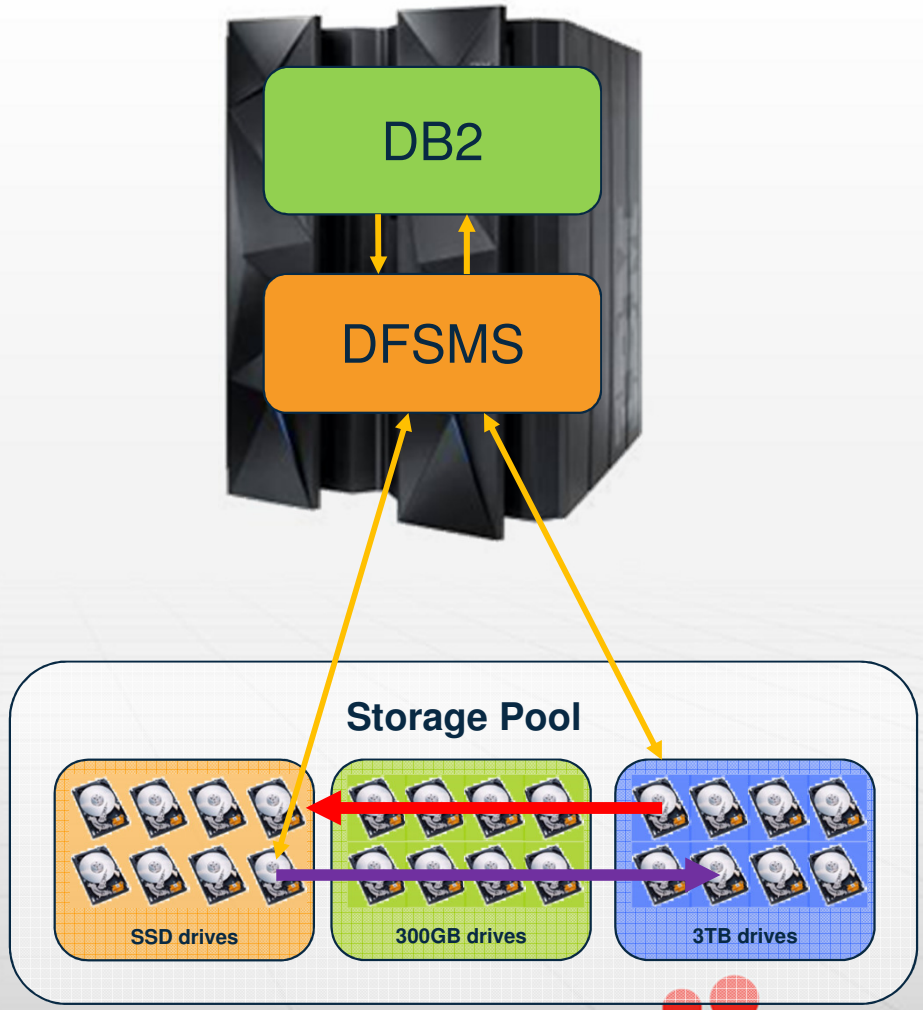

Storage Tiers and Easy Tier

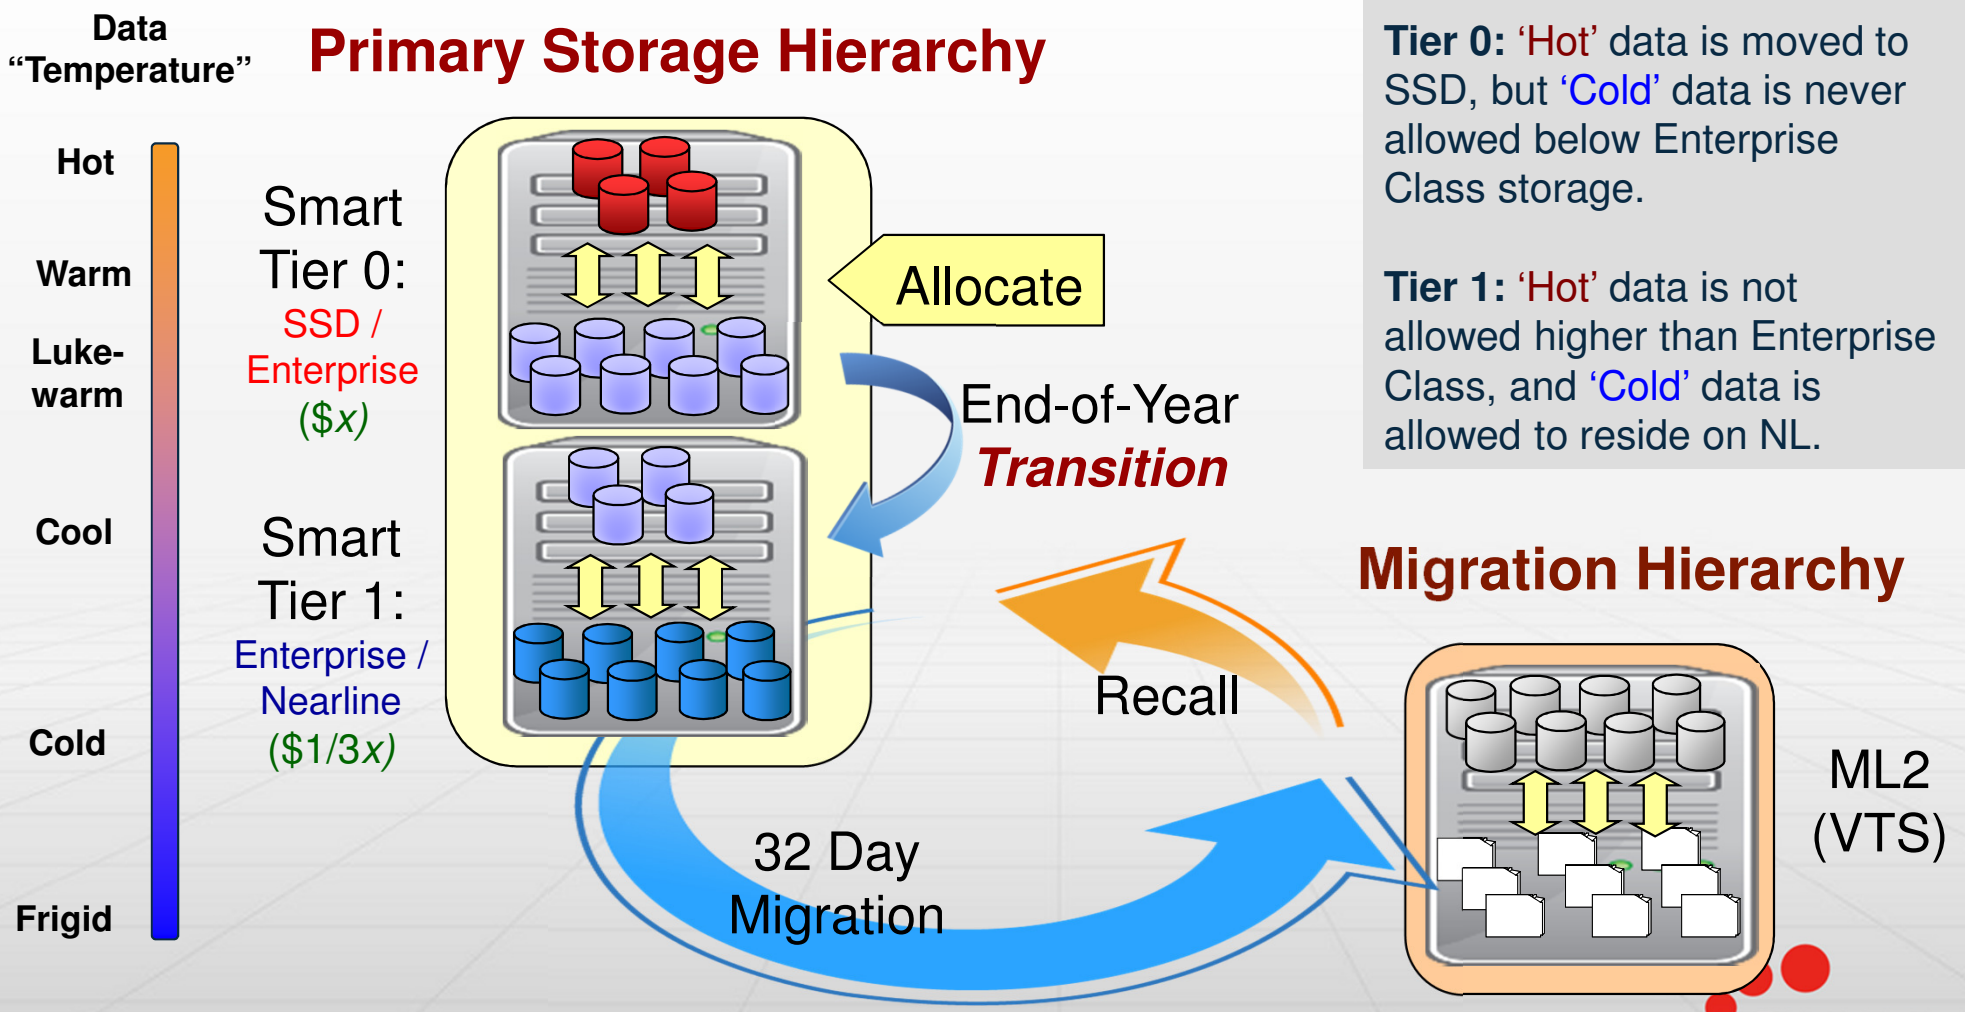

Tier 0: 'Hot' data is moved to SSD, but 'Cold' data is never allowed below Enterprise Class storage.

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

DS8870 Drive Technology

Performance

- Flash – 1.8” in High Performance Flash Enclosure
  - 400 GB drive
- SSD – 2.5” Small Form Factor
  - Latest generation with higher sequential bandwidth
  - 200/400/800/1600GB SSD
- 2.5” Enterprise Class 15K RPM
  - Drive selection traditionally used for OLTP and z/OS
  - 146/300/600GB drives
- 2.5” Enterprise Class 10K RPM
  - Large capacity, much faster than Nearline
  - 600GB and 1.2TB drives
- 3.5” Nearline – 7200RPM Native SAS
  - Extremely high density, direct SAS interface
  - 4TB drives

Drive Selection in an Easy Tier Environment

- 3-5% Flash/SSD, 95-97% Enterprise
  - Provides improved performance compared to single tier solution and enables use of larger Enterprise drives
  - All data guaranteed to have at least enterprise performance
- 10-20% (or more) Flash/SSD, 80-90% Enterprise
  - Provides Flash IOPS and Latency
  - Can be combined with selective pinning of data to Flash if required
- 3-5% Flash/SSD, 25-53% Enterprise, 40-70% NL SAS
  - Provides improved performance and density to a single tier solution
  - Significant reduction in environmental costs
- 20-60% Enterprise, 40-80% NL SAS
  - Provides reduced costs and comparable overall performance to a single tier Enterprise solution

Example Easy Tier implementations (1)

- Mainframe Easy Tier implementation on DS8870
- Two tier Easy Tier for production with separate pool of 900GB drives for low performance workloads
  - 75% of random workload on SSD
- Three tier Easy Tier for development

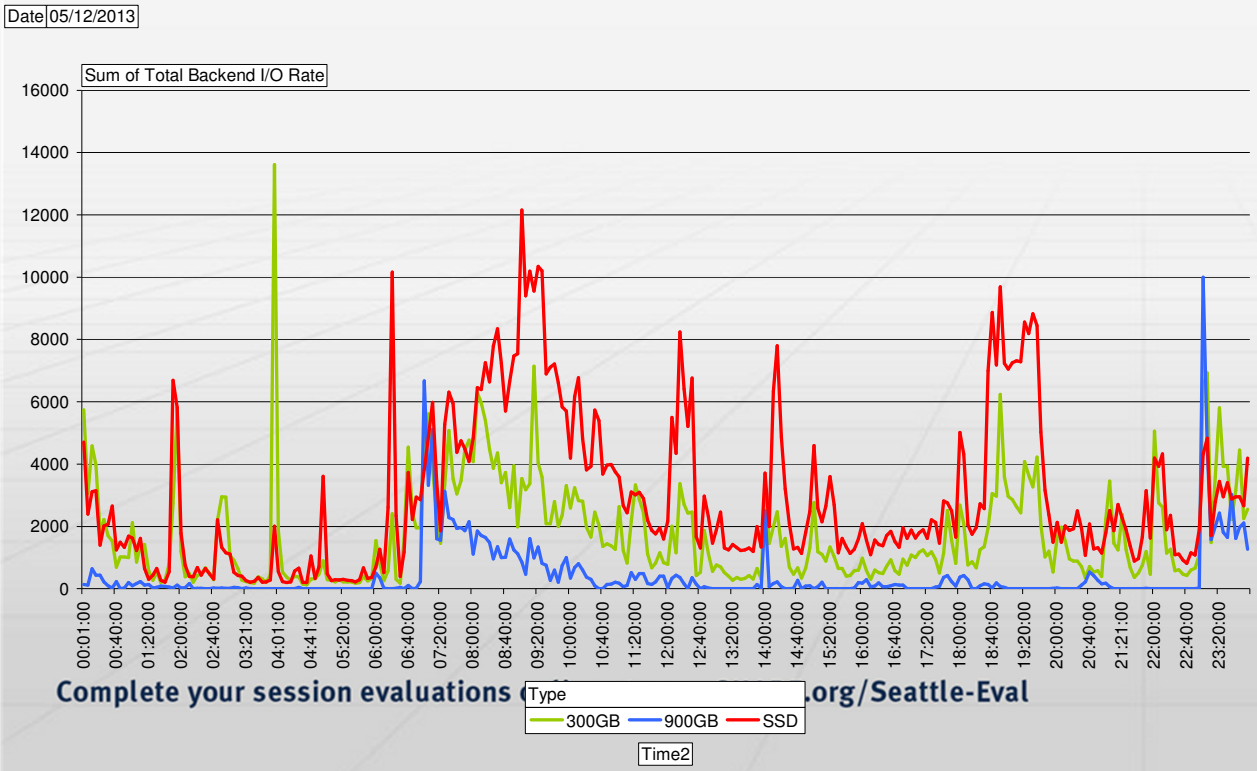

Production

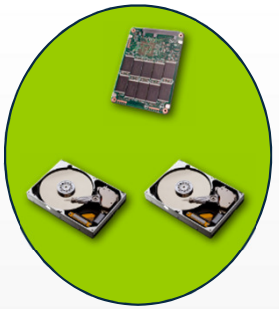

9% 400GB SSD

69% 300GB 15K

22% 900GB 10K

Test/Development

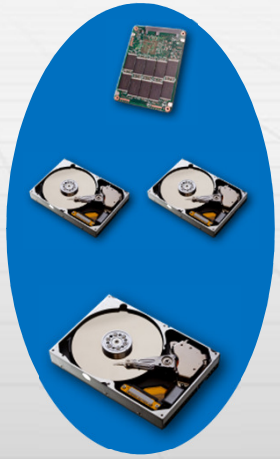

2% 400GB SSD

78% 900GB 10K

20% 3TB 7.2K

Example Easy Tier implementations (2)

- Mainframe Easy Tier implementation on DS8870
- Two tier Easy Tier for production with separate pool of Enterprise/Nearline drives for low performance workloads
  - ~99% of random workload on SSD
- Two tier Easy Tier with Enterprise/Nearline for development

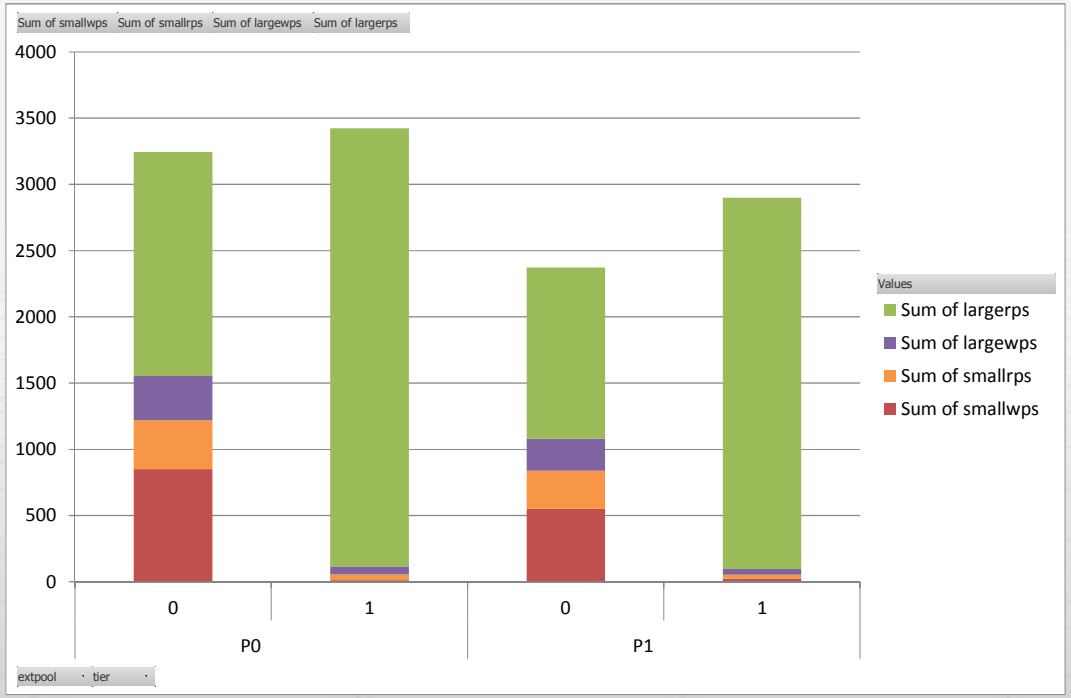

Production

10% 400GB SSD

65% 600GB 10K

5% 600GB 10K

20% 3TB 7.2K

Test/Development

60% 600GB 10K

40% 3TB 7.2K

Easy Tier – Modeling Tools

- Storage Tier Advisor Tool (STAT)
  - Extracts heat data collected by Easy Tier for volumes that are being monitored
- Disk Magic
  - Supports 5 predefined skew levels for prediction with Easy Tier
  - Can utilize detailed Easy Tier data to generate actual client skew chart
  - Utilizes either predefined or actual skew to predict the number of I/Os
    - Higher skew results in a more aggressive sizing
- FLASHDA (z/OS only)
  - Identifies what datasets and devices have the highest accumulated read-only disconnect time
- IBM Tivoli Storage Productivity Center (TPC)

STAT reports review(1)

IBM System Storage® DS8000®

Summary Report

Systemwide Recommendation

System Summary

This report is based on data from Thu Oct 30 11:15:20 2014. Easy Tier has been running continuously since Wed May 21 10:38:34 2014

Storage Tier Advisor Tool version: 9.3.0.0

Storage facility	IBM.2107-SF75CFN60
Total storage pools monitored	2
Total volumes monitored	10452
Total capacity monitored	177156 GiB
Hot data capacity (% of total)	322 GiB (0%)
Data validity	Valid
System state	Latest Warmstart: No Warmstart Latest Failback: No Failback

Storage Pool ID *1	Capacity (GiB)	Configuration	Tier Status*2	Data Management Status*3	
P0	90434	SSD + Enterprise		89932 GiB/100.00%	502 GiB
P1	90434	SSD + Enterprise	*Enterprise is IOPS skewed	87224 GiB/100.00%	3210 GiB

20 Entries Per Page | GO | |< << >> >| Displaying Page 1 of 1

Assigned | Assign In-progress | Easy Tier Managed | Unallocated

STAT reports review(2)

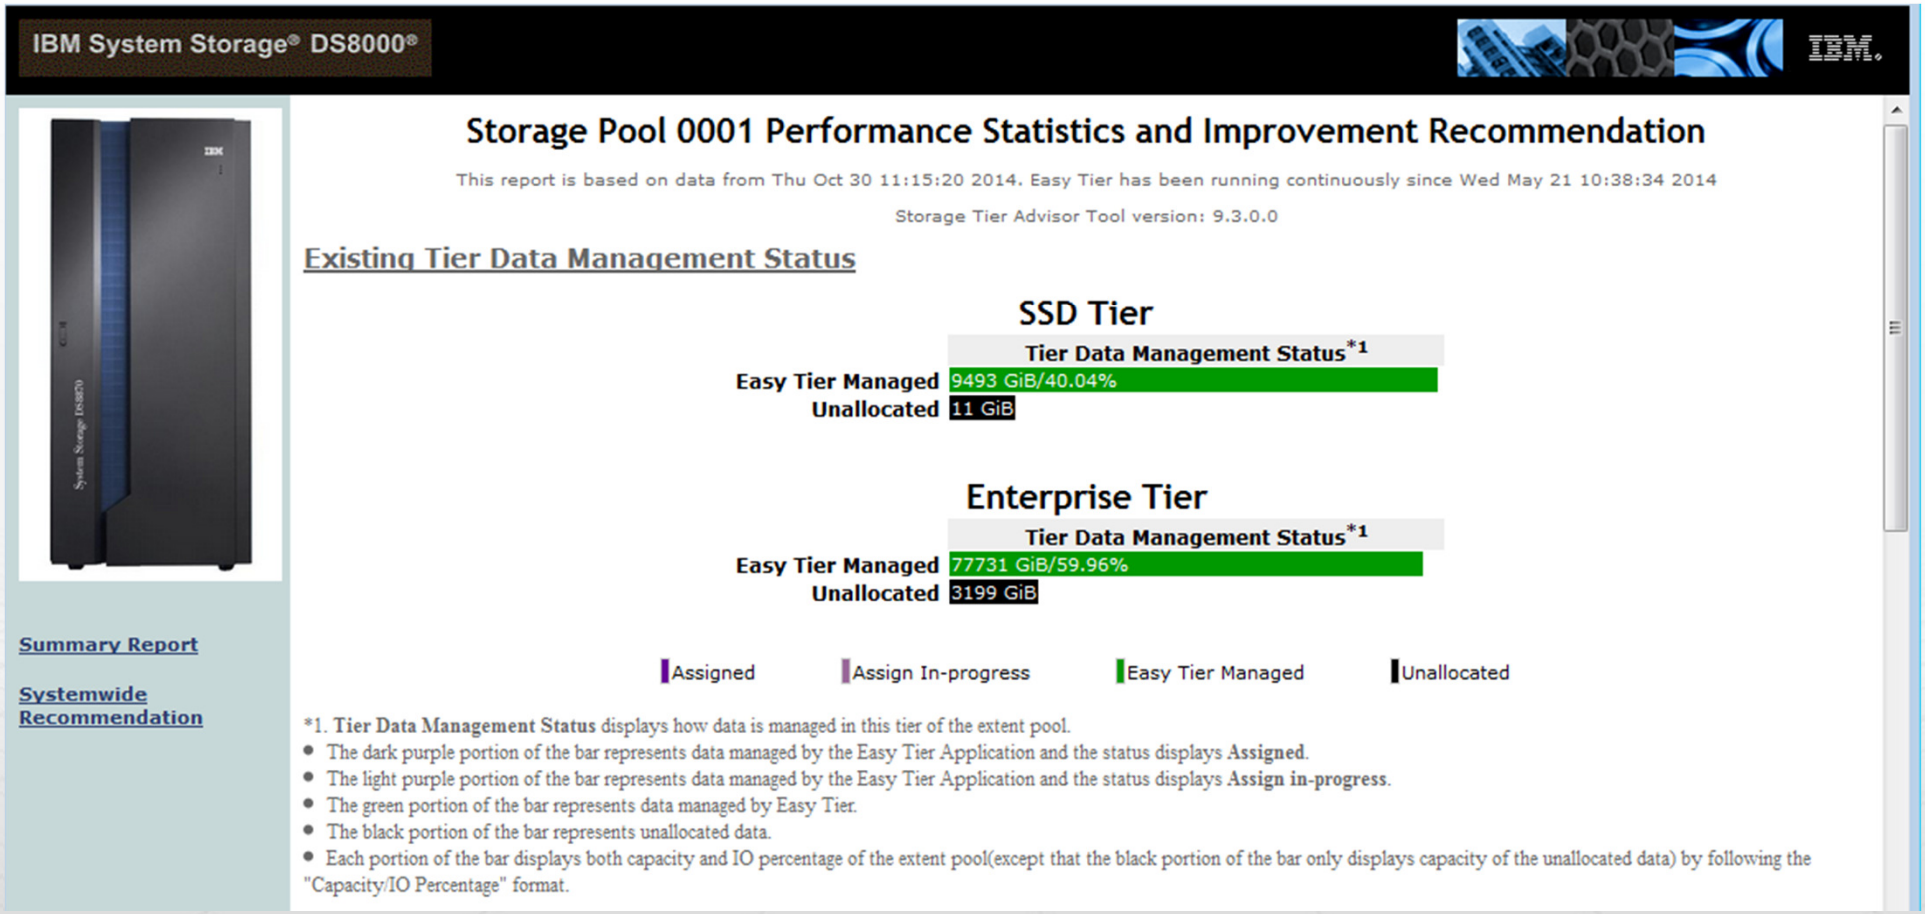

STAT reports review(3)

IBM System Storage® DS8000®

Summary Report

Systemwide Recommendation

Intra-tier Status

SSD Tier(Average Utilization of Rank IOPS is 8%)

Rank ID*2	Storage Pool ID	Rank type	Number of IOPS Threshold Exceeded*3	Utilization of Rank IOPS*4	Projected Utilization of Rank IOPS*5
5	0001	SSD	0	8%	8%
19	0001	SSD	0	8%	8%
33	0001	SSD	0	8% 1%	8% 1%
47	0001	SSD	0	8% 2%	8% 2%

Enterprise Tier(Average Utilization of Rank IOPS is 42%)

Rank ID*2	Storage Pool ID	Rank type	Number of IOPS Threshold Exceeded*3	Utilization of Rank IOPS*4	Projected Utilization of Rank IOPS*5
1	0001	Enterprise	0	42% 6%	42% 2%
3	0001	Enterprise	0	42%	42%
7	0001	Enterprise	0	42%	42%
9	0001	Enterprise	0	42% 3%	42% 2%
11	0001	Enterprise	0	42%	42%
13	0001	Enterprise	0	42%	42%
15	0001	Enterprise	0	42%	42%
21	0001	Enterprise	0	42%	42%
23	0001	Enterprise	0	42%	42%
25	0001	Enterprise	0	42% 4%	42% 2%
29	0001	Enterprise	0	42%	42%
31	0001	Enterprise	0	42% 2%	42% 2%
35	0001	Enterprise	0	42%	42%

STAT reports review(4)

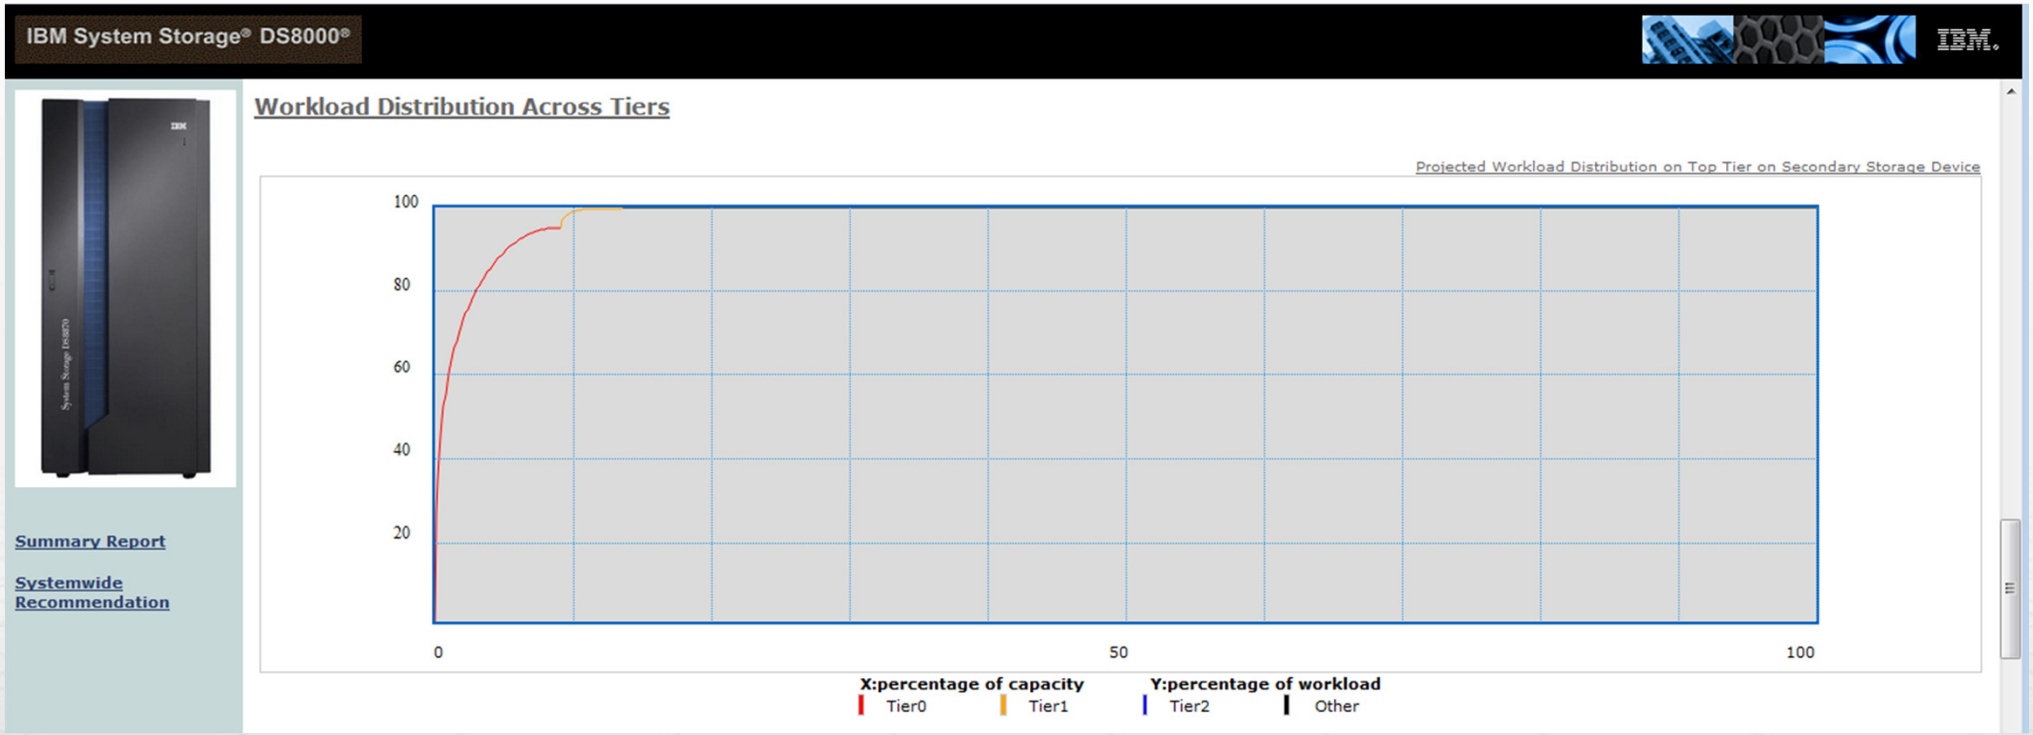

STAT reports review(5)

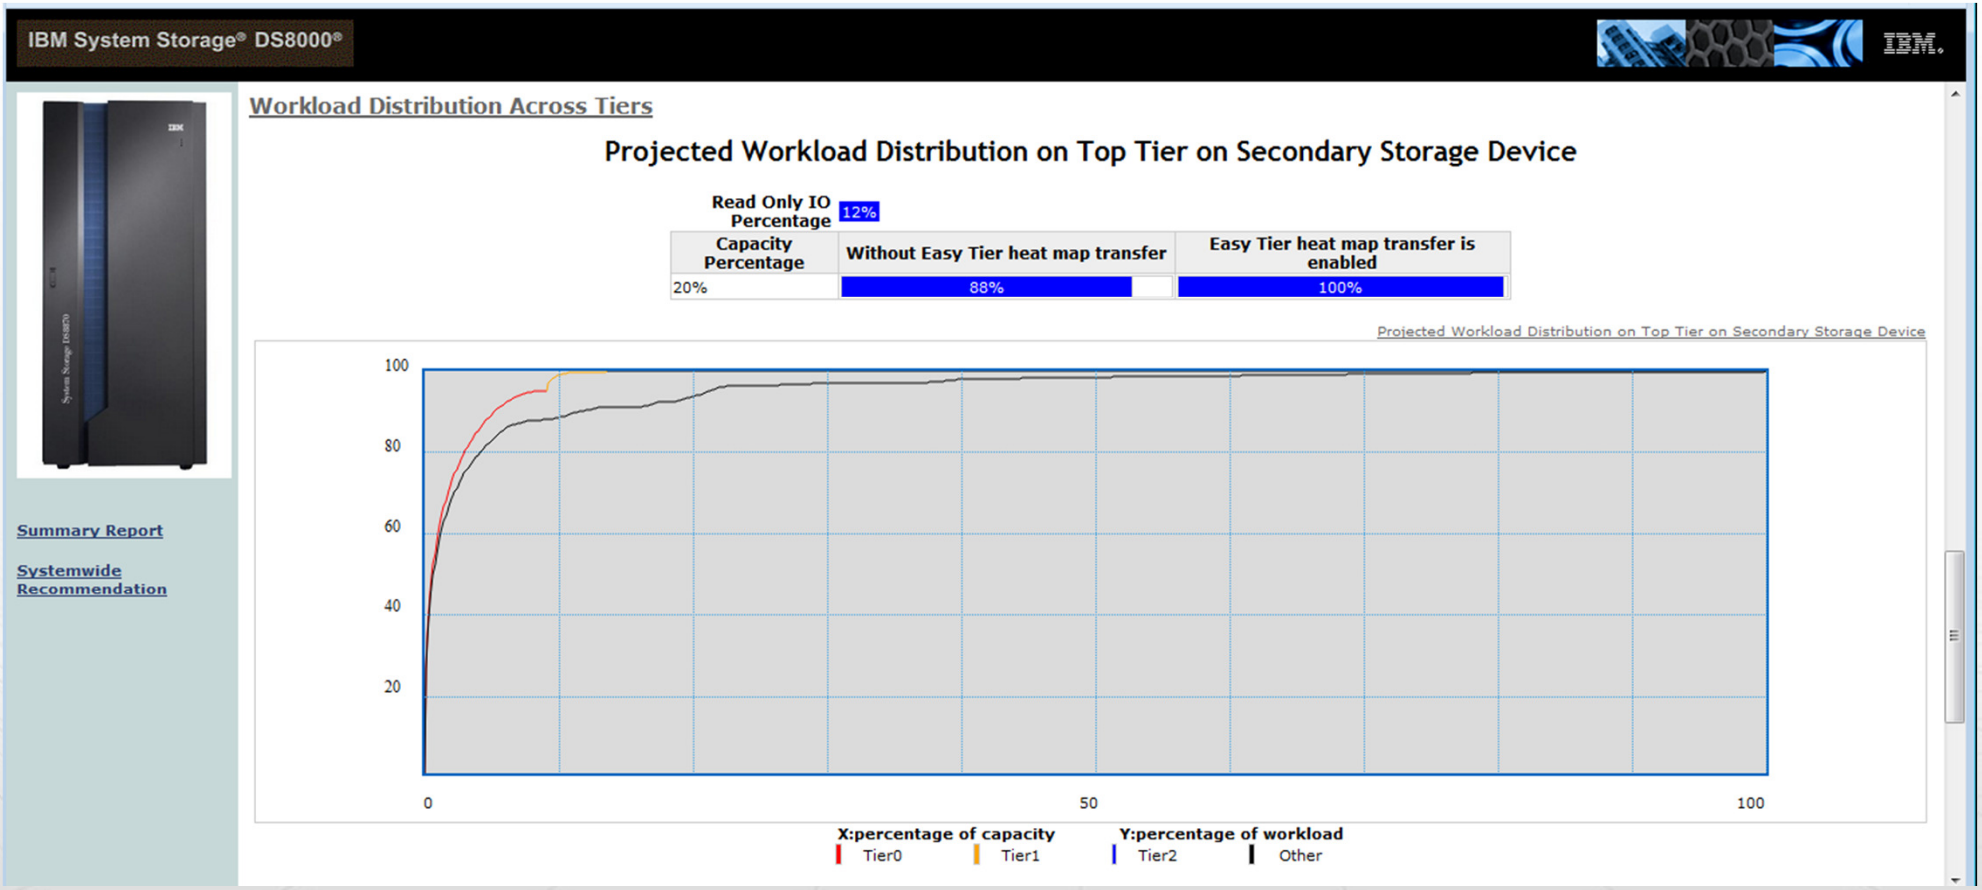

STAT reports review(6)

IBM System Storage® DS8000®

Summary Report

Systemwide Recommendation

Volume Heat Distribution

Volume Heat Distribution

Volume ID *8	Learning Strategy *9	Configured Size *10	IO Percentage of Extent Pool	Tier	Capacity on Tier *11	Heat Distribution *12
0x0100	N/A	1 GiB	0.00%	SSD Tier	0 GiB	
				Enterprise Tier	1 GiB	1 GiB
0x0101	N/A	9 GiB	0.01%	SSD Tier	0 GiB	
				Enterprise Tier	9 GiB	8 GiB, 1 GiB
0x0102	N/A	9 GiB	0.00%	SSD Tier	0 GiB	
				Enterprise Tier	9 GiB	8 GiB, 1 GiB
0x0103	N/A	9 GiB	0.01%	SSD Tier	0 GiB	
				Enterprise Tier	9 GiB	1 GiB, 8 GiB
0x0104	N/A	9 GiB	0.00%	SSD Tier	1 GiB	1 GiB
				Enterprise Tier	8 GiB	5 GiB, 3 GiB
0x0105	N/A	9 GiB	0.02%	SSD Tier	5 GiB	5 GiB
				Enterprise Tier	4 GiB	3 GiB, 1 GiB
0x0106	N/A	9 GiB	0.01%	SSD Tier	0 GiB	
				Enterprise Tier	9 GiB	8 GiB, 1 GiB
0x0107	N/A	9 GiB	0.00%	SSD Tier	3 GiB	3 GiB
				Enterprise Tier	6 GiB	6 GiB
0x0108	N/A	9 GiB	0.00%	SSD Tier	0 GiB	
				Enterprise Tier	9 GiB	9 GiB
0x0109	N/A	9 GiB	0.00%	SSD Tier	0 GiB	
				Enterprise Tier	9 GiB	9 GiB

10 Volumes Per Page GO |< << >> >| Displaying Page 1 of 508

*8. **Volume ID** represents the DS8000 volume ID, which is generated when the volume is created.
*9. **Learning Strategy** displays the specific learning data used by this volume.
- **N/A** indicates that the volume applies learning data based on its own learning results.
- **Based on remote** indicates that the volume applies learning data transferred from a remote storage device.

*10. **Configured Size** displays the configured capacity of the volume.

csv file output from STAT

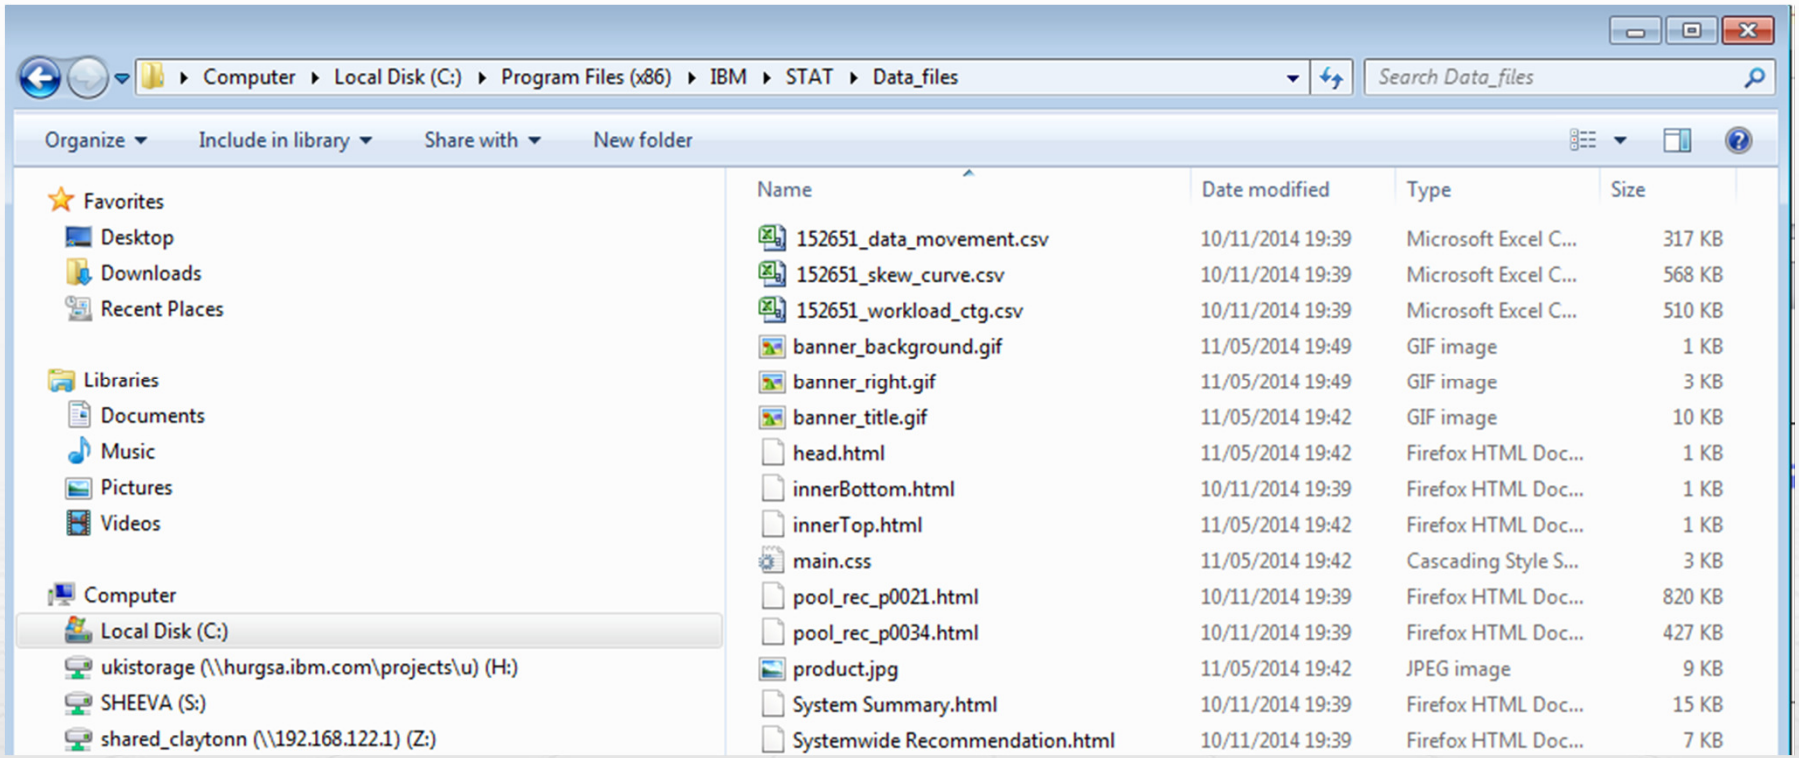

Data movement

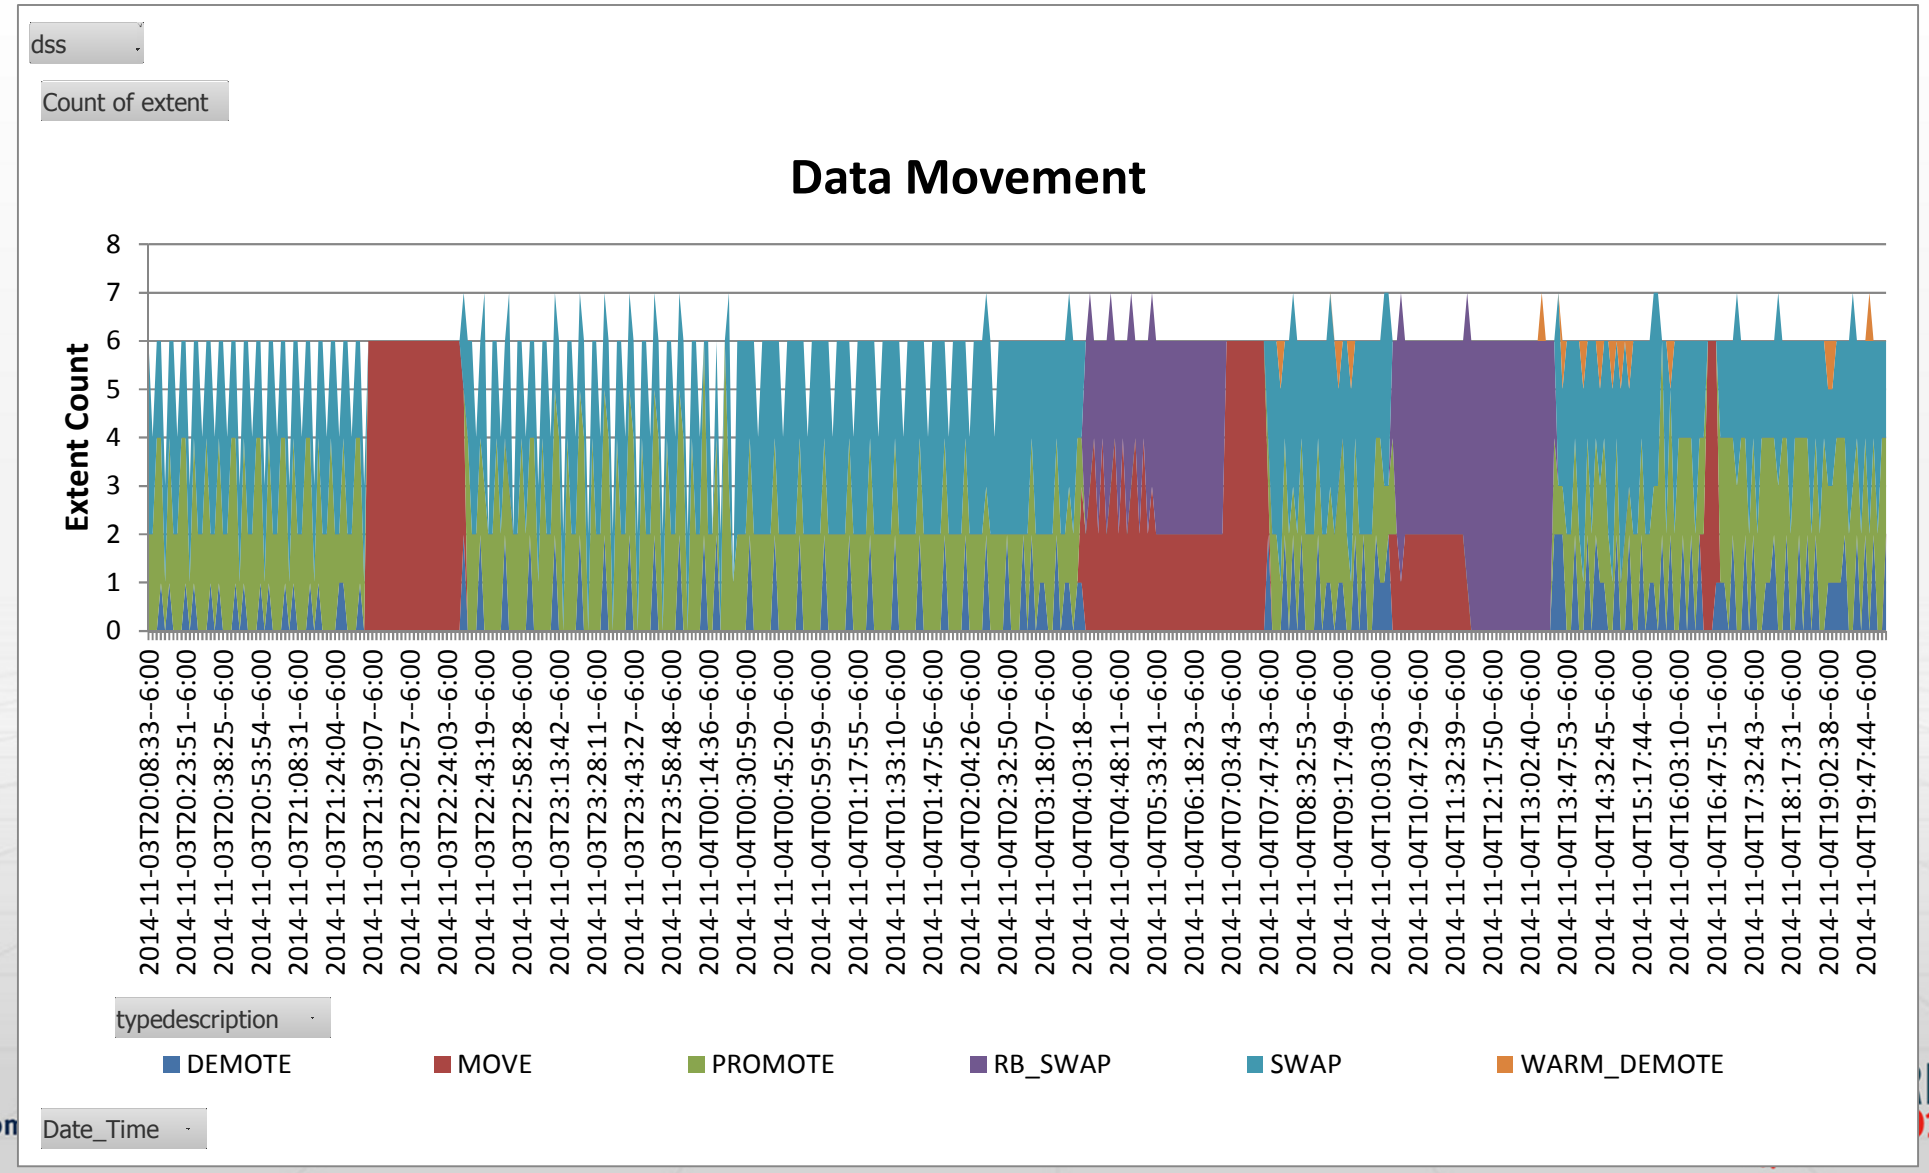

Workload categorisation

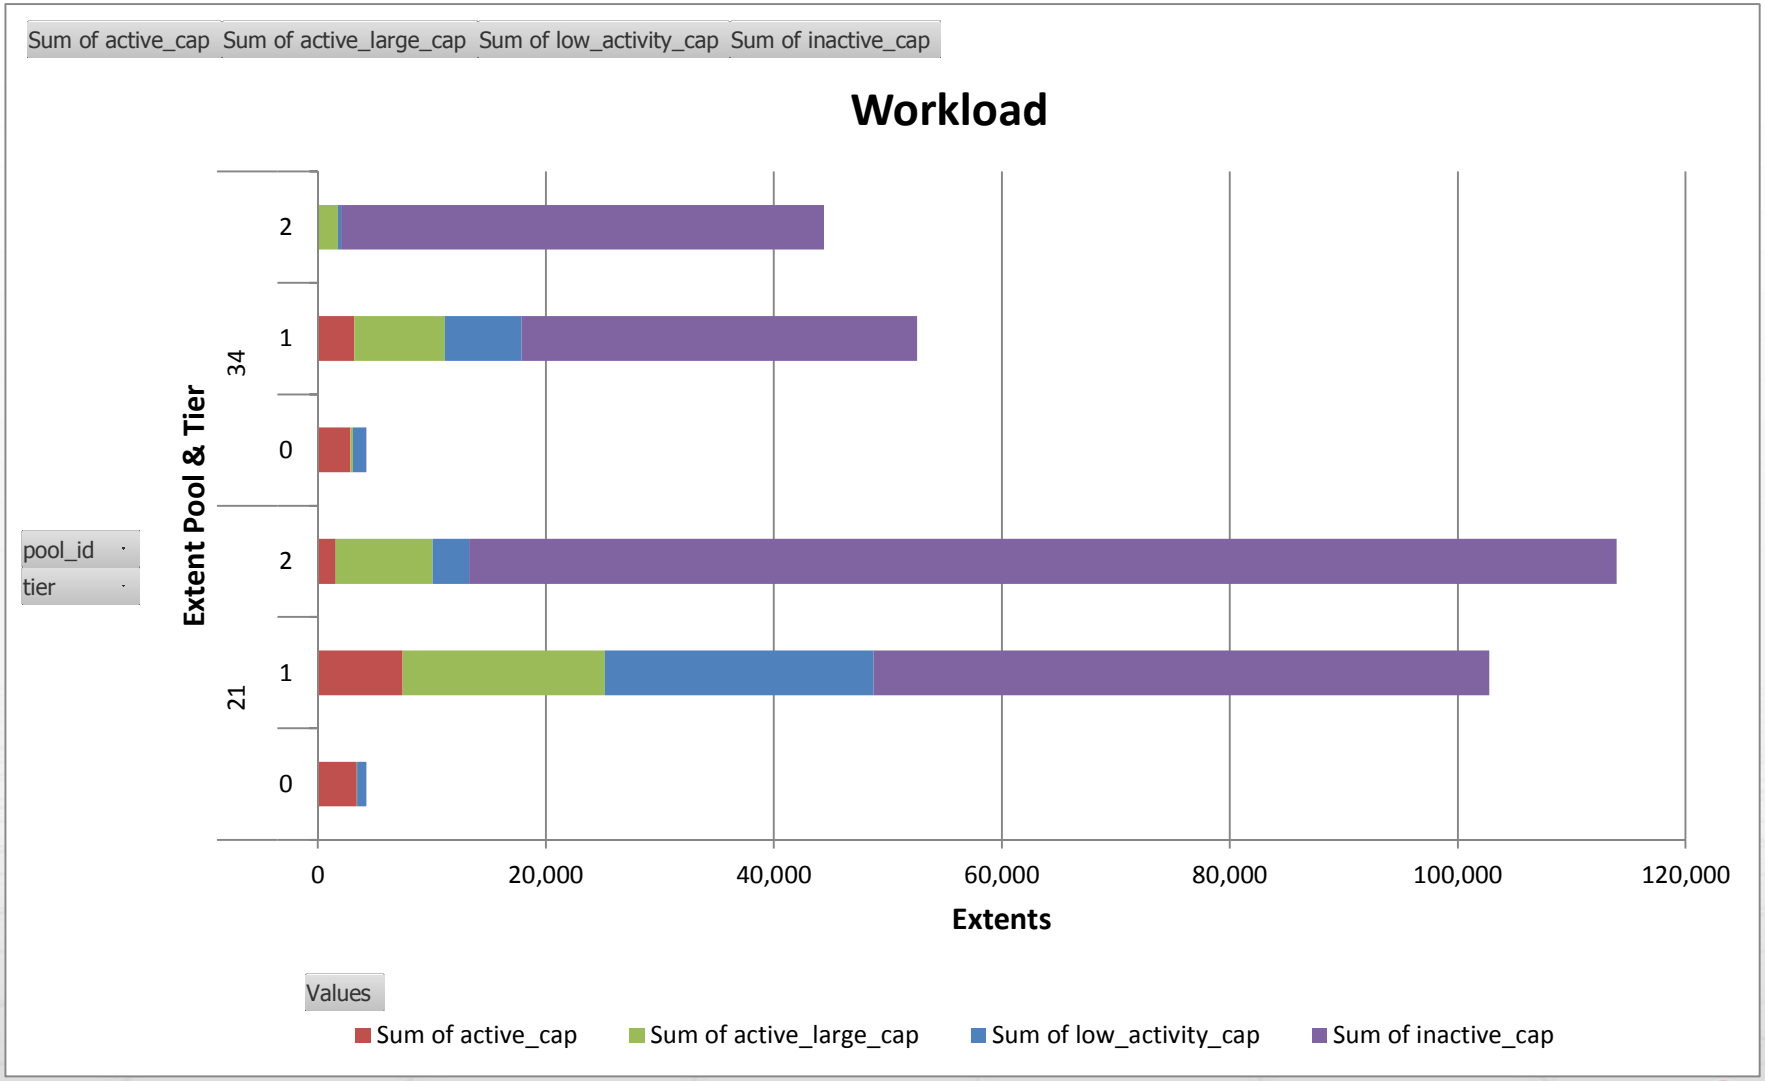

Summary

Actions | Filter

- Tiering Functionality for z/OS
- Easy Tier review
- New Easy Tier functionality
  - Easy Tier Policy and Controls
  - Easy Tier z/OS integration
- Tiered storage configurations for z/OS
- Easy Tier reporting
- Summary