Hot Band Workloads and High End HP Storage Products

Chuck Paridon, HPES Master Storage Performance Architect

SNIA Emerald™ Training

SNIA Emerald Power Efficiency Measurement Specification, for use in EPA ENERGY STAR®

June 24-27, 2013

Hot Band IO Patterns

- The Hot Band workload is comprised of several different IO streams, some of which contain hot spots, or regions of more intense IO demand.
- This results in varying degrees of logical block address re-referencing in certain regions of the overall IO space that can be either contained with the cache of an array, or can be placed on storage devices that deliver a higher IO rate.

Hot Band IO Patterns

Hot Band IO Profile

IO Profile	% of workload	Read/Write Percentage	IO Size (KiB)	Access Pattern	Usable Address Range
Write Stream 1	5	0/100	See Table 12	Sequential	0-100%
Write Stream 2	5	0/100	See Table 12	Sequential	0-100%
Write Stream 3	5	0/100	See Table 12	Sequential	0-100%
Read Stream 1	5	100/0	See Table 12	Sequential	0-100%
Read Stream 2	5	100/0	See Table 12	Sequential	0-100%
Read Stream 3	5	100/0	See Table 12	Sequential	0-100%
Read Stream 4	5	100/0	See Table 12	Sequential	0-100%
Read Stream 5	5	100/0	See Table 12	Sequential	0-100%
Uniform Random	6	50/50	See Table 12	Random	0-100%
Hot Band 1	28	70/30	See Table 12	Random	10 -18%
Hot Band 2	14	70/30	See Table 12	Random	32-40 %
Hot Band 3	7	70/30	See Table 12	Random	55-63 %
Hot Band 4	5	70/30	See Table 12	Random	80-88 %

Hot Bands concentrate 54% of the IO in 32% of the space

Hot Band IO Patterns

- The Hot Band workload when run on High End Storage demonstrates the power/performance advantage of two product features
  - Array Based Cache
  - Storage Tiering
- Although the initial goal was solely cache focused, there is also a benefit of implementing faster tiers in the product such as HP P9500 Smart Tiering

SNIA Green TWG Cache Friendly Performance Comparison Chart Cache and Tiering Speedup on HP P9500 Storage Product

SNIA Green Storage Initiative

	Hot IOPS	Hot RT	Rnd IOPS	Rnd RT	C/WS ratio	Cache Hit Ratio
Small Array	4330	32.8 ms	4 30	33.4 ms	<<	N/A
Large Array Cache Assist only	39900	8.97 ms	18410	22.5 ms	~3%	60%/24%
Large Array (Tiered)	42870	5.77 ms	N/A	N/A	~3%	60%/24%

- The Workload Analysis Process Consists of Two Steps

1. Create a single pool large enough to hold the desired working set
2. Run the Hot Band workload on that pool and use the analysis tools to produce a report guiding the composition of the tier(s)

- The tier construction process involves 2 decision points.

1. Which technology to deploy
2. The capacities of the tier(s)

- The following slides illustrate an example of these activities

SNIA Green TWG Tiered Storage w/Hot Bands
Initial Tier Property Analysis

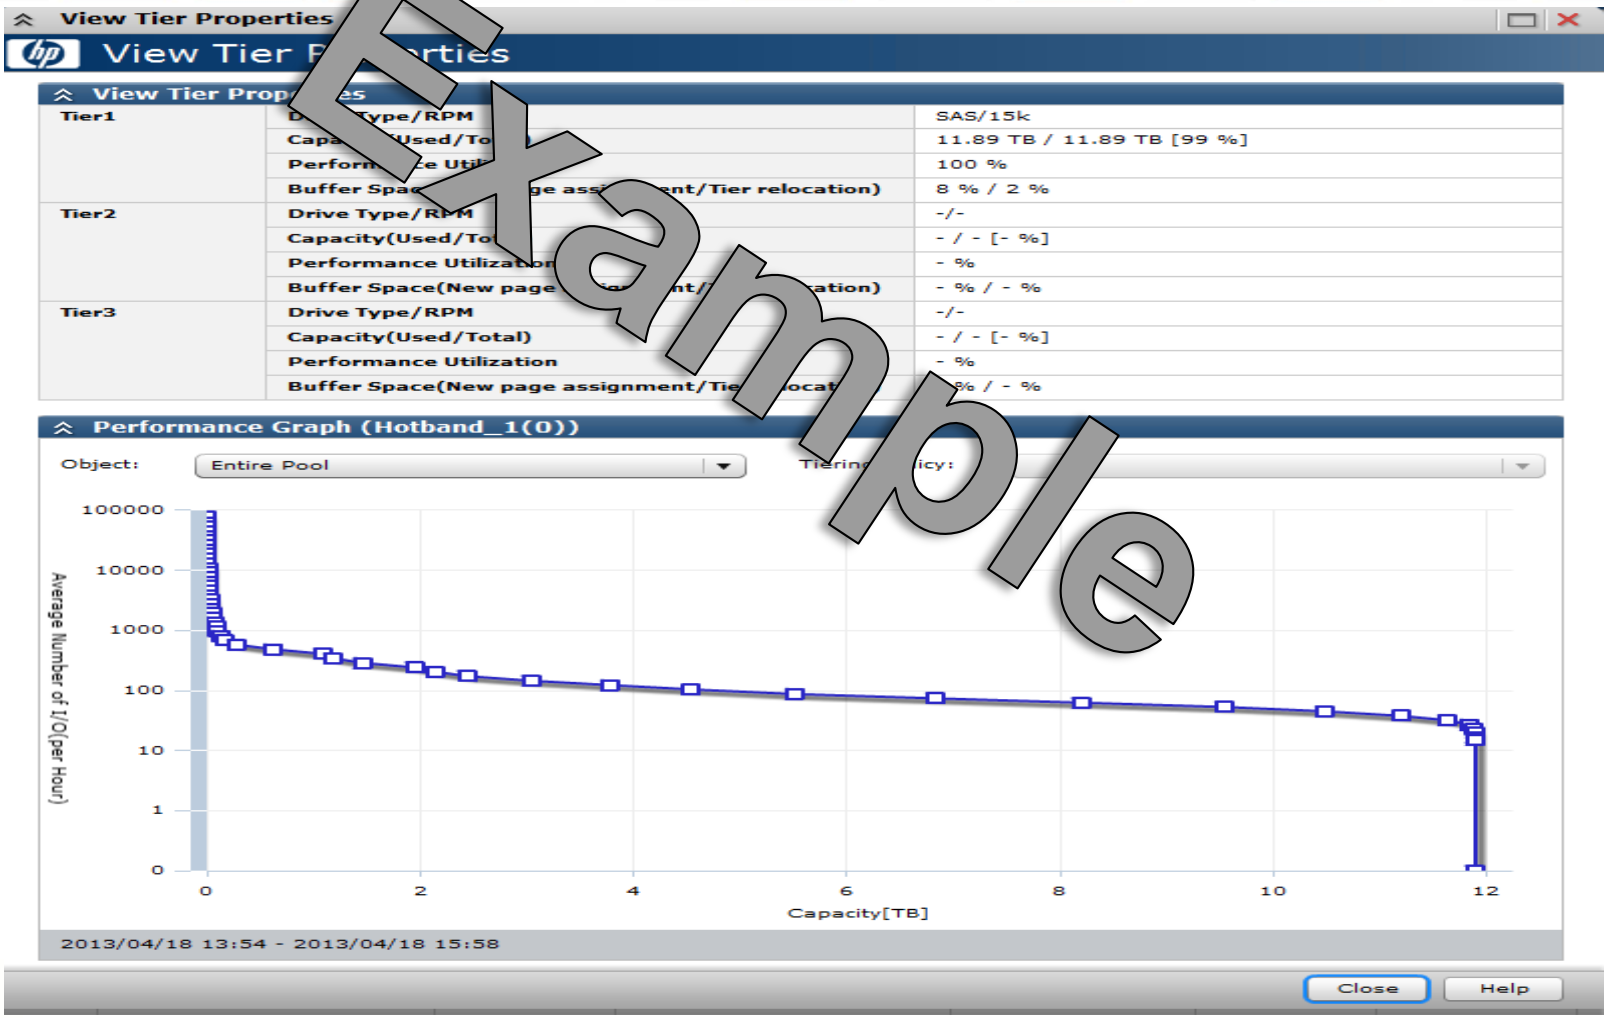

SNIA Green TWG Tiered Storage w/Hot Bands Final Tier Property Analysis

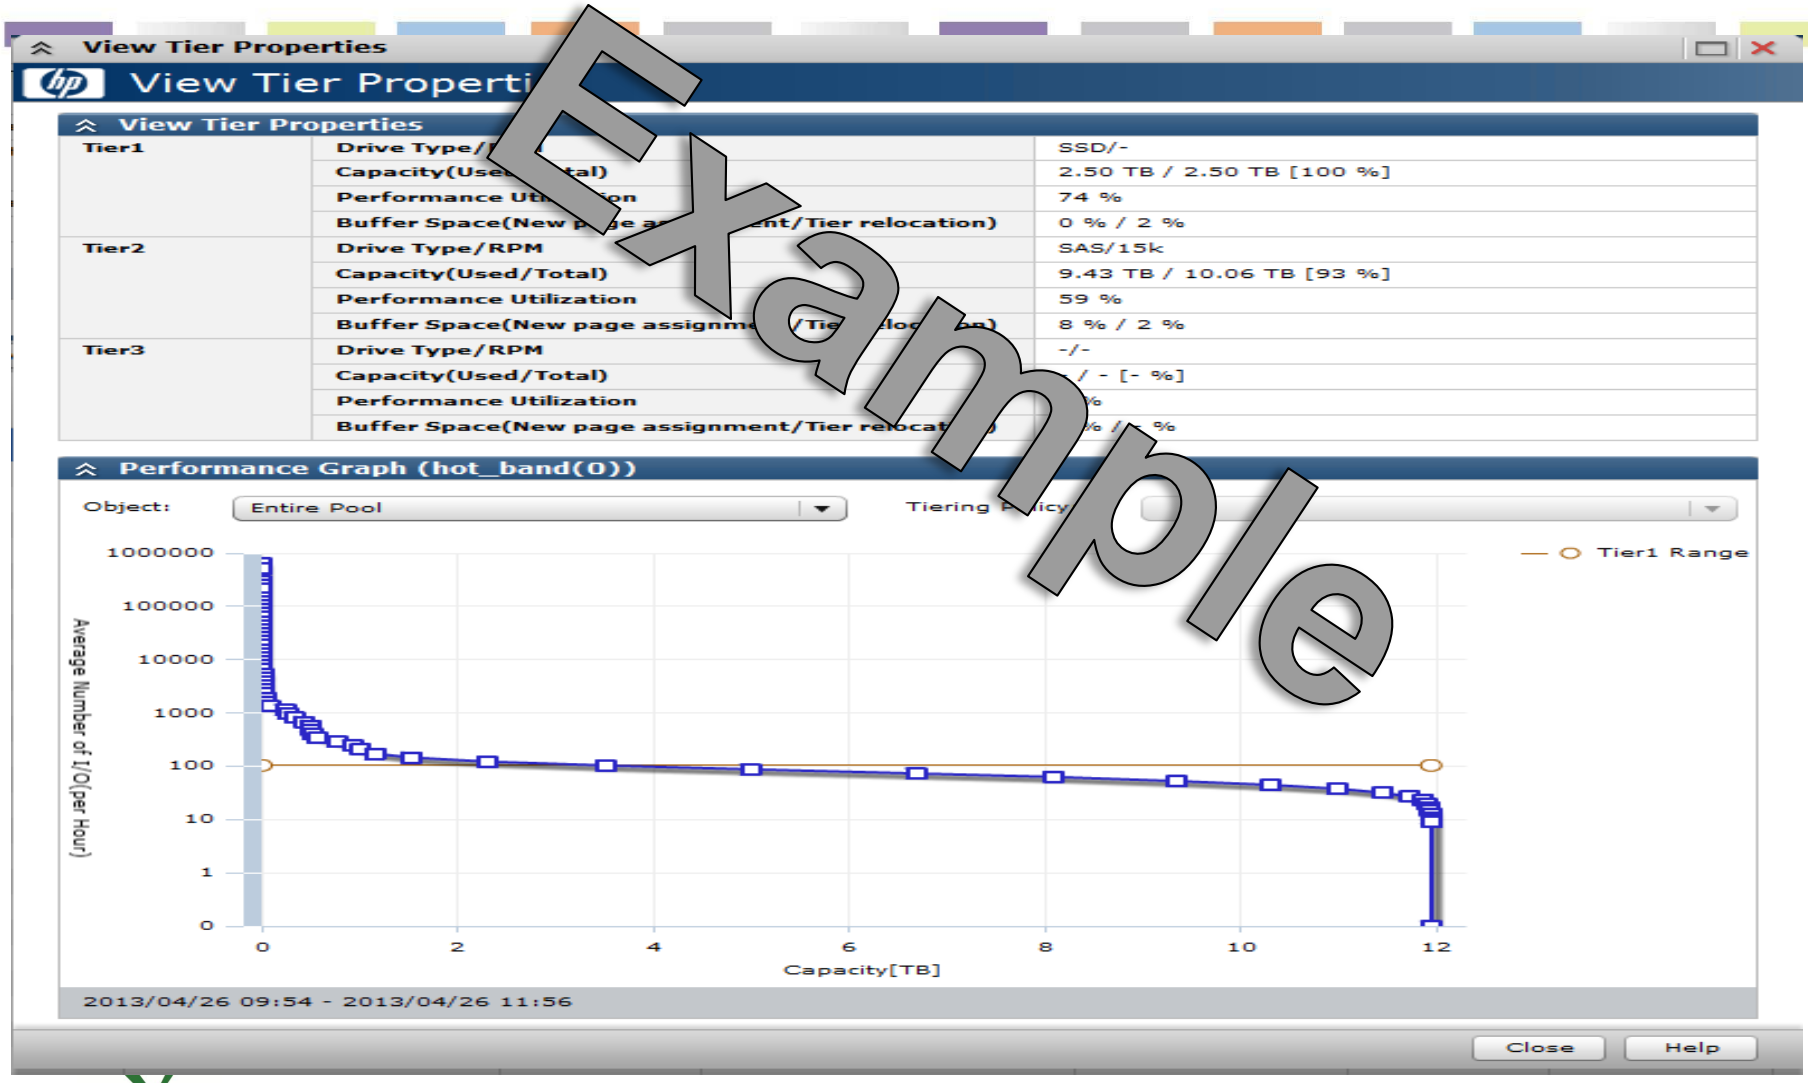

The HP P9500 Power Calculator

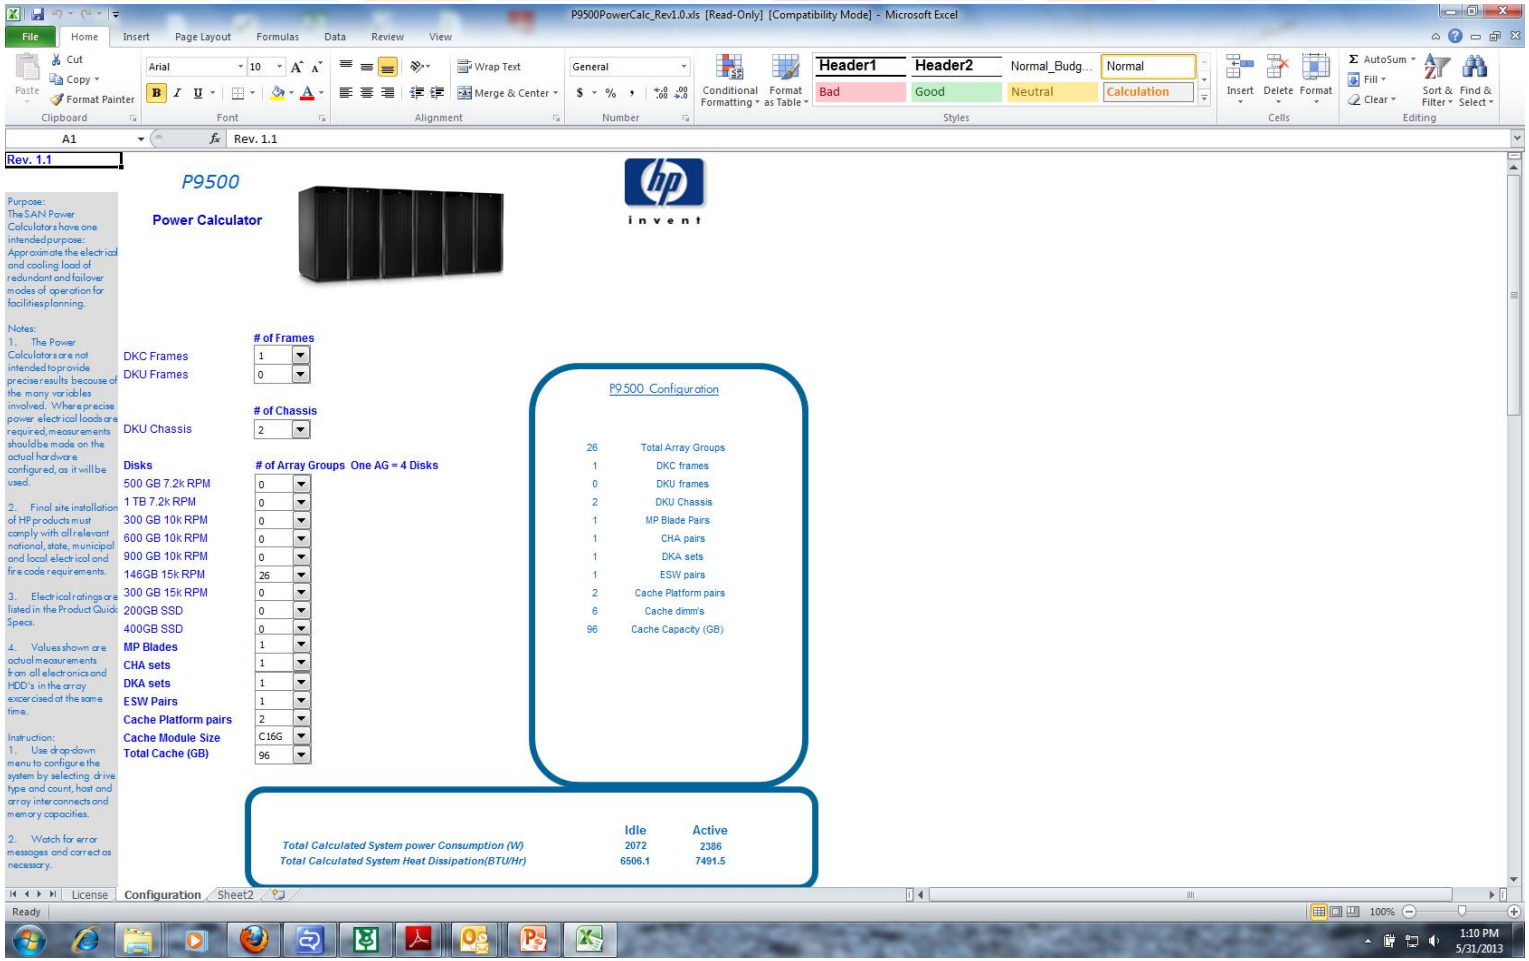

SNIA Green TWG Tiered Storage w/Hot Bands Primary Metric Comparison

Configuration	Tier Type	Power Consumption (Measured)	IOPS	IOPS/Watt
Large Array (Initial)	15k RPM	7491 Watts	18420	2.457
Large Array (Cache Assist)	15k RPM	7491 Watts	39900	5.326
Large Array (Tiered)	15k + SSD	7283 Watts	42870	5.886

SNIA Green TWG Hot Band Workload Conclusions and Observations

- The Hot Band workload is amenable to performance optimization by both storage subsystem cache and the proper deployment of tiered storage.
- As a result of high cache hit rates (~60%) the overall performance contribution of tiered storage is limited.
- In addition to the increase in IO rate (132%), there is also a corresponding decrease in power consumption from the substitution of SSDs in the configuration
- The net effect of these two parameter changes is a 140% improvement in the SNIA primary active metric (IOPS/Watt) of