ENERGY STAR
Data Center Storage Meeting
Initial Insights from Storage Results

November 18, 2015

Key Goals

- Identify submitted dataset issues
- Gain insights from public data
  - Identify any additional insights from data only visible by EPA

Initial High Level Observations

- Issues with submitted data
  - Systems not optimized for workload
  - Missing data / details
- Configurations show meaningful difference
  - Drive type aligned with workload
  - Good results with newer technology (SFF, SSD)
  - No common pattern to drive scaling

Dataset Observations

- Small number of systems submitted
  - 66 total different entries (systems / configs)
- Potential issues with submitted systems / data:
  - Vendor portfolio coverage for US Government sales?
  - Claimed qualification workloads vs.
    - Submitted test data
    - Configured optimization points

Issue: Vendor Portfolio Coverage?

- Some venders submitted wide system set

- Some venders submitted a very small portion of portfolio
  - Does this cover full range of sales opportunity to US Government?
  - Are there other exclusions being used?

Issue: Workload Qualifications

- System claimed qualification under multiple workloads, e.g.: Transition + Capacity
- Same configuration / results used for both workloads

> 8) Capacity Optimization Family Restriction: Since the scope in 2.1.1 only encompasses Online systems, a product family may not be based solely on Capacity Optimized Configurations. Every storage device submitted for qualification under Capacity Optimized Configurations must also include one or more Transaction Optimized Configuration(s) and/or Streaming Optimized Configuration(s) using the same storage device. A Capacity Optimized Configuration may only be submitted as an addition to one (or more) of the other optimizations.

- Each workload claimed should have optimized configuration / dataset for that workload
  - Unlikely same configuration would be optimal for multiple workloads

Issue: Selection of BFF Config

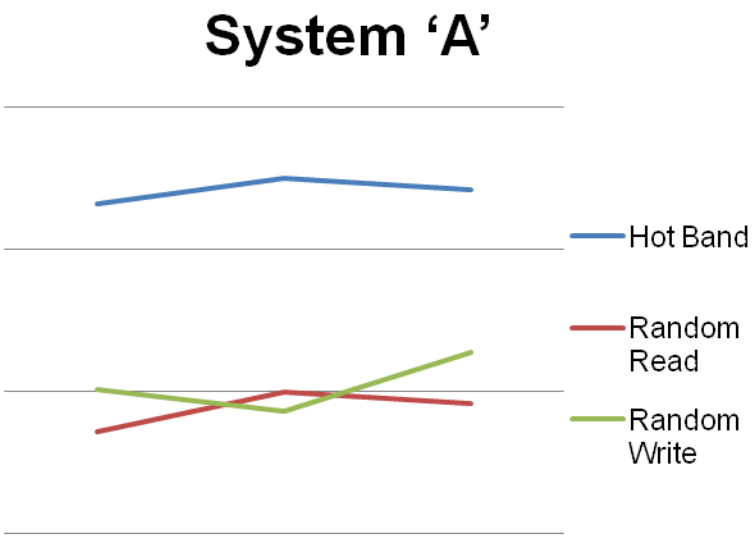

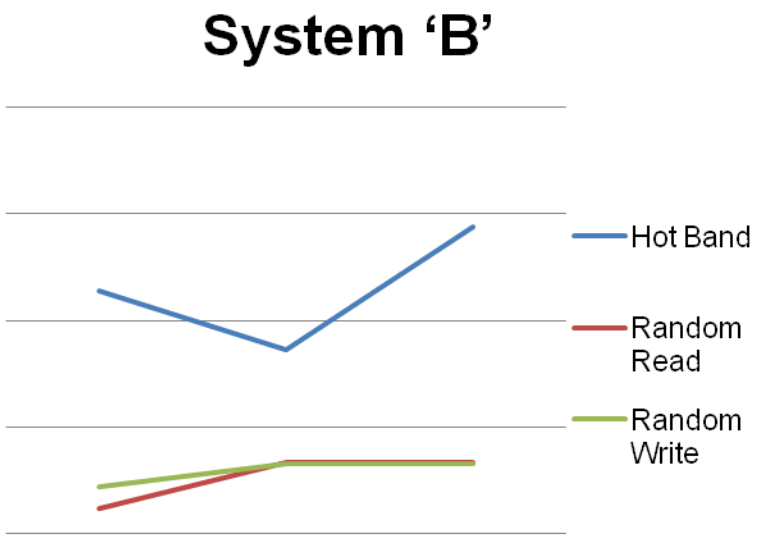

- Table 6 in specification defined workload to be used for BFF optimization
  - i.e. HotBand for Transaction optimized systems
- Some systems clearly have not done this
  - Demonstrated when BFF point is not best work/watt point
  - Side-effect of only submitting one configuration for multiple workloads?

Insights from data?

- Small sample set:
  - Difficult to confirm trends
  - Focusing on Group 3 / transaction systems
    - 35x submissions

- Can we ID a proxy workload?
- System workload efficiency breadth vs. configuration / workload optimization
- Storage device count / scaling trends

Can We ID a Proxy Workload?

- Correlation of RR and RW?
- Many systems show high level of correlation
  - But some do not. . .

Ratio: BFF/RR vs. BFF/RW	
Average	1.30
Median	1.16
Min	0.39
Max	3.16

- Hot band show wider variations.
- Data indicates need for continued evaluation of all three random workloads.

System Workload Breadth

- Half of the configurations exceeded at least one “135% of median” of a transaction workload
  - 8 of 35 configurations exceeded all three

- Good performing systems perform well . . .
  …but not necessarily across all workloads:
  - Best reported Hot-Band system (2x median)
  - Also in top 25% of Random-Read results
  - However bottom third of Random-Write results

Impact of Configuration

- Detailed look at Dell PS6210:
  - Common controller, cache, and PSU
  - Many configurations and optimizations
  - Reported best Hot-Band BFF result of 49.7 IOPS/W

Dell PS6210 Line

	PS6210E	PS6210X	PS6210XV	PS6210XV 3.5"	PS6210S	PS6210XS
Hard Disk Drives	24x 3.5" 7.2K RPM NL-SAS drives	24x 2.5" 10K RPM SAS drives	24x 2.5" 15K RPM SAS drives	24x 3.5" 15K RPM SAS drives	24x 2.5" SSDs	7x 2.5" SSDs + 17x 2.5" 10K RPM SAS drives
Drive Capacities	2TB, 3TB, 4TB	600GB, 900GB, and 1.2TB	300GB	600GB	400GB and 800GB	Combines 400GB or 800GB SSDs and 600GB or 1.2TB 10K RPM SAS drives
System Capacities	Up to 96TB (up to 1.536PB with 16 arrays)	Up to 28.8TB (up to 460.8TB with 16 arrays)	Up to 7.2TB (up to 115.2TB with 16 arrays)	Up to 14.4TB (up to 230.4TB with 16 arrays)	Up to 19.2TB (up to 307.2TB with 16 arrays)	Up to 26TB (up to 416TB with 16 arrays)
Physical	Height: 4U Weight: 47.1 kg (103.8 lb)	Height: 2U Weight: 26.4 kg (58.3 lb)	Height: 2U Weight: 26.4 kg (58.3 lb)	Height: 4U Weight: 47.1 kg (103.8 lb)	Height: 2U Weight: 26.4 kg (58.3 lb)	Height: 2U Weight: 26.4 kg (58.3 lb)
Est. List Price	€ 32.255	€ 31.614	€ 30.544	€ 30.844	€ 102.606 € 150.044	€ 67.685
Hot Band	8.7	35.7	33.1	19.3		49.7
Random Read	4.3	16.7	22	10.9		21.1
Random Write	3.3	18.5	16.2	12.2		7.2
Idle (GB/W)	28.1	12.7	2.7	6.8		17.7

All systems submitted under 'Transaction' workload

Dell PS6210 Line – 135% of Median

	PS6210E	PS6210X	PS6210XV	PS6210XV 3.5"	PS6210S	PS6210XS
Hard Disk Drives	24x 3.5" 7.2K RPM NL-SAS drives	24x 2.5" 10K RPM SAS drives	24x 2.5" 15K RPM SAS drives	24x 3.5" 15K RPM SAS drives	24x 2.5" SSDs	7x 2.5" SSDs + 17x 2.5" 10K RPM SAS drives
Drive Capacities	2TB, 3TB, 4TB	600GB, 900GB, and 1.2TB	300GB	600GB	400GB and 800GB	Combines 400GB or 800GB SSDs and 600GB or 1.2TB 10K RPM SAS drives
System Capacities	Up to 96TB (up to 1.536PB with 16 arrays)	Up to 28.8TB (up to 460.8TB with 16 arrays)	Up to 7.2TB (up to 115.2TB with 16 arrays)	Up to 14.4TB (up to 230.4TB with 16 arrays)	Up to 19.2TB (up to 307.2TB with 16 arrays)	Up to 26TB (up to 416TB with 16 arrays)
Physical	Height: 4U Weight: 47.1 kg (103.8 lb)	Height: 2U Weight: 26.4 kg (58.3 lb)	Height: 2U Weight: 26.4 kg (58.3 lb)	Height: 4U Weight: 47.1 kg (103.8 lb)	Height: 2U Weight: 26.4 kg (58.3 lb)	Height: 2U Weight: 26.4 kg (58.3 lb)
Est. List Price	€ 32.255	€ 31.614	€ 30.544	€ 30.844	€ 102.606 € 150.044	€ 67.685
Hot Band	8.7	35.7	33.1	19.3		49.7
Random Read	4.3	16.7	22	10.9		21.1
Random Write	3.3	18.5	16.2	12.2		7.2
Idle (GB/W)	28.1	12.7	2.7	6.8		17.7

All systems submitted under 'Transaction' workload

Dell PS6201 – What Can We Learn?

- Very high cost for all SSD configuration
- Drive selection impacts results:
  - PS6210E → great GB/W
    … less so active transaction tests
  - PS6210XS→ solid Hot-Band results
    …not so hot Random Writing
  - 2.5” device good overall results
  - Lower power devices (10K vs. 15K) produces good work/watt results
    - Hybrid systems delivers even better results.

Energy Efficiency is Only One Criteria

- PS6210XV w/2.5” devices delivered better active results than w/3.5” devices
  - 3.5” media provided 2x the storage capacity.
  - And much better transaction results than bulk storage PS6210E (1TB) configuration.
- PS6210XS gave great random access and bulk storage results
  - Rather poor pure random-write results
  - Almost 2x cost of non-SSD based systems.

- Final selection of systems consider wide range of qualities and goals.
- Configuration options exist to support different end user needs.

Comparing Controllers

- Like configurations:
  - 24x / 2.5” / 10K / 300GB HDD
  - 2x 700/725W PSU
  - 2x controllers
- Different controllers
  - Positioning within vender line
  - Across venders
- Over 30% delta

Brand Name	Model Name	Trans Optimal Point Hot Band Workload Test (IOPS/W)
DELL	PS4110X	27.3
DELL	PS6100X	30.4
DELL	PS6210X	35.7
NetApp, Inc	E2724	34.4
NetApp, Inc	E5524	34.4

Storage Device Count Scaling

- No alignment between example 4x transaction systems.
- Continued trend of at times rapid decline after optimal point.

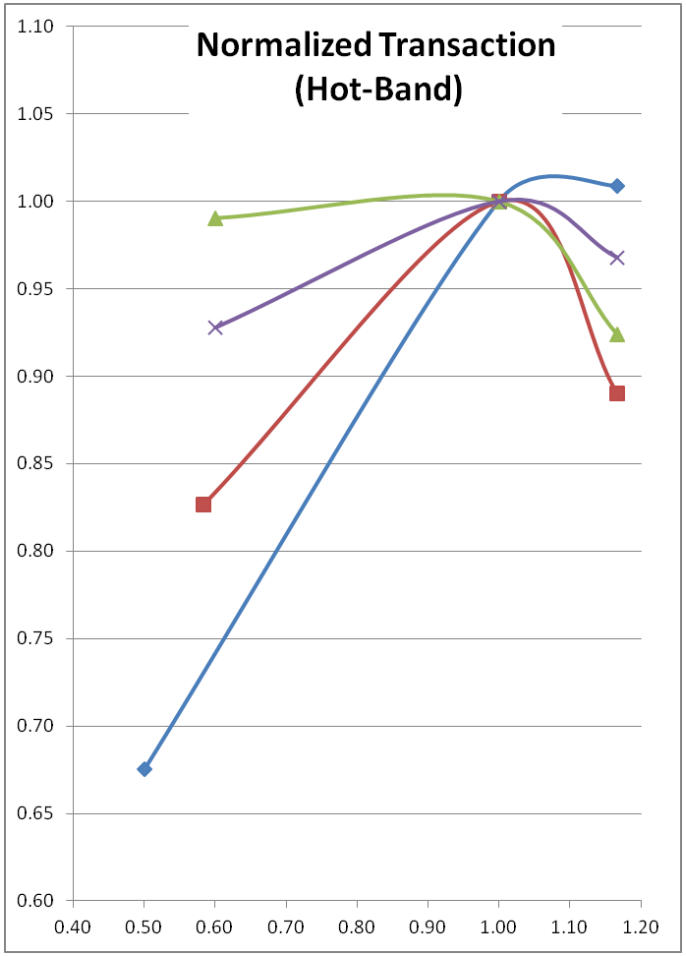

Device count scaling (cont)

- Sequential
  - Again, no identifiable pattern.
- Data issue:
  - Graph clearly shows several systems not optimized for sequential workload.

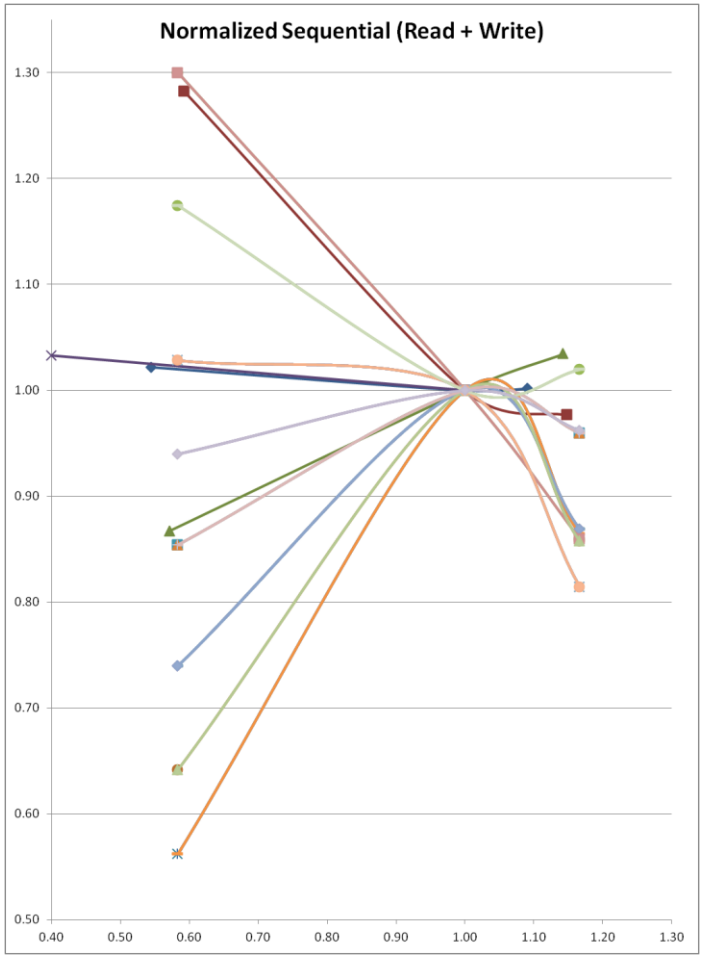

Device count scaling (cont)

- Removed all obvious anomalies
  - Some remaining systems may still be sub-optimized

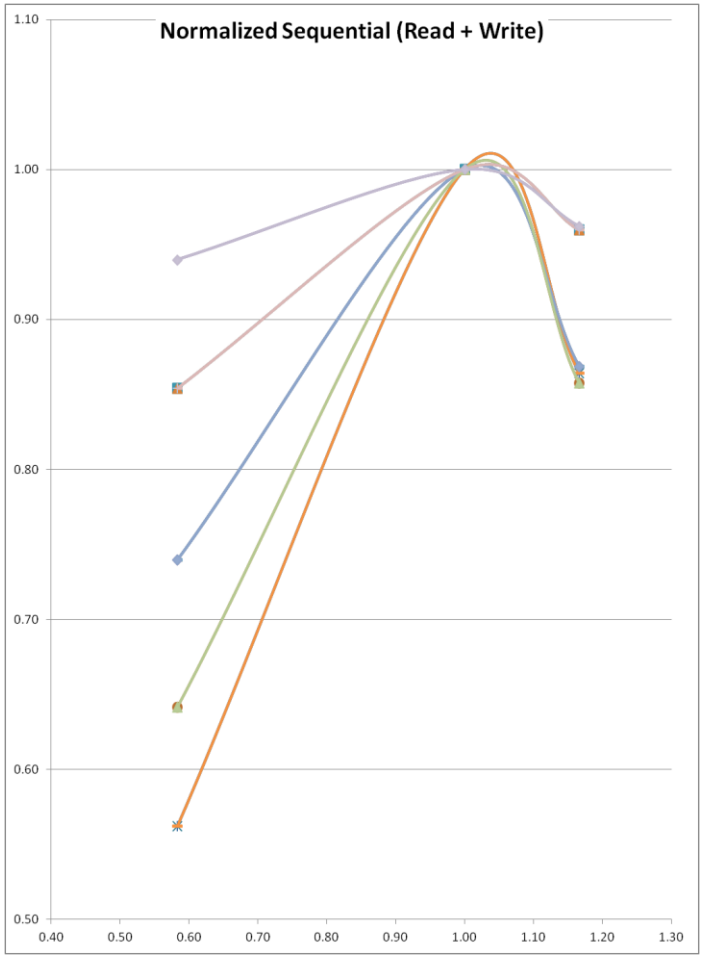

Device Count / Scaling Trends

- Most systems continue to show trend of at times rapid efficiency declines past optimal point
- Configurations smaller than optimal point show variety of slopes
- No strong indication of industry wide scaling trends

Overall Conclusion on Patterns

- To this point, there is no observable general pattern to storage scaling.
- Storage device type has impact on energy work/watt at different workload points
- 2.5” devices show good efficiency
- Hybrid systems can show better efficiencies
  - At high cost
  - Perhaps not for all workloads (random write)

Online 3 vs. 4

- Comparison of Hot-band workloads shows overlap
  - But at bottom portion of Hot-band Scale

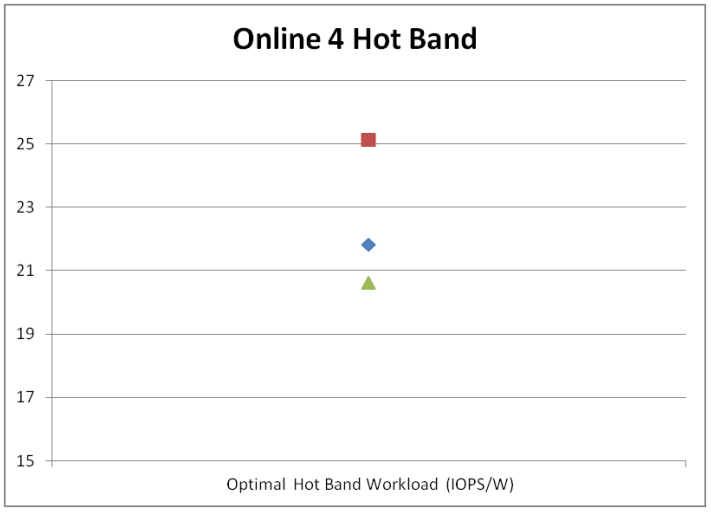

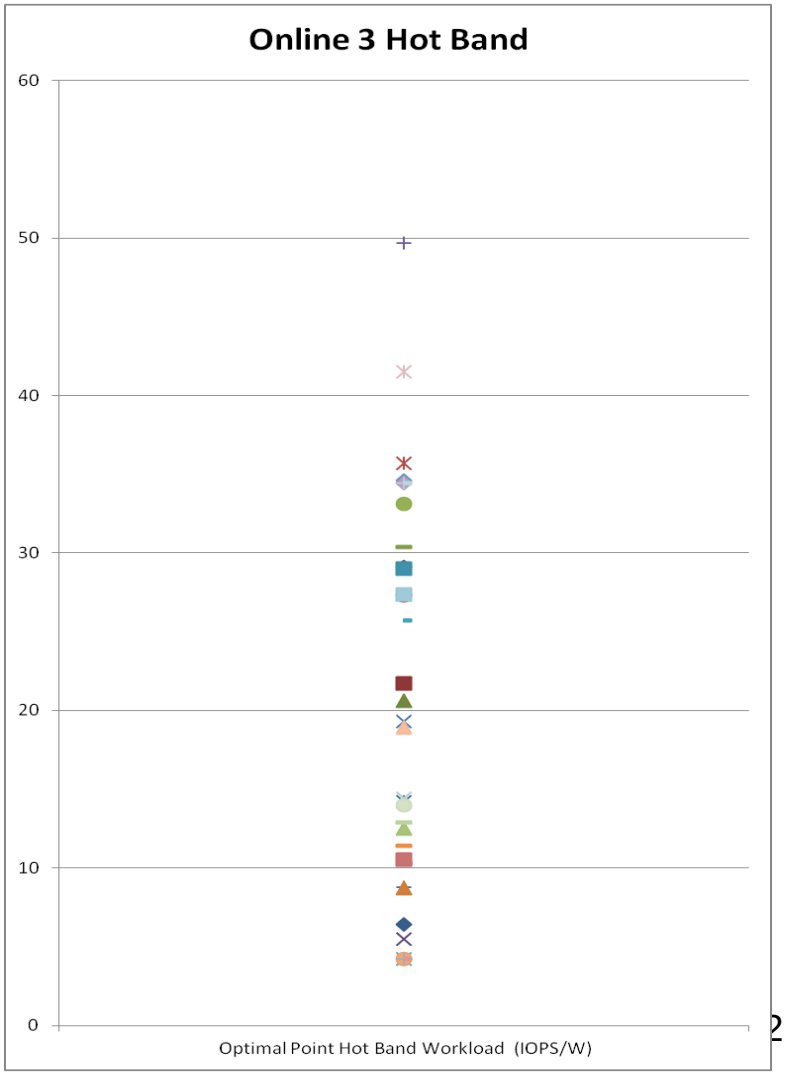

Questions

- Questions around submitted systems
  - Optimization points vs. claimed workloads

- Data set provided minimal insight into actual performance vs. efficiency
  - Are most efficient systems also poor performers?
  - Will such behavior drive up the purchase of a higher quantity ‘efficient’ systems to gain needed performance?