

Apples to Apples, Pears to Pears in SSS Performance Benchmarking



Esther Spanjer

Director, SSD Marketing



The Performance Landscape One Year Later

ME

- Read and Write IOPS Specifications (Iometer* Queue Depth 32)
 - Random 4 KB Reads: Up to 35 K IOPS
 - Random 4 KB Writes:
 - 80 GB - Up to 6.6 K IOPS
 - 160 GB - Up to 8.6 K IOPS
- Bandwidth Performance Specifications
 - Sustained Sequential Read: Up to 250 MB/s
 - Sustained Sequential Write:
 - 80 GB - Up to 70 MB/s
 - 160 GB - Up to 100 MB/s

Performance	
Average Access Time	20-120 microseconds
Sustained Read Throughput	250 MB/sec
Sustained Write Throughput	115 MB/sec
Random IOPS Read Operations	45,000 IO/sec, sustained
Random IOPS Write Operations	16,000 IO/sec, sustained

IOPS?

Block Size?

- Prominent product specifications include:
- Up to 52,000 Sustained Random Read IOPS
 - Up to 17,000 Sustained Random Write IOPS

PEAK/Sustained IOPS - Sector 4KB aligned	(random preconditioned Sustained speed)	
4KB random READ	50K / 50K	50K / 32K
4KB random WRITE	50K / 50K	50K / 11K
8KB random READ	23K / 23K	23K / 23K
8KB random WRITE	28K / 28K	28K / 11K

Random Precondition
Sustain

Sequential read	Up to 250 MB/sec
Sequential write	170 MB/sec

Random or

Sustained

PERFORMANCE	
Sustained data transfer rate	240,000 Mb/s
I/O data transfer rate	300 MB/s

?

Market Segmentation

Client SSD



- Low cost
- C-MLC
- 0-7% over provisioning
- No supercap
- No Enterprise features
- No customization
- Warranty 1-3 yrs

Enterprise SSD




- Higher cost
- E-MLC/SLC
- 28-50% over provisioning
- Supercap support
- Enterprise features
- Customization
- Warranty 5 yrs

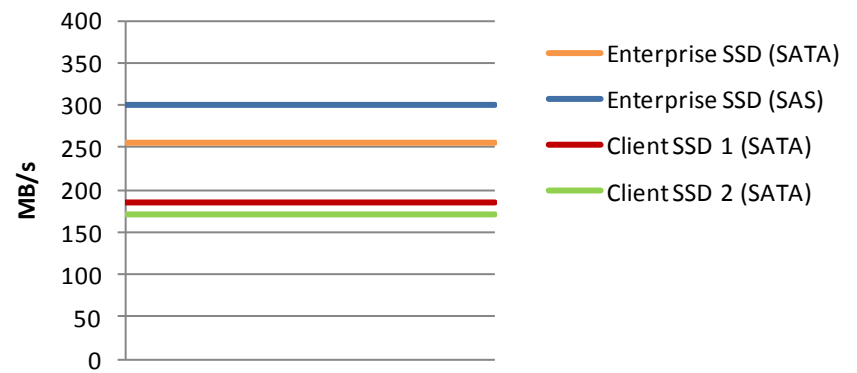


Performance Comparison Enterprise vs. Client SSD


Sequential Read



Sequential Write



Random Read



Random Write



Variables influencing Performance

- Platform
 - Test Hardware (CPU, interface, chipset, etc)
 - Software (OS, drivers)
- SSS Device Architecture
 - Flash geometry, cache, flash management algorithm, etc



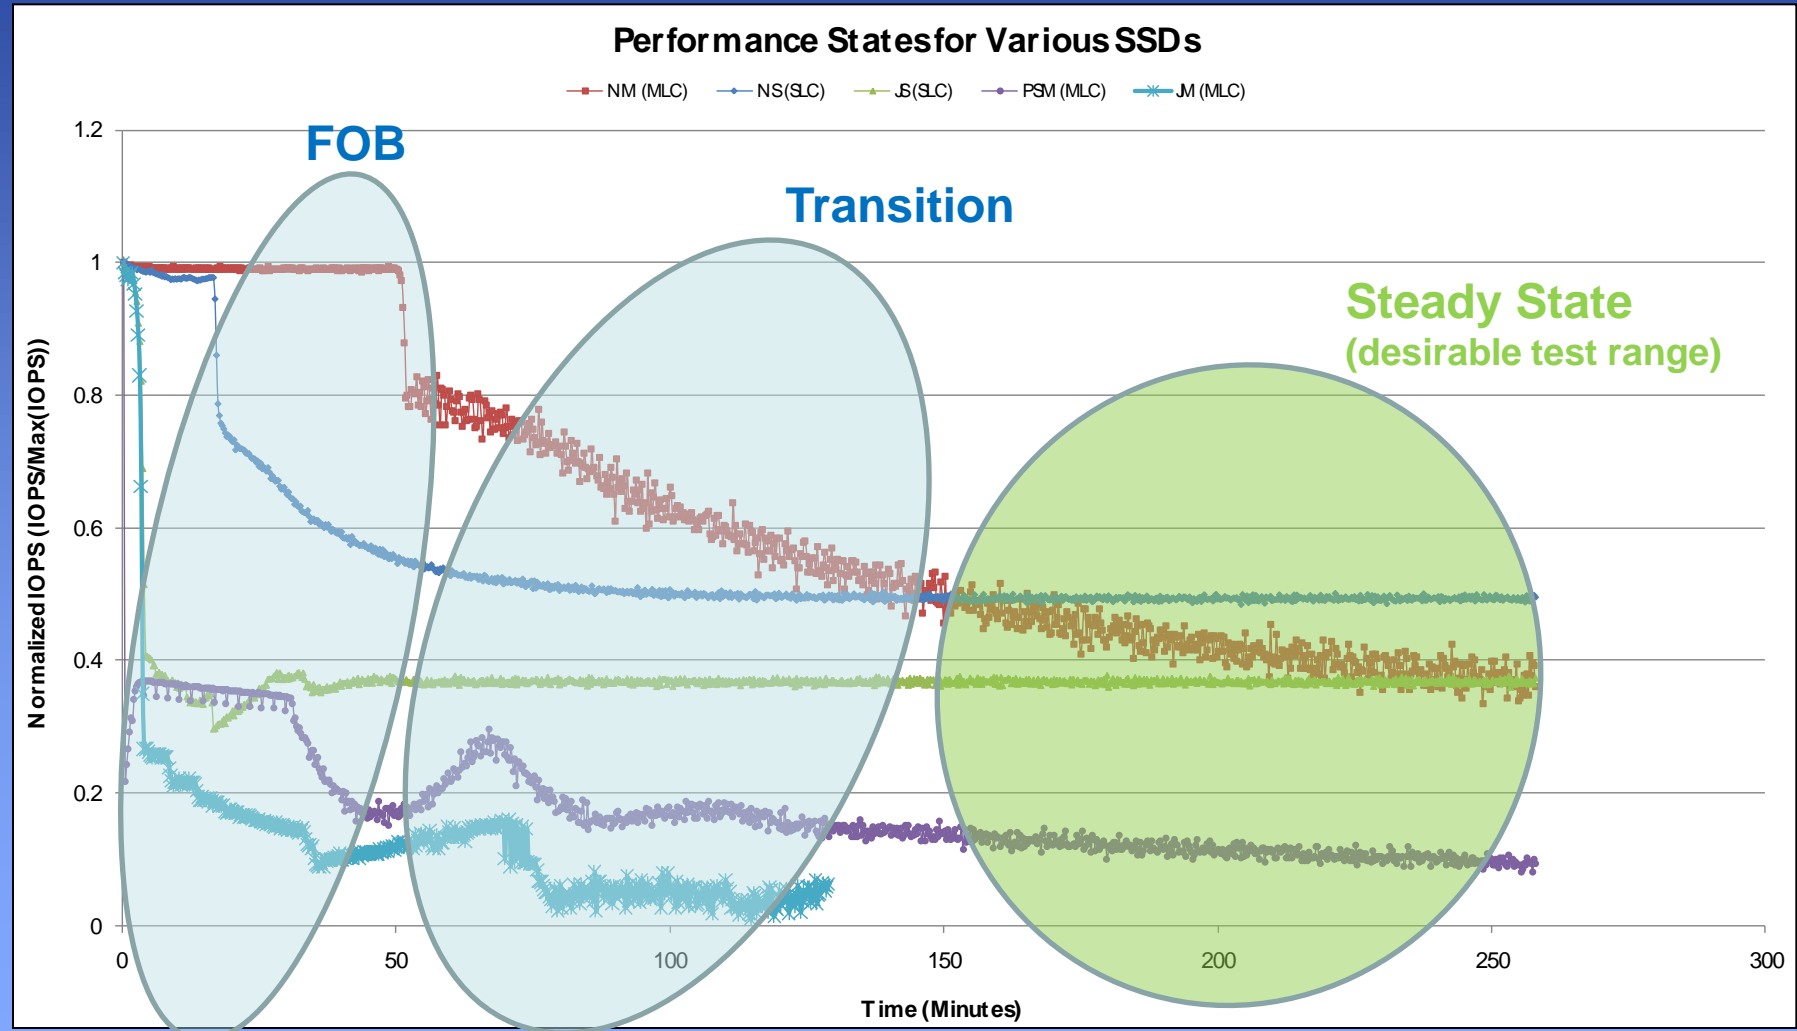


Variables influencing Performance

- Platform
 - Test Hardware (CPU, interface, chipset, etc)
 - Software (OS, drivers)
- SSS Device Architecture
 - Flash geometry, cache, flash management algorithm, etc
- Workload
 1. Write history & preconditioning: State of device before testing

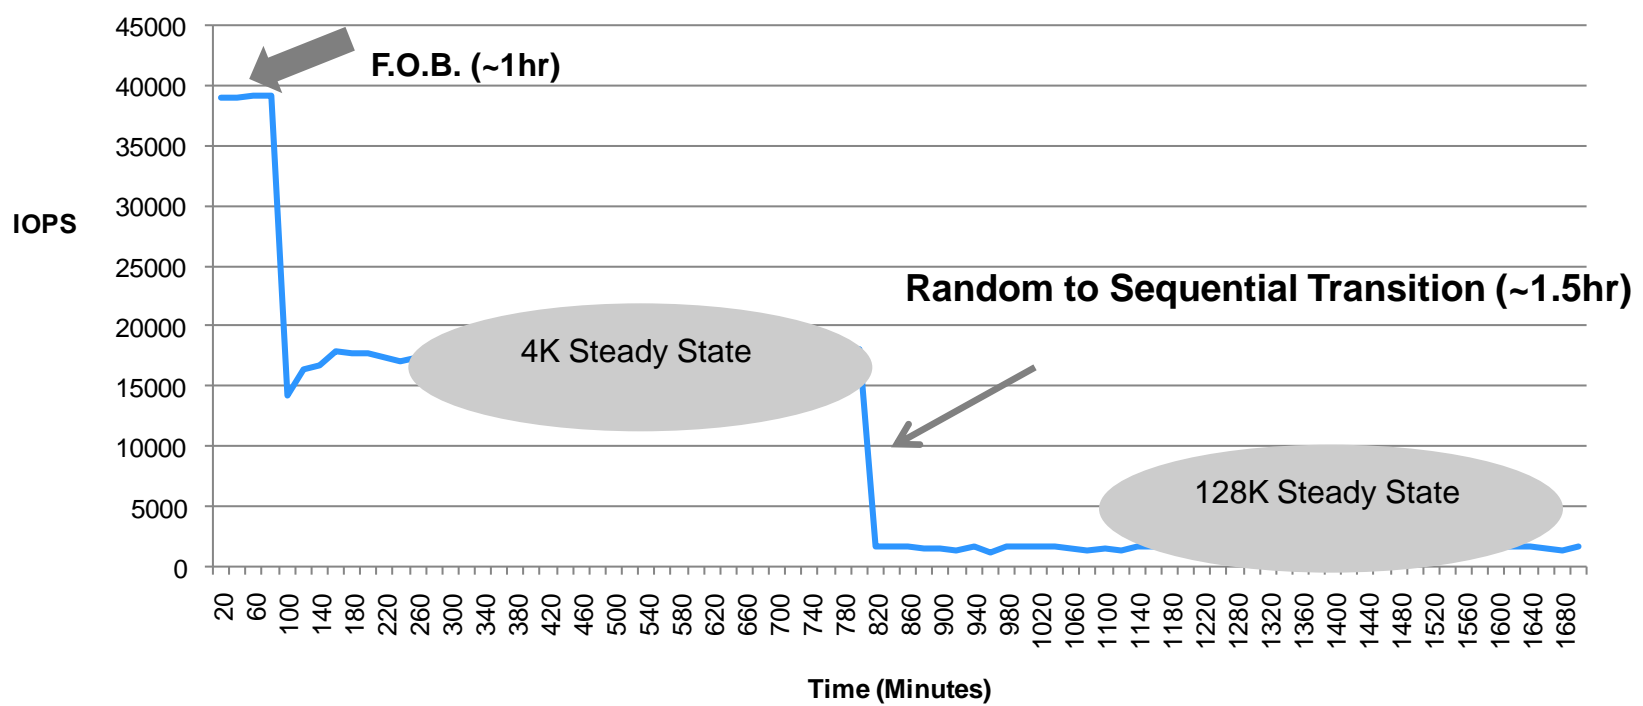


The need for Preconditioning



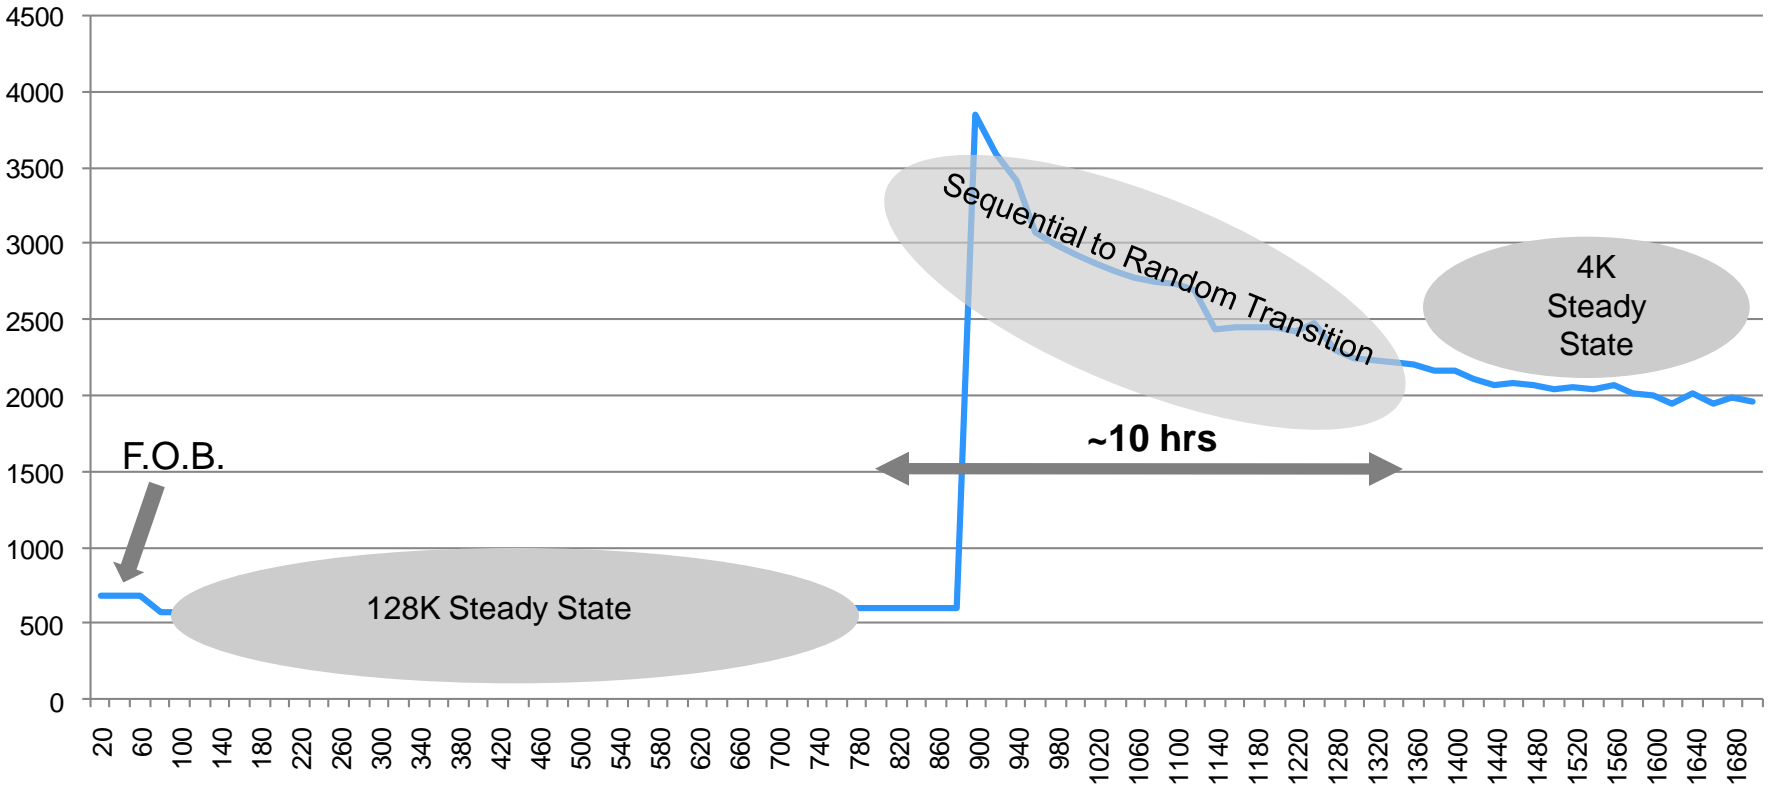
Write History - 1

4K Random to 128K Sequential Transition



Write History - 2

128K Sequential to 4K Random Transition






Variables influencing Performance

- Platform
 - Test Hardware (CPU, interface, chipset, etc)
 - Software (OS, drivers)
- SSS Device Architecture
 - Flash geometry, cache, flash management algorithm, etc
- Workload
 1. Write history & preconditioning: State of device before testing
 2. Workload pattern: Read/write mix, transfer size, sequential/random

Workload Pattern

3D IOPS Surface Profile

- 0.0-500.0
- 500.0-1000.0
- 1000.0-1500.0
- 1500.0-2000.0
- 2000.0-2500.0
- 2500.0-3000.0
- 3000.0-3500.0
- 3500.0-4000.0



Performance depends on

- Read/Write Mix
- Block Size
- Queue Depth (not shown)

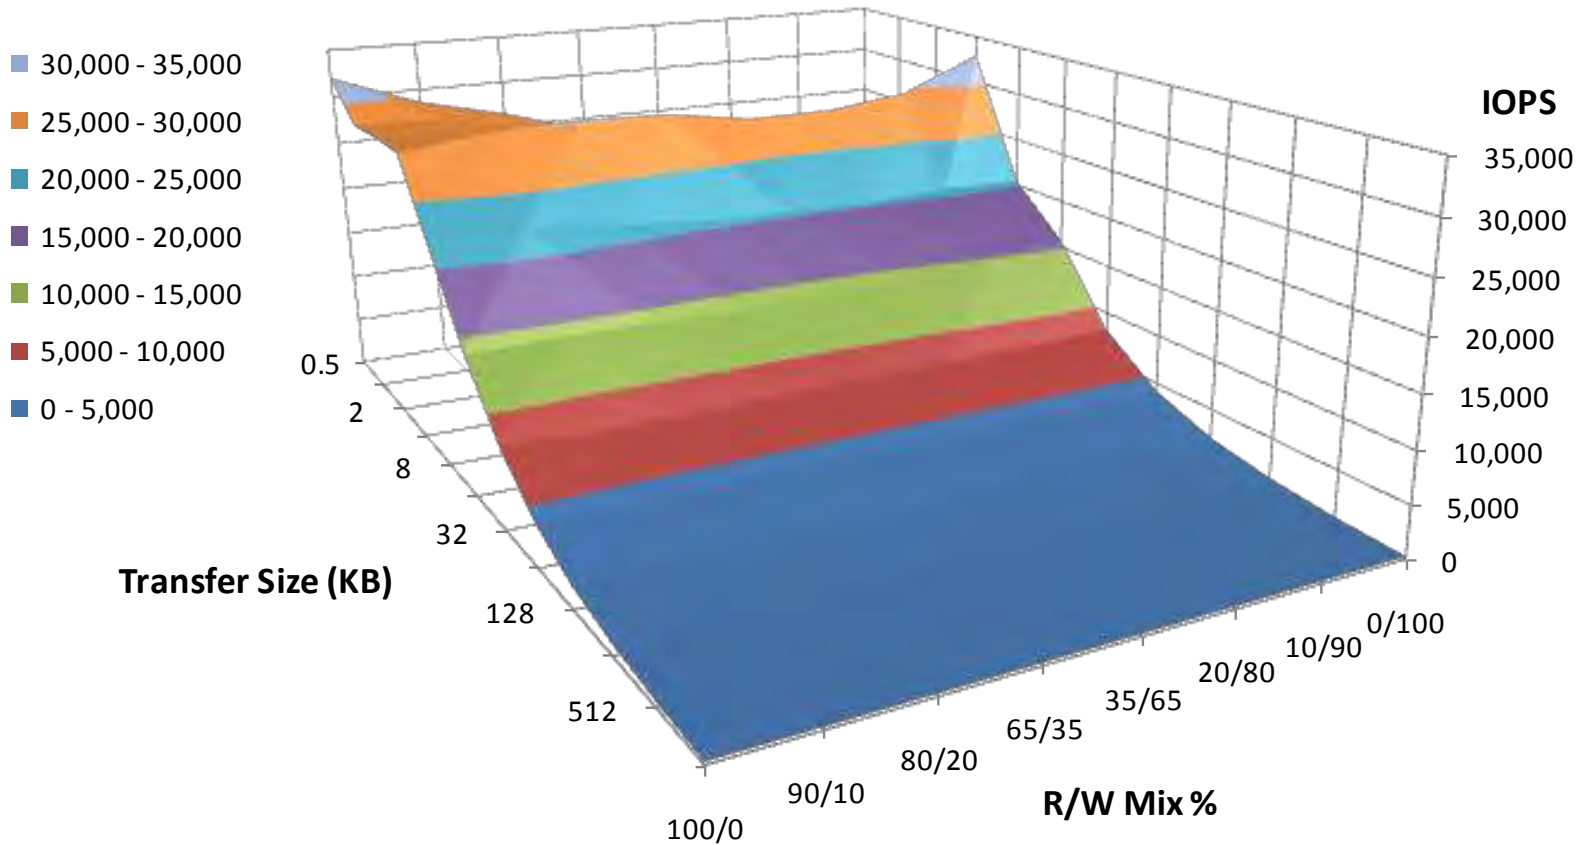


Variables influencing Performance

- Platform
 - Test Hardware (CPU, interface, chipset, etc)
 - Software (OS, drivers)
- SSS Device Architecture
 - Flash geometry, cache, flash management algorithm, etc
- Workload
 1. Write history & preconditioning: State of device before testing
 2. Workload pattern: Read/write mix, transfer size, sequential/random
 3. Data Pattern: The actual bits in the data payload written to the device

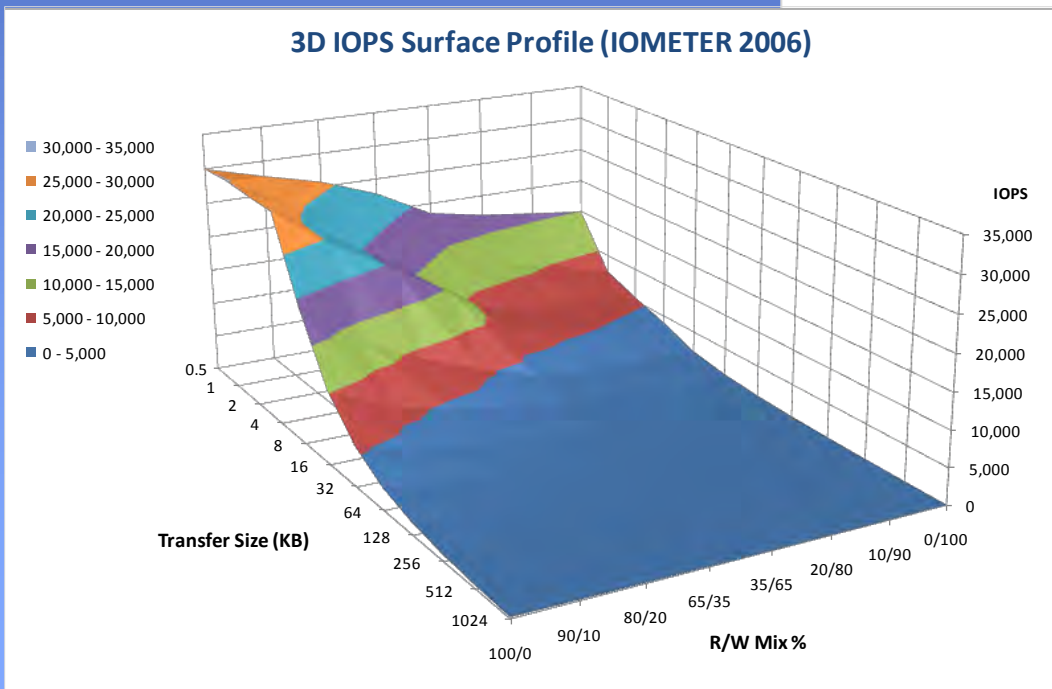

Dependency on data content - 1

3D IOPS Surface Profile (IOMETER 2008)



Dependency on data content - 2

IOMeter 2008 Low Entropy Data Content



IOMeter 2006 High Entropy Data Content



Benchmark Suites

	Test Suite	Client SSD	Enterprise SSD
PCMark	HDD Score, OS and application loading timing, user simulation (surfing web, windows media player, etc)	√	
SysMark	System-level test. Measures performance based on average response time, gives score (0-250)	√	
IOMeter	Sequential/Random performance, workload simulation (file server, web server workload, etc)	√	√
HDTach/ H2benchw	Performance stability, Sequential/Burst performance, Access Time	√	√
HD Tune	Performance stability, Sequential/Burst performance, Access Time		
Everest	Random Access Time (Read/Write)	√	√
VDBench	Workload generator, performance on DAS and NAS		√
Calypso CTS	Device (RAW) level, direct IO synthetic stimulus generator for both client and enterprise	√	√

The Need for Industry Standardization!



- SNIA Technical Working Group (TWG)

• Created in early 2000



- Spec 0.9 open for review now
- Spec 1.0 focus on Enterprise vs Client SSS testing



SNIA Performance Specification

1. Prepare the Device
 - Purge/Erase → put SSD back into “original” state
2. Workload independent preconditioning
 - Write data 2x capacity → bring device to known state
3. Steady State Testing (includes workload based preconditioning)
 - Run Test Loop up until steady state is achieved
 - Performance stays within $\pm 10\%$ margin
4. Test Report
 - Steady state convergence
 - Steady State Verification
 - Performance measurement (2D/3D)



Other Standardization Initiatives

- SSSI Group of SNIA
 - Technical Work Group (TWG) → Performance Benchmark Spec
 - Tech Dev Group → Performance Test Platform
- JEDEC 64.8
 - Specification for SSD endurance measurement
- SSDA
 - Testing of reliability (power cycling, data retention, endurance, etc) and OS compatibility (Windows 7)

Summary

- SSS Performance is dependent on many variables
- Comparing vendors is not trivial → industry standard required
- SNIA Performance Specs allows apples to apples comparison
 - Spec for review at http://www.snia.org/tech_activities/publicreview
 - Send your feedback to ssstwg@snia.org





Thank you!