Education

Introducing and Validating SNIA SSS Performance Test Suite

Esther Spanjer, SMART Modular

SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

Abstract

- SSS Performance Benchmarking Learning Objectives
  - Get a good understanding of the various parameters that influence the performance characteristics of SSDs
  - Get a full understanding of the proposed SNIA Performance Measurement Specification
  - Provide step-by-step guidance on how to set up a test benchmark that enables comparison among the various SSS devices

Definition of SSS

Education SNIA

SSS = Solid State Storage

Traditional hard disk drive

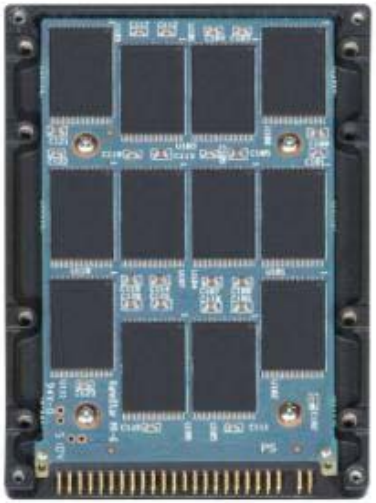

Solid state hard drive

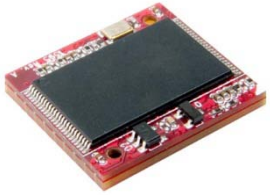

Flashy fists fly as OCZ and DDRdrive row over SSD performance

"dirty tricks used to drastically inflate IOPS results“…

" the egregious disparity between promised performance and reality”…

"This really isn’t an 'apples to apples' comparison”…

“you can't use an X1 dragster on the open road but you can use a Ferrari Vertex 2 EX”…

http://www.theregister.co.uk/2011/01/14/ocz_and_ddrdrive_performance_row/

The Performance Landscape
One Year Later

Education SNIA

- 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 6.7 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

MB/s or Mb/s?

Performance

Average Access Time	20-120 microseconds
Sustained Read Throughput	2[illegible] MBytes/sec
Sustained Write Throughput	115 Mbytes/sec
Random IOPS Read Operations	45,000 IO/sec, sustained
Random IOPS Write Operations	16,000 IO/sec, sustained

IOPS?

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

Block Size?

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

Sustained Speed?

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

Random or Sustained?

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

Up to?

Market Segmentation

Education SNIA

Client SSD

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

Enterprise SSD

- Higher cost
- E-MLC/SLC
- 28-50% over provisioning
- Backup power circuit
- Enterprise features
- Customization
- Warranty 5 yrs

Performance Comparison
Enterprise vs. Client SSD

Education SNIA

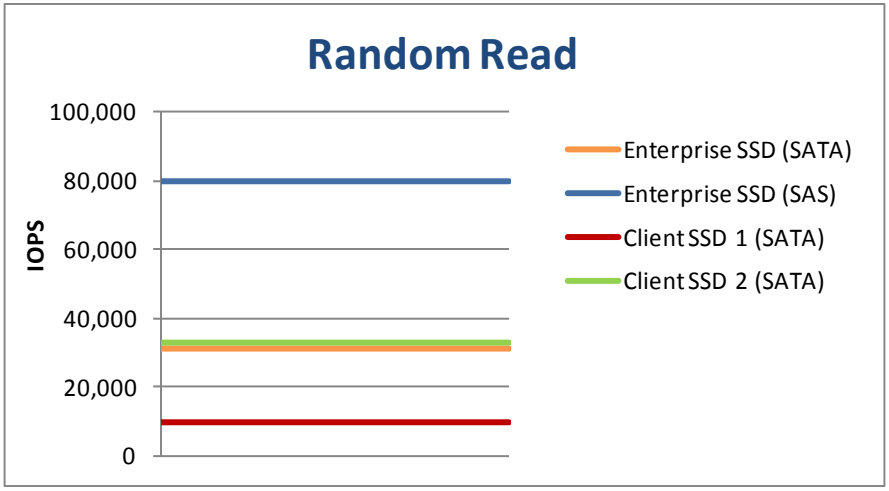

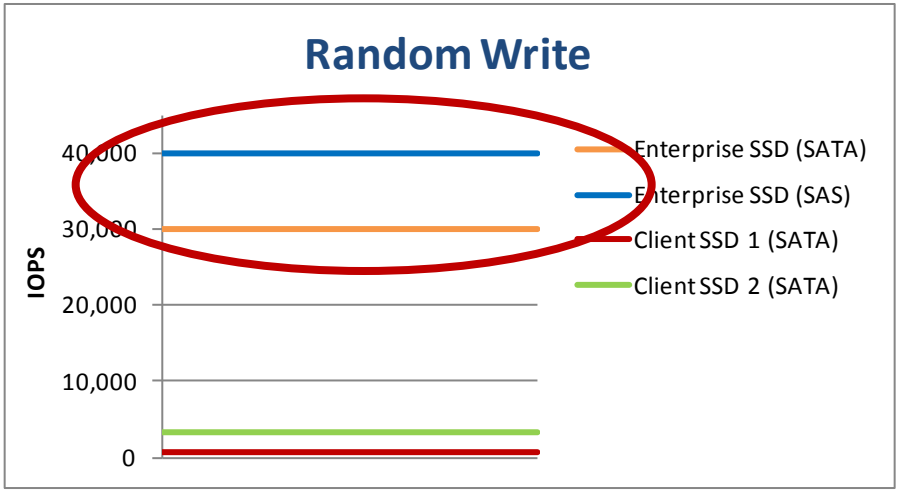

Variables influencing Performance

- 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

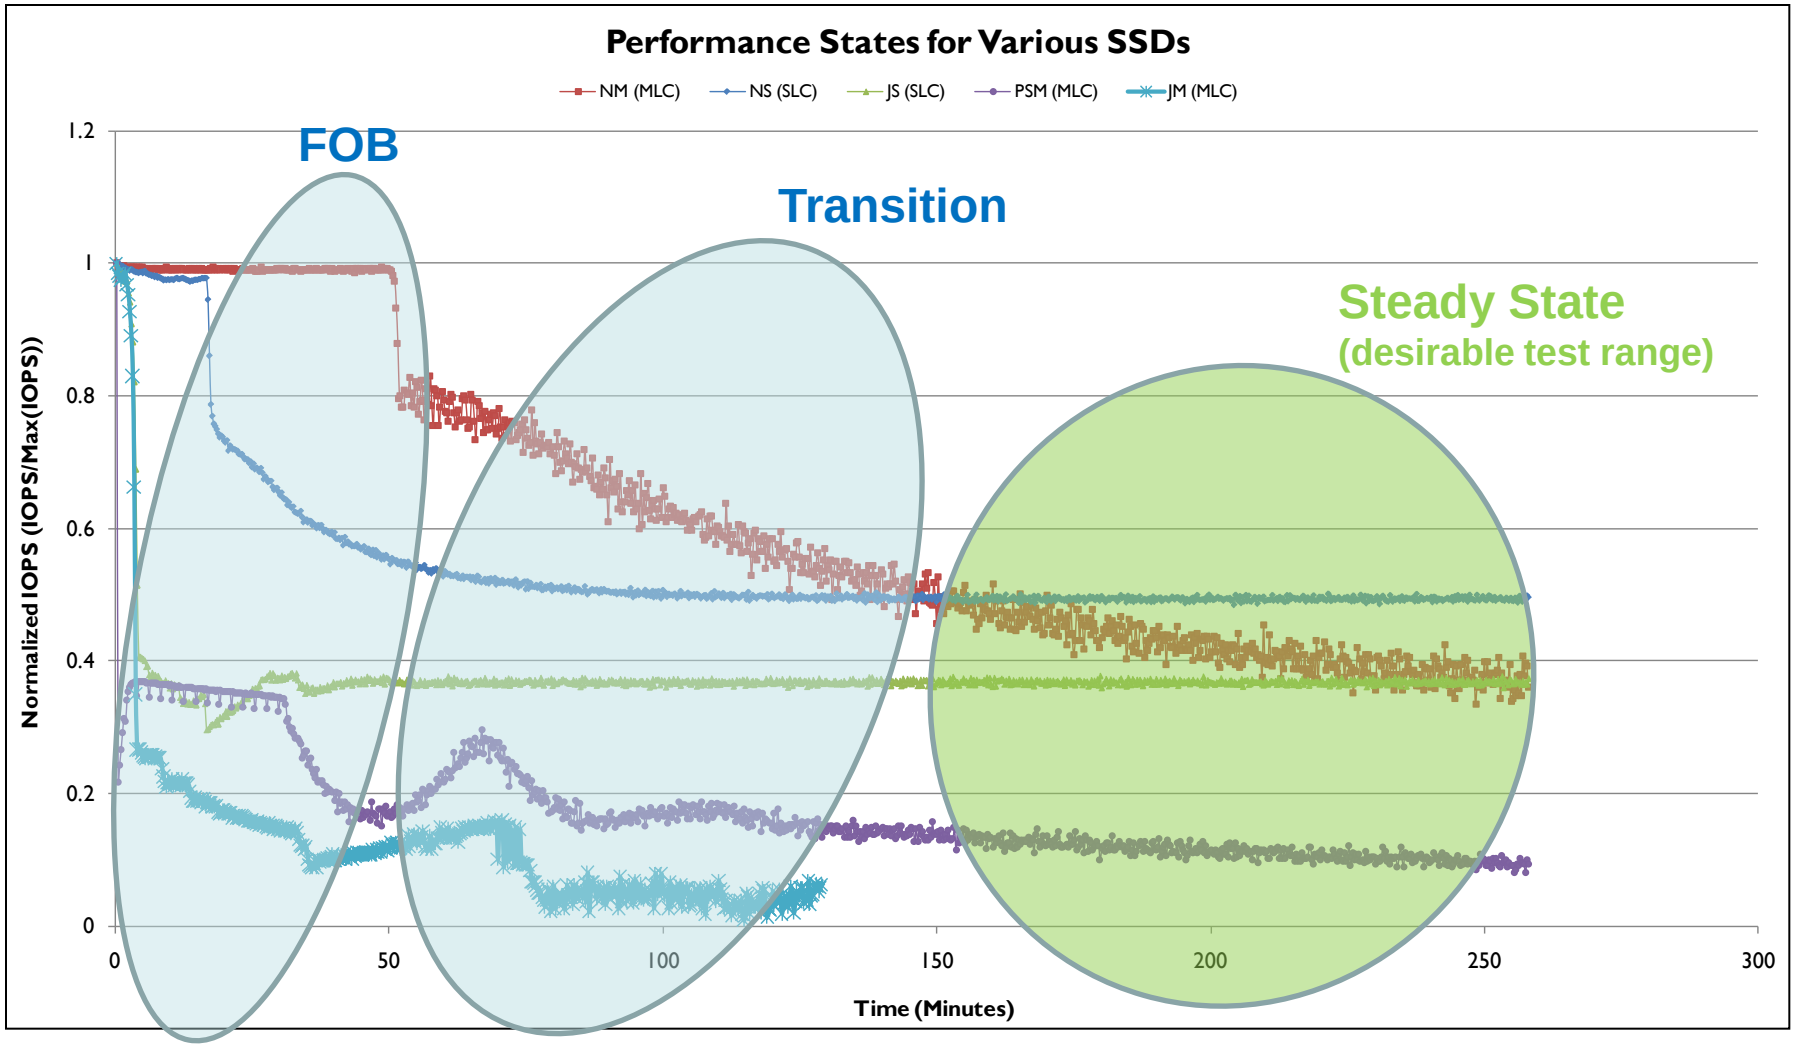

Write History - I

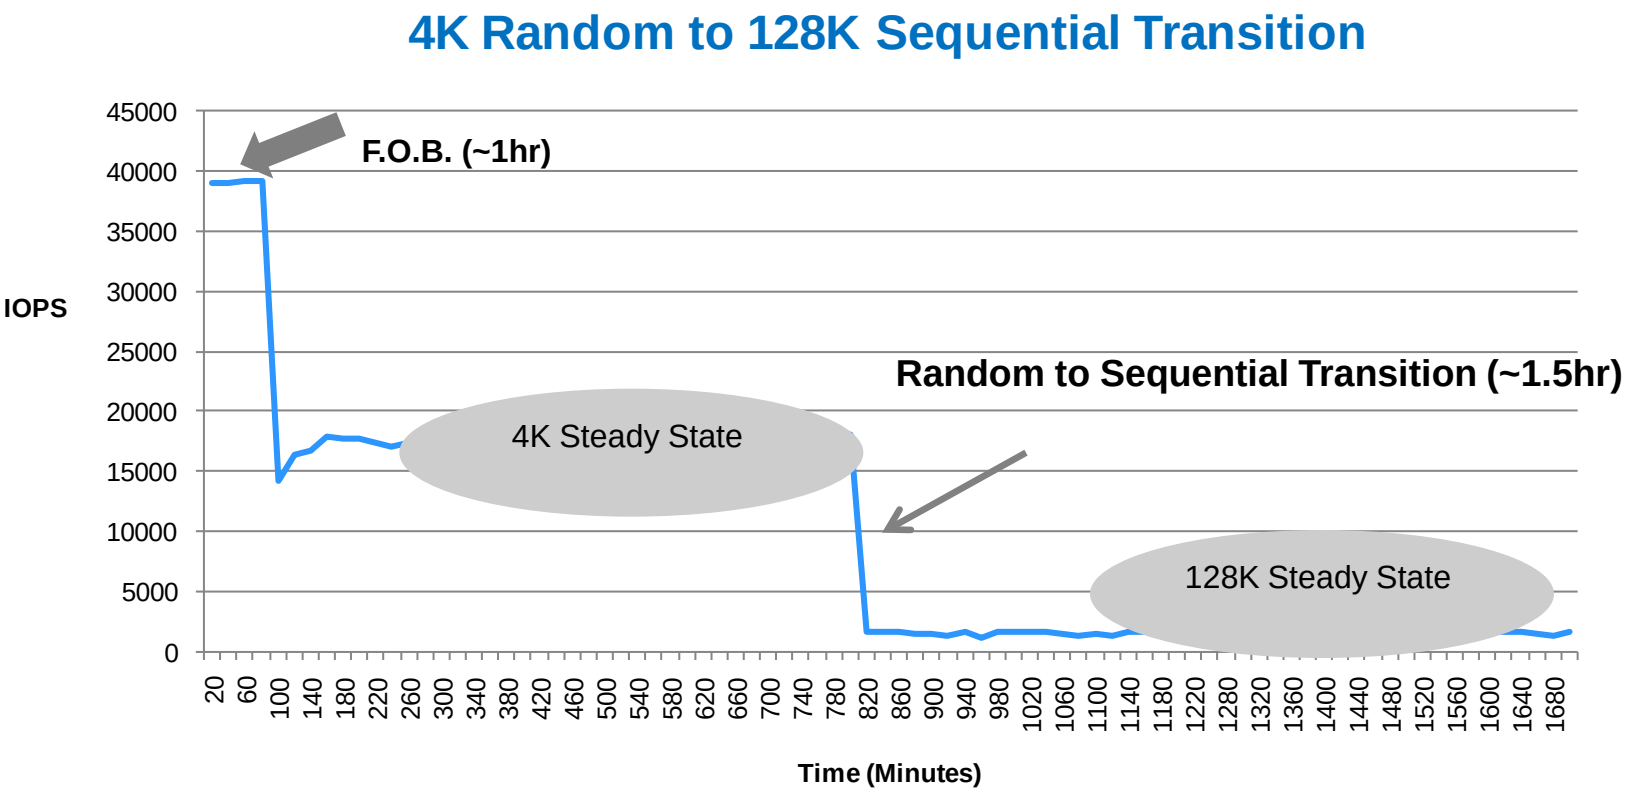

Write History - 2

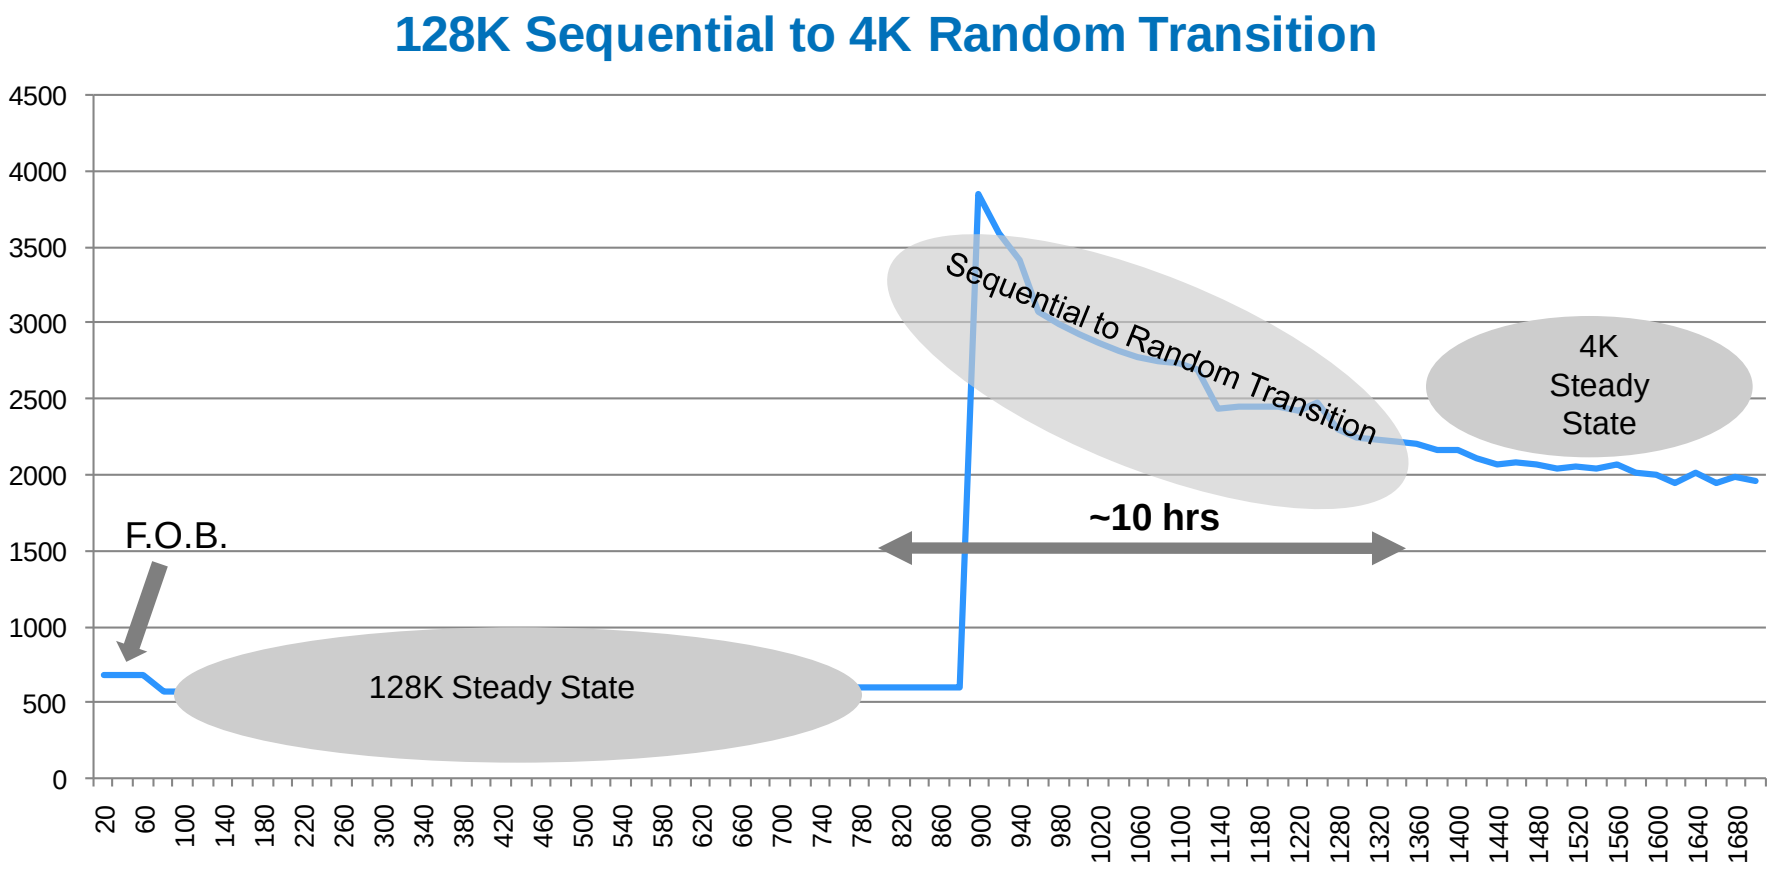

Variables influencing Performance

- 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

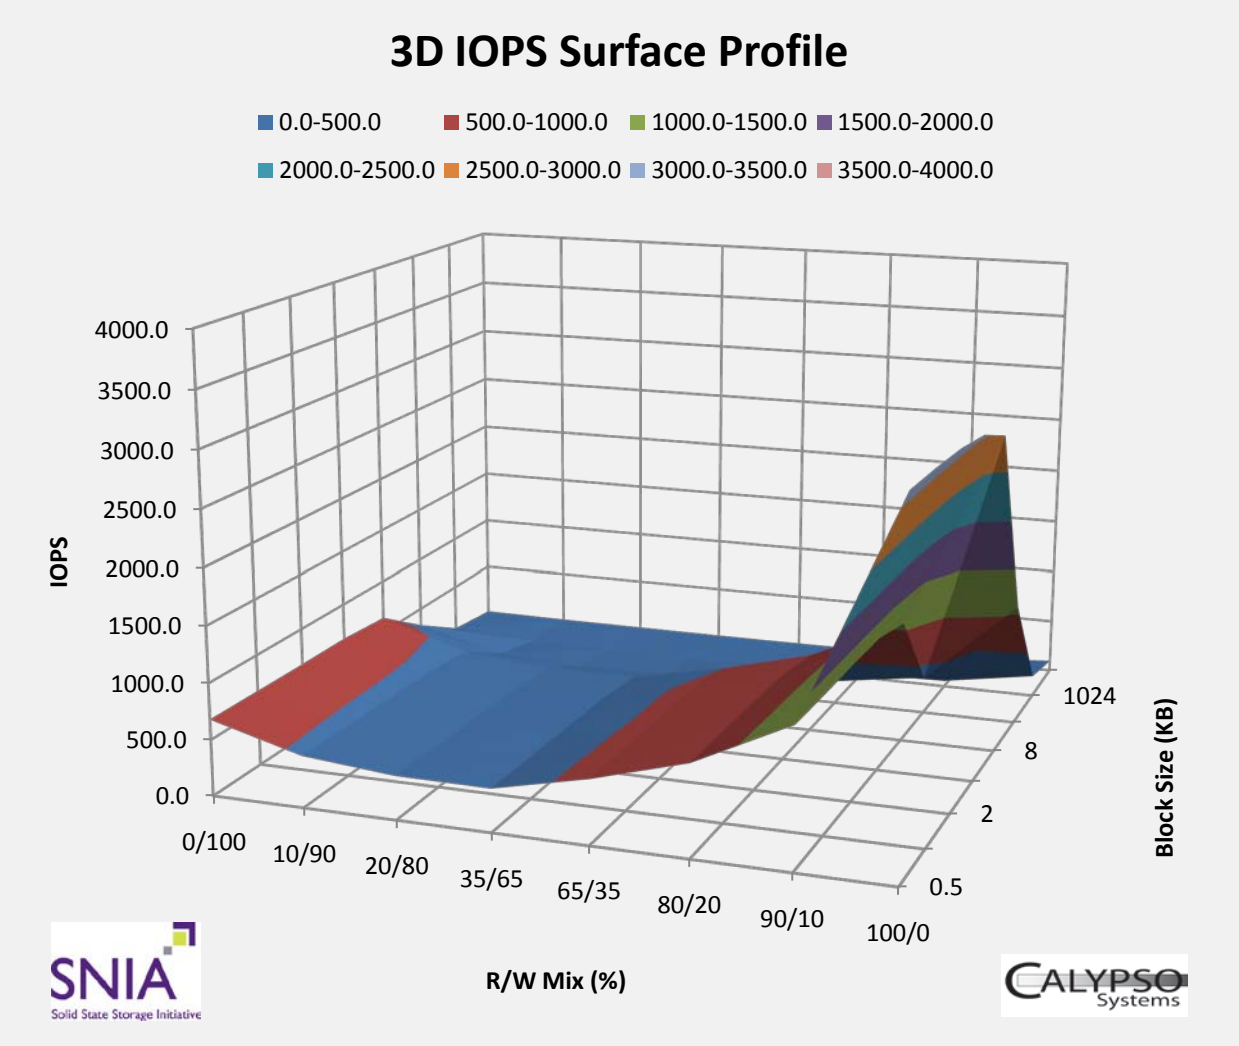

Performance depends on

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

Note: Shown 3D IOPS image courtesy of Calypso Systems

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 - I

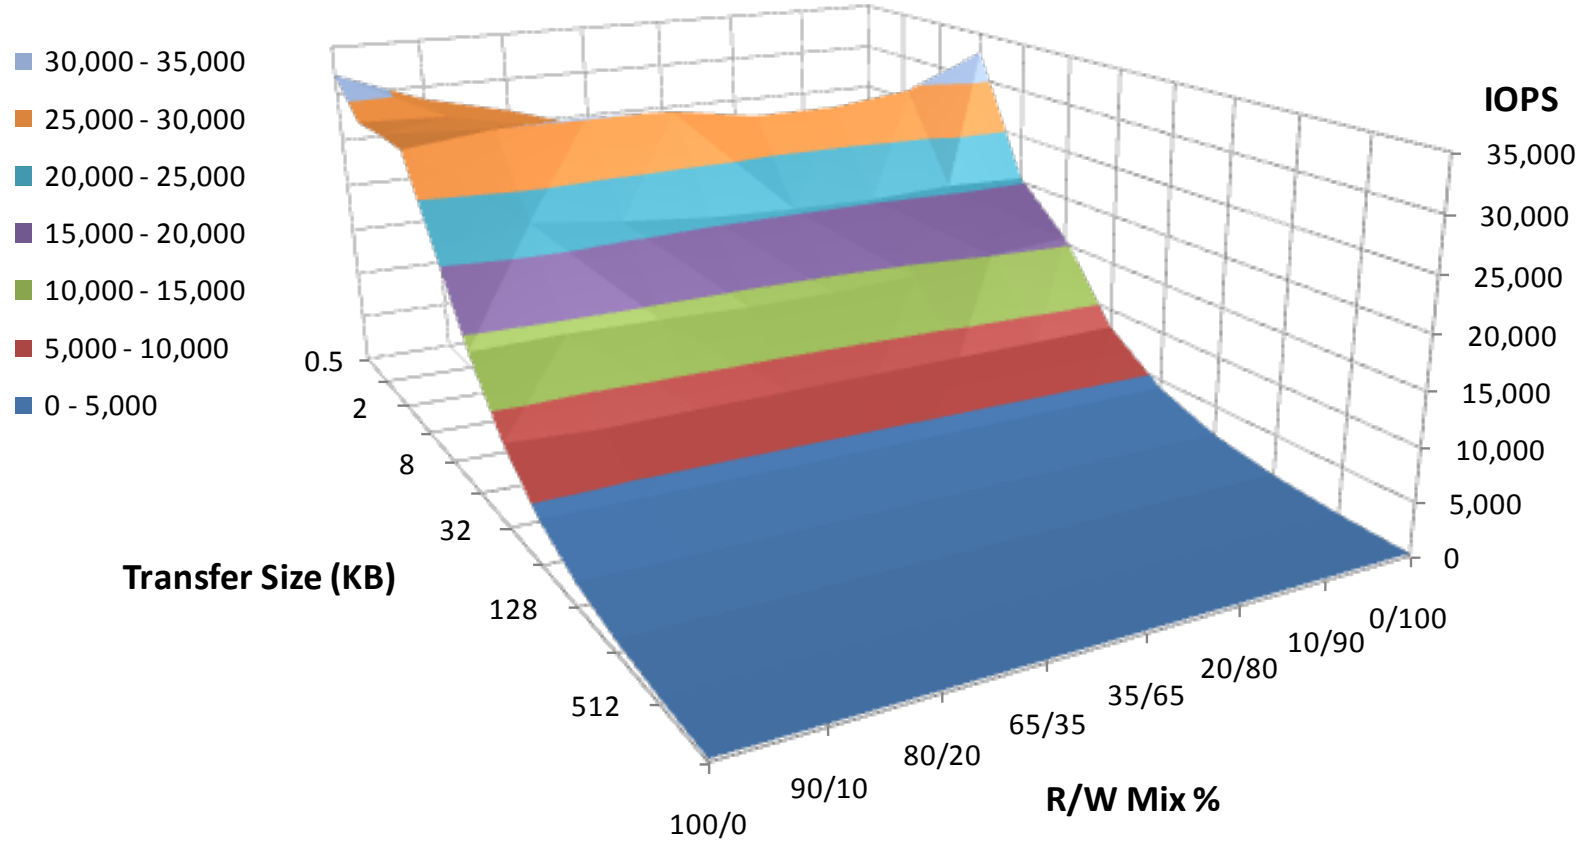

Dependency on data content - 2

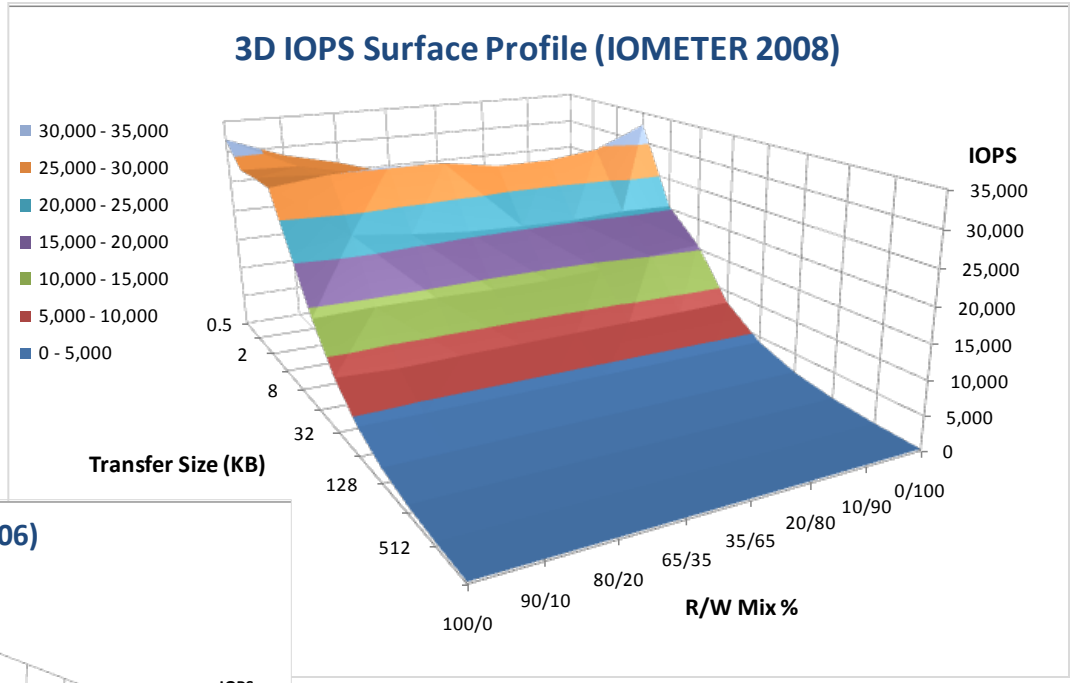

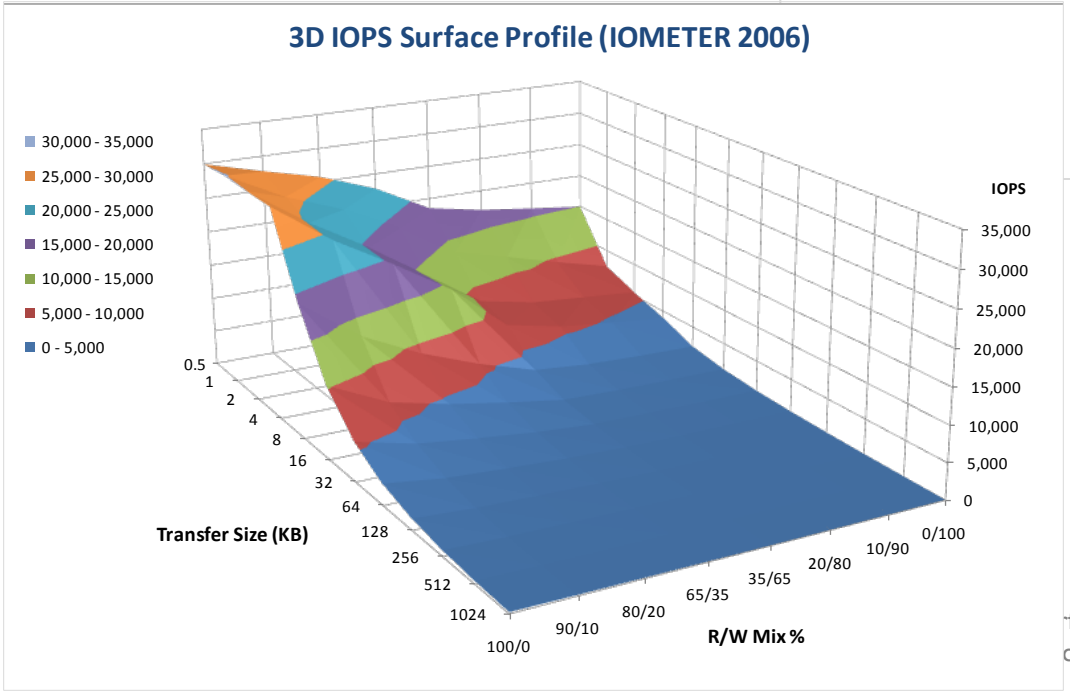

Benchmark Suites

	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)
  -
- S
  c
  -
  -
  -
- S
  - Future: Split of test specification into separate Enterprise and Client categories

ers

effort

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 ± 10% margin
4. Test Report
   - Steady state convergence
   - Steady State Verification
   - Performance measurement (2D/3D)

Preconditioning

- Preconditioning is key to get repeatable results
- Preconditioning needed to get drive in Steady State, after which performance can be measured

Preconditioning

- Preconditioning is key to get repeatable results
- Preconditioning needed to get drive in Steady State, after which performance can be measured
- Two types of preconditioning
  - Workload independent – write 2x capacity with 128KB sequential writes
  - Workload dependent – run workload itself until steady state is achieved

Steady State

- Measurement window is interval for last 5 measured rounds (i.e. test loops) that show steady sate results
- Steady State is
  - Variation of y measuremen window is wi 20% of avera
  - Trending of y within measuremen window is w 10% of avera

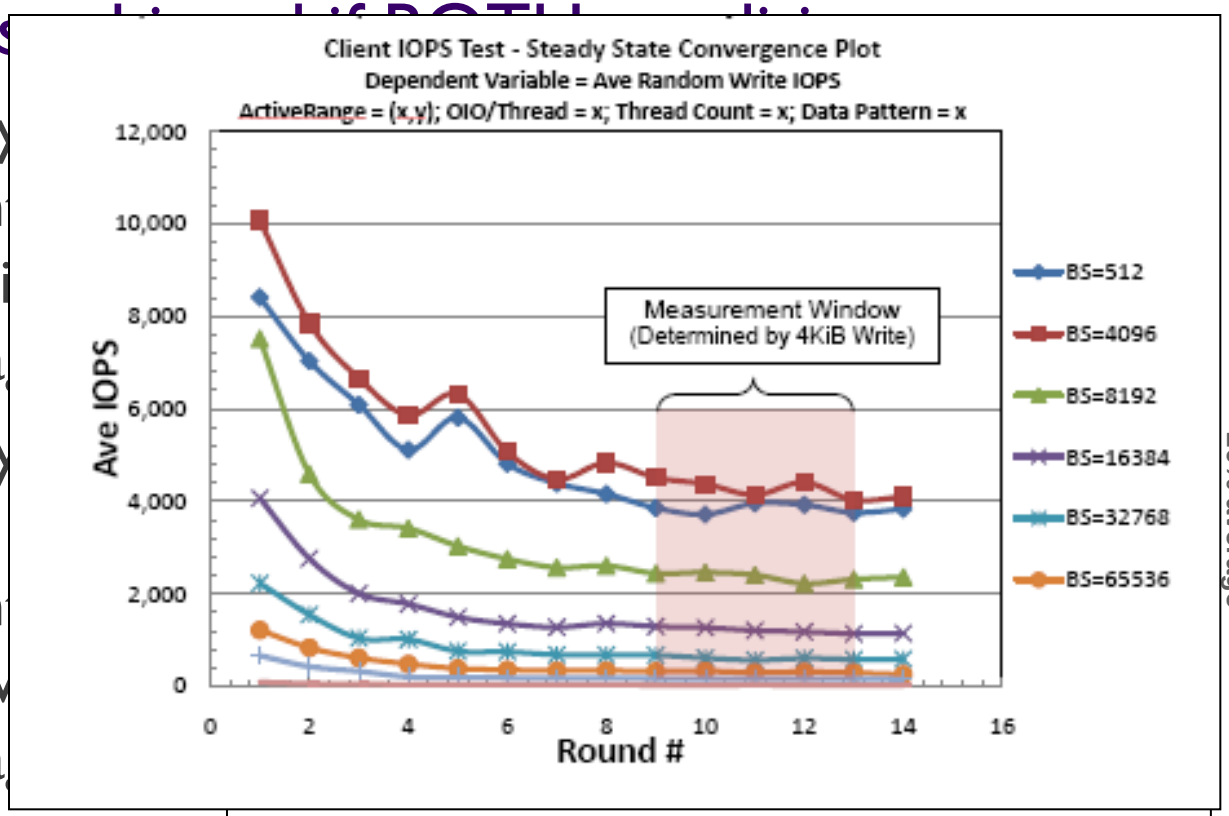

Performance Workloads & Tests

Client Test

- Random IOPS
  - 100/0, 95/5, 65/35, 50/50, 35/65, 5/95, 0/100
  - 1024K, 128K, 64K, 32K, 16K, 8K, 4K, 0.5K
- Sequential MB/s
  - 100/0, 0/100
  - 1024K
- Latency (random access)
  - 100/0, 65/35, 0/100
  - 8K, 4K, 0.5K

Enterprise Test

- Random IOPS
  - 100/0, 95/5, 65/35, 50/50, 35/65, 5/95, 0/100
  - 1024K, 128K, 64K, 32K, 16K, 8K, 4K, 0.5K
- Sequential MB/s
  - 100/0, 0/100
  - 1024K, 64K, 8K, 4K, 0.5K
- Latency (random access)
  - 100/0, 65/35, 0/100
  - 8K, 4K, 0.5K

Version 1.0 will focus further on differentiating Client vs. Enterprise

Standard Reporting - 1

Steady State Convergence Plot

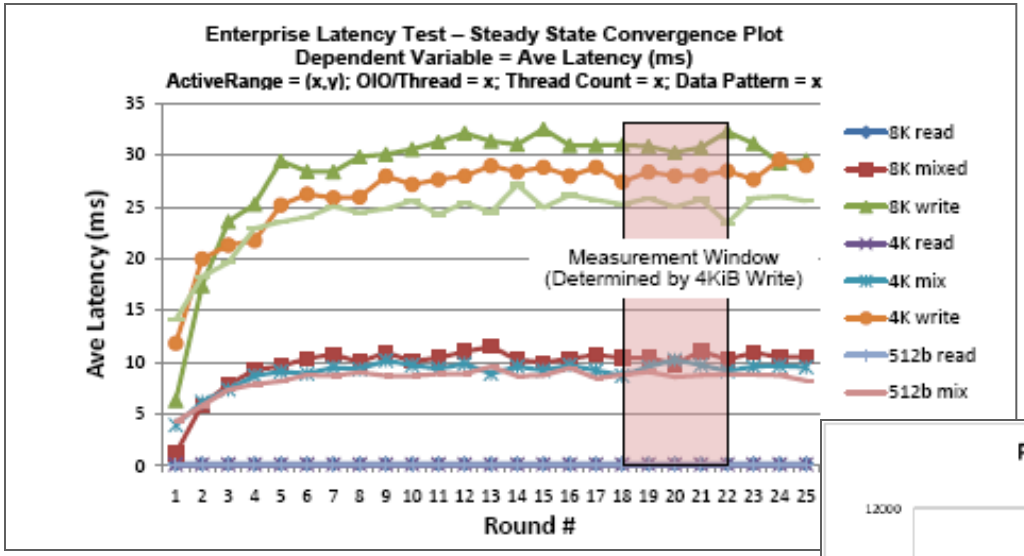

Steady State Verification

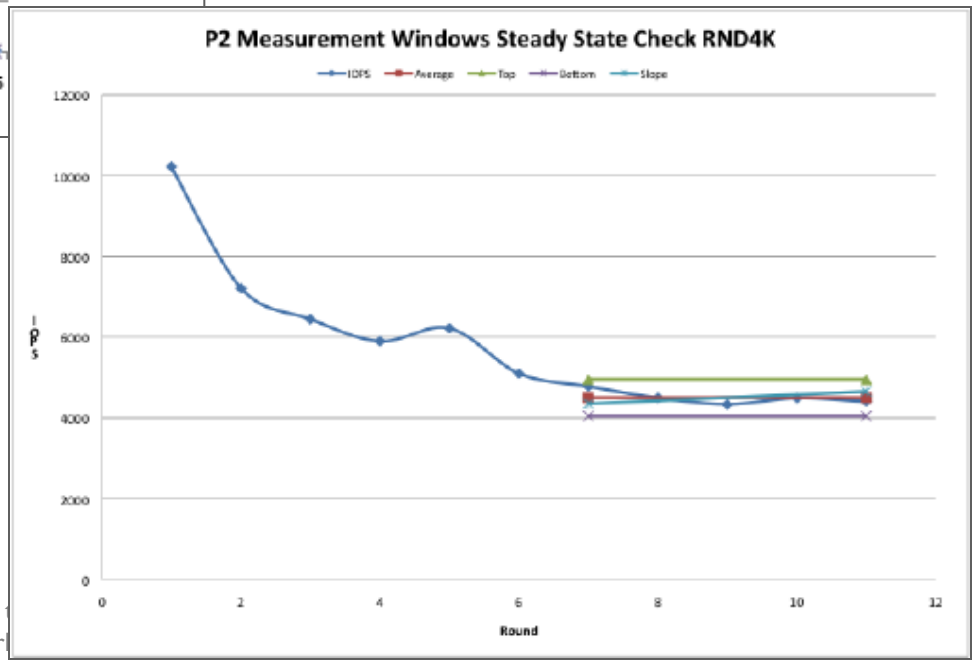

Standard Reporting - 2

Performance Measurements 3D

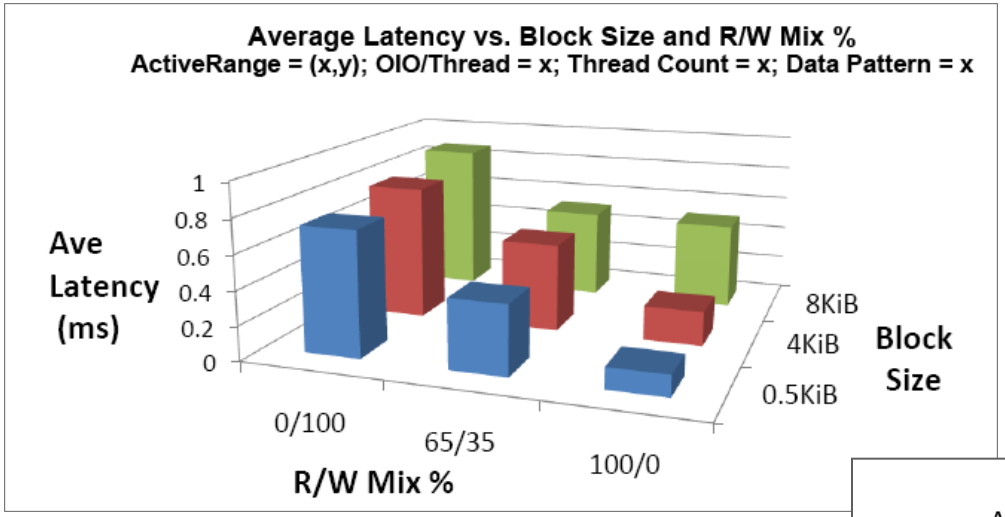

Performance Measurements 2D

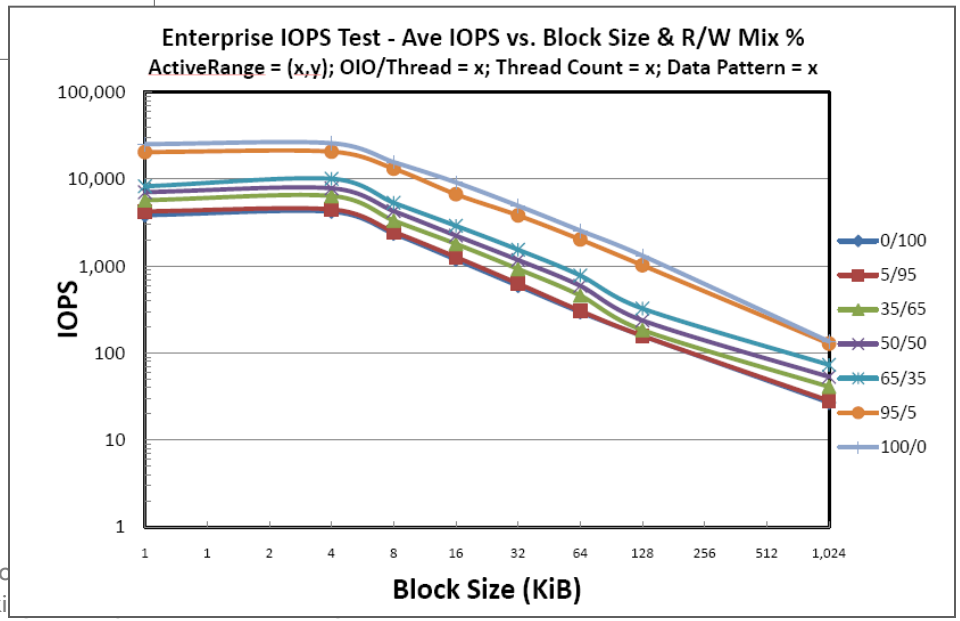

Sample Test Report included

Informative Annex A – Sample Test Report
Summary Report Page

Solid State Storage Performance Test Specification (PTS)				Rev.	0.7
				Page 1	of 5
Device Under Test (DUT)	ABC Co.	SSS TWG PTS Summary Report	SNIA SSS TWG	SNIA	
Model No.:	ABC123	Test Specification:	Test Run Date:	Apr 3-23 2010	
Form Factor:	2.5"		Report Date:	June 01, 2010	
NAND Capacity:	256 GB MLC	SNIA SSS TWG PTS v 0.9	Test Sponsor:	Calypso Systems	
DUT Interface:	SATAII, SAS HBA		Auditor Name:	N/A	

Testing Summary: Tests Run

Test		Preparation			Test Loop Parameters	
		Purge Type	Workload Independent Preconditioning	Active Range %	QD / TC / OIO	Data Pattern
8.1	Client IOPS	Secure Erase	2X 128K SEQ	100%	QD 16 / TC 4 / OIO 64	RND
8.2	Client IOPS OPT – AR 10%	Secure Erase	2X 128K SEQ	10%	QD 16 / TC 4 / OIO 64	RND
8.3	Client IOPS OPT– File Data	Secure Erase	2X 128K SEQ	100%	QD 16 / TC 4 / OIO 64	NON RND File as Data Pattern
9.1	Client Throughput					
10.1	Client Latency					

General Device Description

Device Under Test (DUT)		System Hardware Configuration		System Software Configuration	
Manufacturer	ABC Co.	System Mfgr	Calypso Systems, Inc.	Operating Sys	Linux CentOS 5.4
Model No.	ABC123	Model No.	RTP 2.0		
Serial No.	123.xxx.ffff	Motherboard	Intel 5520HC	Test SW Tool(s)	Calypso CTSv6.5
Firmware Rev No.	Fff.hhhh.abc.123	Chassis	Intel SC5560DP		
User Capacity	256 GB	CPU Type	Intel 3.2GHz W5580	Other SWTool(s)	
Interface/Speed	6Gb/s SATAII	No. CPUs	Single		
Form Factor	2.5"	DRAM Type	1333MHz DDR3 ECC		
Media Type	MLC	DRAM Amt	12 GB		
Major Features:		DUT I/F	SAS HBA		
NCQ:	YES	SAS HBA	LSI 6Gb/s 9212-4e4i		
Hot Plug:	YES	SATAII	ICH10HR		
Sanitize Support:	NO	PCI-e	Gen 2 (8) lane		
Other 1:		Boot HDD	160 GB 7200RPM		
Other 2:		Optical Drive			

Informative Annex A – Sample Test Report
Client IOPS REQUIRED - Report Page

Solid State Storage Performance Test Specification (PTS)							Rev.	0.7
							Page 2	of 5
Device Under Test (DUT)	ABC Co.	8.1 Client IOPS Test				SNIA SSS TWG	SNIA	
Key Set Up Data		DUT Preparation		Test Loop Parameters			Steady State	
DUT	256GB MLC	Purge Done:	Yes	Required:	Data Pattern	random	Convergence	YES
DUT I/F	SAS HBA	Purge Type:	Security Erase				Rounds	7-11
Test HW	RTP 2.0			Tester Choice:	OIO/Thread	16	Active Range	
OS	CentOS 5.4	Workload Independent Preconditioning	2X 128K Sequential Wrt		Thread Count	4	Required:	100%
Test SW	CTSv6.5						Optional:	N/A

8.1.1 Steady State Convergence Plot – All Block Sizes

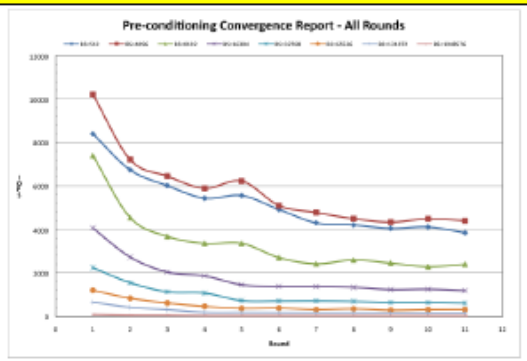

8.1.2 Steady State Convergence Plot – 4K Block Sizes

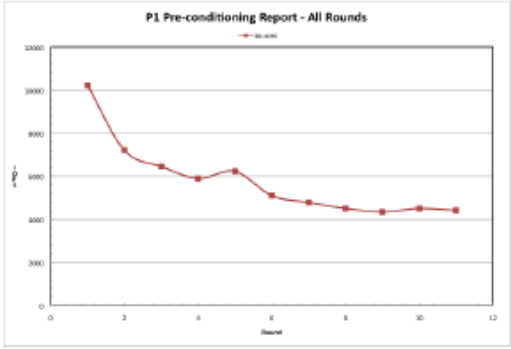

Other Standardization Initiatives

- 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

Q&A / Feedback

- Please send any questions or comments on this presentation to SNIA: tracksolidstate@snia.org

Many thanks to the following individuals for their contributions to this tutorial.

- SNIA Education Committee

David Landsman,
Easen Ho
Eden Kim
Neal Ekker,
Dan Le