



Education

# **Apples to Apples, Pears to Pears in SSS performance Benchmarking**

Esther Spanjer, SMART Modular

- 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.**

## ➤ 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


## SSS = Solid State Storage



Traditional hard disk drive



Solid state hard drive



# The Performance Landscape

- Bandwidth Performance Specifications
  - Sustained Sequential Read: Up to 250 MB/s
  - Sustained Sequential Write: Up to 70 MB/s
- Read and Write IOPS specifications (Iometer Queue Depth 32)
  - Random 4 KB Reads: Up to 35 K IOPS
  - Random 4 KB Writes: Up to 3 K IOPS

## Performance

Average Access Time	20-120 microseconds
Sustained Read Throughput	220 Mbytes/sec
Sustained Write Throughput	115 Mbytes/sec
Random IOPS Read Operations	42,000 IOPS sustained
Random IOPS Write Operations	16,000 IOPS sustained

**IOPS?**

**Sustained?**

Like other high-performance, enterprise-class flash drives, [redacted] drive is being positioned as an alternative to traditional 15,000-rpm serial-attached SCSI (SAS) hard disk drives. [redacted] said its new "performance-optimized" enterprise SSD can process IOPS (input/output per second) more than 10 times faster than the fastest SAS hard drive, with a sequential read rate of 230MB/sec. and a sequential write rate of 180MB/sec.

**10 times faster than HDD?**

**Block Size?**

**Random?**

**Up to?**

- Increased High Speed Performance
- 800 Megabit Full Duplex Burst Rate
  - Up to 230 MB/sec Sustained Rate
  - Up to 35,000 IOPS
  - 30 to 100 usec Access Time

Table 3: Sustained and Random Read/Write Performance

Parameter	Value
Burst Read	150 MBytes/sec <sup>4</sup>
Burst Write	150 MBytes/sec
Sustained Read	Up to 115 MBytes/sec
Sustained Write	Up to 75 MBytes/sec
Random Read <sup>5</sup>	6100 Input/Output Operations Per Second (IOPS)
Random Write	400 IOPS
Random 67% Read, 33% Write	1120 IOPS

<sup>4</sup> One megabyte, or MByte, equals 1,048,576 bytes.

<sup>5</sup> Random performance values are based on 4 KByte transfers.

## ➤ Platform

- ◆ Test Hardware (CPU, interface, chipset, etc)
- ◆ Software (OS, drivers)

## ➤ SSS Device Architecture


- ◆ Flash geometry, cache, flash management algorithm, etc

## ➤ Workload

- ◆ Write history (Terabytes written, % spares)
- ◆ Random, sequential, read/write mix, etc
- ◆ Preconditioning (Random, sequential, transfer size, etc)
- ◆ Data content
- ◆ TRIM command

# The 3 dimensions of SSS performance

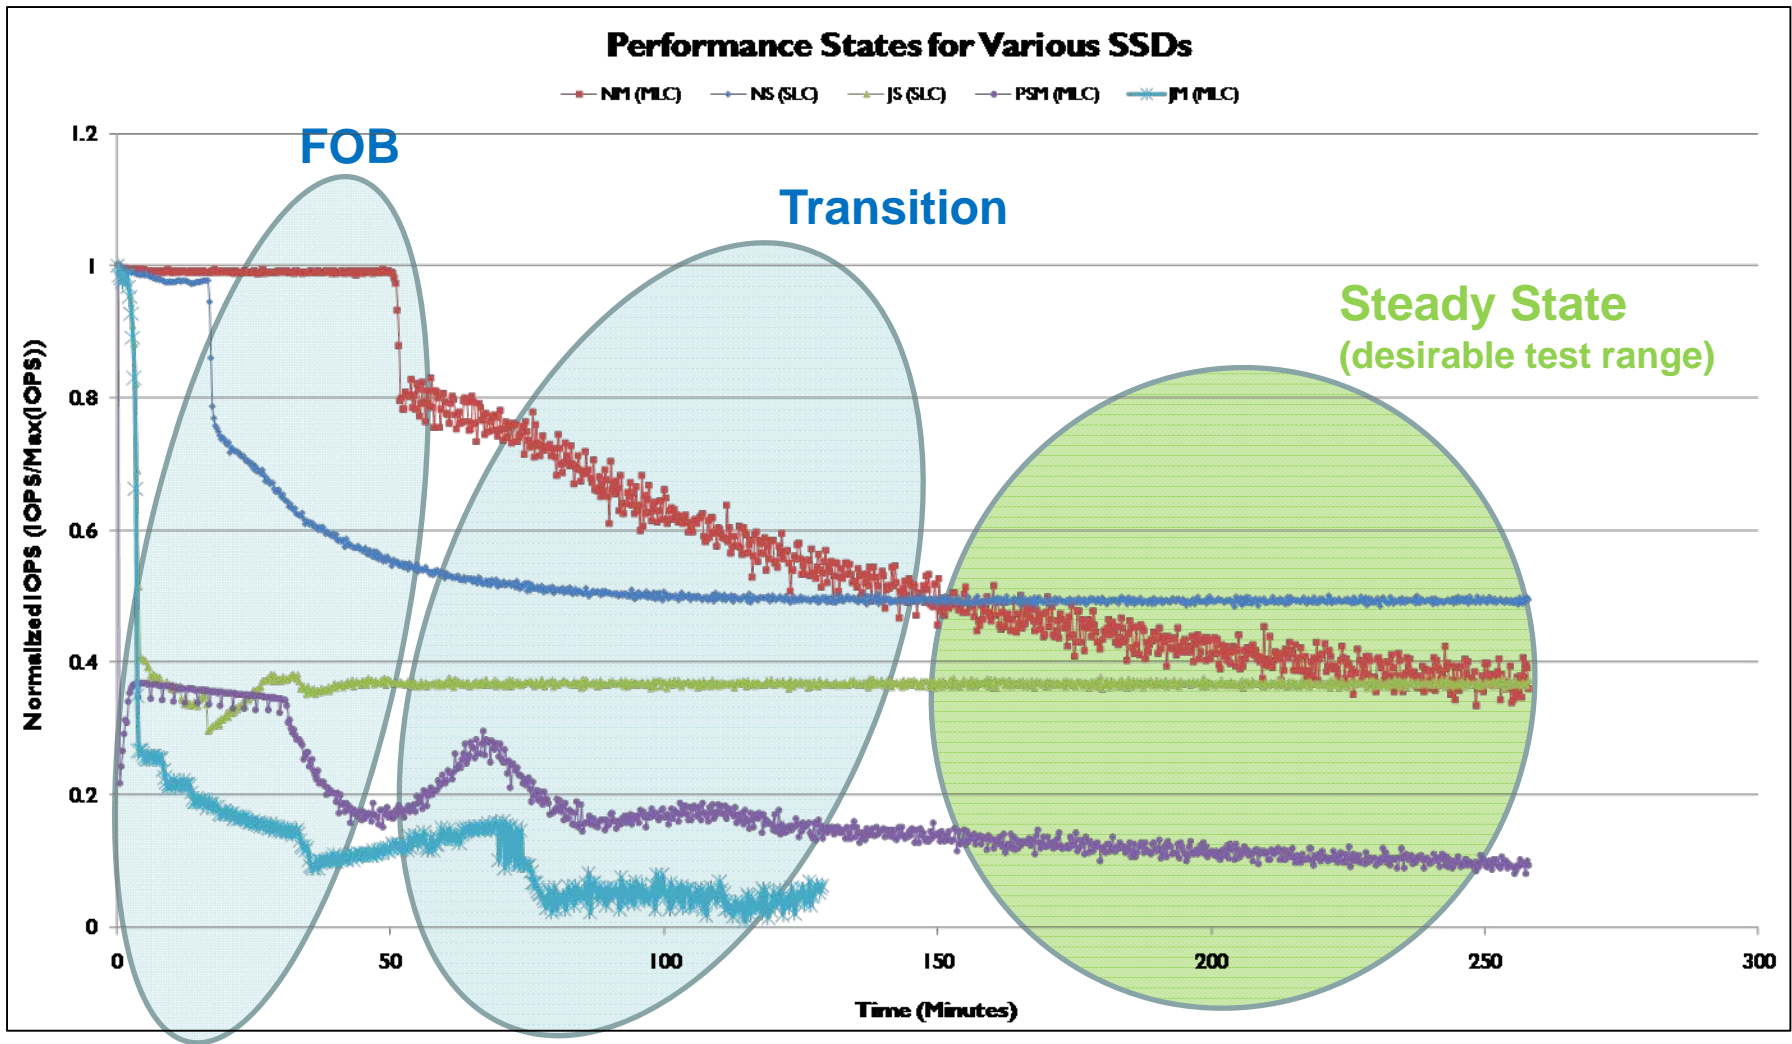
## 3D IOPS Surface Profile



SSS performance depends on

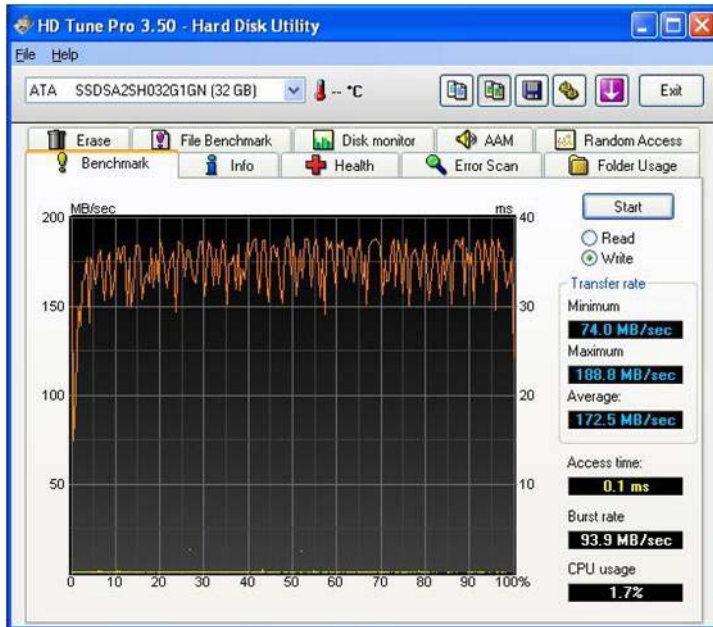
- Read/Write Mix
- Transfer Size
- Queue Depth (not shown)

# Performance States

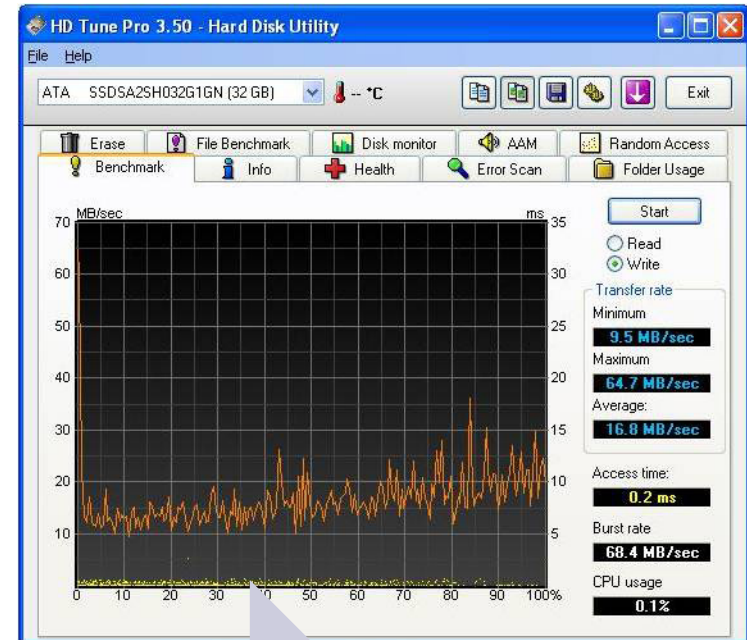




# Workload Dependency - 1



SSD Mfg A

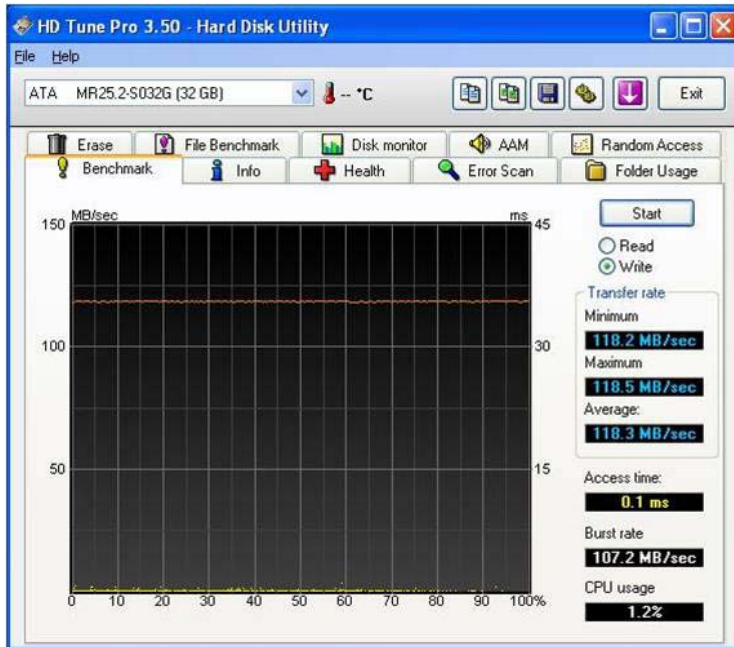


Step 1: HD Tune Pro  
Sequential Write Test

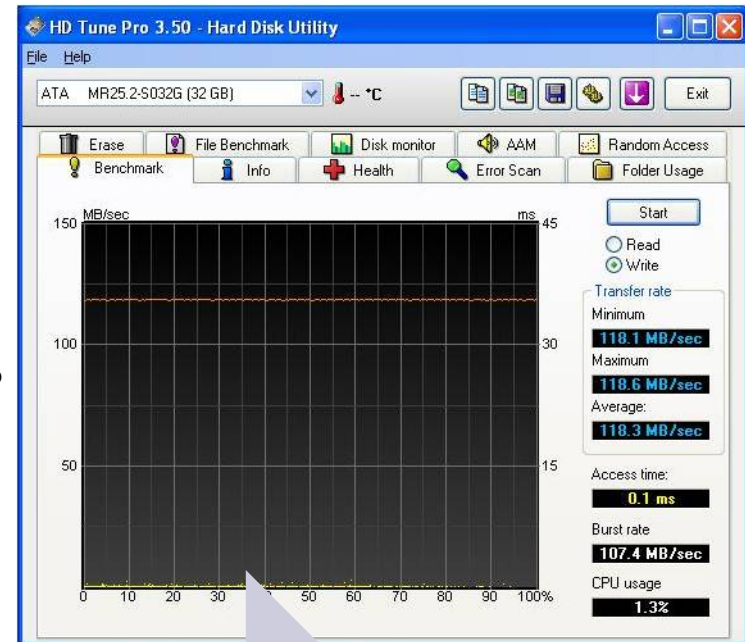
Step 2: IOMeter  
Random Write Test

Step 3: HD Tune Pro Sequential  
Write Test

# Workload dependency - 2



SSD Mfg B



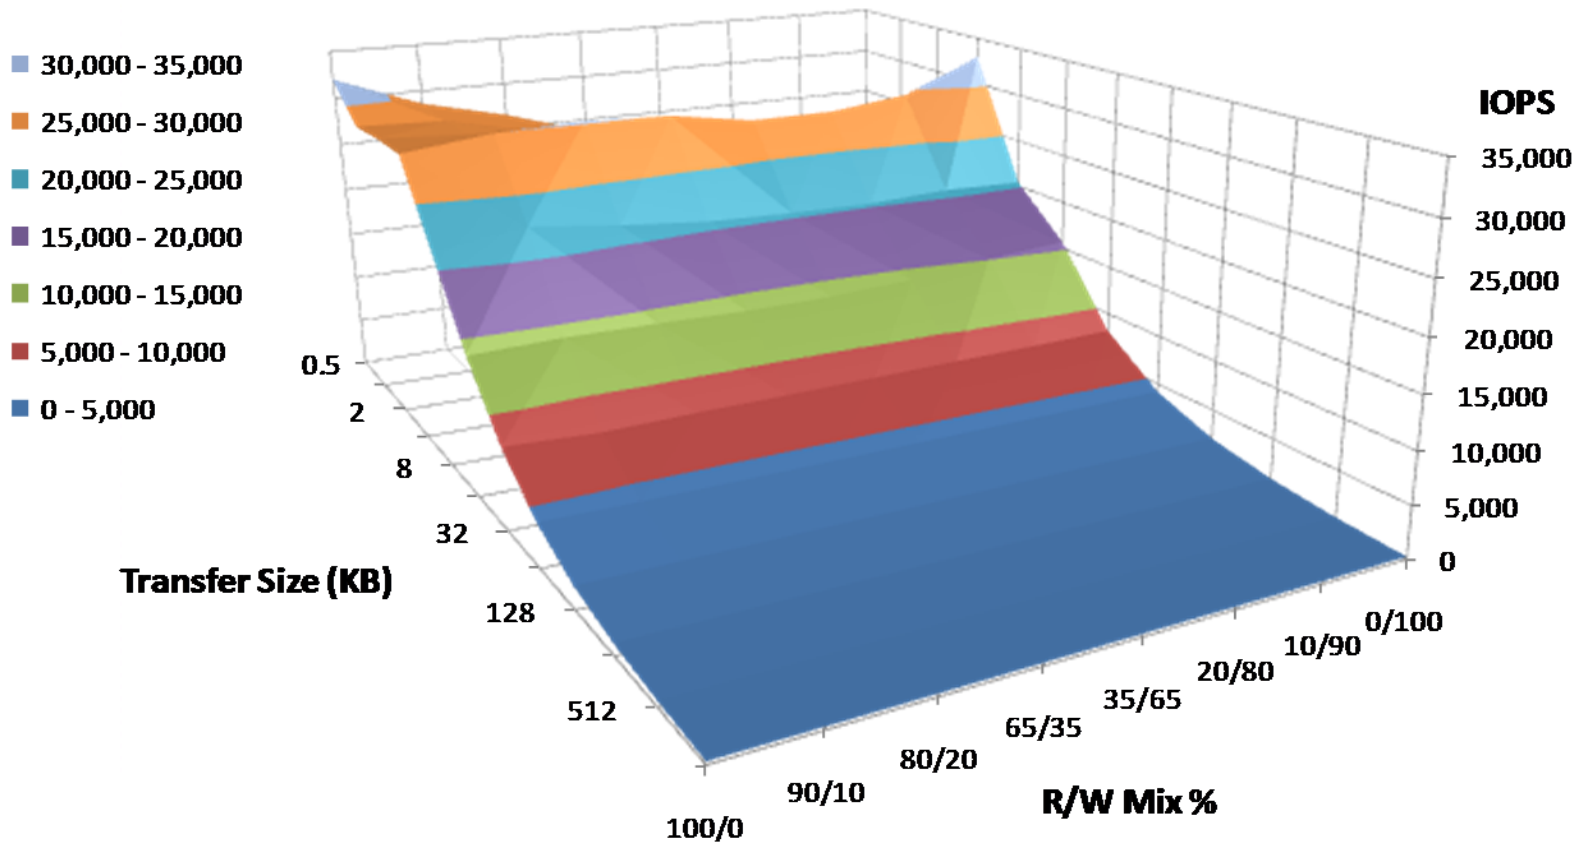
Step 1: HD Tune Pro  
Sequential Write Test

Step 2: IOMeter  
Random Write Test

Step 3: HD Tune Pro Sequential  
Write Test

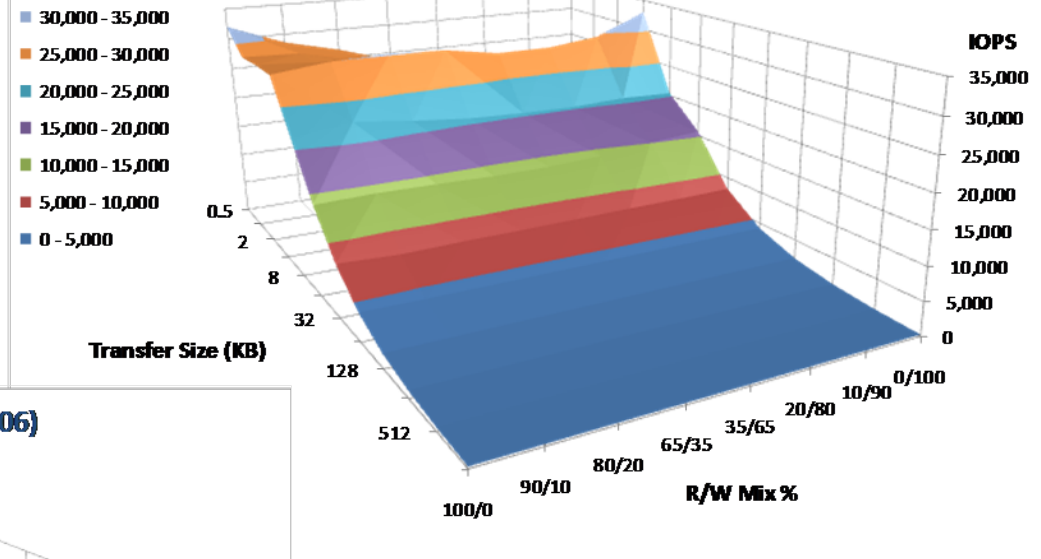
# Dependency on data content

## 3D IOPS Surface Profile (IOMETER 2008)

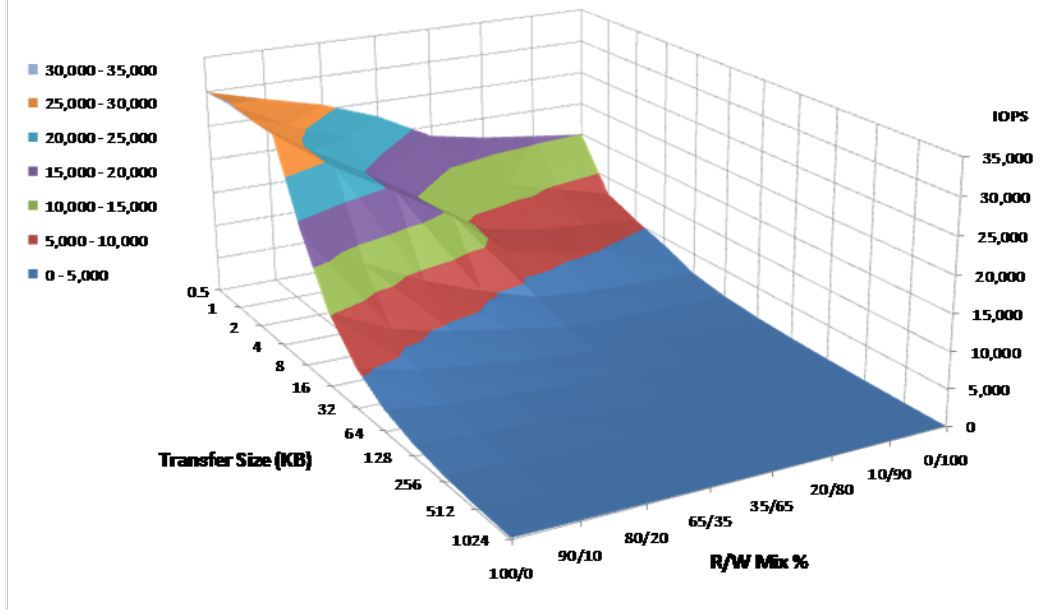


# Dependency on data content

### 3D IOPS Surface Profile (IOMETER 2008)



### 3D IOPS Surface Profile (IOMETER 2006)



# Benchmark Suites

	Test Suite	Commercial/ 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 response time and throughput score (0-250)		
IOMeter	System level test. Measures performance based on response time and throughput score (0-250)		✓
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		✓

**Type of preconditioning  
 and order of  
 benchmarks can  
 influence results**

- SNIA Technical Working Group (TWG) created in

ea

- Sp  
tc  
St



cedures  
of Solid



- ◆ **Repeatable** – Easy comparison between devices

- ◆ **Practical** – Complete with reasonable time and effort

- ◆ **Accessible** – open specification, 3<sup>rd</sup> party validation

## ➤ Setup and Methodology

- ◆ Purge
- ◆ Preconditioning
- ◆ Test Parameters

## ➤ Workloads (synthetic)

- ◆ Client – IOPS, Throughput, Latency
- ◆ Enterprise – IOPS, Throughput, Latency

## ➤ Reporting

- ◆ Show convergence to steady state
- ◆ Show performance results during steady state

## ➤ Performance Test Platform in development by SSSI Tech Dev Group (2H10)

- V0.x available for public review/comment
  - ◆ [www.snia.com](http://www.snia.com).....
- Your Feedback is crucial!
  - ◆ Do we test the right things? Do we need others?
  - ◆ Are the reports useful?
  - ◆ Are the procedures clear?
  - ◆ Does this truly give us apples-to-apples performance comparison?
- Performance Test Platform
  - ◆ Working in SSSI; please join us



# Performance Benchmark Steps

## 1. Prepare the Device

- ◆ Purge/Secure erase → put SSS back into “original” state

## 2. Precondition the Device

- ◆ Write data 2x capacity → bring device to known state

## 3. Steady State Testing

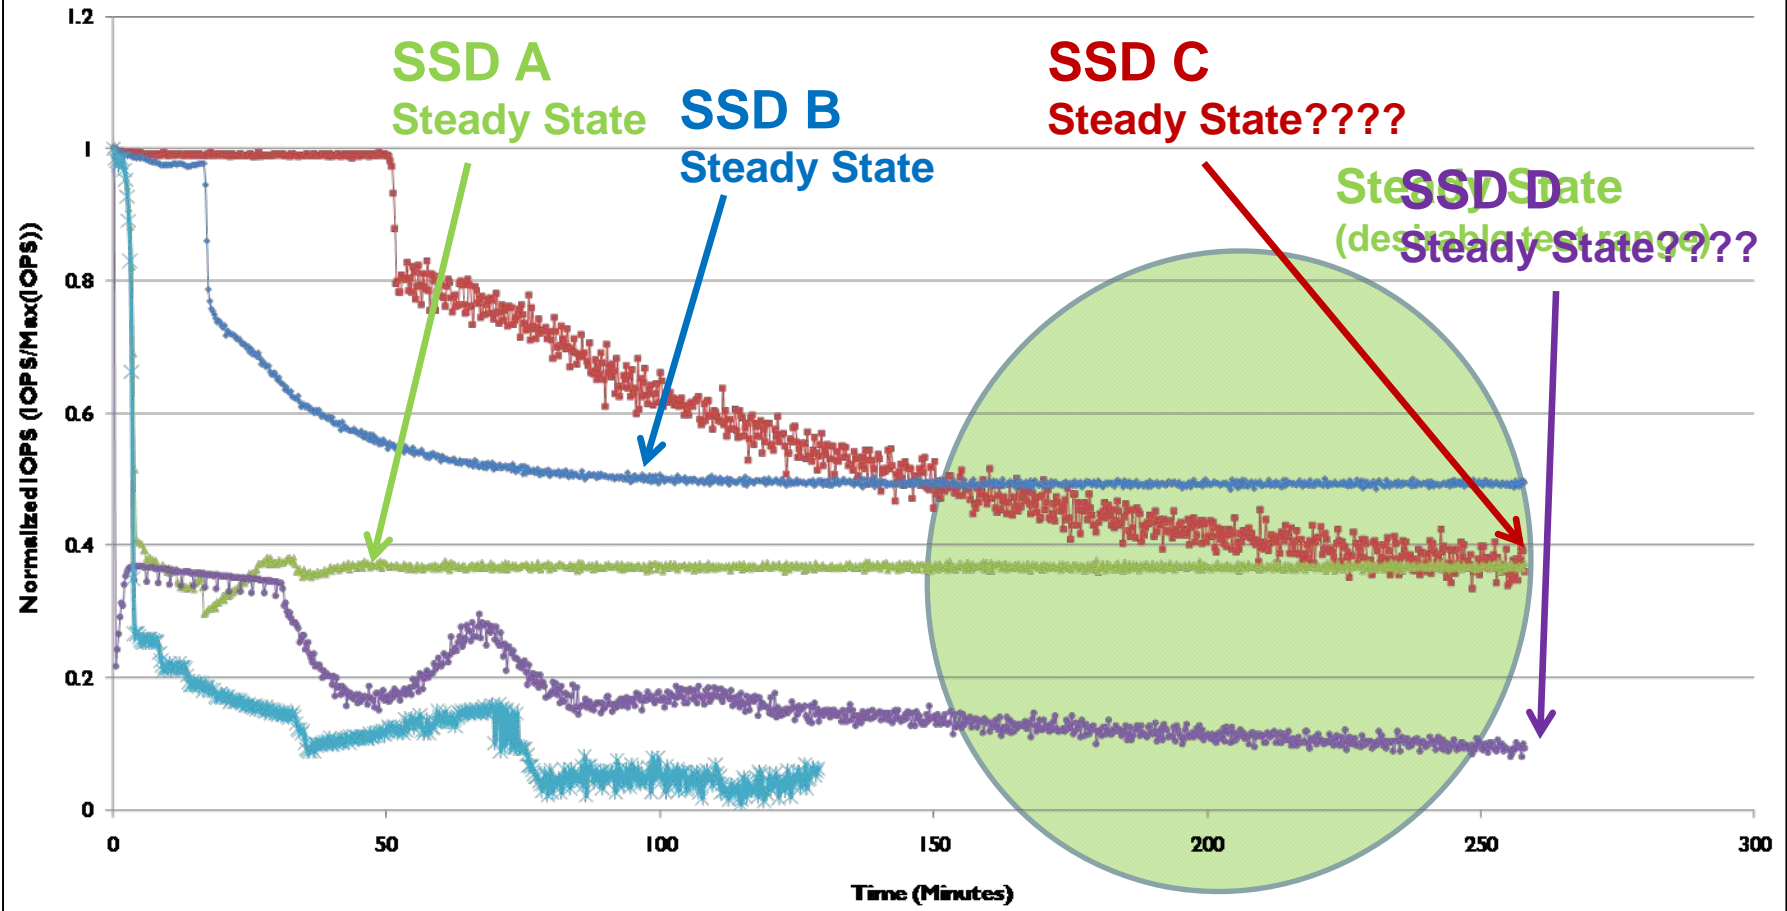
- ◆ Run Test Loop up until steady state is achieved
- ◆ Performance stays within  $\pm 10\%$  margin in last 5 test loops

## 4. Test Report

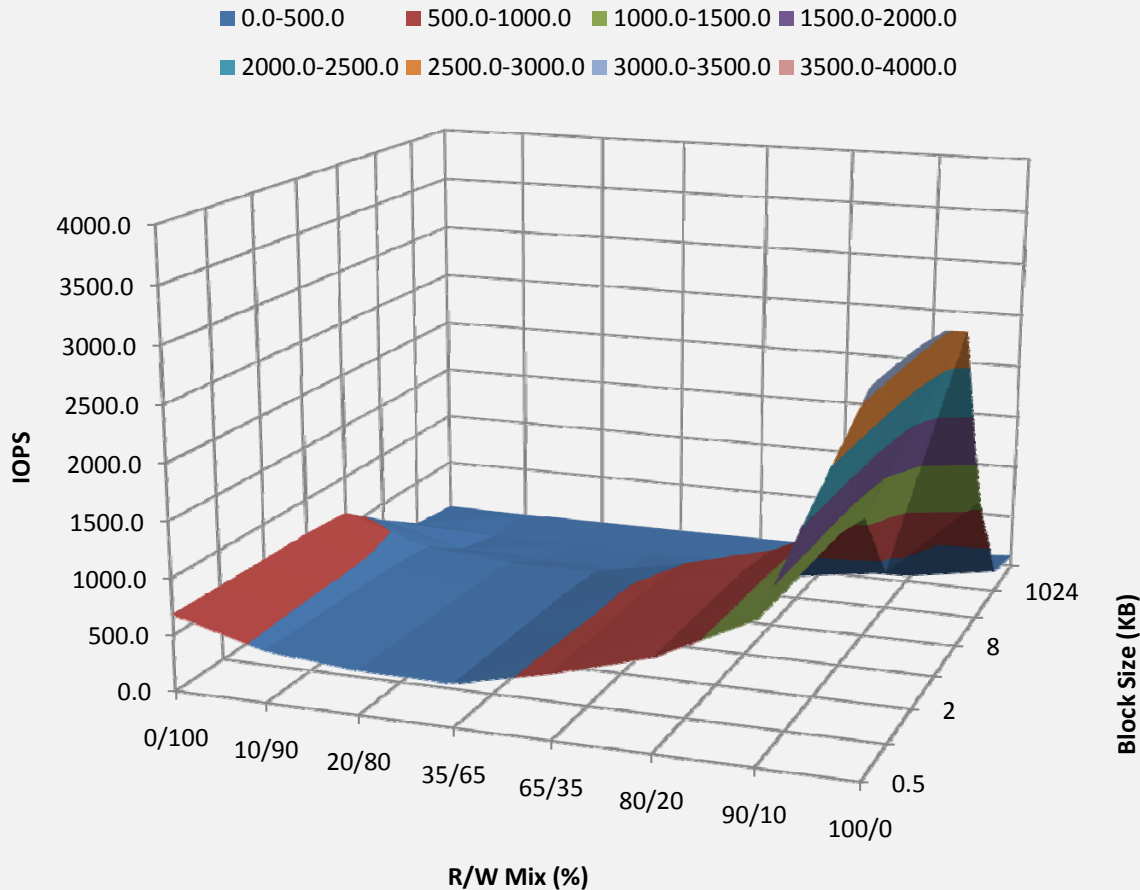
- ◆ Show convergence to steady state
- ◆ Show 3D test results

# Convergence to Steady State

Steady State reached if performance swing is within 10% margin



## 3D IOPS Surface Profile



Indicates steady state performance

- Various Block Size
- Various Read/Write mixes

## ➤ 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)

- 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/forums/sssi>
  - ◆ Send your feedback to [ssstwg@snia.org](mailto:ssstwg@snia.org)

➤ **GET INVOLVED!!**



- Please send any questions or comments on this presentation to SNIA: [tracksolidstate@snia.org](mailto: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**