Emerging SSD Performance Tests

Measuring the Performance of Solid State Storage Devices

Easen Ho

CTO - Calypso Systems, Inc.

Table of Contents

- Why is SSD Performance Testing Difficult?
- SNIA SSS Performance Test Specification
  - V1.0: Device-Centric Test
  - Basic Performance Tests
- Some Emerging Tests
  - Demand Intensity
  - Response Time Histogram
  - Cross Stimulus Recovery
  - Idle Recovery

Why is SSD Performance Testing Difficult?

Two Reasons

- *NAND-based SSDs are complicated*
- *Many ways to test performance*

NAND-based SSDs Are Complicated

Write Saturation: Continuous RND/4K W from FOB

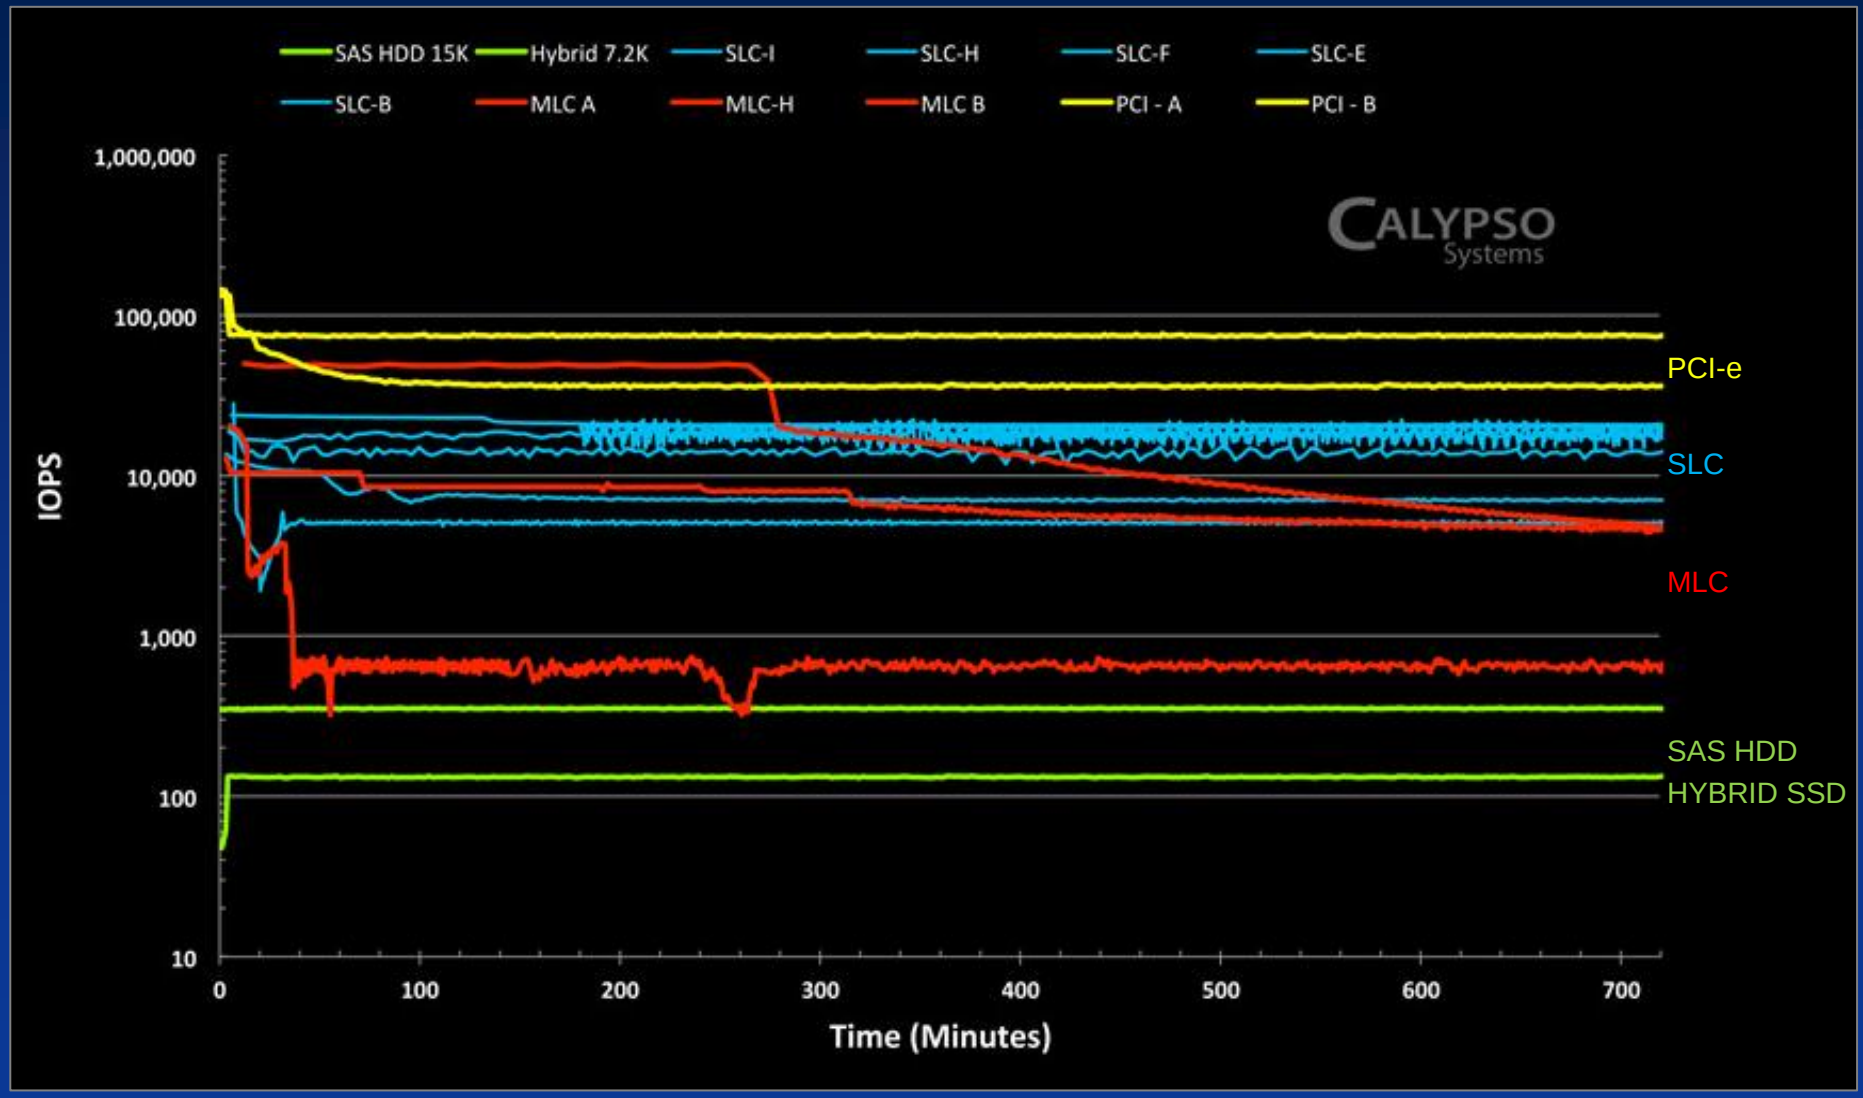

NAND-based SSDs Are Complicated

Items Impacting SSD Performance

Write History	PC Active Range	Test Active Range	Data Content
• What was previously written	• Where data was previously written • Trim effects	• Where data is currently written	• What is the nature of the data

Access Pattern	Demand Intensity	Throttling	?
• Manner in which data is being accessed	• How hard the application is driving the device	• How fast is data being written vs warranty	

Many Ways to Test Performance

Many benchmark tests being used

- 100% Real-World Test
- Trace-Based Capture & Playback
- Trace-Based Playback
  - AnandTech Storage Bench, Drivebench…
- Scripted Application Playback
  - Sysmark…
- Synthetic Stimulus *With or Without Constraints*
  - SNIA PTS, Calypso CTS, SPC-1/2, IOMeter, ATTO, CrystalDiskMark…

Performance Testing

Inherently Involves the Host System

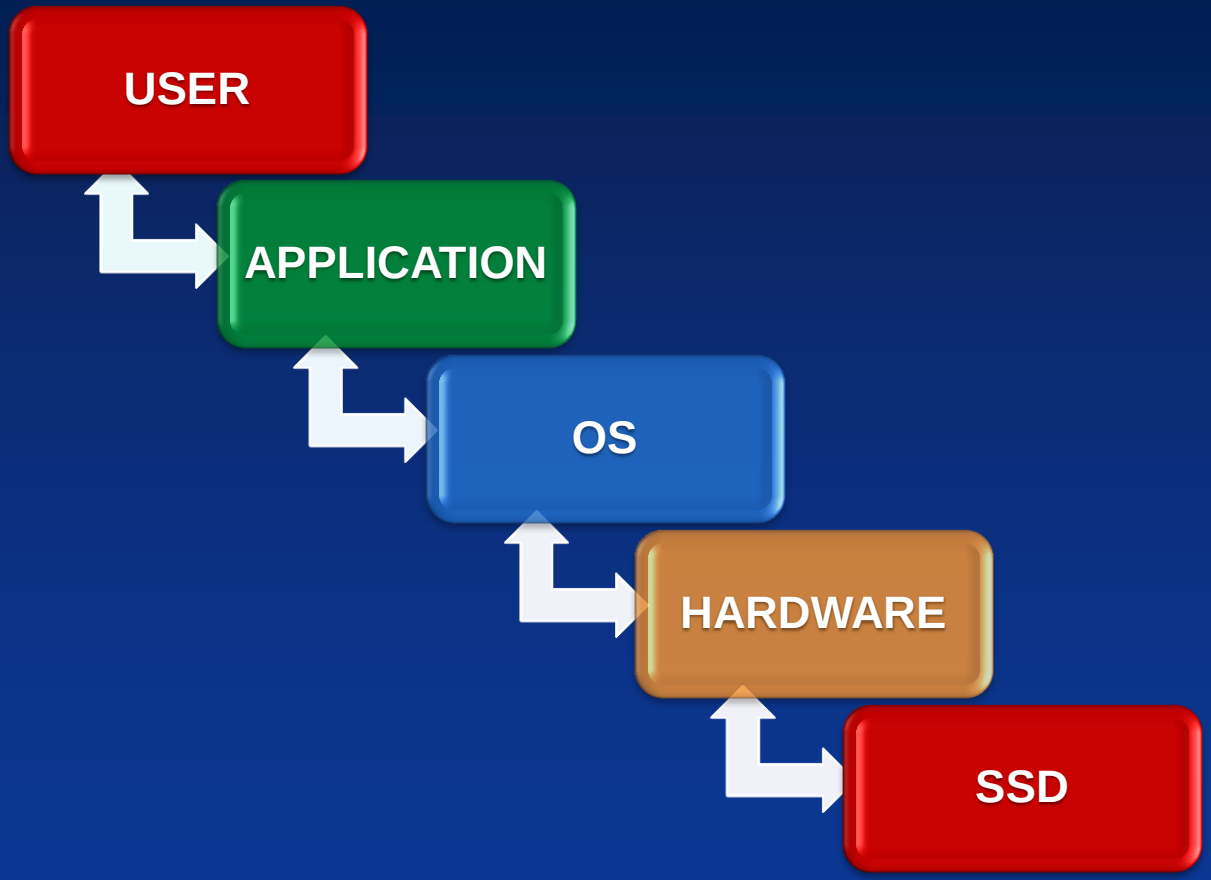

** Different approaches have different levels of involvement*

Trace-Based Capture and Playback

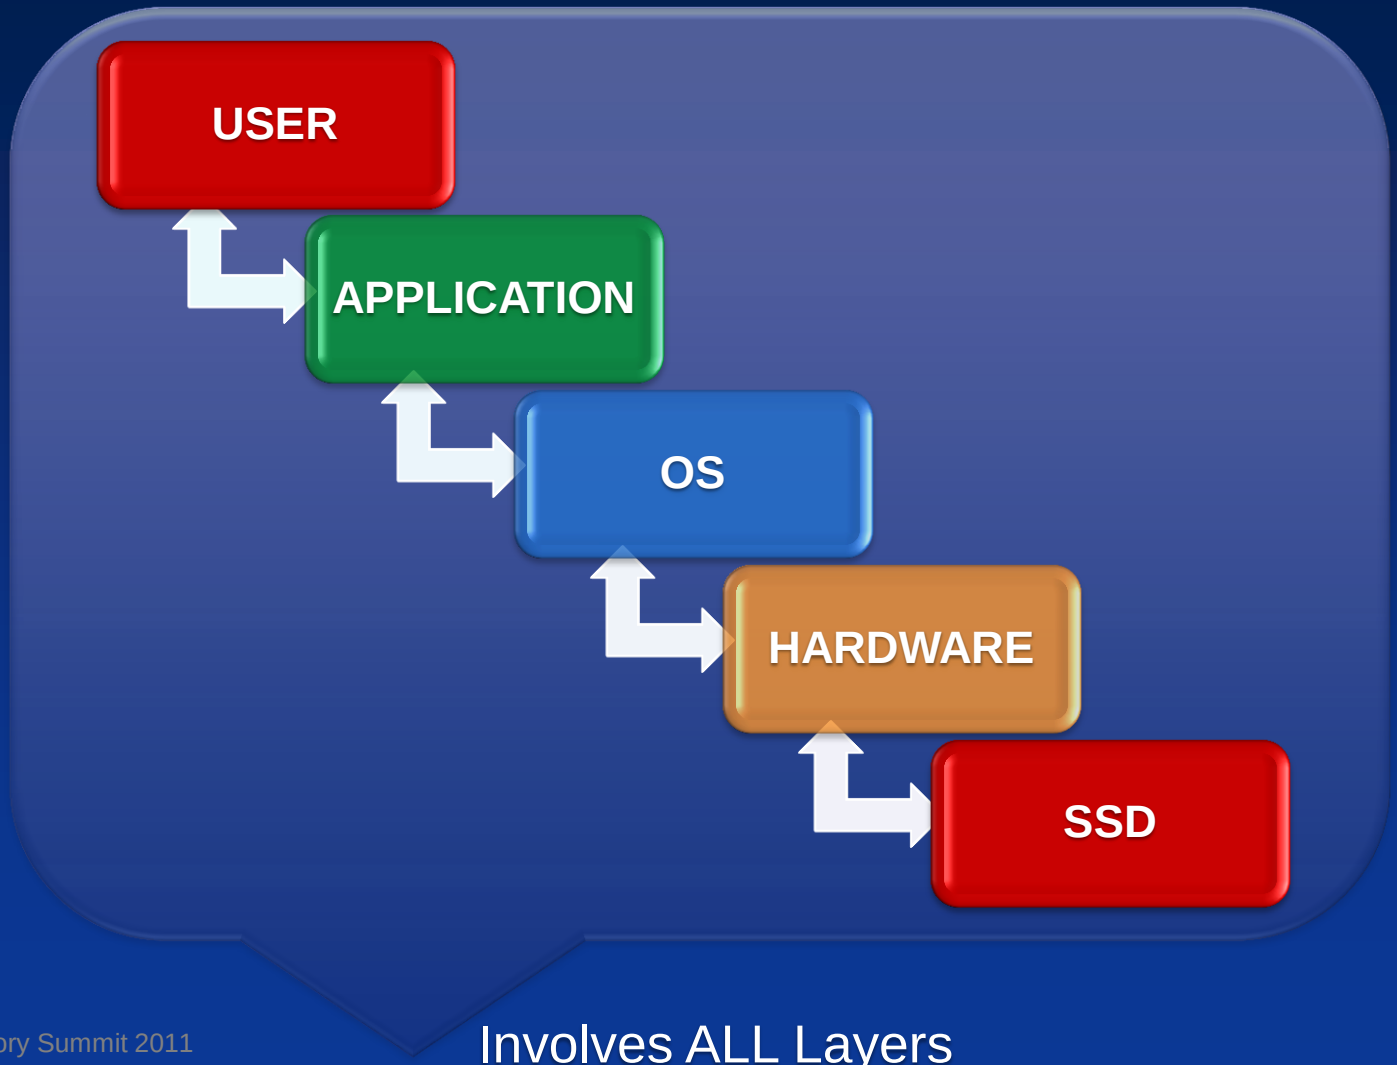

Involves ALL Layers

Traced-Based Playback
Scripted Application Playback
Synthetic Stimulus

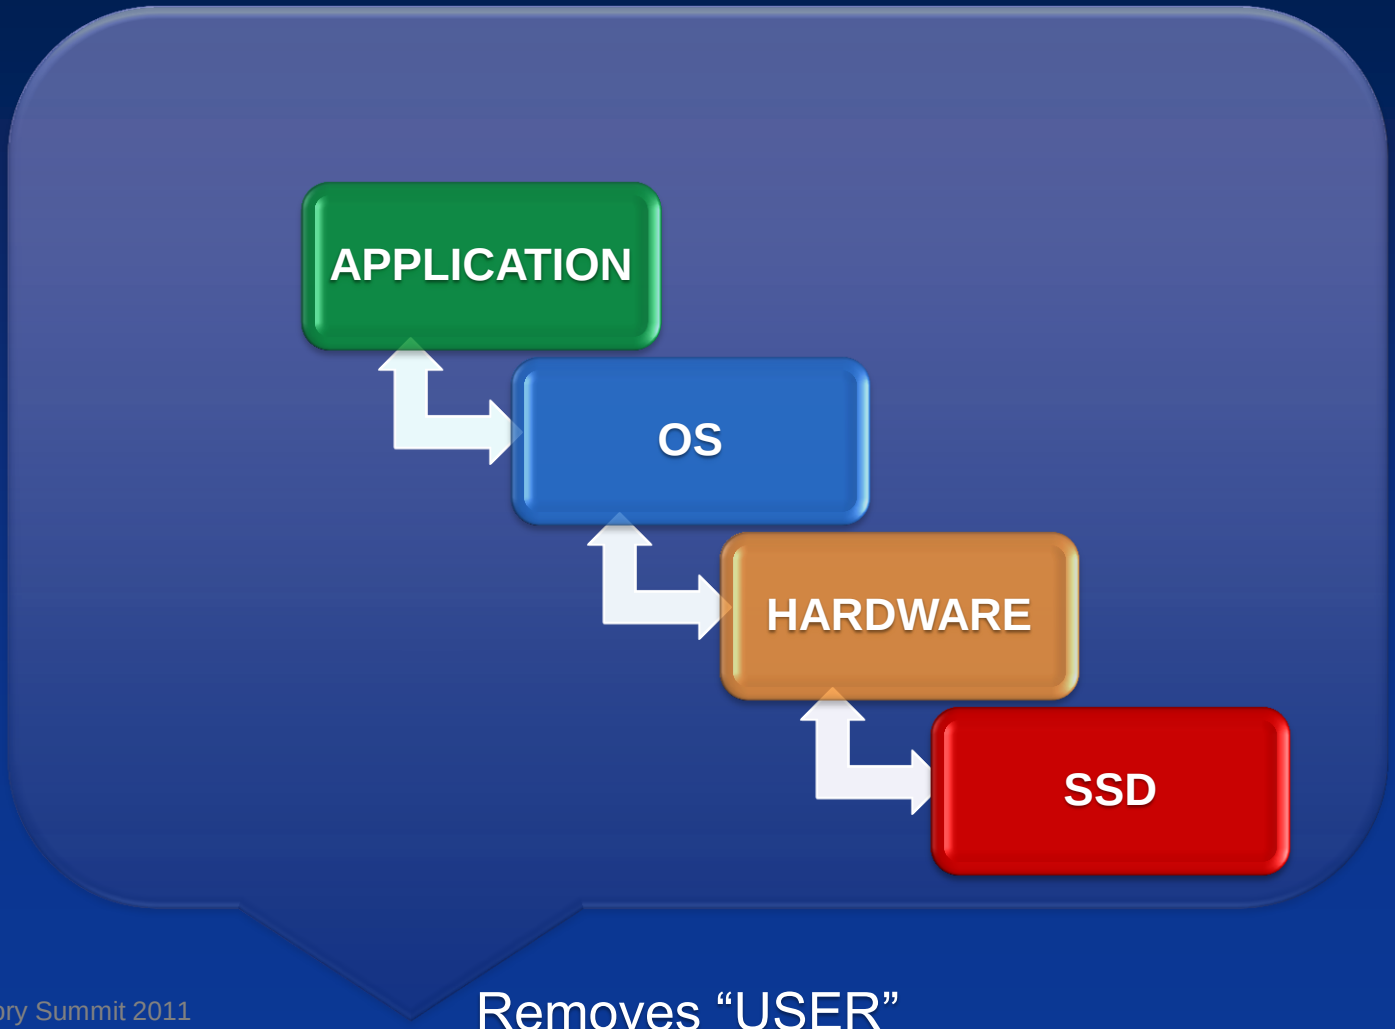

Removes “USER”

State of the Art in SSD Testing

- Trace Based Capture/Playback Solution
  - Ultimately, the Ideal Solution
  - SNIA TWG / member companies working on Trace-Based approach
  - However, many issues still exist

- Advantages of Device Based Synthetic Solution
  - Provides wealth of information on "DEVICE CHARACTERISTICS"
  - Easier to construct a testing environment that minimizes Effects of OS, Apps & Drivers on Device Performance

Synthetic Stimulus focused on "Device Testing" – e.g. SNIA PTS 1.0 / CTS 6.5

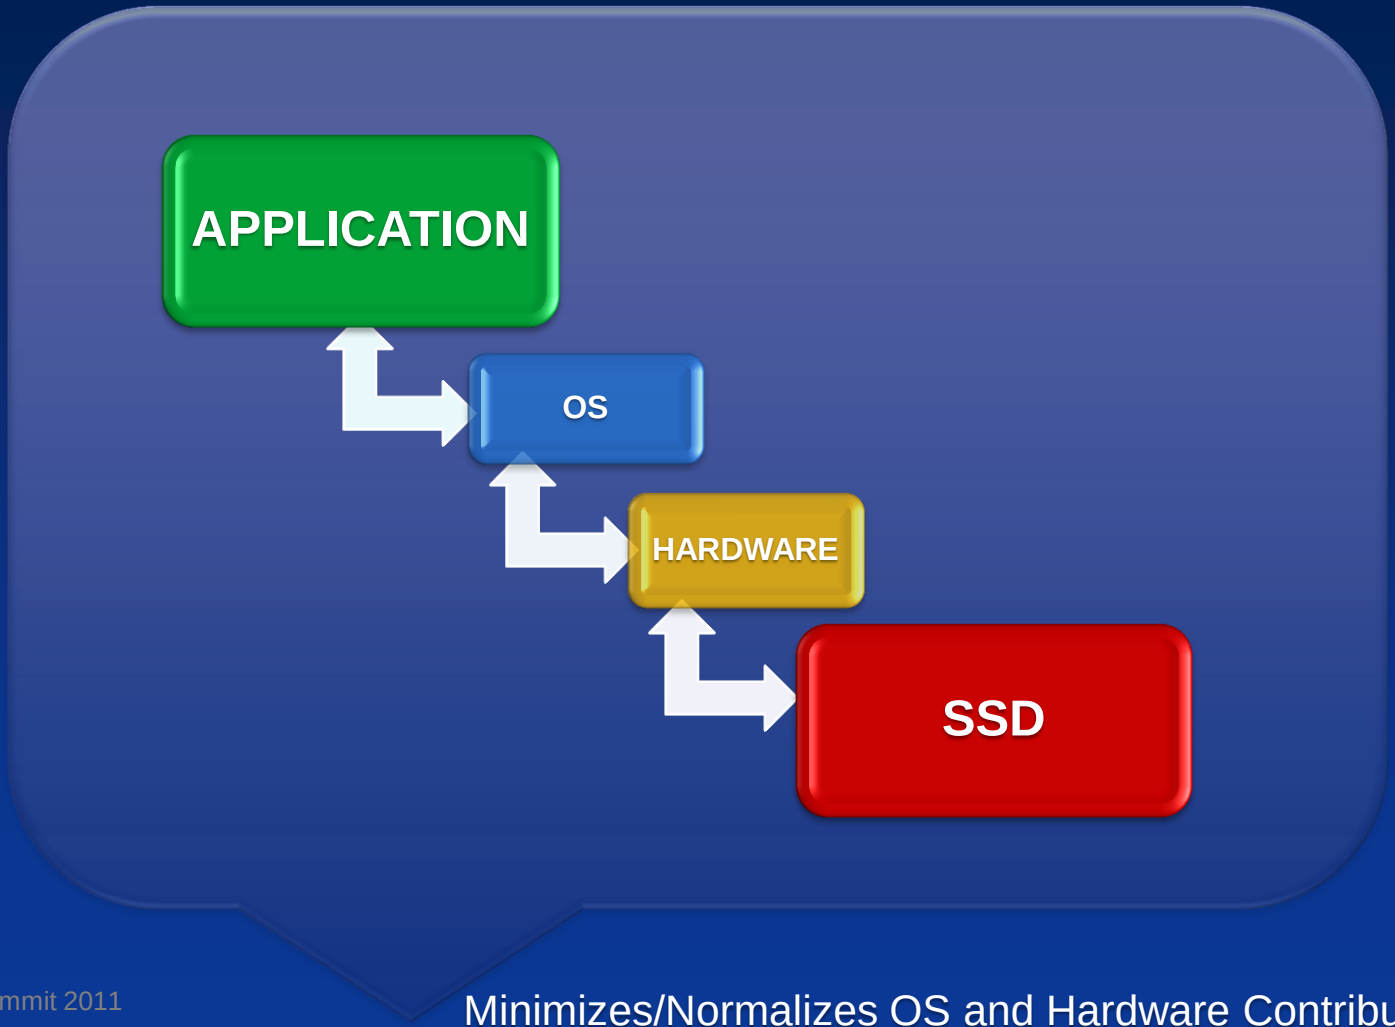

Minimizes/Normalizes OS and Hardware Contribution

Flash Memory SUMMIT

Why "Device Level" Performance Benchmarking?

Throughput?

Synthetic Benchmarks?

IOPS?

IOPS/W?

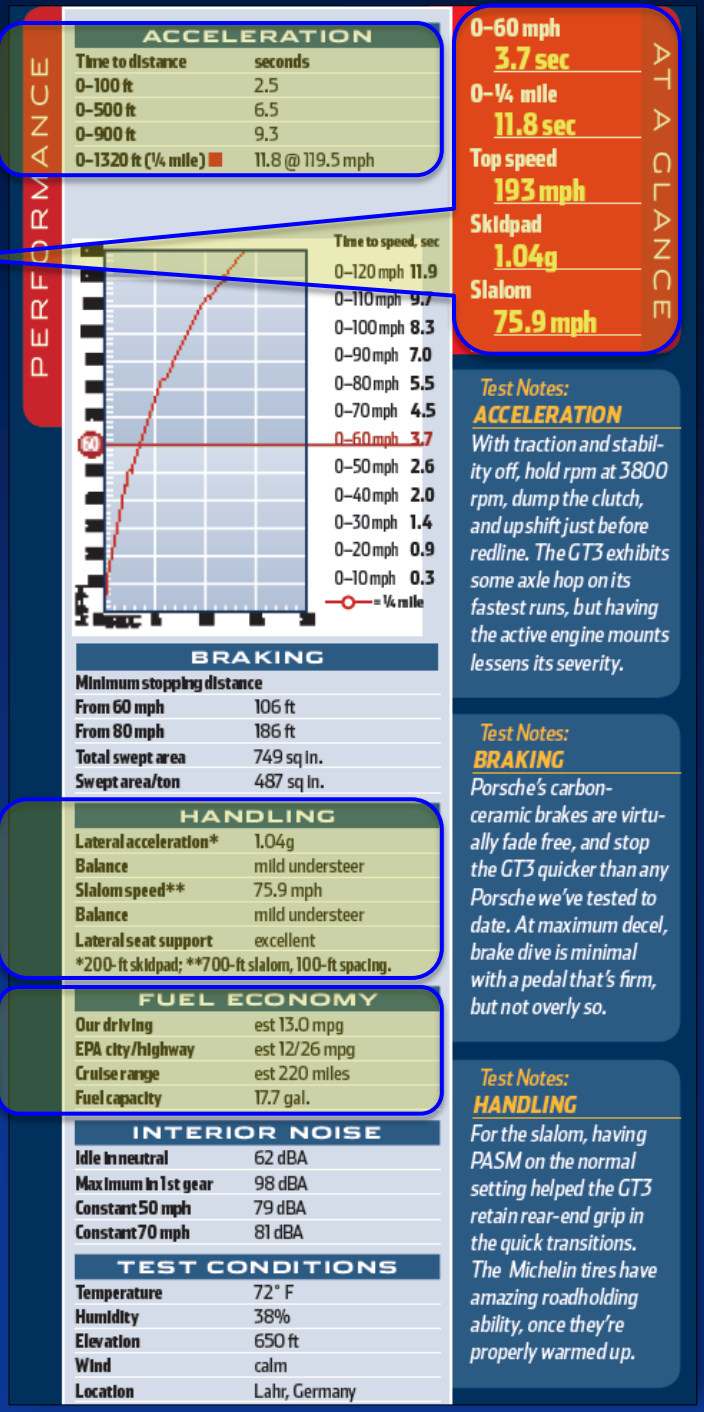

PERFORMANCE

ACCELERATION

Time to distance	seconds
0–100 ft	2.5
0–500 ft	6.5
0–900 ft	9.3
0–1320 ft (¼ mile)	11.8 @ 119.5 mph

Time to speed, sec

0–120 mph	11.9
0–110 mph	9.7
0–100 mph	8.3
0–90 mph	7.0
0–80 mph	5.5
0–70 mph	4.5
0–60 mph	3.7
0–50 mph	2.6
0–40 mph	2.0
0–30 mph	1.4
0–20 mph	0.9
0–10 mph	0.3

= ¼ mile

BRAKING

Minimum stopping distance	
From 60 mph	106 ft
From 80 mph	186 ft
Total swept area	749 sq in.
Swept area/ton	487 sq in.

HANDLING

Lateral acceleration*	1.04g
Balance	mild understeer
Slalom speed**	75.9 mph
Balance	mild understeer
Lateral seat support	excellent

*200-ft skidpad; **700-ft slalom, 100-ft spacing.

FUEL ECONOMY

Our driving	est 13.0 mpg
EPA city/highway	est 12/26 mpg
Cruise range	est 220 miles
Fuel capacity	17.7 gal.

INTERIOR NOISE

Idle in neutral	62 dBA
Maximum in 1st gear	98 dBA
Constant 50 mph	79 dBA
Constant 70 mph	81 dBA

TEST CONDITIONS

Temperature	72° F
Humidity	38%
Elevation	650 ft
Wind	calm
Location	Lahr, Germany

AT A GLANCE

0–60 mph	3.7 sec
0–¼ mile	11.8 sec
Top speed	193 mph
Skidpad	1.04g
Slalom	75.9 mph

Test Notes: **ACCELERATION**

With traction and stability off, hold rpm at 3800 rpm, dump the clutch, and upshift just before redline. The GT3 exhibits some axle hop on its fastest runs, but having the active engine mounts lessens its severity.

Test Notes: **BRAKING**

Porsche's carbon-ceramic brakes are virtually fade free, and stop the GT3 quicker than any Porsche we've tested to date. At maximum decel, brake dive is minimal with a pedal that's firm, but not overly so.

Test Notes: **HANDLING**

For the slalom, having PASM on the normal setting helped the GT3 retain rear-end grip in the quick transitions. The Michelin tires have amazing roadholding ability, once they're properly warmed up.

SNIA's PTS: Purpose

"…This Specification defines a set of
device level tests
and methodologies to enable
comparative testing
of Solid State Storage (SSS) devices."

- Performance Test Specification v1.0 – Section 1.1

SNIA PTS 1.0

- Enterprise 1.0
  - Write Saturation, IOPS, Throughput, Latency
- Client 1.0
  - Limited PC and Test AR
  - IOPS, Throughput, Latency

"Performance" *beyond IOPS, TP….*

There are other more subtle, but equally important performance metrics.

EXAMPLE:

DEMAND INTENSITY

How hard can I drive the SSD while maintaining reasonable response times?

RESPONSE TIME STATISTICS

How well-behaved is the response time statistics?

CROSS STIMULUS RECOVERY

How does the SSD respond to change between sustained stimulus?

IDLE TIME RESPONSE

How does the SSD do during IO idle time?

IOPS Demand Intensity

RND/4K

Purpose		Determines how the DUT responds to increasing demand from Host

Test Setup	Preconditioning	SNIA E-PTS IOPS to Steady State
	Test	Vary Total Outstanding IO (TOIO) Measure RND4K IOPS Average Response Time (ART) for one minutes each IO

RND/4K Read Demand Intensity
(E-SLC)

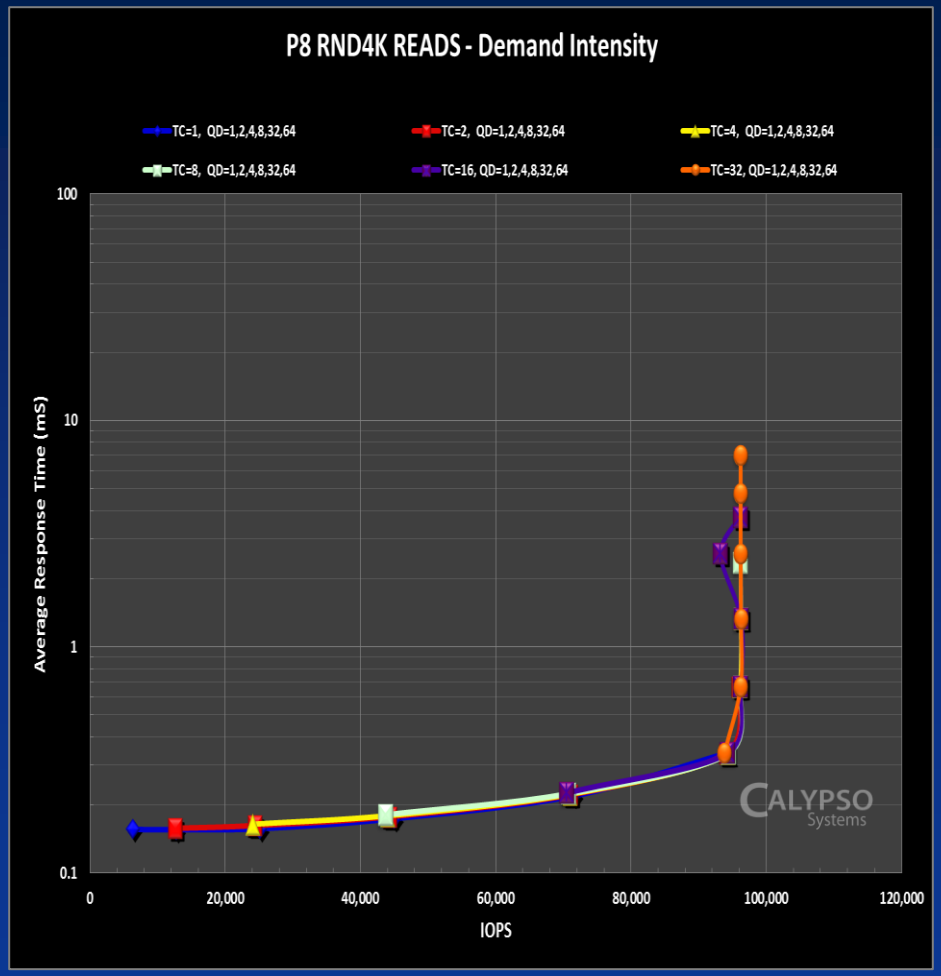

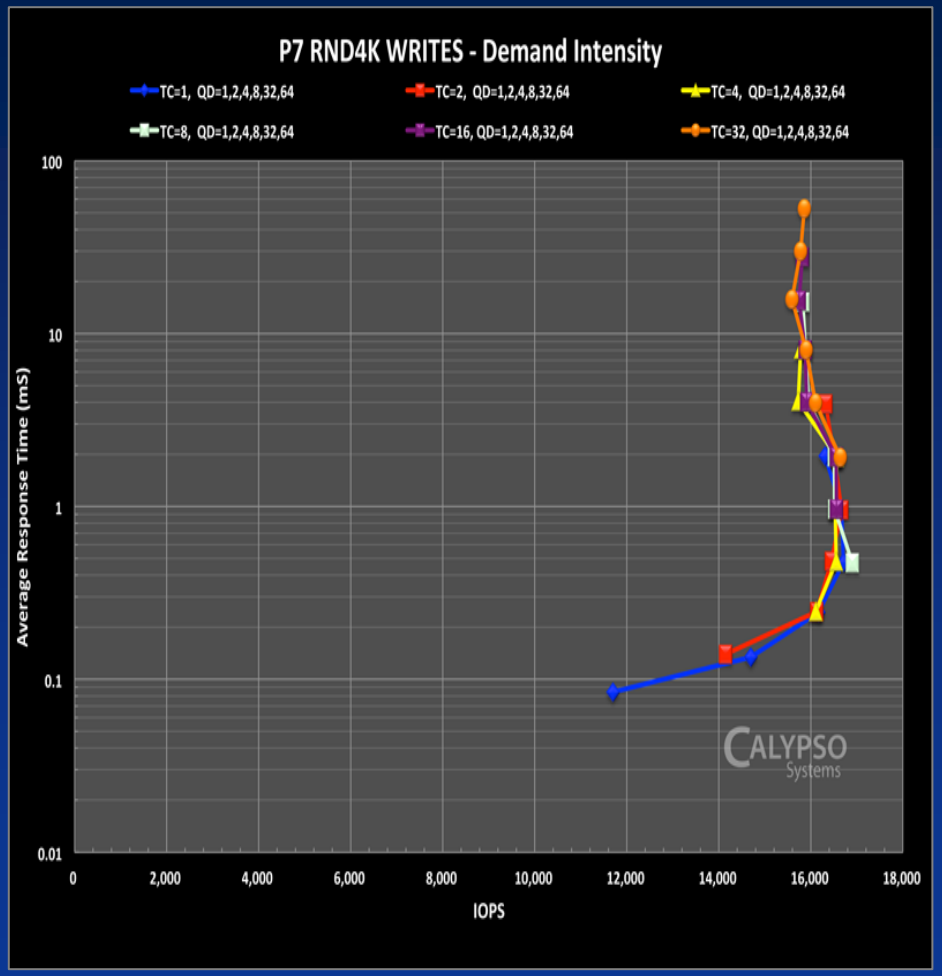

RND/4K Read Demand Intensity
(C-MLC)

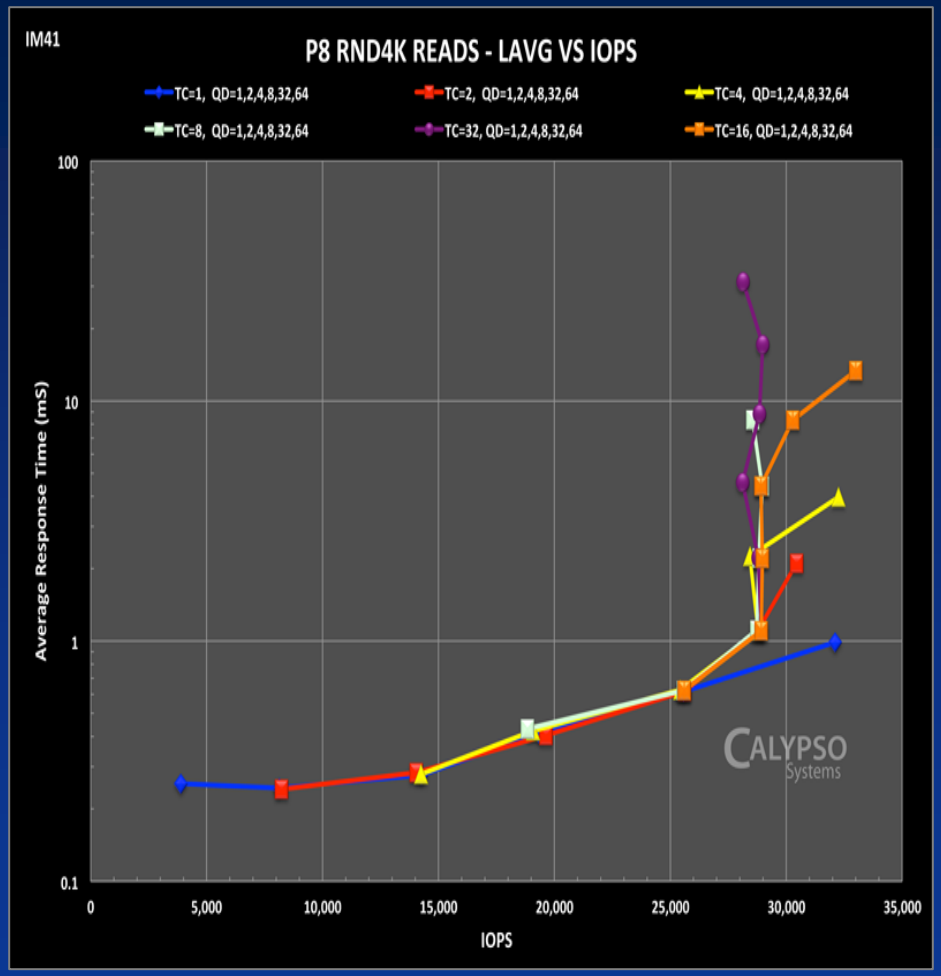

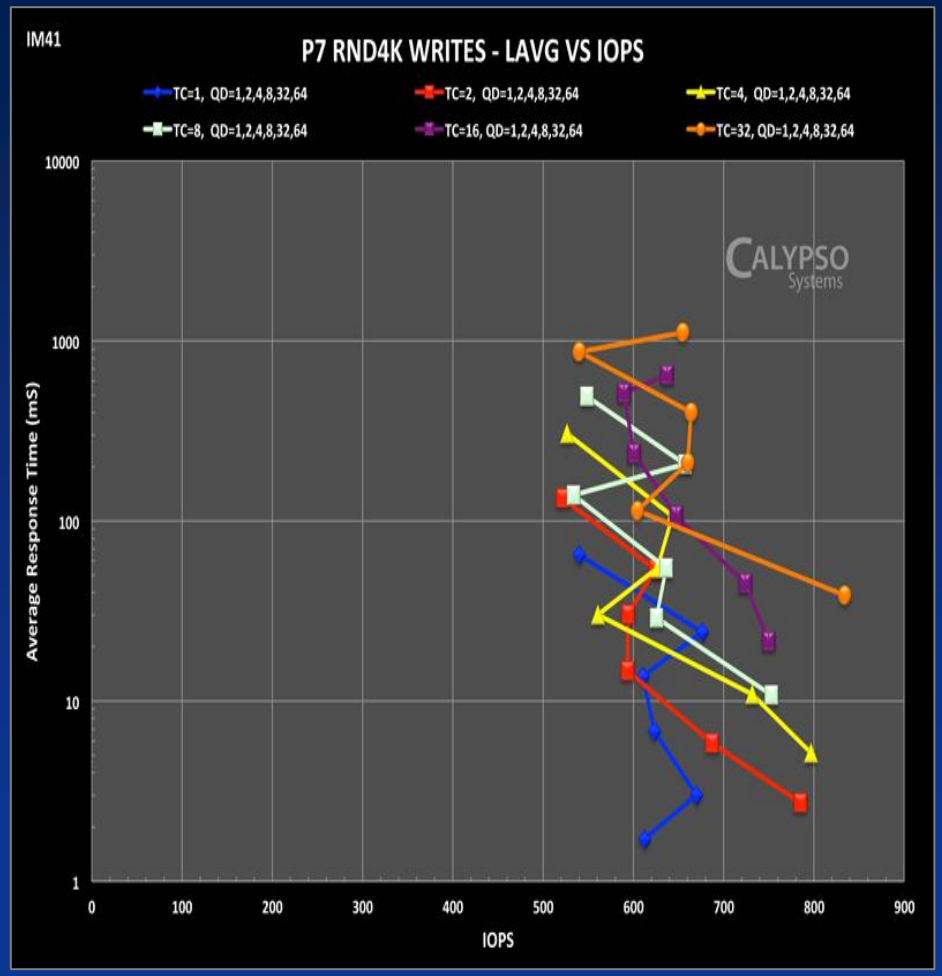

Response Time Histogram

RND/4K

Purpose		Investigate the DUT's response time statistic
Test Setup	Preconditioning	Purge, RND/4K Writes
	Test	Capture very IO's completion time within a specified time duration and sort into specified bins

Response Time Histogram

(10 Min at RND/4K SS)

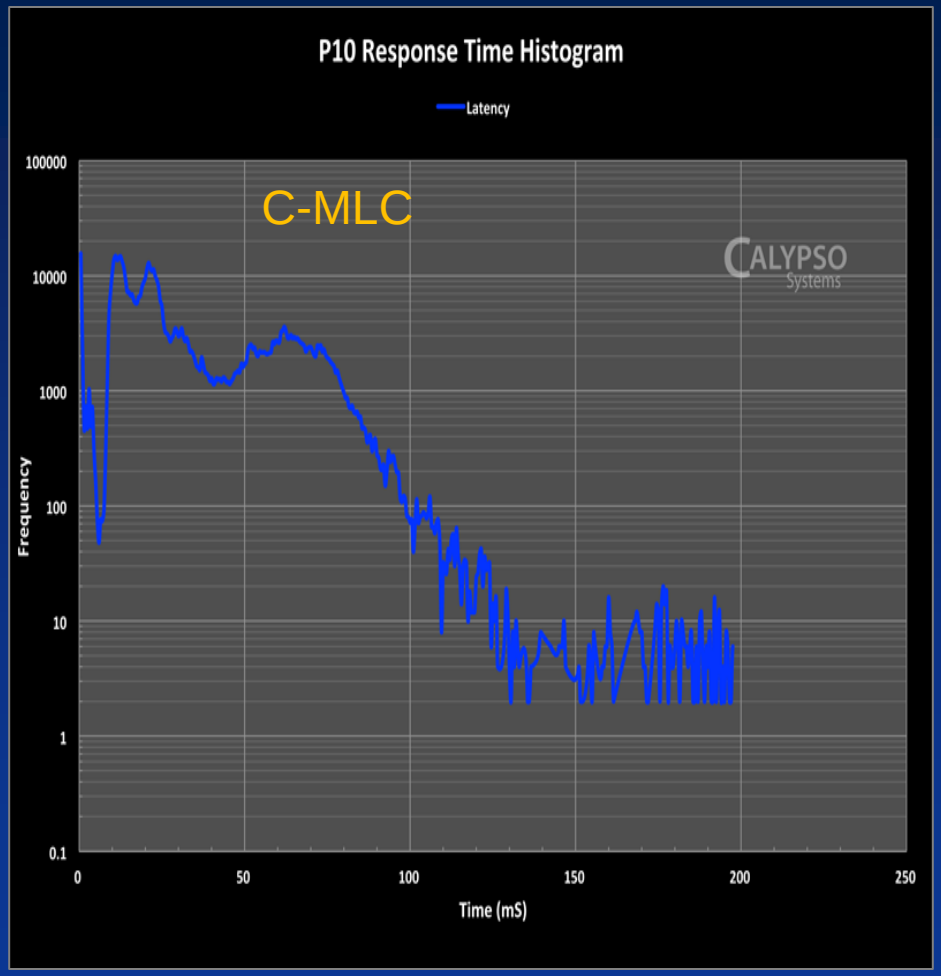

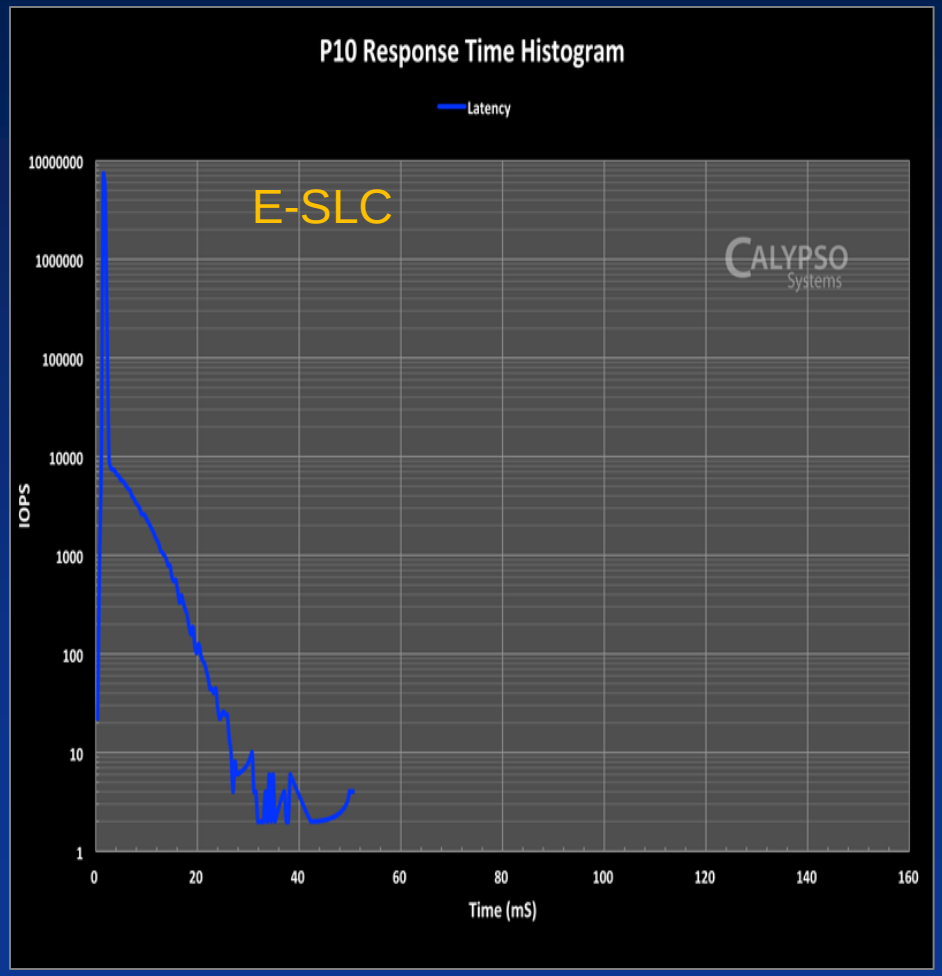

Cross Stimulus Response

SEQ/128K→RND/4K→SEQ/128K

Purpose		Determines how the DUT responds to sudden switch in stimulus to and from sustained large block sequential writes to small block random writes

Test Setup	Preconditioning	PURGE
	Test	Write SEQ/128K Write RND/4K Write SEQ/128K Capture IOPS and ART

Cross Stimulus Recovery:

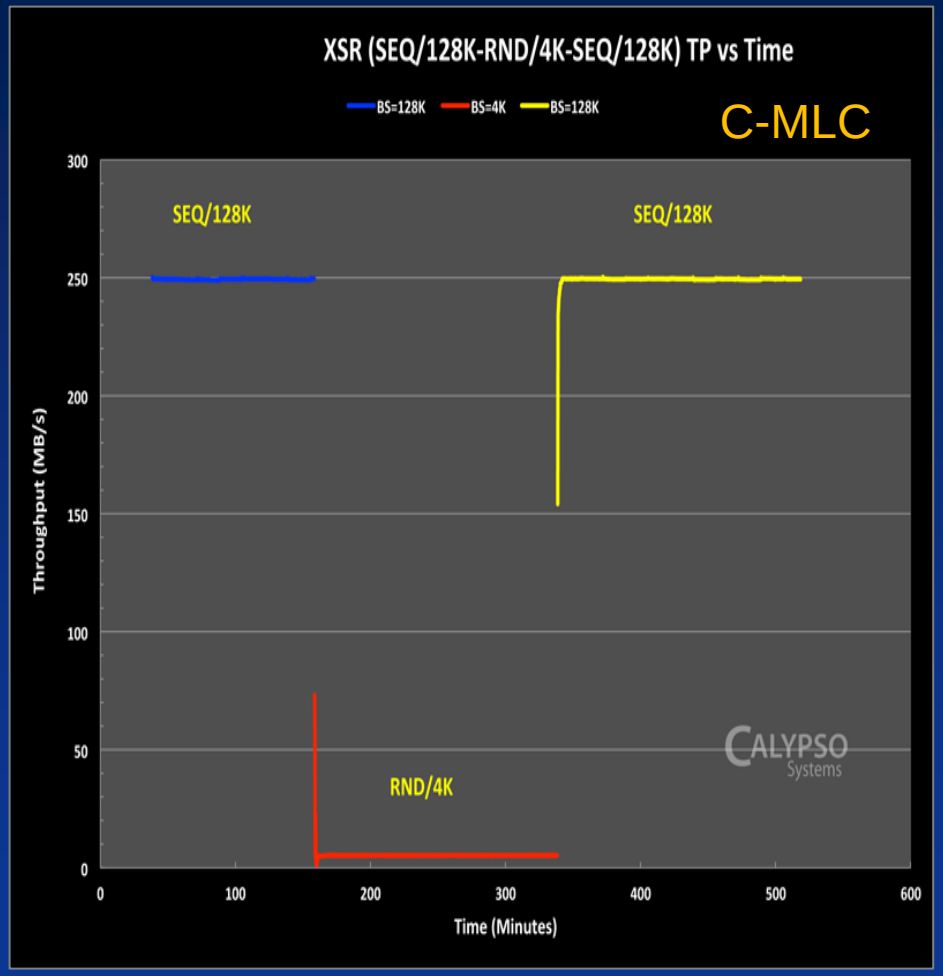

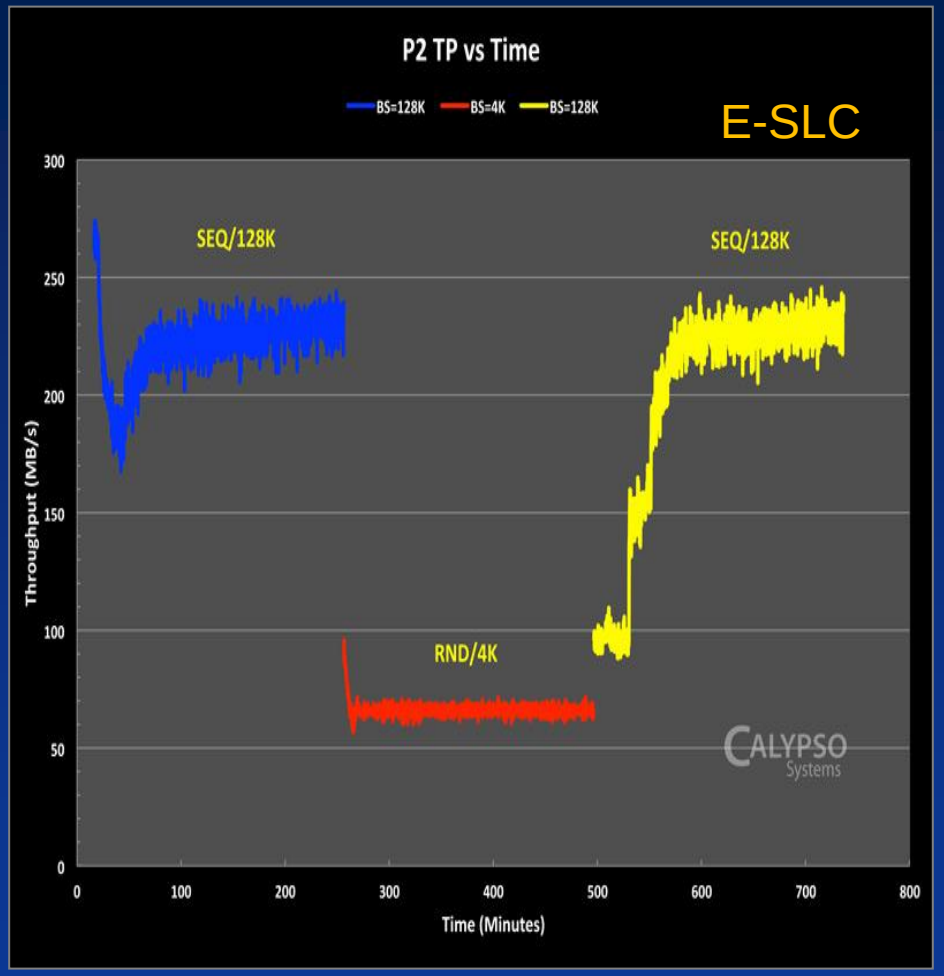

Cross Stimulus Recovery:

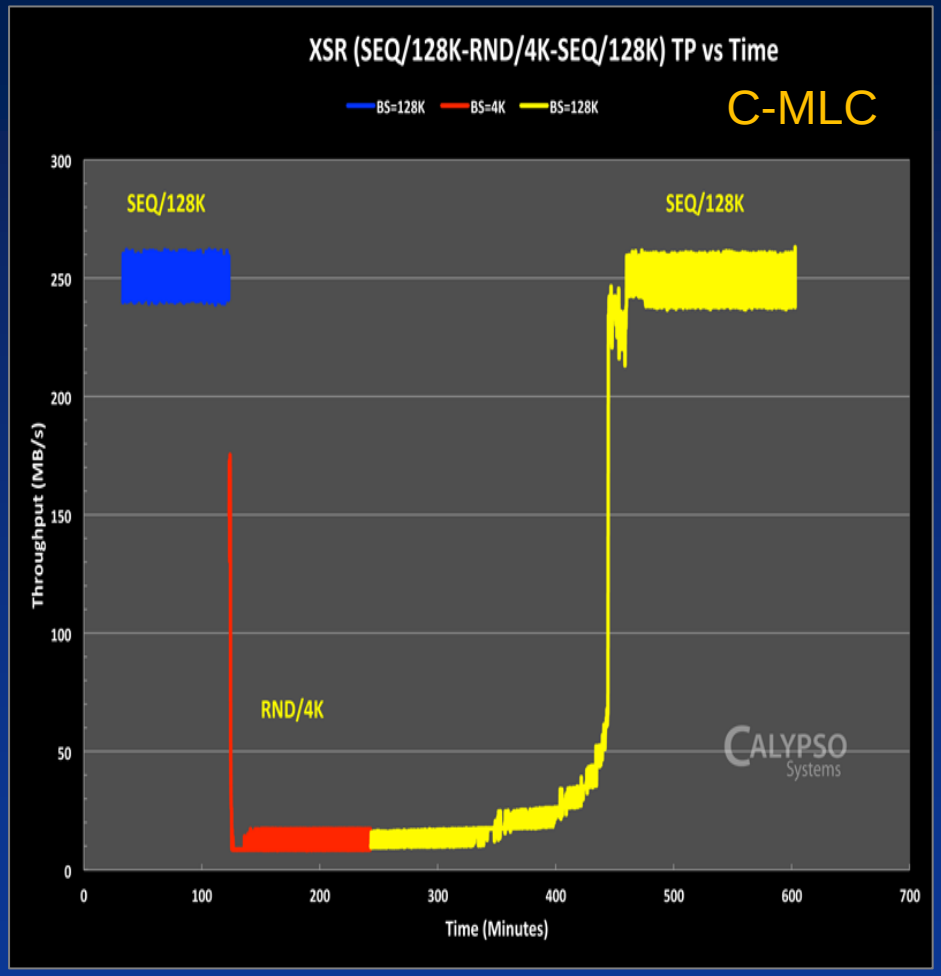

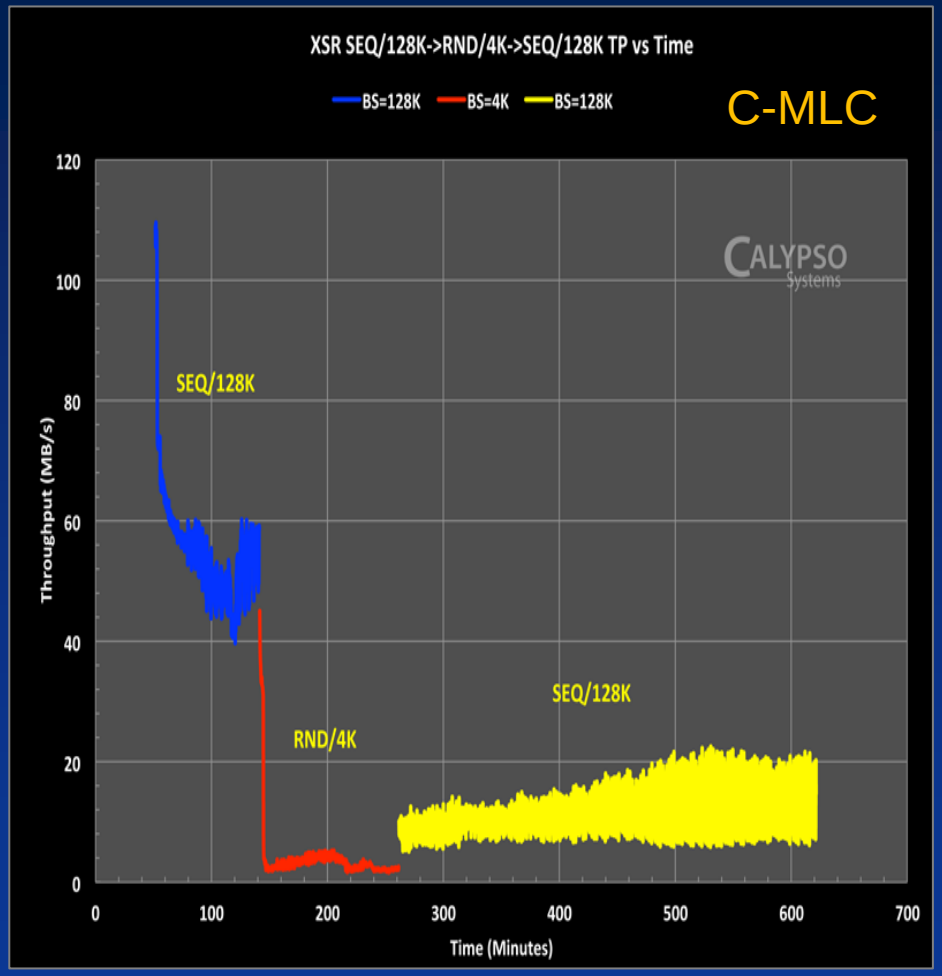

Idle Recovery

RND/4K Write Sat. with Wait States

Purpose		Determines how the DUT responds to Host IO idle period amidst continuous IO
Test Setup	Preconditioning	SNIA E-PTS IOPS till Steady State
	Test	Cease IOs for a various amount of time between segments of continuous RND/4K writes: Wait State 0: Write/Idle=5s/0s Wait State 1: Write/Idle=5s/5s Wait State 5: Write/Idle=5s/25s

Write Saturation with Wait State

RND/4K

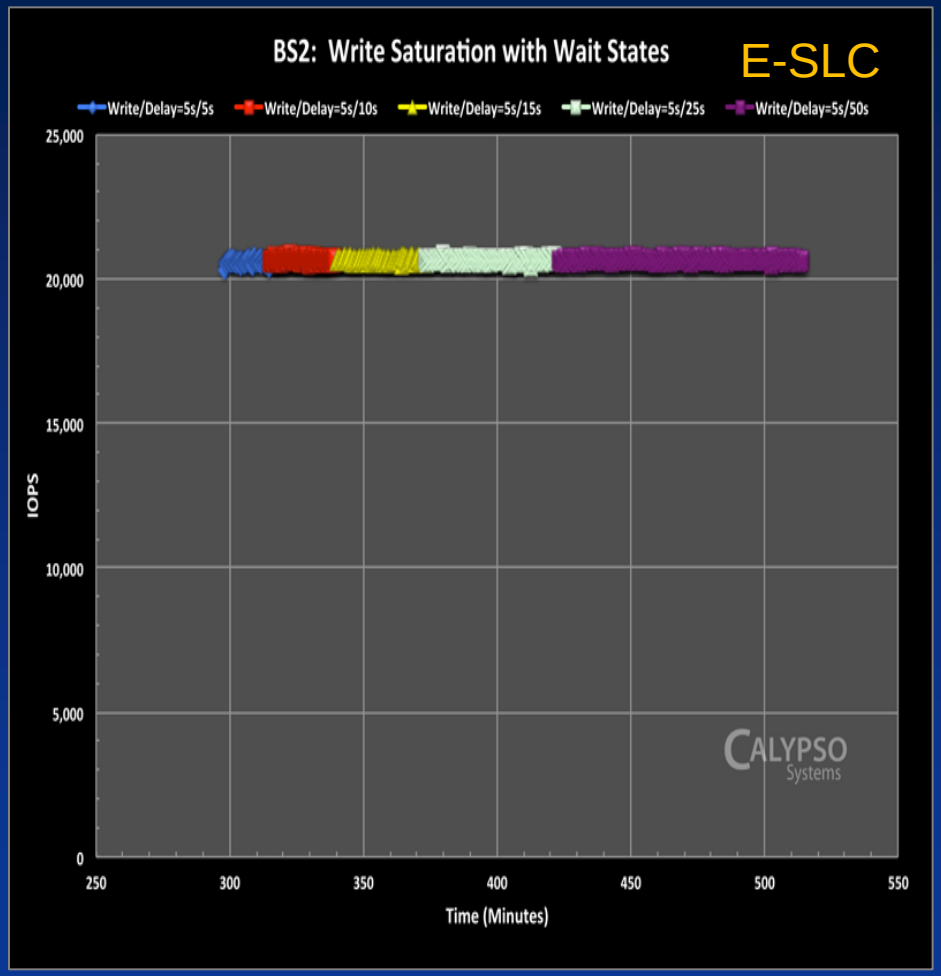

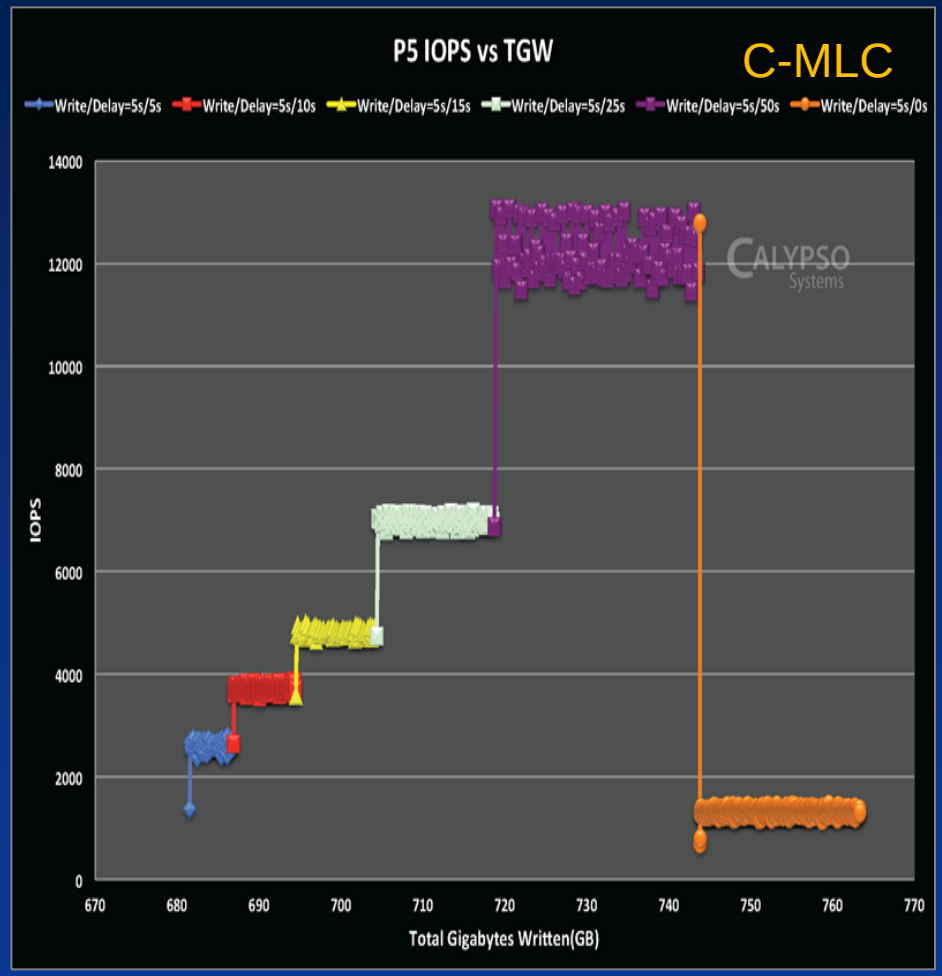

Other Tests Under Development…

- Energy Efficiency
- Data Compressibility
- Enterprise Composite Workload
- Active Range Restrictions
- Active Amount Restrictions
- …

Thank You

Easen Ho

CTO - Calypso Systems, Inc.

eho@calypsotesters.com