



Measuring SSD Performance by Storage Class

Performance Comparison by Storage Class Update
August 2014

Eden Kim, CEO
Calypso Systems, Inc.
Chair, SNIA SSSI Technical Development Committee
Chair, SNIA Solid State Storage Technical Working Group



NAND Flash Performance Issues


SSD NAND Flash Performance

Performance changes over time

Performance is affected by the SSD “write history”

Performance depends on the type of workload

Measurements are affected by the test platform hardware & software





The Solution: Industry Standard Specification

SNIA Solid Storage Storage Performance Test Specification

Defines a standardized test methodology for comparative performance test of NAND Flash based SSD Devices

Specific Pre-conditioning and Steady State Determination Criteria

Standard “Basic Benchmark” tests for IOPS, Throughput, Latency & Write Saturation

SNIA “Advanced” tests focusing on Response Time saturation and changes in sustained stimulus

Defines a normalized Reference Test Platform

Lists Standard Results Reporting Format



Solid State Storage (SSS) Performance Test Specification (PTS) Enterprise Version 1.1

Abstract: This document describes a solid state storage device-level test methodology, test suite and reporting format intended to provide an accurate, repeatable and reliable comparison of NAND Flash-based solid state storage products of various form factors and interfaces used in Client and Enterprise applications.

This document has been released and approved by the SNIA. The SNIA believes that the ideas, methodologies and technologies described in this document accurately represent the SNIA goals and are appropriate for widespread distribution. Suggestion for revision should be directed to <http://www.snia.org/feedback/>.

SNIA Technical Position

September 4, 2013

Note: The Calypso RTP/CTS is the official SSSI PTS Reference Test Platform

TAKE AWAY

Tested to SSS PTS on SSSI RTP 3.0
Performance overlaps categories
Drives are optimized for
workload type

PTS 1.1 “Basic Benchmark” tests reported by Calypso

August 2014 Summary Performance Comparison by Storage Class

| Storage Class | | | IOPS FOB PTS WSAT - T4Q32 | IOPS Steady State PTS IOPS - T2Q16 / T4Q32 | | | | Bandwidth PTS Throughput - T1Q32 | | Response Time PTS Latency - T1Q1 | |
|------------------------|-------------|----------------------------|------------------------------|---|----------------------|--------------------|-----------------------|-------------------------------------|------------------------|-------------------------------------|-------------|
| Category | Device Type | Capacity | RND 4KiB 100% W | RND 4KiB 100% W | RND 4KiB 65:35 RW | RND 4KiB 100% R | SEQ 1024KiB 100% W | SEQ 1024KiB 100% R | RND 4KiB 100% W Ave | RND 4KiB 100% W Max | |
| HDD & SSHD | | | | | | | | | | | |
| 1 | SSHD | 7,200 RPM 2.5" SATA Hybrid | 500 GB | 134 | 134 | 131 | 148 | 107 MB/s | 103 MB/s | 18.54 mSec | 40.63 mSec |
| 2 | SAS HDD | 15,000 RPM 3.5" SAS HDD | 80 GB | 350 | 340 | 398 | 401 | 84 MB/s | 90 MB/s | 55.39 mSec | 97.28 mSec |
| CLIENT SSDs | | | | | | | | | | | |
| 3 | mSATA | mSATA 1.8" MLC | 128 GB | 45,743 | 1,359 | 1,926 | 36,517 | 187 MB/s | 533 MB/s | 0.74 mSec | 543.41 mSec |
| 4 | M.2 x2 | M.2 x2 2280 MLC | 512 GB | 61,506 | 4,185 | 9,532 | 71,282 | 455 MB/s | 535 MB/s | 0.29 mSec | 24.99 mSec |
| 5 | SATA Client | SATAIII 2.5" MLC | 200 GB | 54,788 | 33,583 | 50,708 | 63,640 | 367 MB/s | 480 MB/s | 0.06 mSec | 11.95 mSec |
| ENTERPRISE SSDs | | | | | | | | | | | |
| 6 | SATA 6Gb/s | SATA 6Gb/s 2.5" eMLC | 800 GB | 57,422 | 39,561 | 46,072 | 70,604 | 454 MB/s | 504 MB/s | 0.05 mSec | 0.22 mSec |
| 7 | SAS 12Gb/s | SAS 12Gb/s 2.5" MLC | 800 GB | 97,950 | 41,516 | 72,342 | 145,407 | 448 MB/s | 973 MB/s | 0.05 mSec | 11.84 mSec |
| 8 | SFF 8639 | SFF 8639 4 lane 2.5" MLC | 700 GB | 149,512 | 44,872 | 166,002 | 397,564 | 564 MB/s | 1,698 MB/s | 0.01 mSec | 0.38 mSec |
| 9 | PCIe 8 Lane | PCIe 8 Lane Edge Card MLC | 1400 GB | 159,926 | 87,419 | 236,227 | 742,674 | 614 MB/s | 2,673 MB/s | 0.01 mSec | 0.56 mSec |

All measurements taken on the RTP 3.0 CTS 6.5 Reference Test Platform pursuant to the SNIA PTS-E 1.1.

© August 2014 Calypso Systems, Inc.

NOTE: Thread and Queue settings for PTS IOPS are T2Q16 for HDD/SSHD & Client SSDs and T4Q32 for Enterprise SSDs.

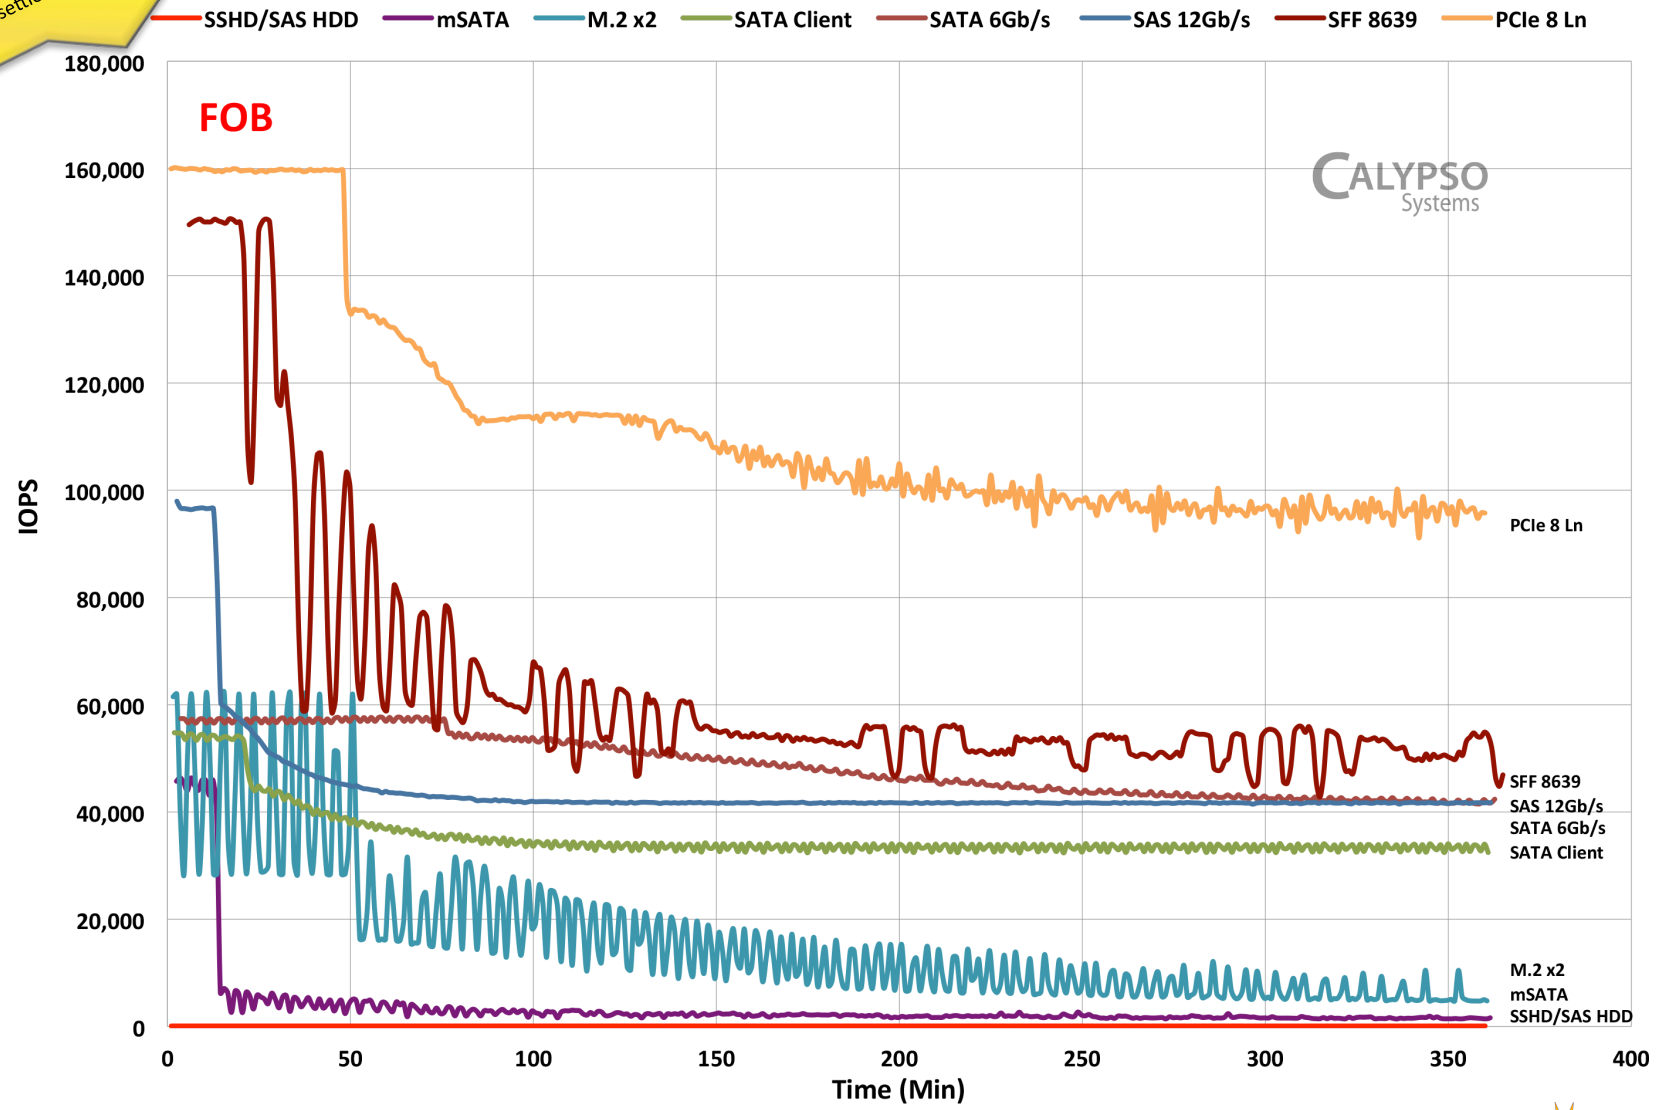
CALYPSO
Systems

Note: Average and Maximum Response Times are reported for the PTS Latency test (at T1Q1 and a lower IOPS level than the PTS IOPS test). Additional Response Time “Confidence Levels” are important and can be observed in Response Time Histograms



TAKE AWAY
 All SSDs show a peak FOB that settles to a steady state

WSAT RND 4K 100% Writes - IOPS vs Time



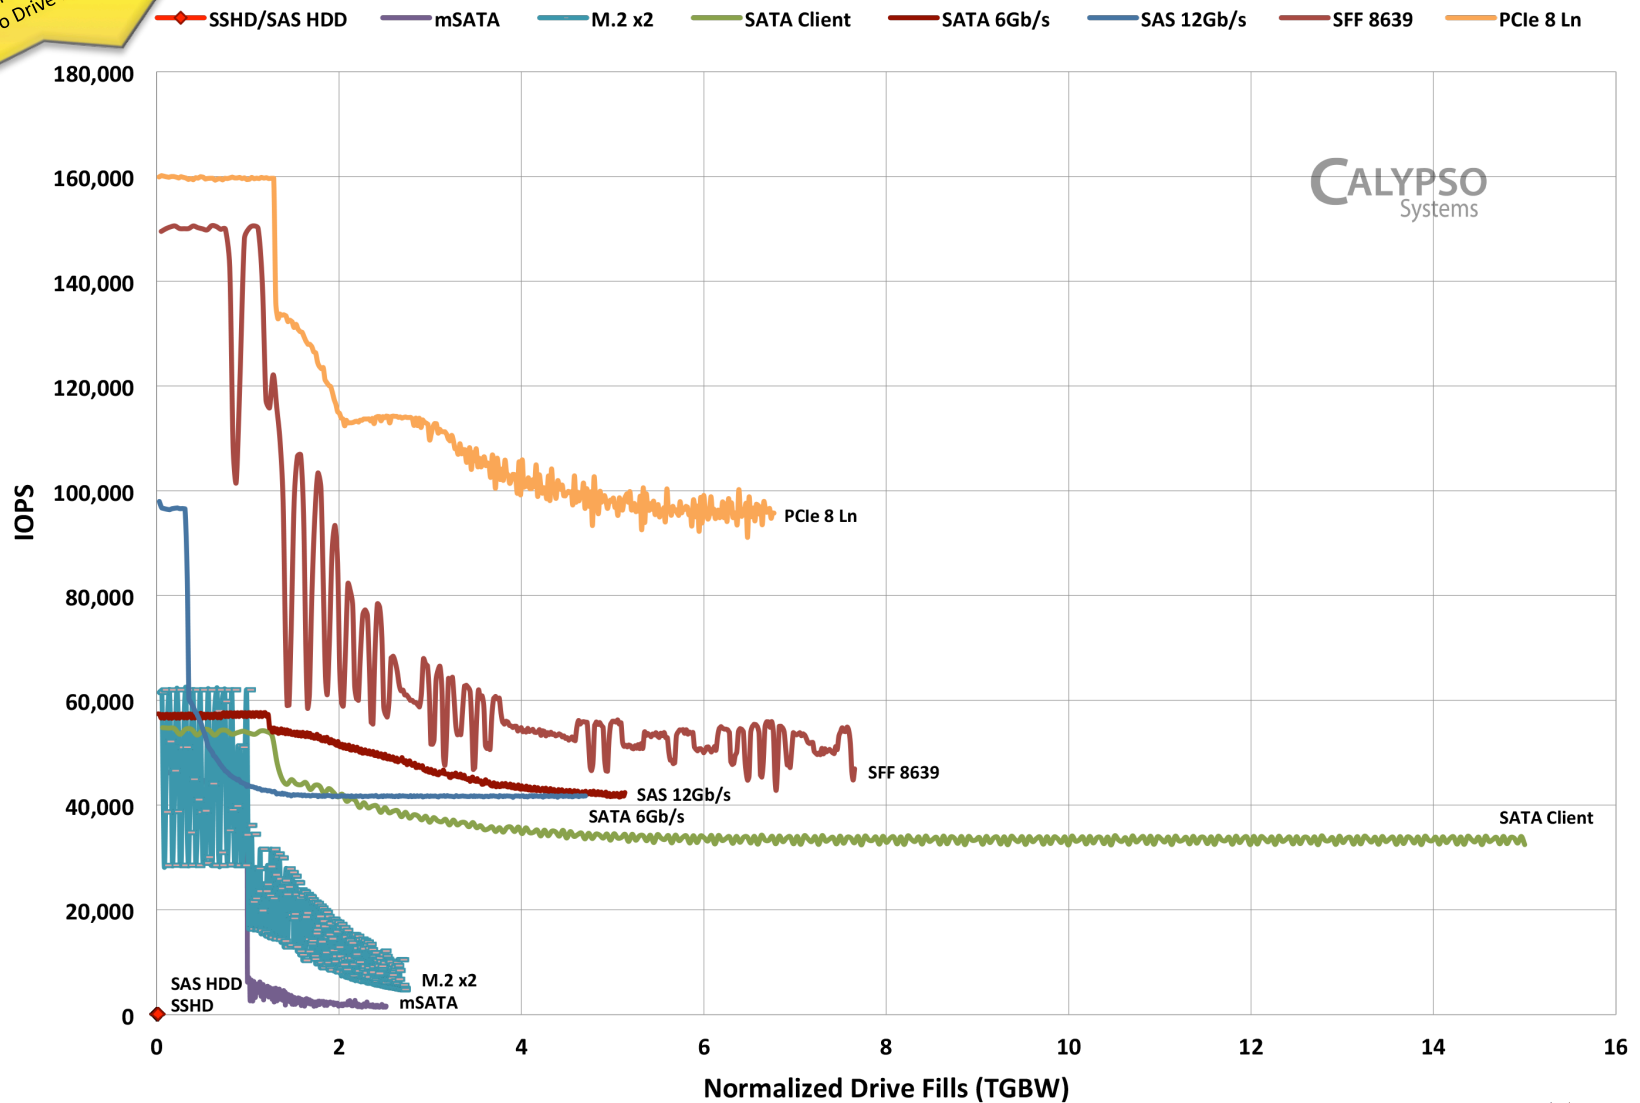
Note: For SSDs, WSAT shows FOB peak immediately after a Device PURGE. Here, RND 4K 100%W are applied for 6 hours.



TAKE AWAY

To measure steady state, be sure to pre-condition with at least two Drive Fills before WDPC.

RND 4K WSAT 6 Hours - IOPS v TGBW in Drive Fills




CALYPSO
Systems

Note: IOPS v TGBW shows how many "Drive Fills" occur by the end of the 6 hour test. Faster drives (higher IOPS) with lower capacities will show more Drive Fills. Slower drives (lower IOPS) with larger capacities will show fewer Drive Fills.

TAKE AWAY
 RND 4K IOPS from PTS IOPS test
 IOPS Increase and Response
 Times decrease through classes

PTS IOPS Test: IOPS v Resp Times for RND 4K RW0 / RW100

■ IOPS RND 4K W ■ IOPS RND 4K R
—■— Ave Response Times W -◆- Ave Response Times R




CALYPSO
Systems

Note: Above IOPS & Response Times are measured during the PTS IOPS test (at T2Q16 or T4Q32) which is optimized for higher IOPS. This differs from IOPS & Response Times measured during the PTS Latency test (at T1Q1) which is optimized for lower Response Times.



TAKE AWAY
 SEQ 1024K from PTS TP test
 PTS Steady State for TP loop
 Bandwidth Increases with lane
 speed and number of lanes

PTS TP Test: Bandwidth v Ave Resp Times - RW0 / RW100 T1Q32



Note: Response Times are for SEQ 1024K measured during the Throughput test (at T1Q32).
 Bandwidth increases as lane speeds and number of lanes increase. Drives can be optimized for higher Bandwidth and lower IOPS.



TAKE AWAY


Latency tests are optimized for Response Time measurement with Response Times at T1Q1

10000

PTS Latency Test: RND 4K T1Q1: Ave v Max Response Time

■ RND 4K W Ave Response Time

■ RND 4K W Max Response Time



CALYPSO
Systems

Note: Response Times are for SEQ 1024K measured during the Throughput test (at T1Q32).
Bandwidth increases as lane speeds and number of lanes increase.

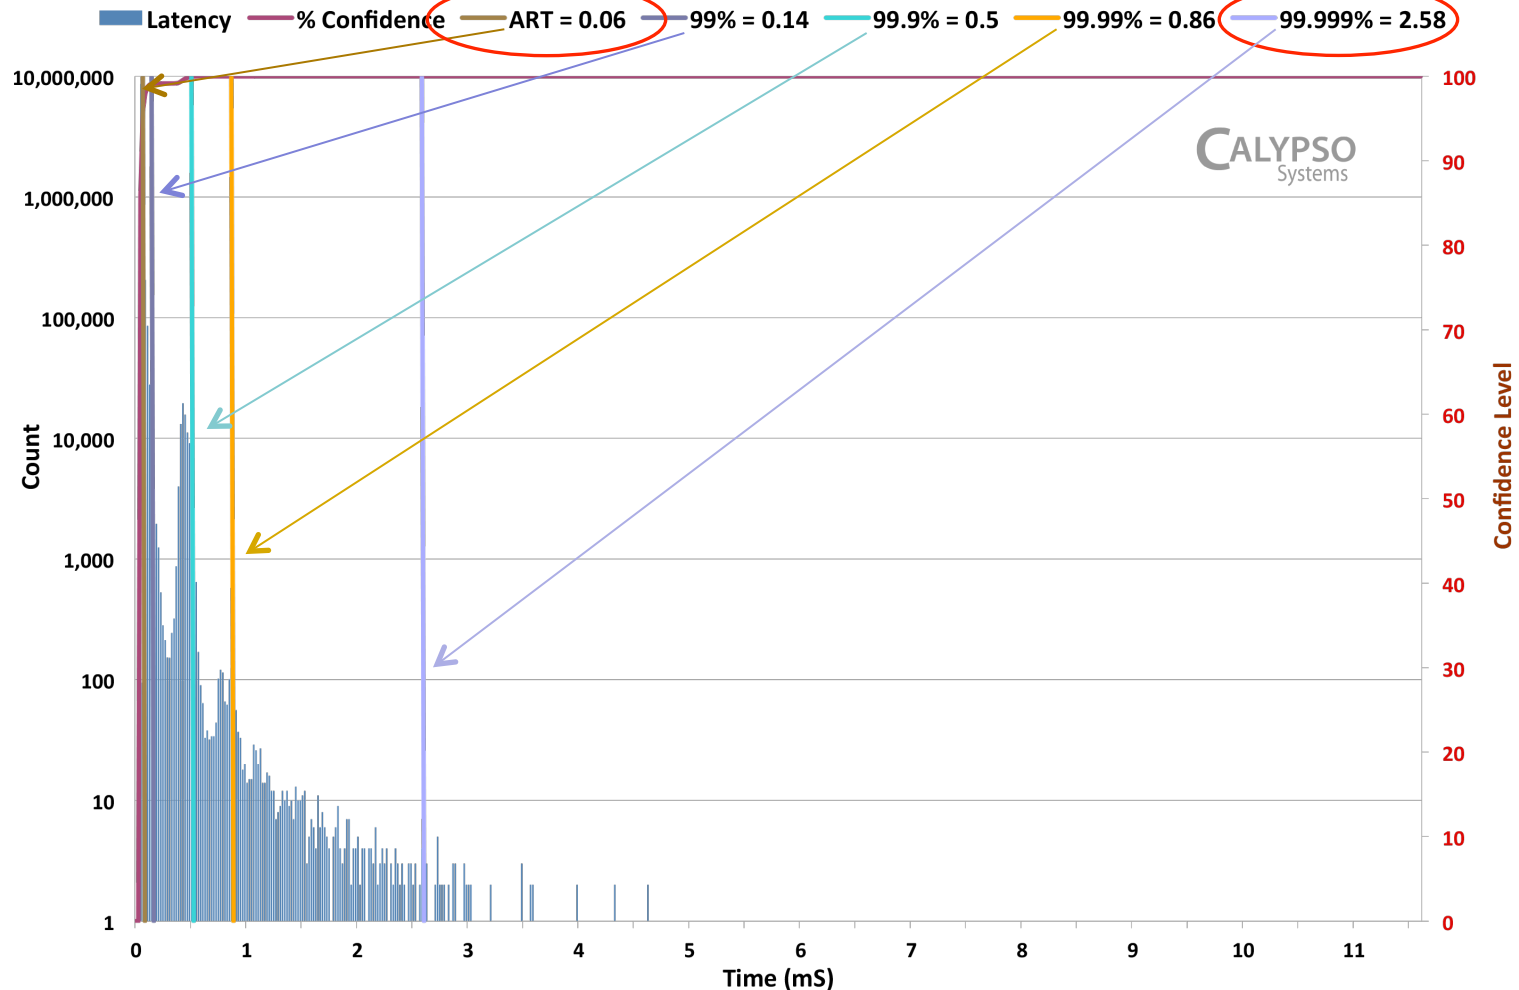
TAKE AWAY

Histograms show the frequency & distribution of Response Times by count & time bin

PTS Latency Test Histogram for RND 4K RW0 at T1Q1

How long does it take for 99.999% of total IOs to occur?

Client SATA: Confidence Level Plot. IOPS=16,775, 65.53 MB/s, **MRT=11.95 mS**



Note: Response Times Histograms show the distribution & frequency of all of the IO response times during the measurement period. In addition to Average and Maximum Response Times, levels are shown for 99%, 99.9%, 99.99% and 99.999% IO completion (or number of '9's' Confidence Level - aka "Quality of Service" or "QoS")

All testing conducted by Calypso on the RTP/CTS pursuant to the SNIA Solid State Storage Performance Test Specification-E v 1.1
For more details, go to www.calypsotesters.com/news




TAKE AWAY
 Confidence Plots or "QoS" show how many IOs are dropped at various response time levels

Comparing PTS Latency Test Response Time Histograms

Bars = Response Time Levels; Tan line = IOPS at T1Q1

Confidence Level Plot Compare (CLPC)
 Response Time Quality of Service (Q o S) - RND 4K RW0 T1Q1



Note: RT Histograms for each device are compared above with RT Bars for ART, MRT and RT Confidence % levels. Dotted red line shows the Client SSD from the previous slide. Response Times & IOPS measured during the PTS Latency test (at T1Q1) and differ from Response Times & IOPS measured during the PTS IOPS test (at T2Q16 or T4Q32)

All testing conducted by Calypso on the RTP/CTS pursuant to the SNIA Solid State Storage Performance Test Specification-E v 1.1
 For more details, go to www.calypso testers.com/news





What's Next?

SSSI Technical works include:

Enterprise Workloads

Investigation of Synthetic Enterprise workloads

Establishing test methodologies and templates for Enterprise Application workloads

Client Composite Workloads

Investigation of Client User workloads

Workload IO Capture Program (WIOCP) tools to download actual user workload statistics

Synthesizing IO statistics into a Client Composite Workload for ranking Client SSDs

Standard Hardware for Reference Test Platform

Qualification of 12Gb/s HBAs, SFF 8639 2.5" platforms, M.2 interposer boards

Client dev sleep motherboards, Power tests

Moving up the IO Stack

Investigation of testing RAID, ARRAY and File System level tests



Interested?
Come join the work.



snia.org/forums/sssi

감사합니다 Natick

Grazie

Danke Ευχαριστίες Dalu

Thank You

Köszönöm

Tack

Спасибо Dank Gracias

谢谢

Merci

Seé
ありがとう

Obrigado

For more information, contact Calypso Systems, Inc.

info@calypsotesters.com

www.calypsotesters.com