SNIA
Solid State Storage Initiative

Understanding SSD Performance Using the SNIA SSS Performance Test Specification

How to Use the PTS Results to Evaluate & Compare SSD Performance

January 2012

SSSI Members and Authors:
Easen Ho, Calypso Systems, Inc.
Eden Kim, Calypso Systems, Inc.
Chuck Paridon, Hewlett Packard Corporation
Doug Rollins, Micron Technology
Tom West, hyperI/O LLC

Table of Contents

The Solid State Storage Initiative

About SNIA

The Storage Networking Industry Association (SNIA) is a not-for-profit global organization made up of some 400-member companies and 7,000 individuals spanning virtually the entire storage industry. SNIA's mission is to lead the storage industry worldwide in developing and promoting standards, technologies, and educational services to empower organizations in the management of information. To this end, SNIA is uniquely committed to delivering standards, education, and services that will propel open storage networking solutions into the broader market. For additional information, visit the SNIA web site at http://www.snia.org.

About SNIA Solid State Storage Performance Test Specification (PTS)

This white paper is based on testing done pursuant to the SNIA Solid State Storage Performance Test Specifications (PTS) rev 1.0 for Client (PTS-C) and Enterprise (PTS-E) applications. All testing was conducted by Calypso Systems, Inc., a certified SSSI PTS Test Lab, using the SNIA compliant Calypso Reference Test Platform (RTP 2.0) and CTS 6.5 test software (see Test Platform Requirements in PTS-E and PTS-C). Test results and this white paper can be downloaded from the SNIA SSSI website at www.snia.org/forums/sssi/pts.

1. Introduction

A commonly asked question for NAND Flash based SSDs is: "which SSD performs best?" Invariably, the informed answer is "it depends" – this is due to several factors inherent in NAND-based SSDs.

Device Level Factors:

- Was the test done at the file system level or at the device level?
- How was the drive treated before the test started? Was it preconditioned? If so, how?
- Did the test sequence ensure the drive had reached steady state before the results were captured?
- How much data was written and where was it written to during the test?
- What data pattern was tested?

System Level Factors:

- What test platform was used to test the SSD?
- What hardware and software package was used?
- Is the HBA bandwidth sufficient?

Architectural Factors:

- What is the type of NAND Flash?
- Is the drive's targeted use for high write workloads, or high read workloads?
- Is the drive's apparent performance designed to meet other criteria such as warranty issues?

This white paper will focus on evaluating and comparing SSD performance using the SNIA PTS Specification and the SNIA Standard Reporting Format test reports. As an aid to better understanding the terminology in this white paper, please see the SSS Glossary at www.snia.org/forums/sssi/knowledge/education.

2. SSD Performance States

All NAND based SSDs exhibit at least three distinct performance states, Fresh-Out-of-Box ("FOB"), Transition and Steady State.

Fresh-Out-of-Box ("FOB")

The condition of a new/unused Solid State Storage device when first received from the manufacturer. Typically the storage cells on the device will have few or no program/erase ("P/E") cycles applied to them when the device is in this state (the exception would be any P/E cycling done at the factory as part of the manufacturing process), or the device has been returned to this state using standard methods such as ATA Security Erase, SCSI Format Unit, or other proprietary methods. This device is ready to have data stored (that is, all storage elements are pre-erased).

Transition

This is a performance state where the device's performance is changing as it goes from one state to another. For example, a typical small block write performance of an SSD will start out very high in the FOB state. After a certain amount of the same stimulus, the SSD would then reach a state where the performance becomes relatively time-invariant. The period of time between the FOB state and this relatively time-invariant state is called the Transition State.

Steady State

The condition under which most of the transient performance behavior (i.e., the Transition State) has ceased is called "Steady State." Steady state performance is typically reflected in a relatively small change in performance over a relatively large timeframe and is specifically defined in the PTS (hereinafter referred to as "Steady State").

As they are written, most SSDs migrate through these performance states sequentially: FOB → Transition → Steady State. Because Steady State (SS) most accurately reflects the SSD's performance in long-term use for a specific IO activity type, the most desirable region in which performance is measured is the Steady State. For example, most SSD Transition States are measured in hours and are generally very short compared to service life of the drive.

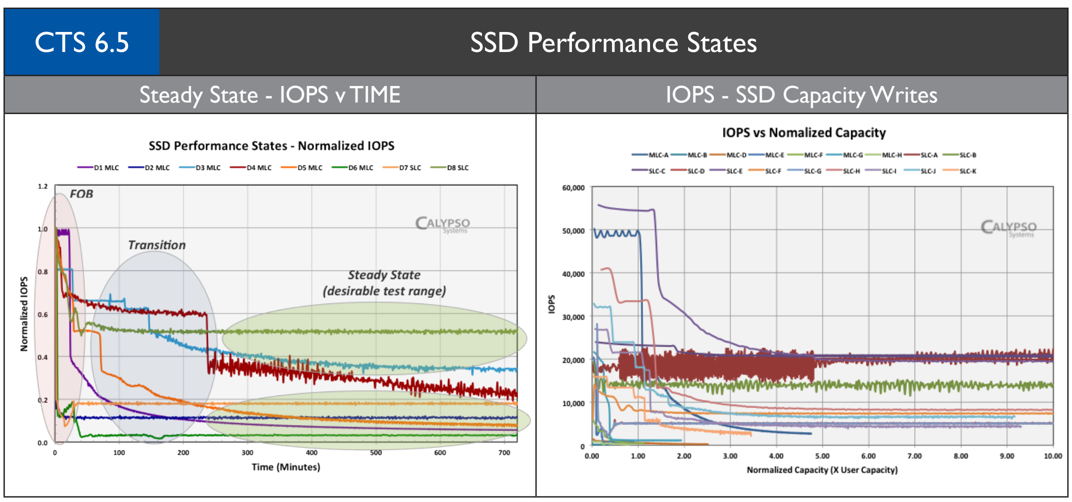

Figure I. SSD Performance States

Factors Affecting SSD Performance

SSD performance is highly dependent on three primary factors, Write History (including host idle time), Measured Workload and Hardware/Software Environment. The ***exact same SSD*** can produce dramatically different results depending on these factors.

Write History and Host Idle Time - a "Fresh-Out-of-Box" (FOB) SSD with no write history will initially show very high peak performance that will dramatically decrease as more data is written to the SSD. Similarly, if an SSD has a buffer enabled (such as a DRAM buffer) and the device is idle for a given amount of time, the SSD may actively migrate the data contained in the DRAM buffer to the NAND. If the host interface remains idle, this newly cleared DRAM is available to store incoming data when the host begins writing again. This will produce a brief period of very high performance that dwindles as the buffer fills.

In Figure 1 IOPS v SSD Capacity Writes, several SSDs have been written (the number of drive fills is on the x-axis) and their performance (in IOPS) plotted on the y-axis. Despite the curves being slightly different, all are approximately the same shape and all of the drives shown exhibit the same performance fall-off. The area on the extreme right of the above performance plots is the Steady State region.

Measured Workload - In addition to the amount of data written, the type of data (transfer size, degree of randomness, etc.) can also affect an SSD's performance.

One might expect that, if a drive were written into Steady State with large, sequential transfers, then the stimulus changed to small block random transfers (again, writing enough data to get the drive into Steady State), then finally changed again back to the large block sequential transfers – that the two regions of sequential transfers would exhibit the same performance.

For many drives, this phenomenon – also known as Cross Stimulus Recovery – can be seen in the examples below. For these plots, Time is shown on the x-axis and Throughput in MB/s is shown on the y-axis.

CTS 6.5 — Cross Stimulus Recovery: SEQ 128KiB - RND 4KiB- SEQ 128KiB

MLC-A - 256 GB

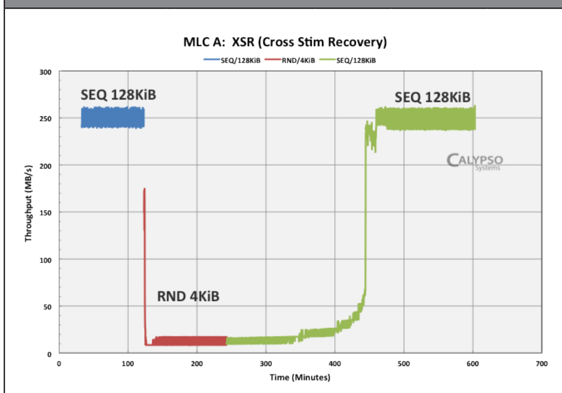

MLC-B - 160 GB

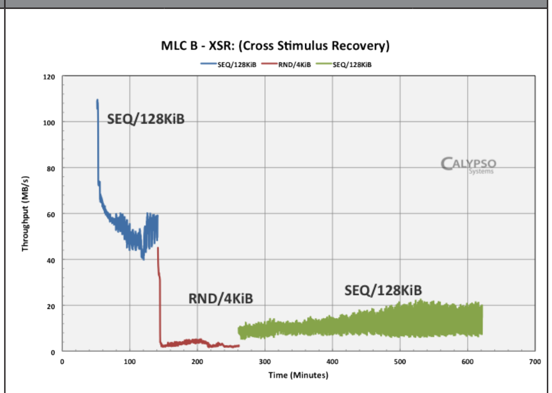

MB/Sec	MLC-A Measurement Period	Start	After RND 4KiB	End	MLC-B Measurement Period	Start	After RND 4KiB	End
	SEQ 128KiB W	250	5	250	SEQ 128KiB W	120 - 60	10	10 - 20

Figure 2. Cross Stimulus Recovery

For SSD MLC-A, as expected, the small random transfers (shown in red and labeled "RND 4KiB") offer lower overall Throughput. Also, as expected, the two regions of large block sequential transfers (shown in blue at the left and green at the right and labeled SEQ 128KiB) offer the same throughput after an initial recovery period.

However, this is not universally true. The second example of SSD MLC-B was written and measured identically. MLC-B never recovers the initial performance level seen in the sequential transfers during this test.

Hardware/Software Environment – The test platform itself can also affect results. Operating system, file system, hardware, driver, and software differences can all influence SSD performance as each is a significant part of the "storage stack." In order to minimize these effects, the host test platform should be designed to affect performance as minimally as possible. For example, it is important that the host test platform has sufficient bandwidth and sufficient host processing resources to generate the necessary IO loads.

3. Creating a Standard - PTS Specification

The SNIA Solid State Storage Performance Test Specification 1.0 (SSS PTS or PTS) has been published in two versions to address the markedly different requirements of the Enterprise and Client use cases. To ensure that test results can be easily compared, the PTS also includes a SNIA Standardized Reporting Format for concise and simple disclosure of required reporting information and easy comparison of SSD performance test results.

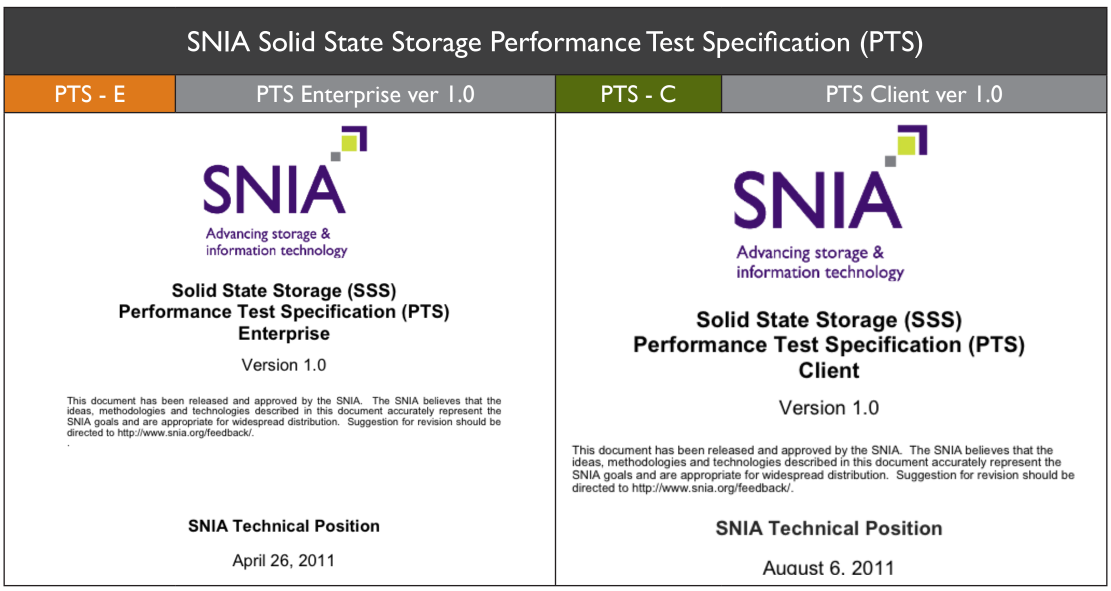

Figure 3. PTS Specification PTS-C & PTS-E

The Performance Test Specification (PTS):

- is based on ***synthetic device level tests***
- Prescribes a standard ***preconditioning*** methodology (to normalize the effect of write history and get the SSD to a steady state)
- Specifies ***test workloads*** (to clearly specify access pattern, data pattern, access range and footprint restrictions)
- Lists requirements for the ***test platform*** (to ensure the hardware/software environment generates sufficient IOs as specified)
- Provides a ***standardized PTS Report Format*** that allows easy reference to key test set-up information and test results

The purpose of this white paper is to introduce the SNIA PTS Report Format, highlight the organization of the informational header, present the key data results of the PTS Enterprise (PTS-E) and PTS Client (PTS-C) tests, and to discuss the evaluation and interpretation of the data contained in the PTS Report.

4. What is a "Synthetic Device Level Test"?

Synthetic Device Level Test

Synthetic Device Level testing, in the context of the PTS, refers to the use of a known and repeatable test stimulus targeted directly at the Block IO devices themselves (as opposed, for instance, to a File System Level Test) while using a particular set of parameters to completely describe the test stimulus. These parameters include:

- Read/Write Mix (relative amount of read versus write IOs)
- Block Size (data transfer size of each IO)
- Data Pattern (related to the data content of the IO)

Additional parameters provide restrictions regarding which Logical Block Addresses (LBAs) are allowed to be accessed within the device:

- ActiveRange (the range of LBAs allowed to be used)
- ActiveRange Amount (the sum of the capacity referenced by the LBAs that are accessed during a test)
- ActiveRange segmentation (the distribution and size of contiguous, equal-sized LBA ranges (or segments) within the ActiveRange).

The PTS utilizes these parameters to specify how the IOs are to be issued from the test application to the Device Under Test (DUT). Note: "ActiveRange and ActiveRange Amount" are terms defined in the PTS.

The test operator should be careful to ensure that the software test tools and test environment (i.e., the test hardware, the operating system and its associated drivers) do not become the limiting factor in obtaining accurate measurements.

File System Level Test

Device Level testing differs from testing at the File System Level. File System Level testing generally involves directly issuing, through a file system, specific file IO operations targeted towards the device to be tested. SSD performance testing at the File System Level introduces additional variables that can affect the tests and the corresponding performance measurements. These variables and effects are related to the application/test software itself, the various components/drivers and their interaction within the OS software stack, the particulars of each file system, and the underlying computer hardware platform. See the Hardware/ Software Stack graphic below.

Some of the specific variables that can impact SSD performance testing at the File System Level as well as application IO performance in general are:

Caching. A single file IO operation issued by an application may result in no physical device access at all due to various caching strategies that may be implemented at the OS or driver level.

Fragmentation. A single file IO operation issued by an application may require multiple IO operations to the physical device due to file fragmentation. Furthermore, the drivers can also split or coalesce IO commands which can result in the loss of 1-to-1 correspondence between the originating IO operation and the physical device access.

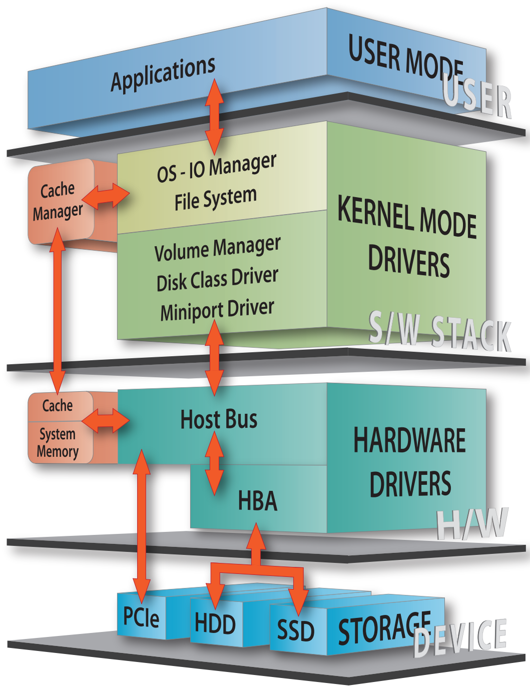

Figure 4. Hardware Software Stack

Timing. Various timing considerations can have a notable impact upon the manner in which IO operations traverse the OS software stack. For instance, while several applications can each be performing sequential access IO operations to their respective files, these concurrent IO operations can be observed to arrive in a more random access pattern at the lower-level disk drivers and other components (due to system task switching, intervening file system metadata IO operations, etc.).

User Workloads. A primary interest for many, if not most, end users when comparing SSD performance is to determine and substantiate the performance benefits that can be gained while operating within their specific computing environments using their particular applications of interest. However, the range and diversity of applications that are available, along with the particular manner in which they are actually used, can introduce a significant set of factors that can impact application IO performance.

In sum, the "synthetic device level testing" provided by the PTS enables the use of a standardized procedure for uncovering the native performance capabilities of different SSDs. In turn, an understanding of such intrinsic performance capabilities can be an important and even fundamental factor when seeking to address or improve application IO performance.

5. What About User Workload Characteristics?

A popular interest and goal of end users is to properly and prudently match the performance needs of their particular applications to their specific storage purchases, especially in a cost-effective manner. As a result, there can be a propensity towards attempting to directly map (i.e., correlate) the advertised/reported performance metrics of storage devices (e.g., IOPS, MB/s, etc.) to the presumed workload characteristics of their applications. This can result in inaccurate assumptions about the workload characteristics.

As noted within the prior section, the IO activity that stems from applications is subject to a variety of variables and effects as the IO operations traverse the OS software stack. This can be confirmed by collecting empirical IO operation performance metrics from the application perspective, and moreover at various key points within the OS software stack.

Furthermore, such mapping of user workload characteristics to the "speeds and feeds" performance of devices is predicated upon the extent to which these determined workload characteristics in fact accurately reflect the actual IO activity of the particular applications of interest in normal, everyday usage.

Various "rules of thumb" can provide some general guidance in this regard. Nevertheless, the common caveats of "your mileage may vary" and "it depends" are often the final advice and caution for these "rules of thumb."

Overall, careful attention should be given to determining and understanding pertinent workload characteristics just as careful attention be given to gathering the PTS performance measurements. The value of the PTS performance measurements can be further enhanced by their greater relevancy to actual user workload characteristics.

6. Understanding PTS Reports

With an understanding of typical user or target workloads, the reader / test sponsor can now analyze and use the various PTS Test Reports. Each PTS test has test conditions, IO parameters, access patterns and metrics designed to implement the workloads associated with each test.

The PTS prescribes different preconditioning and test ranges to differentiate between Enterprise and Client workloads. For example, Enterprise workloads, which are typified by 24/7 continuous use, precondition to 100% of the device capacity (LBAs) and apply the test stimulus to the entire LBA range. Client workloads, by contrast, are preconditioned to a limited LBA range (75% or 100%), in part to account for the impact of the SSD related TRIM command, and limit test ActiveRange Amount (8 GiB or 16 GiB) to reflect smaller active data footprints empirically observed in Client workloads.

Further refinement can be gleaned from the test settings and parameter settings such as outstanding IOs (OIO) as measured by the total Thread Count (TC) and the Queue Depth (QD), which indicates the number of outstanding IOs per thread. Generally, Enterprise SSDs will be optimized to a larger OIO count whereas Client SSDs are designed to function optimally with fewer OIOs. Furthermore, determining optimal OIO settings will depend on the hardware, OS, application software, as well as IO specifications such as read/write ratio, sequential/random ratio, block size, etc.

Finally, when evaluating the test results it is important to note the test platform used to gather the test results (hardware, OS and test software). The PTS requires disclosure of test hardware (CPU, RAM), device bus interface (SAS, SATA, 6Gb/s, 3GB/s), test system manufacturer (motherboard & HBA card vendor or test system house vendor), OS and test software used. The PTS lists both the hardware and software tool requirements as well as a recommended Reference Test Platform (RTP) that was used to develop and validate the PTS.

This multitude of data, test settings and environment are managed in the PTS Test Format set forth in the PTS Enterprise & Client Specifications.

7. PTS Reports

The PTS 1.0 Enterprise (PTS-E) and Client (PTS-C) Specifications set forth required and/or optional tests, test set up conditions and test parameters. The test sponsor ***must*** run the ***required*** tests and ***may*** run ***optional*** tests.

The PTS-E has four required tests: Write Saturation (WSAT), IOPS, Throughput (TP) and Latency (LAT). The PTS-C has three required tests: IOPS, TP and LAT. The test sponsor may elect to run additional optional tests (such as WSAT for PTS-C) as well as use optional test set up conditions and parameters ***in addition to*** those set forth in the PTS as required. Any optional tests or settings shall be clearly reported with the test results.

One report is required for each test run. Use of the PTS Report format ensures that all of the required test settings, parameter disclosures and metrics are uniform and complete. All PTS report pages must contain a Report Header that discloses required reporting, SSD and administrative data. The SNIA Standard PTS Report Format is provided for this purpose and is attached as Annex A to both the PTS-E and PTS-C Specifications.

The PTS Report format consists of a Summary Page and detailed Reporting Pages for each test run. The Summary Page will present summary test set up, device under test (DUT) information and other required and optional information that may not be on each individual Report Page.

Each Report Page must have a Report Header that contains the specific test set up and conditions that pertain to the test results presented on the individual Report Page. Examples of the individual Report Summary Page and individual Report Headers are listed below.

Note: PTS 1.0 Modifications. The tests and Reports contained in this white paper reflect PTS 1.0 tests and certain modifications thereto that have been approved by the SNIA SSS TWG for distribution in the upcoming PTS revision 1.1. Examples of such modifications include:

- *PTS-E TP test Block Sizes have been reduced from five Block Sizes to two Block Sizes (128KiB, 1024KiB) and the two Block Sizes are to be run in separate independent tests with separate test reports.*

Note: Abbreviations. The PTS uses certain abbreviations and conventions in headers and reports for reporting convenience such as:

- *Use of RND for random and SEQ for sequential*
- *Use of R, W and R/W for READ, WRITE and READ/WRITE, respectively.*
- *Use of TP and LAT for Throughput and Latency, and AVE, MAX and MIN for Average, Maximum and Minimum.*

Note: KiB v KB. Block Sizes used and reported in the PTS are in units of KiB (i.e., "kibibytes" where 1 kibibyte equals 1024 bytes) rather than KB (i.e., "kilobytes" where 1 kilobyte equals 1000 bytes). This usage of KiB is in accordance with the device Logical Block/Sector sizes, which generally are either 512 bytes (0.5 KiB) or 4096 bytes (4 KiB). Also note, however, that the manufacturer-stated user storage capacity of a device is typically reported in units of GB (i.e., "gigabytes" where 1 gigabyte equals 1 000 000 000 bytes) rather than GiB (i.e., "gibibytes" where 1 gibibyte equals 1 073 741 824 bytes).

8. PTS Report: Summary Pages & Report Headers

Summary Report Pages

A "Summary Report Page – Individual Test" Informative example is listed in Annex A to the PTS. This Summary Page is useful to list key set up and test parameter data that applies to the particular test run which may not be able to be reported on each individual Report Page Header.

Additional Informative "Summary Report Page - All Tests" can be produced that summarize the key test setup data and conditions for all tests run under the relevant PTS. Examples of both a "Summary Report Page - All Tests" (that contains MLC-A WSAT, IOPS, TP and LAT test information) and an Individual "Summary Report Page – Individual Test" (in this case MLC-A IOPS) are reproduced below.

Summary Report Pages - PTS

PTS - C | Summary Report Page - All Tests

MLC-A Test Report
Summary Report Page – All Tests

SNIA Solid State Storage Performance Test Specification (PTS)						Rev.	PTS 1.0
						Page	1 of 26
Device Under Test (DUT)	MLC-A	SNIA SSS PTS Summary Report		Calypso Systems, Inc.		CALYPSO Systems	
DEVICE INFORMATION		TEST HARDWARE PLATFORM		TEST SOFTWARE		REPORT DATE	
SERIAL NO.	0000-0000-FFFF	SYSTEM	Calypso RTP 2.0	SYS OS	CENT OS 5.6	Report	06DEC11
FIRMWARE REV.	BFO1	Motherboard/cpu	Intel 5520HC / W5580	SW TOOL	Calypso CTS 6.5	Test Run	01NOV – 04DEC11
USER CAPACITY	MLC 256 GB	RAM	12GB ECC DDR3	SW Rev	1.19.10	Test Sponsor	Calypso
DEVICE INTERFACE	6 Gb/s SATAII	Device Interface	LSA 9212-e 6Gb/s HBA	Release	Nov. 2011	Auditor	N/A

Testing Summary: Tests Run

PTS-C	TEST	Purge	DP	OIO	WIPC		WDPC		STEADY STATE	
					PC AR	TEST AR	AR AMT	SEGMENTS	WORKLOAD	TIME/GB
7.0	WSAT - OPTIONAL	Security Erase	RND	TC 1 QD 16	100%	100%	N/A	N/A	RND 4KiB W	24 Hrs 1.9 TB

PTS-C	TEST	Purge	DP	OIO	WIPC		WDPC		STEADY STATE	
					PC AR	TEST AR	AR AMT	SEGMENTS	WORKLOAD	ROUNDS
8.0	IOPS - REQUIRED	Security Erase	RND	TC 2 QD 16	100%	100%	16 GiB	2048	IOPS LOOP	2 - 6

PTS-C	TEST	Purge	DP	OIO	WIPC		WDPC		STEADY STATE	
					PC AR	TEST AR	AR AMT	SEGMENTS	WORKLOAD	ROUNDS
9.0	THROUGHPUT - REQUIRED	Security Erase	RND	TC 32 QD 32	100%	100%	16 GiB	2048	SEQ 1024KiB	1 - 5

PTS-C	TEST	Purge	DP	OIO	WIPC		WDPC		STEADY STATE	
					PC AR	TEST AR	AR AMT	SEGMENTS	WORKLOAD	ROUNDS
10.0	LATENCY - REQUIRED	Security Erase	RND	TC 1 QD 1	100%	100%	16 GiB	2048	LAT LOOP	4 – 8

Test Sponsor – Special Notes

ITEM	NOTATION	COMMENTS

PTS - C | Summary Report Page - IOPS Test

MLC-A Test Report
Summary Report Page - IOPS

SNIA Solid State Storage Performance Test Specification (PTS)						Rev.	PTS 1.0
						Page	7 of 26
Device Under Test (DUT)	MLC-A	SNIA SSS PTS Summary Report		Calypso Systems, Inc.		CALYPSO Systems	
DEVICE INFORMATION		TEST HARDWARE PLATFORM		TEST SOFTWARE		REPORT DATE	
SERIAL NO.	0000-0000-FFFF	SYSTEM	Calypso RTP 2.0	SYS OS	CENT OS 5.6	Report	06DEC11
FIRMWARE REV.	BFO1	Motherboard/cpu	Intel 5520HC / W5580	SW TOOL	Calypso CTS 6.5	Test Run	01NOV – 04DEC11
USER CAPACITY	MLC 256 GB	RAM	12GB ECC DDR3	SW Rev	1.19.10	Test Sponsor	Calypso
DEVICE INTERFACE	6 Gb/s SATAII	Device Interface	LSA 9212-e 6Gb/s HBA	Release	Nov. 2011	Auditor	N/A

Test Description

Purpose	To measure RND IOPS matrix using different BS and R/W Mixes
Test Outline	PURGE, then apply preconditioning until Steady State is achieved according to the SNIA PTS
Preconditioning	PURGE followed by SNIA PTS prescribed WIPC & WDPC

Test Set Up

PTS-C	TEST	Purge	DP	OIO	WIPC		WDPC		STEADY STATE	
					PC AR	TEST AR	AR AMT	SEGMENTS	WORKLOAD	ROUNDS
8.0	IOPS - REQUIRED	Security Erase	RND	TC 2 QD 16	100%	100%	16 GiB	2048	IOPS LOOP	2 - 6

Select Performance Data

RND 4KiB W	RND 4KiB R	RND 8KiB W	RND 8KiB R
3,147	29,876	1,584	21,723

Test Sponsor – Special Notes

ITEM	NOTATION	COMMENTS

Figure 5. PTS Report Summary Pages

Report Page Header

PTS Report Headers are Normative and are required for PTS Reporting. A PTS Report Header is listed on each page of the PTS Report and contains a concise summary of key reporting requirements and informational data. Critical information is contained in the header and an experienced PTS Report reader will refer to the header as a reference or a "check list" when reviewing data or comparing different SSD PTS reports. Key information is organized temporally in shaded boxes across the bottom half of the header.

PTS - C	SNIA PTS REPORT HEADER
	MLC-A 256 GB

Test Run Date:	11/14/2011 12:39 AM	**Report Run Date:**	11/21/2011 04:12 PM

Client IOPS (REQUIRED) - Report Page

SNIA SSS TWG: Solid State Storage Performance Test Specification (PTS)		Rev.	**PTS-C 1.0**
		Page	**1 of 6**

Device Under Test (DUT)	**VENDOR: ABC CO.**	**SSD MODEL NO: MLC-A 256 GB**	**TEST SPONSOR**	CALYPSO Systems

Serial No.	0000-0000-FFFF	**DUT Preparation**		**Test Loop Parameters**		**Steady State**	
Firmware Rev	BF01	**Purge**	Security Erase	**REQUIRED:**		**Convergence**	YES
Capacity	256 GB	**Pre-Conditioning**		**Data Pattern**	RND	**Rounds**	1-5
NAND Type	MLC	**Workload Independent**	2X SEQ/128KiB	**Tester's Choice:**		**PC AR**	100%
Device I/F	6 Gb/s SATA			**OIO/Thread (QD)**	8	**AR AMOUNT**	16 GiB
Test Platform	RTP 2.0 CTS 6.5	**Workload Dep.**	Full IOPS Loop	**Thread Count (TC)**	1	**AR Segments**	2048

Figure 6. PTS Report Header

General Information

General Information about the ***test administration*** (Test Run Date, Report Date, PTS Test Run, PTS revision number) and ***SSD test*** (SSD Vendor, Test Reported, and Test Sponsor) is set forth in the white area in the top half of the Header block. Here, the reader can identify the Vendor, SSD Model Number, test sponsor and the test run date, and the date the PTS Report was generated.

Test Environment

The salmon shaded left hand box sets forth the key Device Under Test ***(DUT) information***: Serial Number, Firmware Revision, Capacity, NAND type, Device Interface and Test Platform. Here, the reader can note the bandwidth of the DUT interface (which can limit SSD performance), in this case 6Gb/s SAS or 3Gb/s SATA and the Test Platform hardware and software tools used.

DUT Preparation

The blue shaded middle left box sets forth ***DUT Preparation*** information related to the preconditioning regime of the PTS test. DUT Preparation identifies the type of ***PURGE*** applied at the beginning of each test - either Security Erase for ATA, Format Unit for SCSI, or other proprietary PURGE command that meets the requirements defined in the PTS. PURGE is required in order to reset the virtual mapping look-up tables and ensure that all NAND cells are programmed to a state "as if no writes had occurred." The purpose of PURGE is to "reset" the effect of the write history for the impending test.

The ***Preconditioning*** section identifies Workload Independent Preconditioning ***(WIPC)*** - in this case 2X (i.e., twice) the user capacity in Sequential (SEQ) 128KiB Writes - and the Workload Dependent Preconditioning ***(WDPC)*** - in this case, use of the full IOPS Test Loop consisting of 7 R/W mixes and 8 Block Sizes.

Note: WIPC writes 2X the user capacity in SEQ 128KiB blocks to quickly touch all the LBAs in the Test ActiveRange to expedite convergence to test Steady State. WDPC immediately follows WIPC (with no delay) to begin Steady State Convergence rounds by writing "test loop data" in order to avoid the effects of cross stimulus on performance (in this case writing preconditioning data that is different, in terms of data transfer size and degree of randomness, than the test loop data).

Test Loop Parameters

The green ***Test Loop Parameters*** box contains required parameter disclosures - in this case, for the IOPS test. ***Data Pattern*** refers to the Random or non-Random data content of the test workload (not to be confused with the "Access Pattern" of Random or Sequential R/W mixes and Block Sizes).

Note: The Random or non-Random data content pattern is important to note for SSD controller architectures that may optimize performance for non-Random data content patterns.

The user selected ***OIO setting*** must also be disclosed - in this case a Thread Count of 1 and OIO/Thread (i.e., Queue Depth) of 8 results in a total OIO of 8.

Note: At a general level, the SSD must have an adequate OIO count to generate enough stimuli to measure its maximum performance. For example, an OIO of one (1 TC and 1 QD) may starve a DUT and yield artificially low maximum IOPS. On one hand, some SSDs may not be designed to handle a very large number of threads, such as some Client SSDs, and may see a decrease in maximum IOPS when the TC exceeds some number. On the other hand, Enterprise drives prevalent in multi-user or virtualized machine (VM) environments are generally designed for higher OIO and TC.

Steady State

The purple block on the right hand side presents summary ***Steady State information***: whether Steady State was reached (yes or no) and the ***number of Rounds*** measured (a minimum of 5 Rounds that meet the Steady State requirement OR a total of 25 Rounds).

Note: A "Round" refers to, for example, one pass through the IOPS loop of 7 R/W Mixes by 8 Block Sizes, or 56 one-minute tests for a "Round" duration of 56 minutes.

Additional Preconditioning parameters are set forth, as applicable, including the ActiveRange, which will be either a percentage of LBAs (100% for PTS-E) or a PC AR Amount (8GiB or 16GiB for PTS-C), and AR Segments (e.g. 2048 for PTS-C).

Note: The minimum test time to run a PTS IOPS test is the time required for PURGE and WIPC (perhaps one hour) and at least 5 Steady State Rounds, each Round taking 56 minutes, or about 6 total hours minimum for a single IOPS test. If the test software does not programmatically determine Steady State, up to 25 Rounds (or 25 hours plus the approximate 1 hour PURGE / WIPC) may need to be taken using post processing and manual inspection to ascertain the five Round Steady State Window. See section 11 Steady State Measurement Window Calculations.

ActiveRange sets forth the ***Test ActiveRange***, or LBAs written to, in the test loop. The PTS sets forth a required LBA range (or ranges) and allows the test sponsor to conduct optional additional tests (in addition to the required test ActiveRange or ranges) at different ActiveRange settings that must be disclosed (within the "optional" setting box).

Note: PTS-C IOPS requires four IOPS runs. There are two LBA Active Ranges (75%, 100%) and two Test AR Amounts (8GiB, 16GiB) that make up the 2x2 matrix.

9. Steady State Convergence

Each of the Steady State tests (IOPS, TP and LAT) begin with a Steady State Convergence Plot. Figure 7, PTS-C IOPS Page 1, "Steady State Convergence Plot" tracks various variables of interest for each of the Block Sizes across the total test Rounds measured.

PTS - C | **STEADY STATE CONVERGENCE**

MLC-K 240 GB

Test Run Date:		10/26/2011 01:14 PM		**Report Run Date:**		10/27/2011 09:37 AM		
Client IOPS (REQUIRED) - Report Page								
SNIA SSS TWG: Solid State Storage Performance Test Specification (PTS)							Rev.	**PTS-C 1.0**
							Page	**1 of 6**
Device Under Test (DUT)		**VENDOR: ABC CO.**		**SSD MODEL NO: MLC-K 240 GB**		**TEST SPONSOR**	CALYPSO Systems	
Serial No.	0000-0011-FFFF	**DUT Preparation**		**Test Loop Parameters**			**Steady State**	
Firmware Rev	BBFO	**Purge**	Security Erase	**REQUIRED:**			**Convergence**	YES
Capacity	240 GB	**Pre-Conditioning**		**Data Pattern**		RND	**Rounds**	14-18
NAND Type	MLC	**Workload Independent**	2X SEQ/128KiB	**Tester's Choice:**			**PC AR**	100%
Device I/F	6 Gb/s SATA			**OIO/Thread (QD)**		16	**AR AMOUNT**	16 GiB
Test Platform	RTP 2.0 CTS 6.5	**Workload Dep.**	Full IOPS Loop	**Thread Count (TC)**		2	**AR Segments**	2048
Client Steady State Convergence Plot – All Block Sizes								

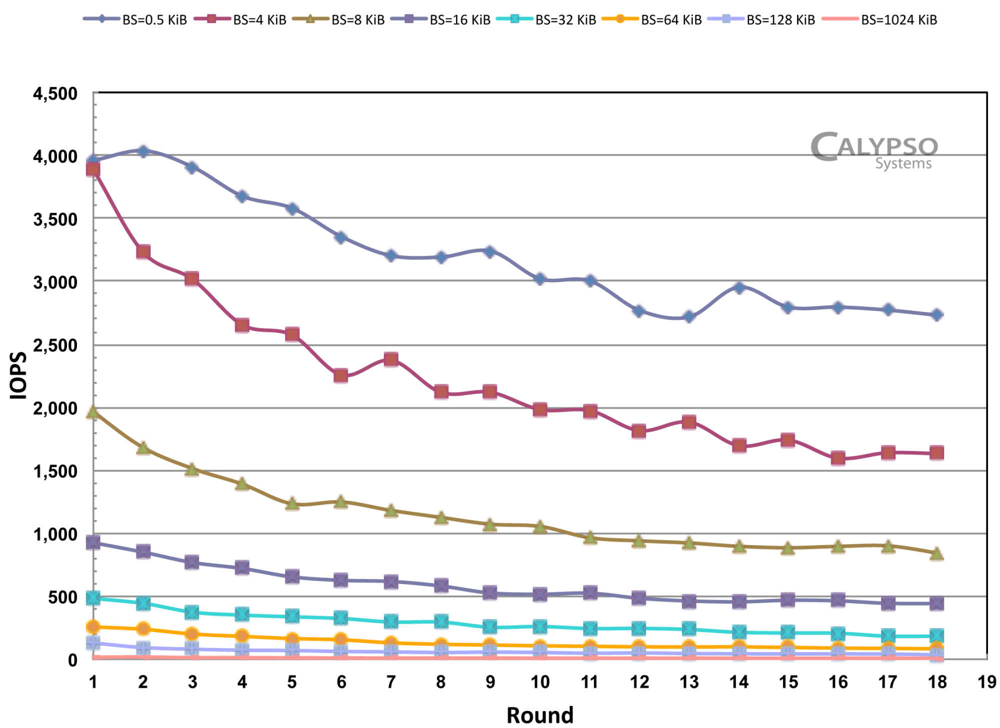

Figure 7. Steady State Convergence Plot IOPS

10. Steady State Measurement Window

The PTS Steady State test reports (IOPS, TP and LAT) present two Steady State Measurement plots that demonstrate adherence to the Steady State criteria. If Steady State was achieved, this plot shows the tracking variable (which can be different for different PTS tests) as a function of all Rounds, including the 5 Rounds that have been determined to be at Steady State.

In the PTS-C IOPS example below (Figure 8) the RND 4KiB IOPS meets the PTS Steady State criteria in Rounds 14-18. This information is reflected both on the plot (as the last five Rounds where various SS metrics have been plotted along with the tracking variable) as well as in the Header under "Steady State Convergence - Rounds" (purple box Figure 7).

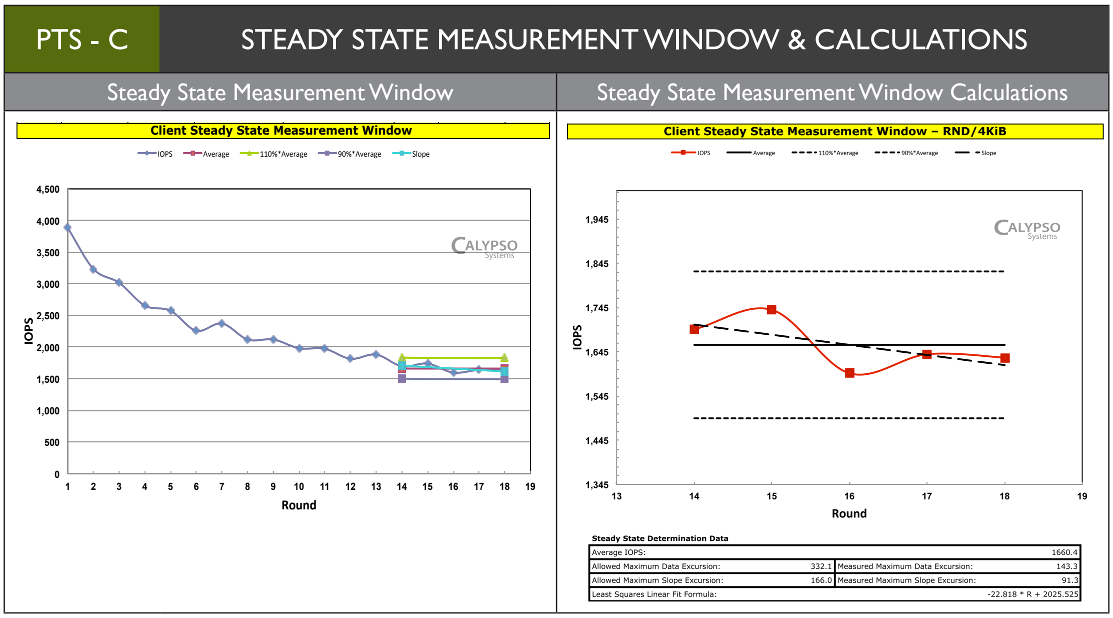

Figure 8 & 9. Steady State Window & Measurement Calculations

11. Steady State Measurement Window Calculations

The PTS Steady State test reports also provide details of the Steady State region and show the tracking variable versus Rounds for the 5 Rounds that meet the PTS Steady State criteria.

The IOPS Steady State plot in Figure 9 above expands rounds 14-18 and shows the RND 4KiB tracking variable (shown as red data markers with curve fit lines), the average of the five IOPS Rounds (black solid line), the least mean squares linear fit (long black dash lines), and the data excursion bounding lines at 110% and 90% of the average IOPS (short black dash lines). Detailed Steady State determination parameters are also listed at the bottom of the page.

The corresponding plots for TP and LAT will plot the tracking variable (TP in MB/S and Latency in mSec - where "mSec" is milliseconds) versus the 5 Rounds within the Steady State Measurement Window.

12. Test Data Results Reporting - Tables, Plots, Charts

For the Steady State test reports for IOPS, TP and LAT, the detailed Steady State Convergence and Measurement plots are followed by various data tables, plots and charts associated with each test. Detailed examples are presented under each test section below. These charts are excerpted from the SNIA PTS Report pages and are reproduced, on occasion, without the Headers for presentation clarity.

PTS-E WSAT

The Write Saturation (WSAT) test is intended to show evolution of continuous RND 4KiB W performance over Time and Total GiB Written (TGBW). The first WSAT plot is IOPS (y linear scale) v Time (x linear scale) while subsequent plots show IOPS v TGBW (x linear scale). The WSAT test specifies that the RND 4KiB Writes be performed from FOB over a 24 hour period or 4X the user capacity, whichever is achieved first. The test begins with a PURGE followed by continuous RND 4KiB Writes.

Note: Test time may be increased for comparison or plotting clarity.

PTS - E	WSAT PLOTS - 100 GB SLC			
	IOPS v TIME		IOPS v TGBW	

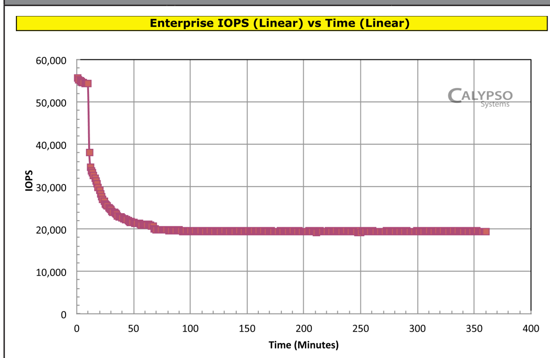

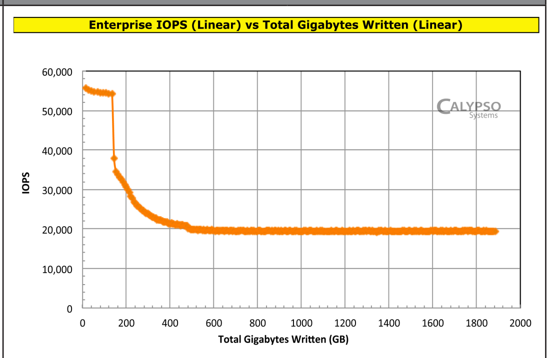

RND 4KiB W IOPS	PEAK FOB RND 4KiB W	55,677	STEADY STATE RND 4KiB W	19,415
	Time to Steady State	70 Minutes	TGBW to Steady State	500 GB

Figure 10. WSAT Plots Time v TGBW

Figure 10 above shows the IOPS evolution from FOB as a function of Time (left panel) as well as TGBW (right panel). Near Time or TGBW=0, the RND 4KiB performance right after PURGE show high IOPS (~56,000 IOPS) that is only sustained momentarily. Then the performance goes through a Transition Region and reaches Steady State where the IOPS is fairly time-invariant at around 19,000. In terms of TGBW, this drive's IOPS reaches Steady State at around 500GB, while the enhanced performance region lasted for less than 200GB.

This set of charts allows the user to see where the various vendor claims of "maximum IOPS" occur, and how long such performance is likely to sustain. One can see where the Steady State performance is relative to the claimed "maximum IOPS."

Note: For the WSAT test, the prescription for the test procedure is to either run continuously for 24 hours, or when 4X the user capacity is reached. There is no "Steady State" determination that is used for this test. Steady State is determined by inspection where IOPS are substantially time invariant.

Note: This drive was tested for 1.9 TB for comparison to a different 400 GB SSD with a 4X user capacity of 1.6TB.

PTS-E & PTS-C IOPS

The IOPS test measures Random performance at various Block Sizes from 0.5KiB to1024KiB, using a range of R/W mixes from 0/100 to 100/0. Figure 11, Page 4 of the IOPS test, is a tabular summary of all R/W & BS combinations measured. From this table, the reader can easily select and review the R/W / BS combinations of interest.

PTS - C — IOPS Table - R/W Mix x BS

MLC-A 256 GB

Test Run Date:	11/14/2011 12:39 AM	Report Run Date:	11/21/2011 04:12 PM

Client IOPS (REQUIRED) - Report Page

SNIA SSS TWG: Solid State Storage Performance Test Specification (PTS)				Rev.	PTS-C 1.0
				Page	4 of 6
Device Under Test (DUT)	VENDOR: ABC CO.	SSD MODEL NO: MLC-A 256 GB	TEST SPONSOR	CALYPSO Systems	

Serial No.	0000-0000-FFFF	DUT Preparation		Test Loop Parameters		Steady State	
Firmware Rev	BF01	Purge	Security Erase	REQUIRED:		Convergence	YES
Capacity	256 GB	Pre-Conditioning		Data Pattern	RND	Rounds	1-5
NAND Type	MLC	Workload Independent	2X SEQ/128KiB	Tester's Choice:		PC AR	100%
Device I/F	6 Gb/s SATA			OIO/Thread (QD)	8	AR AMOUNT	16 GiB
Test Platform	RTP 2.0 CTS 6.5	Workload Dep.	Full IOPS Loop	Thread Count (TC)	1	AR Segments	2048

Client IOPS - ALL RW Mix & BS – Tabular Data

Block Size (KiB)	Read / Write Mix % 0/100	5/95	65/35	50/50	35/65	95/5	100/0
0.5	1,122.3	1,162.2	1,654.6	1,965.6	2,717.7	11,970.0	29,860.1
4	3,147.0	2,896.6	3,044.4	3,454.4	3,779.3	13,005.8	29,876.3
8	1,584.9	1,589.7	2,055.0	2,238.9	2,898.1	11,568.2	21,723.1
16	765.8	786.3	1,028.1	1,272.6	1,604.9	6,208.3	12,482.5
32	392.7	401.0	525.8	652.7	963.8	4,129.6	7,011.6
64	196.4	205.9	291.3	352.3	565.4	2,372.7	3,791.5
128	92.5	97.1	139.9	185.4	377.9	1,410.2	2,015.3
1024	16.4	16.5	23.3	27.3	90.8	191.4	266.7

Figure 11. IOPS Table

PTS-E & PTS-C IOPS have different settings. It is important to refer to the Report Header to ascertain which test settings are applicable. The PTS-E IOPS test requires the use of 100% of the user capacity (100% ActiveRange) for testing, i.e., both WIPC and WDPC use the drive's entire available user capacity. In PTS-C IOPS tests, additional information stating which of the required ActiveRange settings (100% or 75%) and ActiveRange Amount (8GiB or 16GiB), along with number of segments used (2048) is required.

PTS - C	IOPS Plots - 256 GB MLC			
2D IOPS - BS x R/W MIX			3D IOPS BAR PLOT - BS x R/W MIX	

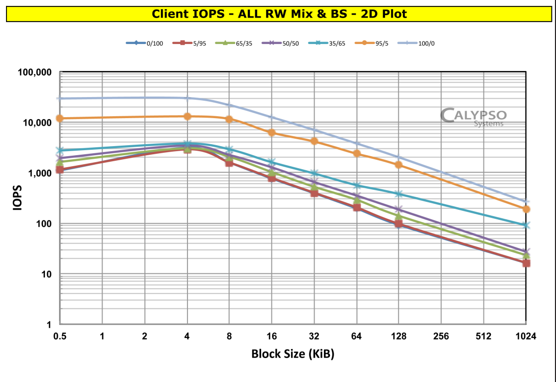

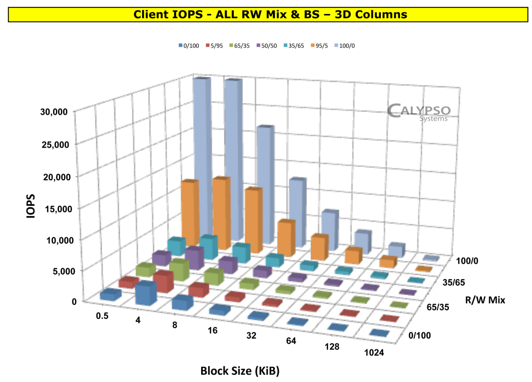

STEADY STATE IOPS	RND 4KiB 100% R	29,876	RND 4KiB 65:35 R/W	3,779
	RND 4KiB 100% W	3,147	RND 128KiB 65:35 R/W	378

Figure 12. IOPS 2D Plot & 3D Bar Plot

PTS-E & PTS-C TP

The Throughput (TP) test is intended to show large block SEQ IOs at Steady State measured in MBs per Second. The PTS-E TP requires 2 Block Sizes, 128KiB and 1024KiB, while the PTS-C TP Test only requires 1024KiB. Two separate TP reports are generated for the PTS-E TP at the 2 Block Sizes. Figure 13 below is the Throughput Table while Figure 14 is the Throughput 2D Bar Plot.

PTS - E	THROUGHPUT Tables - 100 GB SLC			
	SEQ 128 KiB		SEQ 1024 KiB	

SEQ 128 KiB

Test Run Date:	12/03/2011 04:32 PM	Report Run Date:	12/04/2011 10:04 AM					
Enterprise Throughput Test (REQUIRED) - Report Page								
SNIA SSS TWG: Solid State Storage Performance Test Specification (PTS)							Rev.	PTS-E 1.0
							Page	4 of 5
Device Under Test (DUT)		VENDOR: ABC CO.	SSD MODEL NO: SLC-A 100 GB		TEST SPONSOR		CALYPSO Systems	
Serial No.	1111-0000-FFFF	DUT Preparation		Test Loop Parameters			Steady State	
Firmware Rev	BFOA	Purge	Security Erase	REQUIRED:			Convergence	YES
Capacity	100 GB	Pre-Conditioning			Data Pattern	RND	Rounds	4-8
NAND Type	SLC	Workload Independent	2X SEQ/128KiB	Tester's Choice:			PC AR	100%
Device I/F	6 Gb/s SATA				OIO/Thread (QD)	16	AR AMOUNT	100%
Test Platform	RTP 2.0 CTS 6.5	Workload Dep.	SEQ 128KiB		Thread Count (TC)	2	AR Segments	N/A
Enterprise 128KiB Throughput - ALL RW Mix & BS – Tabular Data								

Block Size (KiB)	Read / Write Mix %	
	0/100	100/0
128	144.5	409.3

SEQ 1024 KiB

Test Run Date:	12/04/2011 08:21 AM	Report Run Date:	12/04/2011 10:03 AM					
Enterprise Throughput Test (REQUIRED) - Report Page								
SNIA SSS TWG: Solid State Storage Performance Test Specification (PTS)							Rev.	PTS-E 1.0
							Page	4 of 5
Device Under Test (DUT)		VENDOR: ABC CO.	SSD MODEL NO: SLC-A 100 GB		TEST SPONSOR		CALYPSO Systems	
Serial No.	1111-0000-FFFF	DUT Preparation		Test Loop Parameters			Steady State	
Firmware Rev	BFOA	Purge	Security Erase	REQUIRED:			Convergence	YES
Capacity	100 GB	Pre-Conditioning			Data Pattern	RND	Rounds	4-8
NAND Type	SLC	Workload Independent	2X SEQ/128KiB	Tester's Choice:			PC AR	100%
Device I/F	6 Gb/s SATA				OIO/Thread (QD)	16	AR AMOUNT	100%
Test Platform	RTP 2.0 CTS 6.5	Workload Dep.	SEQ 1024KiB		Thread Count (TC)	2	AR Segments	N/A
Enterprise 1024KiB Throughput - ALL RW Mix & BS – Tabular Data								

Block Size (KiB)	Read / Write Mix %	
	0/100	100/0
1024	157.5	514.3

Throughput MB / Sec	SEQ 128KiB 100% W	SEQ 128KiB 100% R	SEQ 1024KiB 100% W	SEQ 1024KiB 100% R
	145	409	158	514

Figure 13. PTS - E TP Tables

Note: PTS-C TP tests require a total of 4 test runs: ActiveRange 75% and 100%, each with ActiveRange Amount settings of 8GiB and 16GiB at a single Block Size of 1024KiB. The PTS-E TP requires 2 test runs: ActiveRange of 100% for 2 Block Sizes of 128KiB and 1024KiB.

PTS - E	THROUGHPUT Plots - 100 GB SLC			
	SEQ 128 KiB		SEQ 1024 KiB	

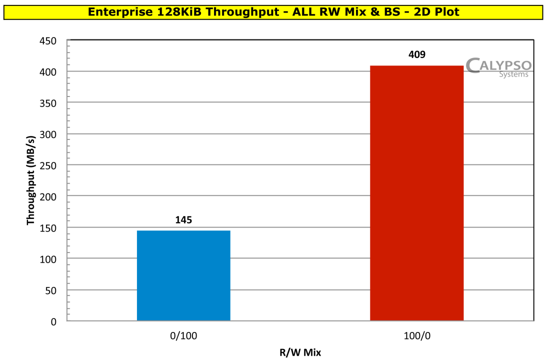

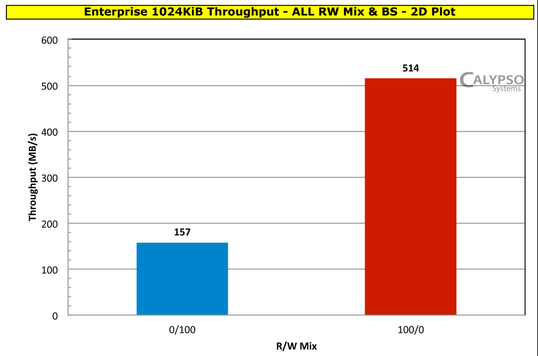

Throughput MB / Sec	SEQ 128KiB 100% W	SEQ 128KiB 100% R	SEQ 1024KiB 100% W	SEQ 1024KiB 100% R
	145	409	157	514

Figure 14. PTS-E TP 2D Plots

PTS-E & PTS-C LAT

The LAT test reports MAX and AVE Response Times with a total outstanding IO setting of 1. Steady State results are measured in mSec. In Figure 15, PTS-C Latency, Steady State was reached after 8 rounds with the SS Window in rounds 4 – 8.

PTS - C	LATENCY Table - R/W Mix x BS
	MLC-A 256 GB

Test Run Date:	11/11/2011 09:53 AM	Report Run Date:	11/15/2011 03:34 PM

Client Latency (REQUIRED) - Report Page

SNIA SSS TWG: Solid State Storage Performance Test Specification (PTS)	Rev.	PTS-C 1.0
	Page	4 of 6

Device Under Test (DUT)	VENDOR: ABC CO.	SSD MODEL NO: MLC-A 256 GB	TEST SPONSOR	CALYPSO Systems

Serial No.	0000-0000-FFFF	DUT Preparation		Test Loop Parameters		Steady State	
Firmware Rev	BF01	Purge	Security Erase	REQUIRED:		Convergence	YES
Capacity	256 GB	Pre-Conditioning		Data Pattern	RND	Rounds	4-8
NAND Type	MLC	Workload Independent	2X SEQ/128KiB	Tester's Choice:		PC AR	100%
Device I/F	6 Gb/s SATA			OIO/Thread (QD)	1	AR AMOUNT	16 GiB
Test Platform	RTP 2.0 CTS 6.5	Workload Dep.	Full Latency Loop	Thread Count (TC)	1	AR Segments	2048

Client - AVE and MAX Response Time - ALL RW Mix & BS – Tabular Data

Average Response Time (ms)

	Read / Write Mix %		
Block Size (KiB)	0/100	65/35	100/0
0.5	0.932	0.449	0.196
4	0.349	0.332	0.203
8	0.537	0.432	0.240

Maximum Response Time (ms)

	Read / Write Mix %		
Block Size (KiB)	0/100	65/35	100/0
0.5	52.1	50.2	1.0
4	51.0	50.0	1.6
8	52.8	53.3	39.5

Figure 15. MLC-A LAT Table

The 3D Bar Plots in Figure 16 show the three block sizes and three R/W mixes for AVE and MAX LAT. PTS-C LAT requires separate reports for ActiveRanges of 75% and 100%, using ActiveRange Amounts of 8GiB and 16GiB. The PTS-E LAT uses Active Amount of 100%. Figure 16 below shows PTS-C MAX and AVE Latency for Block Sizes 0.5KiB, 4KiB and 8KiB using R/W Mixes 0/100, 65/35 and 100/0.

PTS - C	LATENCY Plots - 256 GB MLC
AVE LATENCY	MAX LATENCY

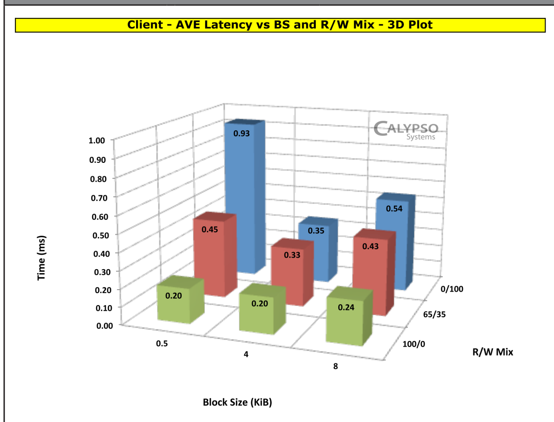

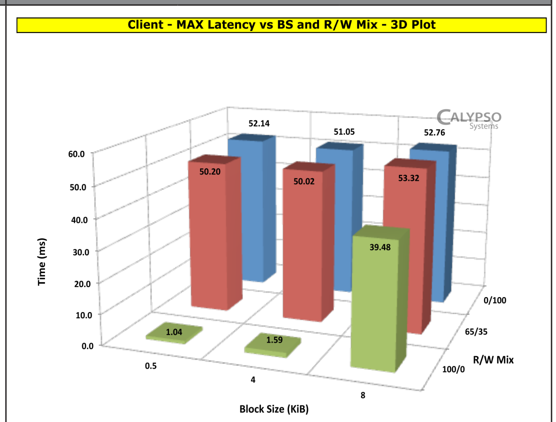

Response Time mSec	RND 4KiB 100%W	RND 4KiB 65:35 R/W	RND 4KiB 100% R	RND 4KiB 100% W	RND 4KiB 65:35 R/W	RND 4KiB 100%R
	0.20	0.33	0.35	1.59	50.02	51.02

Figure 16. MLC-A LAT AVE & MAX

13. Using the PTS to Compare SSD Performance

The PTS Report Format allows for easy comparison of performance. The reader can compare the performance characteristics of SSDs once PTS Reports for WSAT, IOPS, TP and LAT are generated. The same SSD can be tested under varying conditions for comparison or different SSDs can be compared by evaluating PTS test reports.

In this section, several examples are provided that illustrate how the reader or test sponsor may use features of the PTS Report Format to make useful comparisons between test runs, either for different drives (using reports generated by different drives) or for the same drive (with different testing conditions on a single drive).

Steady State Convergence – IOPS Comparison

The Steady State Convergence Plot found in the IOPS Report is useful to the reader because it visualizes a number of important drive characteristics:

1. From a drive that has only been sequentially written, the reader can see how the drive's IOPS evolve as more random data is written to it.

2. Since all of the reporting Block Sizes are represented (including the tracking RND 4KiB), the reader can see at a glance if all the Block Sizes are evolving toward a "Steady State."

3. By focusing on the tracking variable, the reader can get a sense of the quality of Steady State results, for example, if the tracking variable is fluctuating from Round-to-Round, or if the tracking variable shows slowly increasing or decreasing trends.

PTS - C	STEADY STATE CONVERGENCE - MLC IOPS
MLC-A 256 GB	MLC-B 160 GB

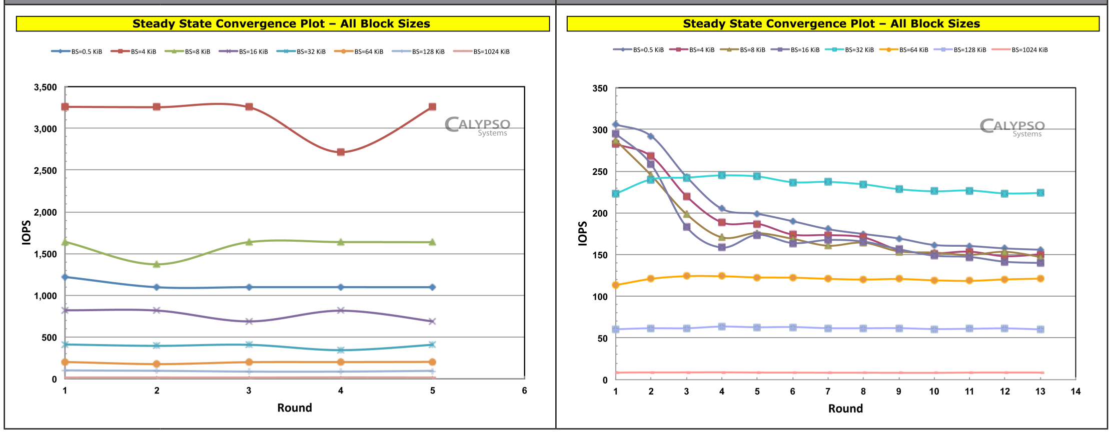

Figure 17. Steady State Convergence Plot - MLC IOPS

In general, SLC SSDs tend to present faster and more stable performance, with less performance difference between the beginning and the end of the Steady State Convergence Plot, and generally takes fewer rounds to reach Steady State.

Figure 17 above shows the Steady State Convergence Plot for two client-class MLC drives. The Steady State for MLC-A is reached between Round 1 and 5 with little difference between maximum and Steady State IOPS. For MLC-B, Steady State is reached between Round 9 and 13 with large relative difference between the maximum and Steady State IOPS for small Block Sizes. In contrast, enterprise-class SLC drives below (Figure 18) show significantly higher IOPS performance overall, with smaller differences between maximum and Steady State IOPS.

PTS - E	STEADY STATE CONVERGENCE - SLC IOPS
SLC-A 100 GB	SLC-B 100 GB

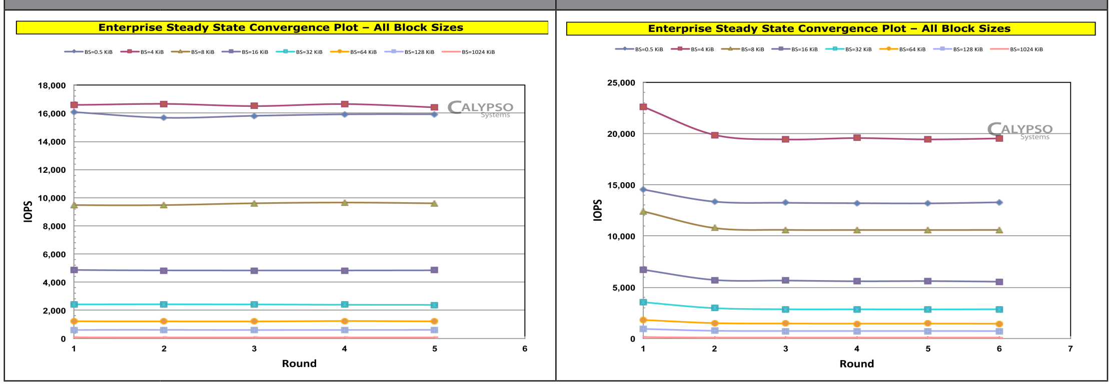

Figure 18. Steady State Convergence Plot - SLC IOPS

Steady State Convergence –AVE LAT W Comparison

Latency SS Convergence Plots shows both the AVE and MAX latencies (in mSec) versus Rounds. The AVE Latency plots are similar to IOPS Steady State Convergence plots in that one can, at a glance, get a feel for overall latency trends for all of the Block Sizes. Figure 19 below presents two different Steady State Convergence plots for two client-class MLC drives. Note the markedly different behavior between the two drives.

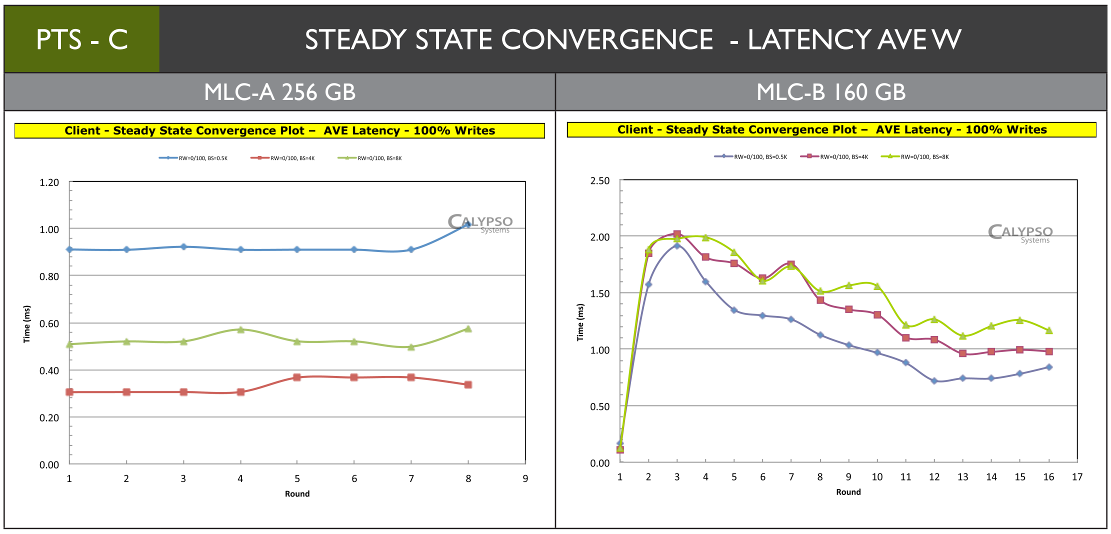

Figure 19. Steady State Convergence Plot - MLC AVE Latency

MAX Latency plots are also useful to spot the existence of processes within the SSD that can sometimes cause spikes in MAX Latency events that may not be seen when looking only at AVE Latencies. Figure 20 below gives two such examples for two client-class MLC drives. The reader should note the larger response times (y-axis) compared to AVE Latency above.

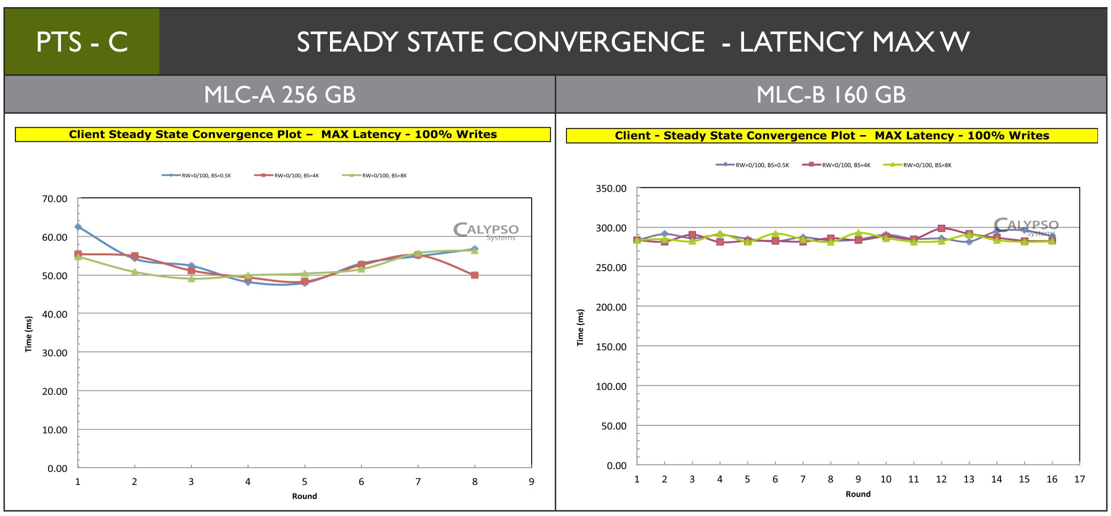

Figure 20. Steady State Convergence Plot - MLC MAX Latency

Steady State Convergence – PTS-E 128K TP W Comparison

SS Convergence Plots for TP track SEQ TP in MB/sec for all Rounds. Each of the TP tests for both PTS-C and PTS-E uses a single block size and is continuously applied to the device from FOB until Steady State.

PTS - E | STEADY STATE CONVERGENCE - TP 128KiB W

SLC-A 100 GB | SLC-B 100 GB

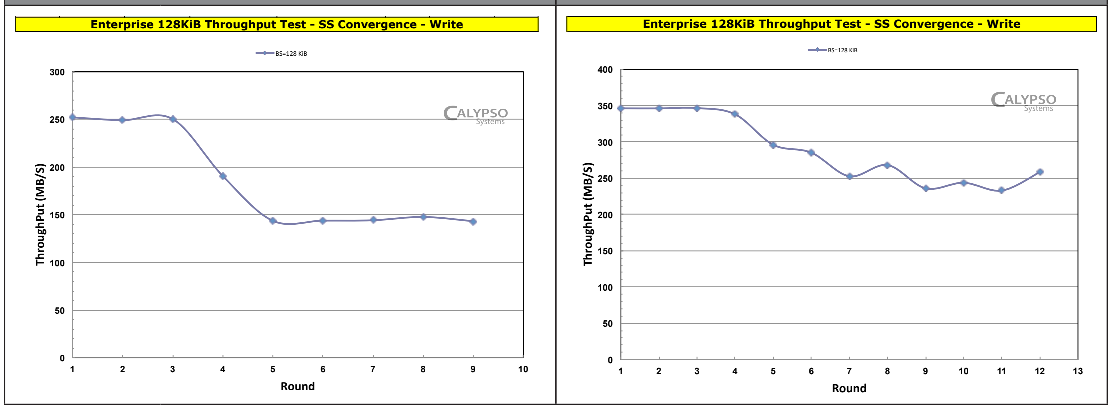

Figure 21. Steady State Convergence Plot - SLC 128KiB Throughput

PTS-E requires running two independent TP tests using both 128KiB and 1024KiB. The PTS-C test requires running a matrix of 75% and 100% ActiveRange settings, each with ActiveRange Amounts of 8GiB and 16GiB. In Figure 21, SLC-A shows peak TP in Rounds 1-3 with a transition to SS in Rounds 4-8 while SLC-B shows a similar peak TP in Rounds 1-3 but with a more gradual progression to SS in Rounds 8-12.

Steady State Measurement Window – IOPS Comparison

The SS Measurement Window plots IOPS and detailed Steady State determination information which allows the reader to verify that Steady State has been achieved according to the PTS Steady State criteria. The reader can also examine the linear curve fit to see if the slope may be caused by IOPS fluctuation or trending of the IOPS.

PTS - C | STEADY STATE WINDOW - IOPS

MLC-A 256 GB | MLC-B 160 GB

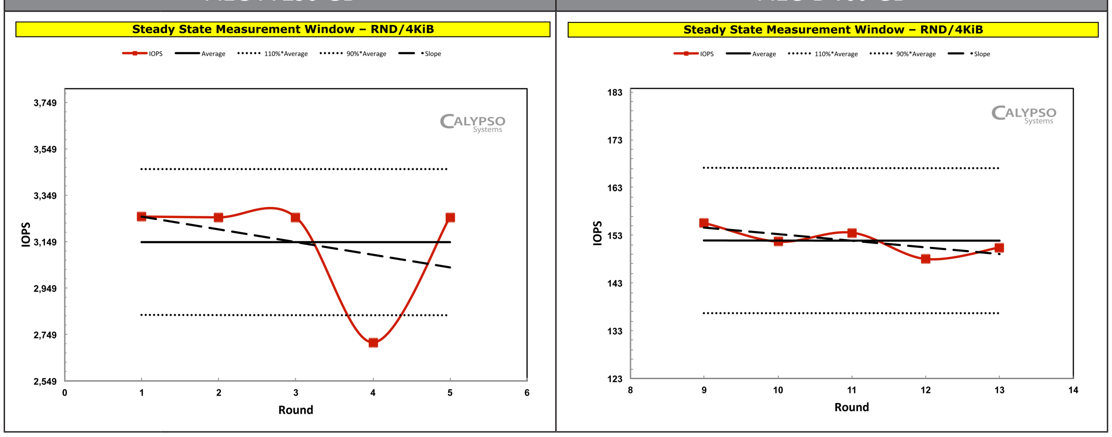

Figure 22. Steady State Measurement Window - MLC IOPS

The dotted bounding Max and Min data excursion lines are shown for illustration purposes – the Steady State criteria calls for the total data excursion to simply fit within this +/ – 10% band. In Figure 22, MLC-A Data point Round 4 dips below the -10% Min data excursion line but still fits within the band defined by +/ -10% of the average. Figure 23 shows SLC drive stability.

PTS - E	STEADY STATE WINDOW - IOPS
SLC-A 100 GB	SLC-B 100 GB

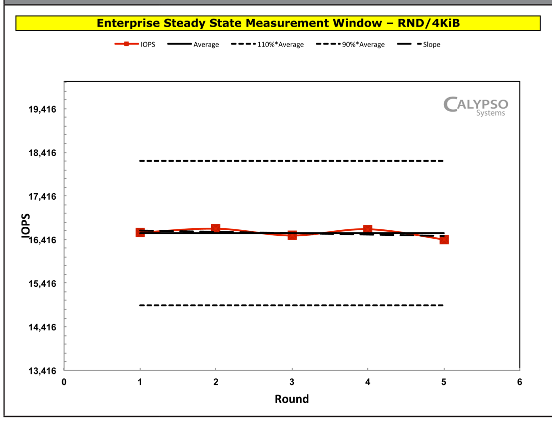

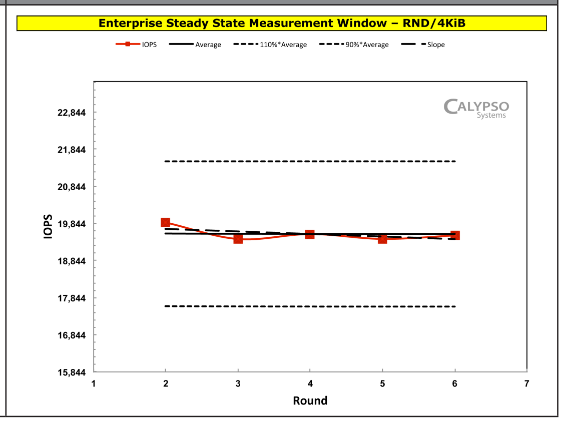

Figure 23. Steady State Measurement Window - SLC IOPS

WSAT – Single SSD Comparison

The WSAT report can be used to quickly evaluate SSD write performance. WSAT shows initial FOB peak performance and IOPS behavior as it evolves over Time and Total GiB Written. The reader can review the transition from FOB to Steady State IOPS, the amount of Time and Total GiB Written during peak FOB performance, the slope and length of the transition zone leading to Steady State IOPS, and the overall performance behavior response to continuous small block RND 4KiB Writes. WSAT plots create a device specific profile that is often discernible among drives and across successive drive releases.

PTS - C	WSAT Plots - MLC 160 GB
IOPS v TIME	IOPS v TGBW

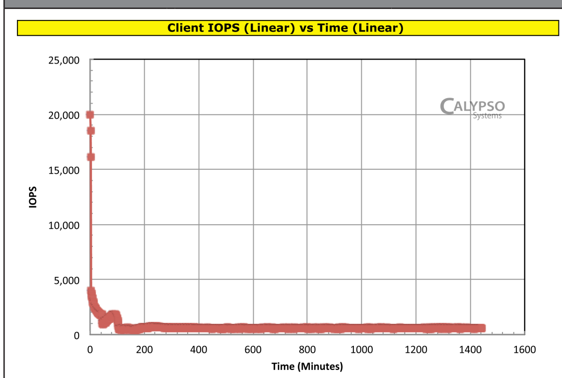

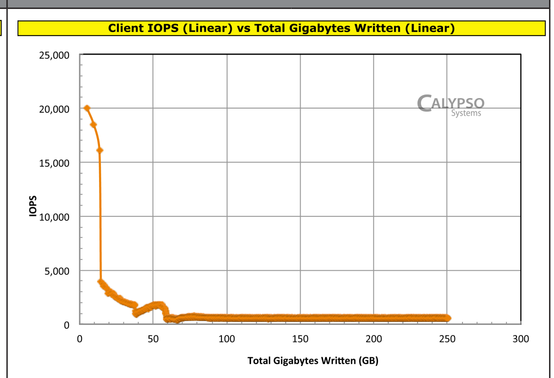

RND 4KiB W IOPS	FOB IOPS	20,364	Steady State IOPS	658
	Time to Steady State	240 Minutes	TGBW to Steady State	75 GB

Figure 24. WSAT MLC-B - Time & TGBW

Note: For a given 24 hour period, a test SSD may not write a full 4 x User Capacity (due to a slow RND 4KiB W speed). Depending upon their transient IO rate period, some WSAT tests are run longer than the required 24 hours or 4 x User Capacity to facilitate comparison between different SSD WSAT characteristics.

PTS - E	WSAT PLOTS - 100 GB SLC		
	IOPS v TIME		IOPS v TGBW

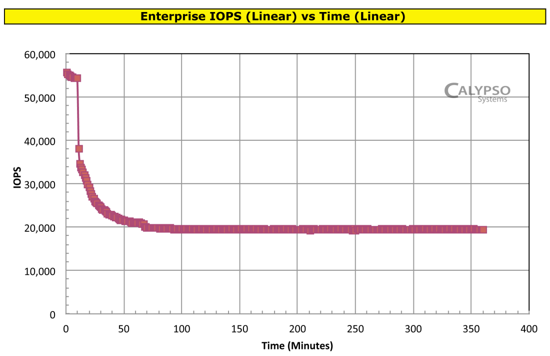

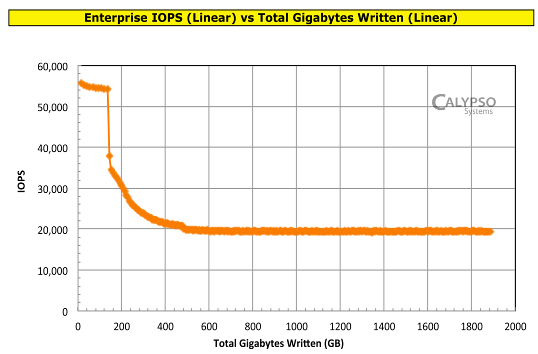

RND 4KiB W IOPS	FOB IOPS	55,677	Steady State IOPS	19,415
	Time to Steady State	70 Minutes	TGBW to Steady State	500 GB

Figure 25. WSAT SLC-B - Time & TGBW

WSAT - SSD Comparisons

One can use WSAT plots to quickly compare the maximum IOPS to vendor-provided metrics and observe how long the drive is able to sustain such maximum IOPS, the rate in which the drives reach Steady State, and the actual performance level of the Steady State IOPS.

PTS - C	WSAT TGBW COMPARISON - MLC	
	MLC-A 256 GB	MLC-B 160 GB

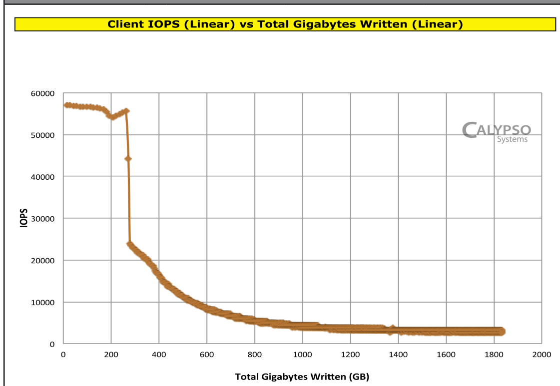

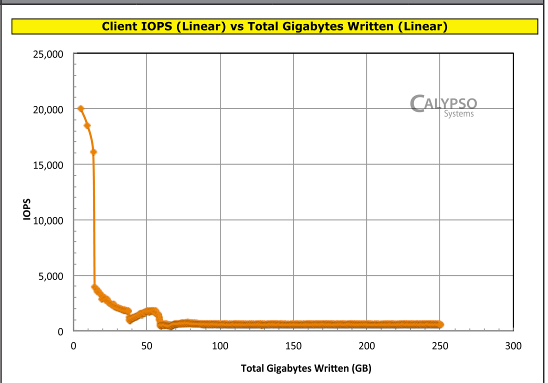

FOB RND 4KiB W	56,896	Steady State RND 4KiB W	2,714	FOB RND 4KiB W	20,364	Steady State RND 4KiB W	658
TIME to Steady State	580 Minutes	TGBW to Steady State	1,280 GB	TIME to Steady State	180 Minutes	TGBW to Steady State	70 GB

Figure 26. WSAT TGBW Comparison - MLC

In Figure 26, SSD MLC-A is able to sustain maximum performance in excess of 200GB, while MLC-B has maximum performance only for a few tens of GBs. The SS IOPS are also markedly different. In Figure 27, SLC SS IOPS are higher and the drop from FOB to SS IOPS is smaller with a shorter Transition State.

PTS - E | **WSAT TGBW COMPARISON - SLC**

SLC-A 100 GB | SLC-B 100 GB

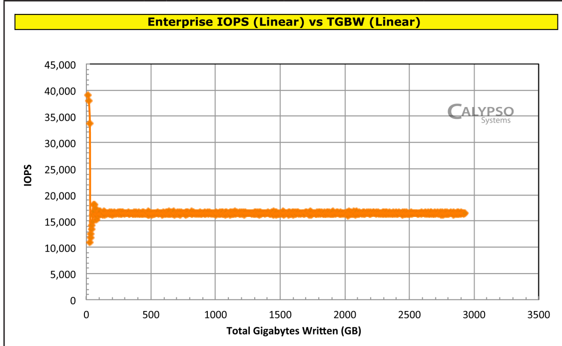

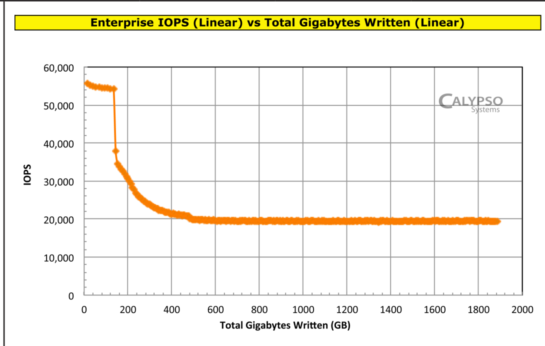

FOB RND 4KiB W	39,092	Steady State RND 4KiB W	16,305	FOB RND 4KiB W	55,677	Steady State RND 4KiB W	19,415
TIME to Steady State	50 Minutes	TGBW to Steady State	150 GB	TIME to Steady State	70 Minutes	TGBW to Steady State	500 GB

Figure 27. WSAT TGBW Comparison - SLC

Note: The test sponsor may elect to re-plot PTS charts to reflect specific metrics of interest, present data for comparison, or for plotting clarity. For example, WSAT TGBW can be plotted against Normalized Capacity i.e. the x-axis is Normalized and expressed as a multiple of the SSD's User Capacity.

IOPS – SSD Comparisons

SS IOPS reports present a 56 element matrix of RND IOs at varying Block Sizes and R/W Mixes. The PTS-C IOPS table allows the reader to quickly select the BS / R/W measurement of interest and to reference the OIO setting, preconditioning rounds to Steady State and data pattern.

PTS - C | **IOPS COMPARISON - MLC**

MLC-A 256 GB | MLC-B 160 GB

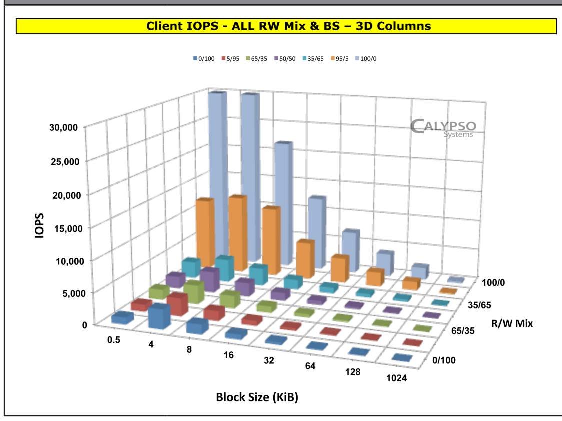

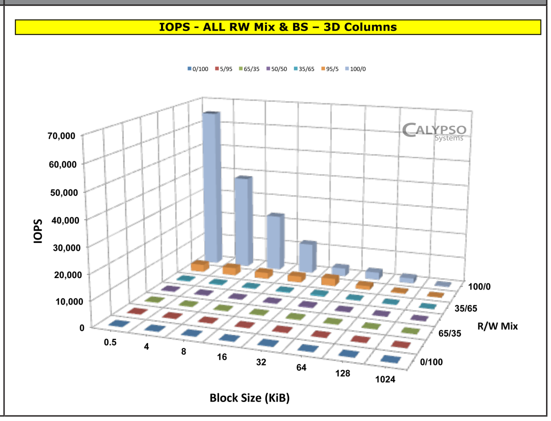

Figure 28. IOPS COMPARISON - MLC

IOPS 3D Bar Plots present a three dimensional representation of the 56 element IOPS BS x R/W matrix which allows the reader to graphically interpret the overall SSD IOPS performance in large and small block RND IOPS. In the IOPS 3D Bar Plot, IOPS are on the y-axis, Block Size on the x-axis and R/W Mix is along the z-axis.

PTS - E	IOPS COMPARISON - SLC
SLC-A 100 GB	SLC-B 100 GB

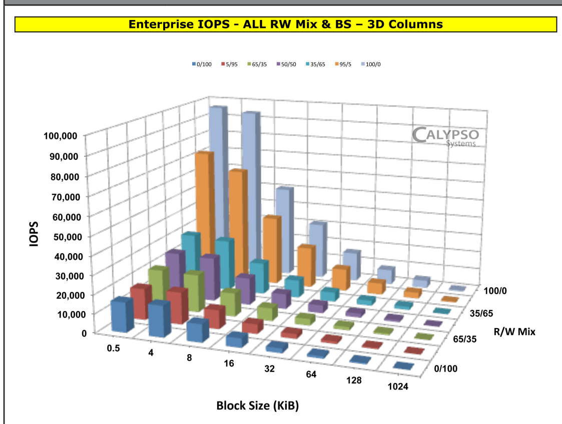

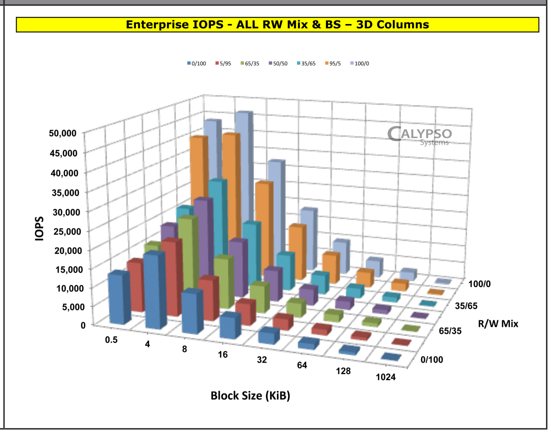

Figure 29. IOPS COMPARISON - SLC

Note: The R/W mix axis in the 3D IOPS plot represents the R/W ratio as categories. The distance between the R/W mixtures are not scaled according to either the Read % or Write %. Thus, the reader is cautioned when attempting to visually interpolate results between the data points given.

The region to note is the left-front side for small block RND W IOPS – the area of key differentiation for most SSDs. The reader should also take note of the IOPS scale when comparing plots. In general, drives achieving good R and W performance parity will show comparable height bars in the "depth or z" direction. The characteristic "waterfall" effect can be observed in the 3D plot going from left to right in Figure 29.

TP – SSD Comparisons

TP-C Reports provide a comparison of large block SEQ Throughput. The PTS TP Tabular data compares large block R and W. Using similar test conditions and using results from two different drives will quickly allow the reader to compare drive performances. Figure 30 shows PTS-C TP at 1024KiB for two client-class MLC drives.

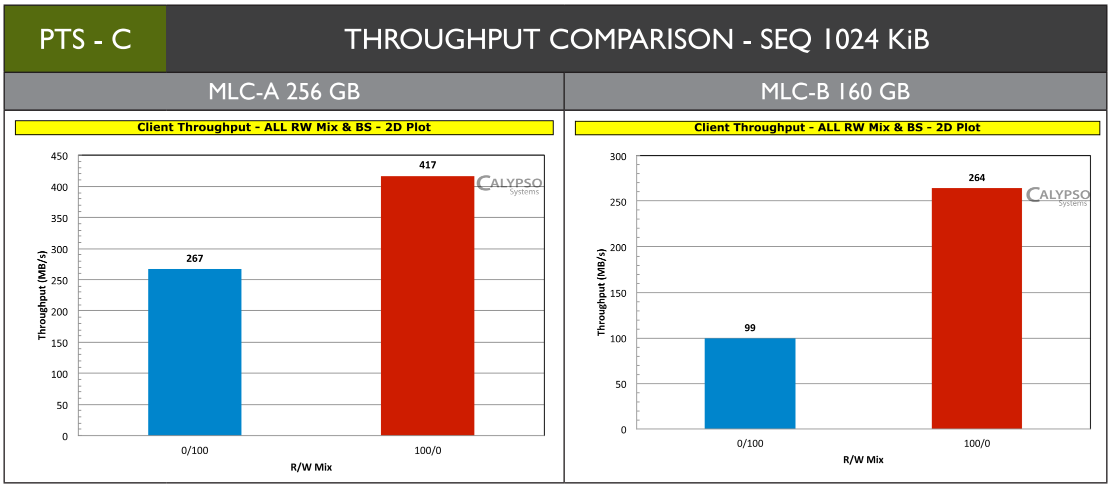

Figure 30. Throughput Comparison - SEQ 1024 KiB MLC

PTS-C TP reports SEQ 1024KiB 100% R and 100% W in MB/sec. The reader and test sponsor should note the System Interface and Device Interface reported in the PTS Report Header to determine if the results are limited by hardware or interface speeds.

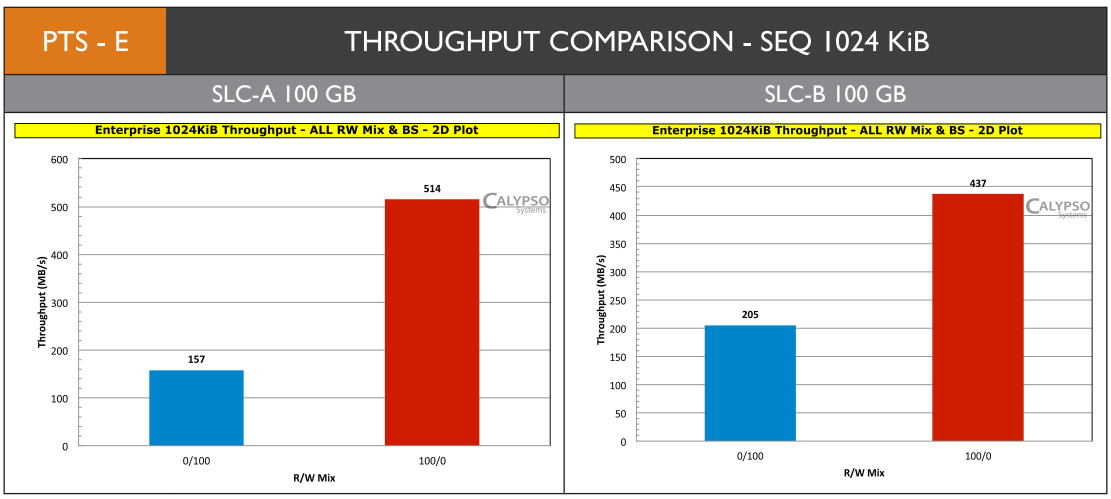

Figure 31. Throughput Comparison - SEQ 1024 KiB SLC

PTS-E TP calls for two separate TP runs using both SEQ 128KiB and SEQ 1024KiB in 100% R and 100% W. An example of PTS-E TP result for SEQ 1024KiB TP test only is presented in Figure 31 above. The test sponsor may combine plots to present results on a single chart if desired.

LAT AVE – SSD Comparisons

Latency Reports present both AVE and MAX Latencies for the three required Block Sizes. The reader can refer to either Latency tabular data or 3D plots. Here, AVE Latency plots are presented. Latency chart data series labels have been added for white paper presentation.

PTS - C	LATENCY COMPARISON - AVE Latency
MLC-A 256 GB	MLC-B 160 GB

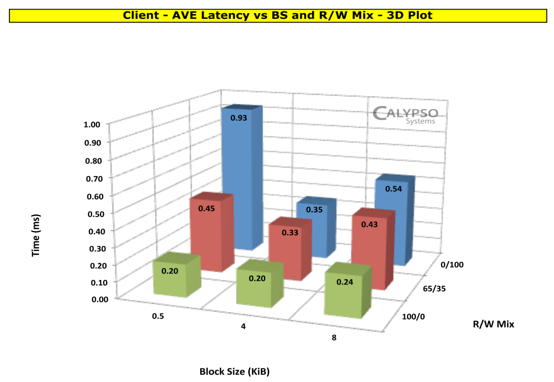

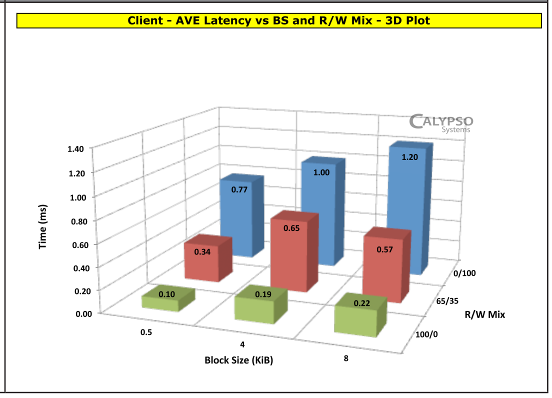

Figure 32. Latency Comparison - MLC AVE LAT

Average Latency shows the average response time for all IOs that are completed within the measurement period of 1 minute at a total OIO of 1. Note that this is the inverse of IOPS when total OIO=1.

PTS - E	LATENCY COMPARISON - AVE Latency
SLC-A 100 GB	SLC-B 100 GB

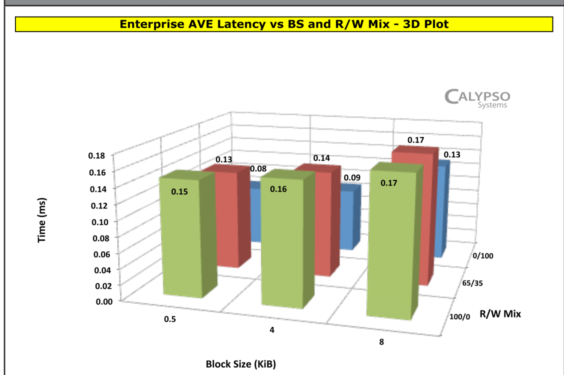

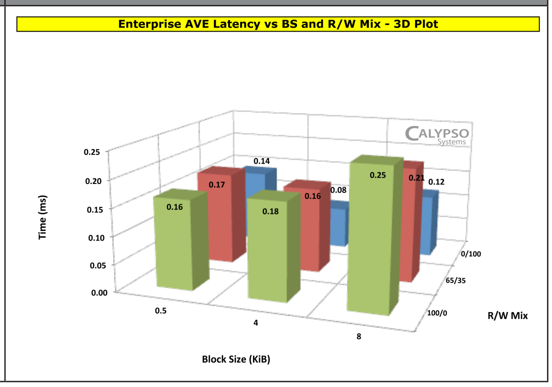

Figure 33. Latency Comparison - SLC AVE LAT

Note: Due to generally higher single OIO IOPS observed in SLC drives, SLC AVE Latencies are correspondingly smaller than MLC AVE Latencies. Again, the reader is cautioned to note the y-axis scale of Response Time when comparing charts.

LAT MAX – SSD Comparisons

Maximum Latency shows the maximum response time for all IOs within the measurement period of 1 minute at a total OIO of 1.

PTS - C | **LATENCY COMPARISON - MAX Latency**

MLC-A 256 GB | MLC-B 160 GB

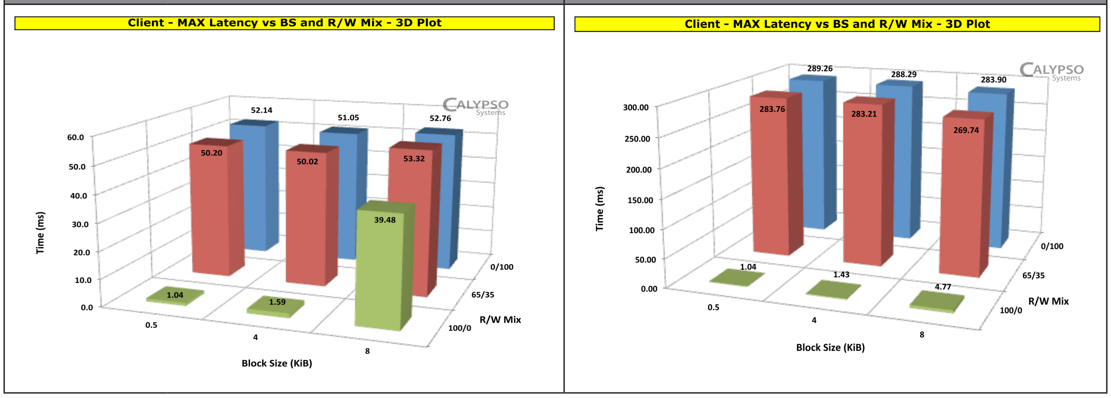

Figure 34. Latency Comparison - MLC MAX LAT

Excessively long MAX Latency Response Times may indicate issues with the drive firmware's ability to handle IOs consistently, e.g. difficulty with handling background task scheduling under continuous load conditions.

PTS - E | **LATENCY COMPARISON - AVE Latency**

SLC-A 100 GB | SLC-B 100 GB

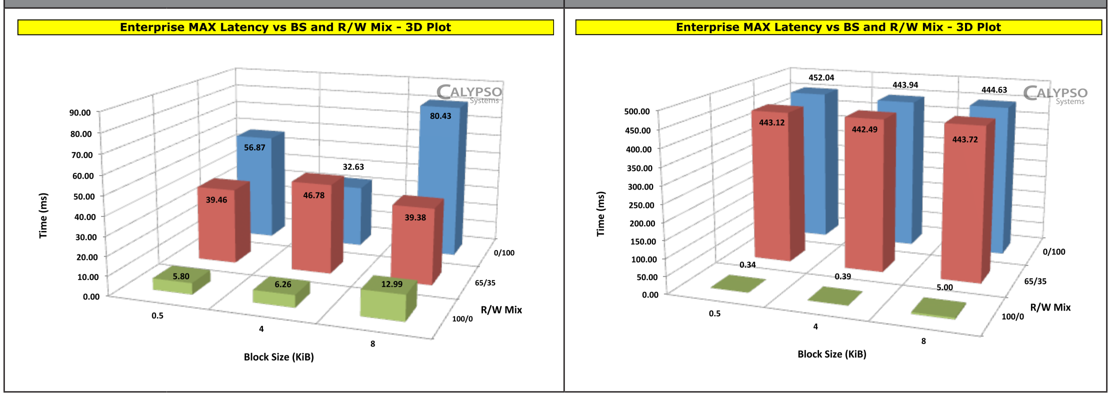

Figure 35. Latency Comparison - SLC MAX LAT

Note that AVE and MAX Latency information provided with the current PTS Latency Tests provides scalar values of AVE and MAX Latency numbers over the observation period. The test sponsor may have an interest in observing the frequency and distribution of each of the individual IO's Response Times within the measurement period. For example, if a drive shows a RND 4KiB W Latency of 5 mSec, one may be interested in knowing what percentage of the total IOs are < 5 mSec.

These, and other tests, are under consideration by the SNIA SSS Technical Working Group (TWG) and may be issued as tests in future versions of the PTS.

14. SSD Test Best Practices

The test and measurement of NAND Flash-based solid state storage performance is highly dependent on the test environment, test stimulus and test methodology. In order to obtain relevant, accurate, reliable and reproducible performance test results, the reader / test sponsor should take care to incorporate "good test procedures" and "SSD Test Best Practices."

While the efficacy of any specific test practice depends on the goals and objectives of the particular test plan, the following SSD Test Best Practices can be utilized in basic test procedures as well as in SSD specific testing.

Basic Test Procedures

Test Hardware & Operating System (OS). Care should be taken in the selection of the test hardware and software to ensure that the test environment does not bottleneck SSD performance nor otheR/ Wise hinder the test stimulus or measured data response.

- Hardware Bottlenecks can occur anywhere within the data/control paths from SSD device interface (e.g. SAS/SATA HBA connection) to the motherboard data bus lanes, RAM and CPU.
- Software Influences can include background processes occurring in the OS, software applications, APIs, and device drivers. To the extent possible, OS background tasks and application software should be terminated and only a single SSD should be tested at a time.
- Test Software Tools are also critical in taking SSD performance measurements. Care should be taken to understand the overhead of the test tools and the effects of the stimulus generator.

Normalized Test Platform. When evaluating and managing the effects of the hardware and software environment on test execution and measurements, the test sponsor should strive to take and compare test results using the SAME test environment - hardware, OS and test tools. By using the identical or equivalent test platform, the impact of the test environment can be normalized across test measurements.

Calibration. Once the test environment is selected, periodic calibration using the same test stimulus / workload on a known device, or the use of a "golden" reference test SSD and test procedure, should be used to ensure the repeatability and reliability of the test measurements.

Test Plan. A good test plan enumerates test objectives, test methodology and selection of tests. This includes establishing the relevance of the test to the test objectives (see "Test Stimulus Workload" below), defining the test baseline, and prescribing the test procedures, number of test samples, test runs and statistical analysis employed.

SSD Specific Testing

Purge. Any SSD test should begin with a device Purge. This white paper has demonstrated the significant effect that write history and workloads have on SSD performance. The test sponsor should ensure that the Test Plan prescribes use of a relevant Purge (Security Erase for SATA, Format Unit for SCSI or other proprietary command specified by the drive controller vendor) to ensure that the drive is put into a state "as if no writes had occurred."

Preconditioning. The effects of preconditioning on SSD performance is well documented here and in other works. Care must be taken to ensure that the preconditioning regime is well defined and targeted to the purpose of the particular SSD performance tests (e.g. see "Block Size Sequencing" below).

Steady State. It has been demonstrated that a single SSD can exhibit many "steady states" depending on the write history, workload and definition of steady state. SSDs have been shown to demonstrate behaviors such as RND and SEQ Write Saturation where performance degrades over time until the device reaches a relatively stable "steady state." The definition of both preconditioning and steady state is a key determinant in any SSD performance metric and measurement.

Demand Intensity. SSD performance under a given workload (e.g. RND 4KiB Writes or other specific Access Pattern of Block Size and Read / Write mix) can change depending on the driving intensity of the host test system as measured by OIO (Outstanding IOs). The test sponsor should take the time to map the target SSD on the test platform to determine the optimal OIO settings for the given test workload access pattern.

Block Size Sequencing. Previous analysis of Cross Stimulus Recovery (see section 2) shows the effect of Block Size Sequence on SSD performance. Care should be taken to ensure that the preconditioning and workload stimulus do not introduce unwanted or unanticipated Block Size Sequence Cross Stimulus effects on the SSD performance.

Test Stimulus Workload. While the use of synthetic device level tests allows the test sponsor to achieve repeatable and reliable test measurements, care must be taken to ensure that the prescribed test stimulus workload is relevant to the characteristics of the test sponsor's targeted user workload.

SSD Test Best Practices

Standardized Methodologies. Employment of standardized test methodologies ensures the test sponsor will benefit from the investigation and development of SSD tests by industry technical working groups and other scientific and academic bodies.

Reference Test Platform (RTP). Use of a Reference Test Platform (as defined in the SNIA SSS PTS) can help normalize the test environment as well as ensure repeatable/reproducible and comparable SSD performance test results.

Standardized Tests. Use of standardized tests (in conjunction with an RTP and use of the PTS) allows for easy comparison of performance between different SSD devices.

Standardized Reporting. It is important to report the Test Environment, Test Settings and Test Measurements in a standardized format. This will ensure that the testing is performed in compliance with stated standards and ensure disclosure of test set-up and specific tests associated with a particular set of test measurements.

Use of SNIA SSS PTS. Use of an industry standard SSD Performance Test methodology, such as the SNIA SSS PTS, allows test sponsors and readers to benefit from the body of industry work undertaken to understand and evaluate NAND Flash-based SSD performance. Test sponsors, end users and SSD vendors can benefit from the uniform prescriptions for SSD performance testing that allow for a quicker comparison and understanding of SSD device performance.

15. Conclusion

The primary purpose of this white paper is to assist the reader when evaluating and comparing SSD performance using the SNIA SSS PTS and SNIA Standard Reporting Format test reports.

Although NAND storage technology is mature, having found its way into everyday life (thumb drives, music players, and the like) it is now migrating into the traditional storage IO stack and enabling performance that to date was unheard of.

Despite the familiar "drop-in replacement" form factors of NAND based SSDs, their performance characteristics are considerably different from those of conventional spinning drives. As a result, new performance testing practices, methodologies, metrics and data consolidation and presentation techniques are required to enable their accurate, objective performance comparison.

First, the performance of these devices is so much greater than traditional storage that the test environments themselves can adversely affect test results. Whether measuring Throughput (Megabytes per Second), Input/Output Operations Per Second (IOPs), or Latency (mSec) the difference relative to rotating drives can be orders of magnitude. Hence, the test platforms themselves require a new level of performance and robustness. This, in turn, places greater requirements on the types of systems, HBAs, operating systems, and stimulus generator and measurement tools.

Second, NAND based SSDs are very "write history" sensitive: the loads to which they have been subjected can have a substantial effect on drive performance – in many cases far more so than the current IO demand. This characteristic requires very precise preconditioning to achieve a true steady state performance measure. Similarly, other parameters such as mis-aligning the IO transfer boundaries can exhibit hysteresis effects which must be allowed to settle out of the measurement interval, necessitating the need to evaluate cross stimulus characteristics.

Finally, an end user's workload is as diverse as the environment in which the SSD is to be placed. For this reason, the PTS facilitates the comparison of SSD performance under a wide variety of workloads and demand intensities. The end user, knowing the attributes of their particular IO profile, can select those test results which best represents their environment and disregard those less relevant.

Editors Note:

The PTS documents may be downloaded at www.snia.org/pts. The reader is also encouraged to visit SNIA TWG portal www.snia.org/publicreview download draft PTS specifications open for public review and to submit comments on drafts.

Further, the reader may visit the SNIA SSSI website at www.snia.org/forums/sssi to download this white paper, to view summary sample PTS results and access other areas of interest. For more information on joining the SSSI, please visit www.snia.org/forums/sssi/about/join.

About the Authors

The authors of this white paper have been intimately involved with the development and execution of the SNIA Solid State Storage Performance Test Specification. This experienced and authoritative group includes members and chairs of key industry and SNIA Technical Groups including:

- SNIA Solid State Storage Technical Working Group (SSS TWG) – Chair and principal members
- SNIA IO Traces, Tools and Analysis Technical Working Group (IOTTA) – Co-Chair
- SNIA Solid State Storage Initiative – Governing Board and members
- SNIA Storage Management Initiative Specification (SMI-S)
- SNIA Green Storage TWG
- SNIA Data Protection and Capacity Optimization TWG (DPCO)
- SNIA Hands on Labs and Tutorials
- Storage Performance Council (SPC)

Representing over 100 years of collective experience in computer mass storage technologies, the authors provide authoritative technical expertise in a wide range of key solid state storage technologies and disciplines including NAND Flash technologies, Enterprise and Client SSD device and system architectures, IO Trace Methodology and software capture and analysis tools and SSD test and measurement technologies, software and test platforms. This group represents some of the key and primary developers of the SNIA SSS PTS specifications.

Eden Kim, CEO Calypso Systems, Inc. www.calypsotesters.com
Chair SSS TWG, SNIA SSSI Governing Board
Eden is Chair of the SNIA Solid State Storage Technical Working Group and a member of the SNIA Solid State Storage Initiative Governing Board. Mr. Kim was recognized in 2010 as the SNIA Outstanding Contributor for his work with the Solid State Storage Initiative and SSS Technical Working Group. Mr. Kim has been Chair of the SSS TWG since 2009 and has ushered the PTS through to publication. Mr. Kim received his BA/JD from the University of CA.

Mr. Kim is also CEO of Calypso Systems, Inc., a solid state storage test and measurement company and developer of the Calypso Reference Test Platform tester and CTS test software - the reference test platform of record and test system used for the development and validation of the SSS PTS specification and used to generate the data contained in this white paper.

Tom West, President, hyperI/O LLC www.hyperIO.com
Co-Chair SNIA IOTTA TWG; Member SSS TWG
Tom has been a co-chair of the SNIA IOTTA TWG since 2005 and led the establishment of the SNIA IOTTA Repository. He is also the author of the Filesystem Performance Profile within the SNIA Storage Management Initiative Specification (SMI-S) and actively participates in the SSS TWG, particularly in IO trace analysis related to the PTS.

Tom is a named inventor in 2 U.S. patents. His storage background includes 15 years at Storage Technology Corp., where he held advisory engineering positions in Systems Engineering and led the design and development of IBM-compatible enterprise mainframe disk storage subsystems. He more recently has been engaged in consulting for SSD firmware development and in the design and development of innovative software tools for measuring and monitoring disk and file IO operation performance.

Chuck Paridon, Storage Performance Architect
HP Storage Data Center Development Unit
SNIA Member: SSS TWG, Green Storage TWG, SPC

Chuck is a 24 year veteran in computer and storage benchmark development and performance analysis. He is currently responsible for the development of Hewlett-Packard Storage Performance collateral in the form of field performance troubleshooting training, efficient storage technology deployment and pre-sales consulting. He is a member of the SNIA Solid State Storage Technical Workgroup, the SNIA Green Storage Technical Workgroup and the Storage Performance Council.

Chuck has provided key insight to the SSS TWG pertaining to storage performance measurement, data collection and reduction techniques in addition to Enterprise applications / use cases and system level deployment of SSD. Chuck has earned a BS in Mechanical Engineering from the University of Akron, Oh. as well as an MS in Computer Science from California State University, Chico.

Doug Rollins, Senior Applications Engineer, Enterprise Solid State Drives,
SNIA Member: SSS TWG, IOTTA, TechDev, DPCO, HOL

Doug Rollins joined Micron Technology in 2009 as an applications engineer with the Enterprise SSD Products group. Prior to joining Micron, Mr. Rollins spent 13 years working in server system, network appliance, and storage platform/data protection design and manufacture.

Mr. Rollins is the named inventor in 13 U.S. patents and has been recognized by both the Storage Networking Industry Association (SNIA) and Intel Corporation for outstanding technical achievement. Mr. Rollins is an active member of several technical groups within SNIA including: the Solid State Storage Initiative and its Technical Working Group; Data Protection and Capacity Optimization; Marketing and Technical Development; Total Cost of Ownership; and the IO Trace Tools Analysis. As co-chair of SNIA's Solid State Storage Initiative's Technical Working Group, Mr. Rollins was instrumental in the early development and validation of SNIA's Solid State Storage Performance Test Specification

Easen Ho, PhD, CTO Calypso Systems, Inc. www.calypsotesters.com
SNIA Member: SSS TWG, SNIA Tutorials

Dr. Ho is the CTO of Calypso Systems, Inc. and has been a principal architect of the recently released SNIA Solid State Storage Performance Test Specification. Dr. Ho has been intimately involved in the development of performance benchmarking for NAND Flash based solid state storage devices. Dr. Ho continues to pursue the most advanced test and measurement methodologies and testing of NAND Flash based devices, development of emerging performance tests for SSDs, and validation of the SSS TWG specifications.

Dr. Ho received his doctorate in laser physics from MIT and his BSEE from the Tokyo Institute of Technology. Dr. Ho previously was founder, CEO and CTO of digital papyrus, inc. a laser optical mass storage technology firm.

Storage Networking Industry Association
425 Market Street, Suite 1020 • San Francisco, CA 94105 • Phone: 415-402-0006 • Fax: 415-402-0009 • www.snia.org