SNIA
Solid State Storage Initiative

Choosing the Right Solid State Storage Device: A Case Study of the SNIA SSSI Performance Test Specification

November 2010

SSSI Members and Authors:
Tom Coughlin, Coughlin Associates
Jim Handy, Objective Analysis
Easen Ho, Calypso Systems
Eden Kim, Calypso Systems

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.

Introduction

The Storage Networking Industry Association's (SNIA) Solid State Storage Initiative (SSSI) has released the first version of its Solid State Storage Performance Test Specification (SSS PTS)[i]. The SSS PTS, combined with a good reference test system, is an important tool for system integrators to use to find the right solid state drive (SSD) for their intended application. This white paper illustrates how the SSS PTS can effectively be used to compare a number of SSDs. The paper describes test results from 17 SSDs currently on the market.

Although SSDs are generally specified to perform much faster than hard disk drives (HDDs), SSD performance usually degrades over time, while HDD performance remains the same over the life of the product. In some cases an SSD can underperform an HDD. Actual SSD performance can vary greatly depending upon the product design (especially the controller and controller firmware) and the application workload.

Calypso Test Systems rigorously tested a variety of generally-available Single Level Cell (SLC) and Multi-Level Cell (MLC) drives to measure the time-dependent behavior and general spread of performance across the industry. This white paper compares a test sample of 17 SLC and MLC SSDs designed for enterprise, gaming, and laptop computer applications.

The data is excerpted from a report written by Coughlin Associates and Objective Analysis[ii].

Realistic Workload Testing of SSDs

The time (and use) dependent performance of an SSD is very different from that of a hard disk drive whose read and write performance is independent of data write activity.

SSD read and write performance changes with use, a fact that became evident in the course of developing the SSS PTS. Write speeds tend to slow as more of the available memory cells are written to, fewer pre-erased blocks become available, and the Logical Block Address (LBA) look-up tables become randomized. Thus the Fresh Out of the Box (FOB) or "purged drive" performance is often much greater than the performance after the SSD has undergone a number of writes. Purging the drive can return an SSD's performance to a level close to that of its FOB state; such a purge is usually performed using a secure erase, sanitize, format unit, or other proprietary command.

For this reason the SSS PTS defines conditions for establishing "steady state" performance. The PTS prescribes a preconditioning methodology that includes both a workload independent preconditioning step by writing sequential 128kB writes sufficient to fill twice the SSD's capacity, followed by a workload dependent preconditioning step that applies a real-world test load to the SSD to attain a steady state measurement window. The range of the LBAs preconditioned and tested can be a variable in the test set-up, depending on the anticipated workload and application.

Another Way to Estimate Steady State

Other workloads that exercise SSDs harder can give hints about when an SSD might reach its steady state during the SSS PTS workload and what the performance drop-off will be once the steady state is reached. One such test is the Write Saturation (WSAT) test. In the WSAT test, a new drive is subjected to continuous 4kB random writes. To determine when and how an SSD will reach a steady state, the WSAT load is presented to the SSD until the IOPS performance settles into a stable range (which is defined as 10 rounds that fall within a 5% maximum data excursion and slope). WSAT does not use the preconditioning methodology of the SSS PTS, but WSAT can be useful for comparing the FOB performance of an SSD with its more time-evolved state. For the SSDs measured in these tests, the steady state performance after WSAT is 45%-95% lower than the purged or FOB performance.

Figure 1 shows how performance changes over time for eight of the test's SSDs from the FOB state under a WSAT workload (continuous 4kB random writes). In this figure the vertical axis shows the maximum write IOPS performance normalized to the FOB performance. Performance decreases significantly for all of the tested SSDs as data is written into them. After a large number of write cycles (typically taking less than 700 minutes), most of the SSDs eventually achieve a WSAT steady state after passing through a transition region, although some take significantly more time than others.

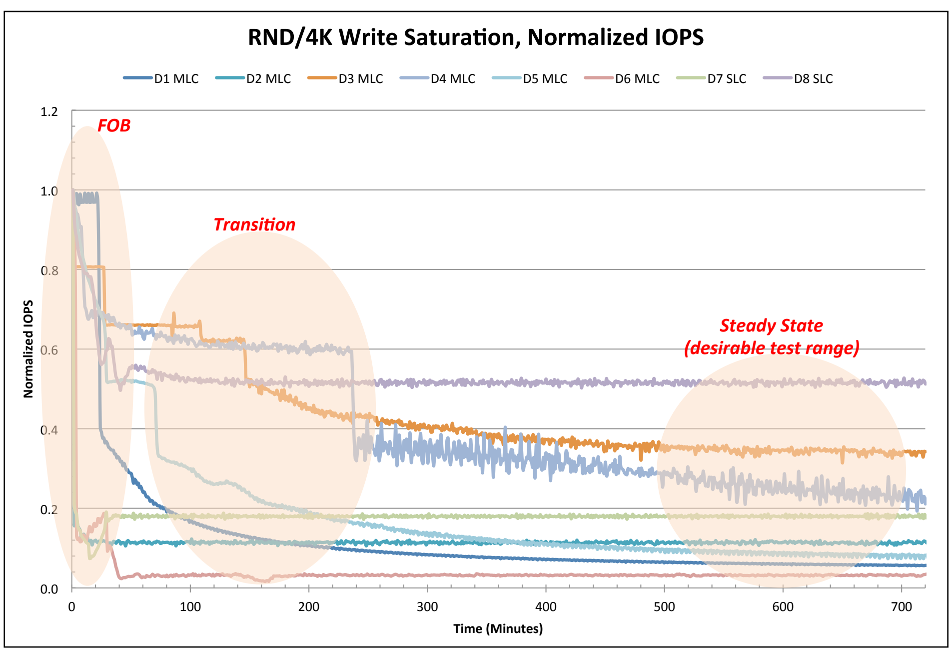

Figure 1. Performance Changes of SSDs Undergoing Continuous Random 4KB Writes

Steady-State Comparisons

The SSS PTS sets forth a standardized preconditioning methodology to ensure measurements are taken during a specifically defined steady state window. It is difficult to determine how any SSD will perform under a given application workload until it has reached its steady state performance, so Steps 1-4 of the SSS PTS procedure described in Table I are used to establish the steady state. Once the steady state has been reached, data can be collected under the desired workload and this data can then be used to generate comparative reports (steps 5 & 6).

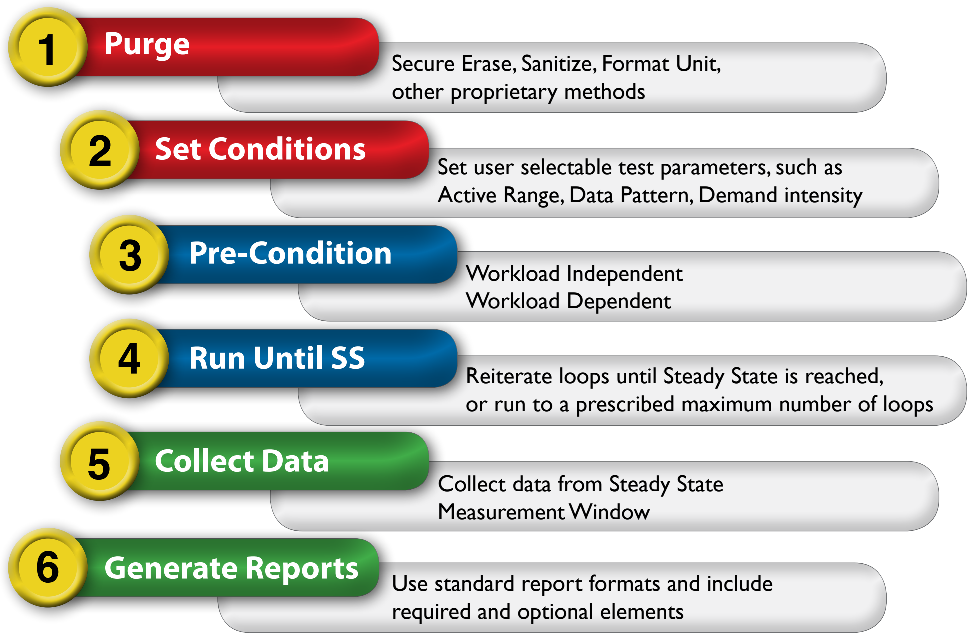

Table I. SSS Performance Test Specifications Test Flow

Figure 2 below plots the steady state performance of nine SLC (blue) and eight MLC (red) SSDs. The data in this diagram was measured using a workload of 35% writes and 65% reads. The third axis represents the various sized random data blocks used for the test – from 0.5 to 1,000 kB. Each row represents a single SLC or MLC SSD, with the last row giving data for the same tests performed on a 15,000 RPM enterprise Serial Attached SCSI (SAS) HDD (yellow).

Block Size Sensitivity

Even after a steady state has been reached, the performance of the various SLC and MLC SSDs varies greatly, with IOPS peaking at certain transfer sizes. In general, the SSDs perform better than the HDD but this is not true under all workloads, especially when large block sizes are used. To help users understand how a certain SSD will work in their own application, the SSS PTS can be adapted to use a variety of alternative workloads to provide valuable data on the best performing SSD for a particular application. This makes the test a valuable tool for storage system designers and integrators.

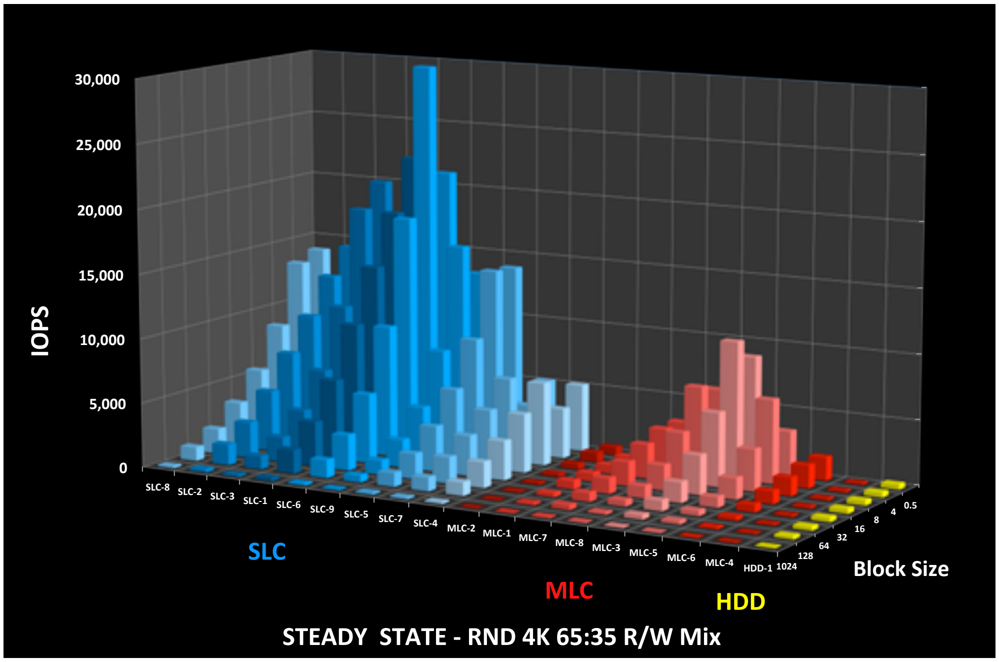

Figure 2. Performance Changes of SSDs Undergoing Continuous Random 4KB Writes

The block size peaks for most of the drives in Figure 2 occur at transfers of 4-8kB. This is a good match for virtual memory systems using variants of the Windows, Linux, and many other operating systems, since virtual memory pages are 4kB. In a system that does not use a virtual memory operating system, particularly a system designed to capture long data transfers, SSD users should look for drives whose performance exceeds that of other SSDs at the longer block length.

Small block random I/O activity is common in OLTP database systems and virtualized computing applications. Large block sequential transactions are found in video on demand (VOD) edge servers.

These SSDs vary significantly from each other largely because there are many "knobs and buttons" to push in order to optimize SSDs for a predicted workload. SSD controller designers must choose from a long list of trade-offs to determine which design approach they will use when designing their SSDs. This makes it imperative for the SSD buyer to compare each SSD's performance to the anticipated workload to guarantee that the best match is made.

The Impact of Different Data Patterns

Block size is not the only factor that impacts SSD performance. Data patterns can also cause an SSD to perform differently than anticipated. Figure 3 shows how IOPS vary as a function of both data patterns and block size for one of the 17 SSDs in the test. As we saw in Figure 2, there is a strong peak at the 4kB block size during database accesses, after which performance decreases as the block size increases.

A random data pattern drives performance below that experienced with either an actual database access or an all-zero database file access and does not peak at any particular block size. The performance at the 4kB peaks with the database loads is much higher than the highest performance level under a random data load. This implies that this particular SSD would yield a higher performance boost in a database application than it might in a more random access application typical of the storage patterns of a highly virtualized server.

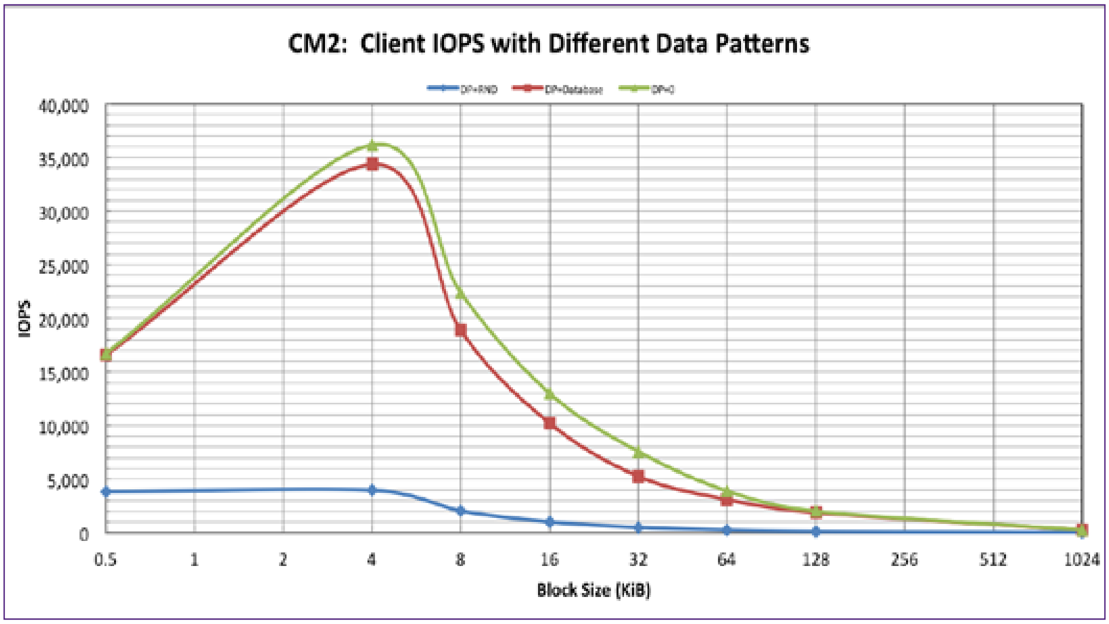

Figure 3. Example Data Dependent Performance of One SSD

A significant volume of similar data can be found in the report from which this white paper has been drawn: ***Putting SSDs to the Test!***[ii]**.**

The SSS PTS supports a wide variety of performance measurements that are useful in designing or integrating storage systems. These include tests with varying amounts of over provisioning, testing with various active region ranges, write saturation results where wait states are allowed, and tests that use different data patterns. In addition, the write and read ratios and block transfer sizes can be modified to match the requirements of the application.

The SSS PTS isolates devices and stimulus while open source benchmark tools do not. It is a very useful tool for testing SSDs to determine the best product for the application. Further, the database of measurements gathered on the Calypso RTP 2.0 and CTS 6.5 test software allows users to compare the same data taken on these 17 SSDs against new data on an "apples-to-apples" basis. All of the data in the tests shown here was collected using the Calypso RTP / CTS. This test system complies with the requirements set forth in the SSS PTS v1.0.

Future PTS Testing

The SSS PTS v1.0 covers only basic drive performance characteristics such as IOPS, throughput, and latency. It serves to introduce key SSD testing procedures such as preconditioning and steady state measurements. It is intended to be a device-level test using synthetic stimulus to allow direct comparison between different drives without taking into account the specific workloads in different user environments. Every system, application, hardware, and user will create a different "typical" workload. Real world workloads are difficult to capture, making it almost impossible to characterize "typical" workloads. Despite this obstacle, it is possible to categorize general workload types.

The goal of the next generation SSS PTS is to identify broad classes of workloads and characterize them (both for enterprise and client environments). For example, client environments could be gaming, light office productivity, Web surfing, desktop publishing, or some other general application type.

Furthermore, there are numerous other performance phenomena specific to Flash-based SSDs that are not normally observed when testing HDDs. The ongoing and future work for the SNIA SSS Technical Work Group (TWG) will be to address many of these key issues, such as:

- Determining how to take particular user workloads into consideration, including access range and data pattern issues.
- Defining a power-efficiency metric, such as IOPS/Watt.
- Investigating performance transients when switching from one sustained stimulus to another.
- Recommending a method for defining application demand intensity.
- Creating detailed analyses of response time statistics beyond average and maximum response times.
- Determining the possibility of one or more "SSD Figure(s) of Merit" for marketing and education purposes using results from the SSS PTS.

The SNIA SSS TWG and Calypso Test Systems are trying to characterize "Task Based" activities and synthesize those attributes to create a test stimulus. Eventually, a library of real world input/output traces may be gathered, but this will take time because it is difficult for both technical and business reasons.

The TWG also is actively seeking input on other tests of significant interest to the SSS industry. The SSS TWG invites new participants from SNIA SSSI member companies to help create these useful test standards that will help designers and users choose SSDs that match their workloads and needs.

More information on the SSS PTS and other solid state storage topics can be found on the SNIA website at: http://www.snia.org/forums/sssi/knowledge/education/.

i Solid State Storage Performance Storage Test Specifications, Neal Ekker, Khaled Amer and Paul Wassenberg, SNIA, 2010.

ii **Putting SSDs to the Test!**, Tom Coughlin, Jim Handy, Easen Ho and Eden Kim, 2010, available from Coughlin Associates (www.tomcoughlin.com) or Objective Analysis (www.Objective-Analysis.com).

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

©Copyright 2010 Storage Networking Industry Association. All Rights Reserved.