Tiered Storage: Enterprise SSHD and Flash SSD – Better Together

Solid State Hybrid Drives (SSHD)

Part of an Enterprise Tiered Storage Strategy
Enterprise SSHD and Flash SSD – Better Together

Performance, Availability, Capacity, Energy and Economic Effectiveness

By Greg Schulz

Founder and Senior Advisory Analyst
Server and StorageIO Group @StorageIO

March 17, 2014

This Industry Trends and Perspectives White Paper is compliments of

Seagate®

www.seagate.com

Tiered Storage: Enterprise SSHD and Flash SSD – Better Together

The best I/O is the one that you do not have to do. The second best I/O or IOP is the one with least impact and most benefit to your applications.

Introduction

The question to ask yourself is not if flash Solid State Device (SSD) technologies are in your future. Instead the questions are when, where, using what, how to configure and related themes. SSD including traditional DRAM and NAND flash-based technologies are like real estate where location matters; however, there are different types of properties to meet various needs. This means leveraging different types of NAND flash SSD technologies in different locations in a complementary and cooperative aka hybrid manner.

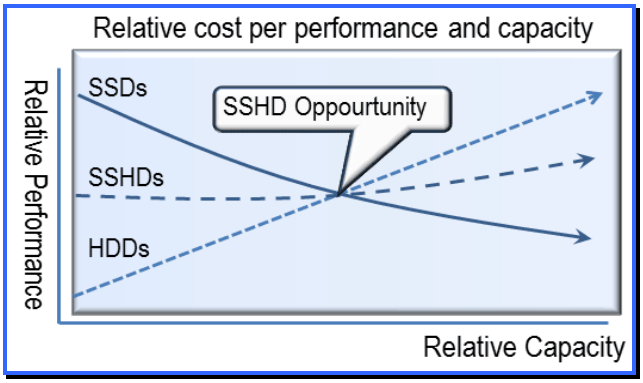

In this StorageIO Industry Trends Perspective thought leadership white paper we look at how enterprise class Solid State Hybrid Drives (SSHD) and how they address current and next generation tiered storage for virtual, cloud, traditional Little and Big Data infrastructure environments.

Background and Common Data Storage Challenges

There is no such thing as an information recession with more data being generated, processed, moved, stored and retained for longer periods. In addition, people and data are living longer as well as getting larger. We have continued growth with traditional little data (databases and traditional applications) along with Big Data and Very Big Data. These and other applications have the common characteristics of needing more performance to support transactions, videos, messages, pages and files served along with bandwidth and reduced latency (response time).

There are economic challenges requiring new outside of the box thinking boosting your return on innovation (the new ROI). Doing more with what you have or less includes consolidating data centers, servers and storage using virtualization. However, consolidation or aggregation also causes aggravation known as bottlenecks. The most common approach is to leverage lower cost, slower and high capacity Hard Disk Drive (HDD) to support storage space consolidation along with various data footprint reduction (DFR) techniques including compression, dedupe and thin provisioning. Another approach is to leverage fast yet expensive lower space capacity NAND flash solid-state devices (SSD) to support IOPs or performance consolidation.

Storage Tiers (device and mediums)

Tier 0 – SSD (DRAM or NAND flash)
Tier 1 – Fast HDD & SSHD
Tier 2 – Slower high capacity HDD
Tier 3 – HDD, Tape or cloud service

Tiered Storage: Enterprise SSHD and Flash SSD – Better Together

Flash SSD in Different Locations that Matter

SSD technology including NAND flash is like real estate in that location matters, as well as a relatively small amount in the right place can have a big impact. SSD can and should be deployed in different locations to meet various workload and data infrastructures demand requirements. This means that SSD are like real estate that comes in different varieties to meet various usage needs one of which is being close to where the IOPs are occurring. For applications and servers, this means caching as the best IO is the one that you do not have to do while the second is the one with least impact.

Using SSD in storage systems provides shared access benefit across multiple servers for both reads and writes. SSHD compliment server and storage system based SSD by placing flash technology close to where actual reads and writes IOPs are performed to be more effective. Leveraging SSD technology inside, SSHD do more work (IOPs) effectively than traditional HDDs.

SSHD Complimenting other Flash-Based SSD Technologies

Enterprise SSHDs are a good complement to other NAND flash Solid State Devices (SSD) including PCIe cards, drives and storage systems as well as tiering and caching software tools. Too often in the IT industry and particularly around data infrastructure related discussions, technologies are positioned as competitive, this vs. that as opposed to how they can complement making each more effective. For example, your applications need more I/O performance (IOPs, bandwidth or lower response time/latency) than what traditional HDDs provide. Thus, you may want SSD type performance for your applications. However, your budget requires that you innovate to stretch it further to support space capacity needs. This is where balancing performance, availability, capacity and economics (along with energy) come into play along with tiered storage devices (mediums).

What is an SSHD?

It is a Hybrid, best of SSD and HDD
Functions and appears as a HDD
Combines NAND flash with an HDD
More performance than HDD
Read and write IO acceleration
More capacity per cost than SSD
Durability duty-cycle of HDD
No special drivers or adapters
Do more with less

The traditional approach has been to use some amount of NAND flash-based SSD located in servers and storage systems along with tiering and caching software for performance, and HDDs for space capacity. SSHD help to close the gap between higher cost per capacity performance based SSD, and lower cost higher capacity slower HDDs. In other words as a Hybrid you get best of benefits from both technologies (performance, availability, capacity, energy and economics). How you leverage SSHD to compliment your SSD deployed in various locations will determine your return on innovation. For example using SSHD as fast disk storage to compliment Flash-Based SSD server or storage system based caching. Keep in mind that fast servers, fast memory, fast SSD also need fast underlying magnetic storage devices to remove complexities, cost and bottlenecks.

Tiered Storage: Enterprise SSHD and Flash SSD – Better Together

Solid State Hybrid Disks (SSHD) Storage Today

Hybrid Hard Disk Drives (HHDD) desktop, workstation or client class devices have been in the market for several years. However those HHDD only helped speed up some read operations where current generation enterprise class SSHD accelerate reads and writes. DRAM cache has been available in enterprise class HDDs for many years; NAND flash SSD technology is the next evolution step to improve performance.

Seagate Enterprise Turbo SSHD

2.5" 6Gbps SAS 600GB device
Faster than 10K and 15K HDDs
Includes read and write acceleration
32GB flash & 128MB DRAM buffers
Low power energy efficient

SSHD and SSD tiered storage – Better together

By placing some amount of DRAM and now NAND flash inside the HDD, the resulting SSHD is able to deliver more performance (read and write). For example, writes can be buffered with persistency to maintain data integrity and application transaction state consistency while reducing response, the result is faster application performance. Another example is that reads can be optimized by increasing read ahead or other predictable activities to provide servers and storage systems with data in advance, so they can reduce the amount of IOPs they may have to do (e.g. help their read and other cache functions).

Collectively these read and write enhancements make SSHD a good option to compliment SSD based cache solutions that rely on underlying magnetic HDD for data storage. An example being VMware VSAN and caching for virtual server including VDI environments where SSD is used for reads with write-thru cache to fast SSHD.

When to use Enterprise SSHD

- ✓Compliment SSD based solutions
- ✓Fast magnetic storage needed
- ✓Replacement several slower HDDs
- ✓Eliminate HDD based short stroking
- ✓Virtual, cloud and physical servers

An example is the Seagate® Enterprise Turbo SSHD that provides up to 600GB of storage space capacity along with faster performance compared to traditional enterprise 15K RPM HDDs. The performance boost comes from read and write optimizations that leverage NAND flash SSD technology integrated within the drive itself. The technical benefit is providing higher performance compared to traditional HDDs at a lower cost vs. standard SSD devices. Table-1 (below) shows how SSHD can be used as a high-performance storage tier, including with SSD based solutions.

In addition to those shown below in Table-1, other candidate applications and workloads include medical imaging, home directors and general file sharing. Other uses include cloud and object storage, physical and virtual machine page and swap files, content distribution networks (CDN) as a high speed magnetic object store complimenting SSD cache and meta data lookup.

Tiered Storage: Enterprise SSHD and Flash SSD – Better Together

Data Protection (Archiving, Backup, BC, DR)	Staging cache buffer area for snapshots, replication or current copies before streaming to other storage tier using fast read/write capabilities. Meta data, index and catalogs benefit from fast reads and writes for faster protection.
Big Data DSS Data Warehouse	Support sequential read-ahead operations and "hot-band" data caching in a cost effective manner using SSHD vs. slower similar capacity size HDDs for Data warehouse, DSS and other analytic environments.
Email, Text and Voice Messaging	Microsoft Exchange and other email journals, mailbox or object repositories can leverage faster read and write I/Os with more space capacity.
OLTP, Database Key Value Stores SQL and NoSQL	Eliminate the need to short stroke HDDs to gain performance, provide more space capacity and IOP performance per device for tables, logs, journals, import/export and scratch, temporary ephemeral storage. Leverage random and sequential read acceleration to compliment server side SSD-based read and write-thru caching. Utilize fast magnetic media for persistent data reducing wear and tear on more costly flash SSD storage devices.
Server Virtualization	Fast disk storage for data stores and virtual disks supporting VMware vSphere/ESXi, Microsoft Hyper-V, KVM, Xen and others. Holding virtual machines such as VMware VMDKs, along with Hyper-V and other hypervisor virtual disks. Compliment virtual server read cache and I/O optimization using SSD as a cache with writes going to fast SSHD. For example VMware V5.5 Virtual SAN host disk groups use SSD as a read cache and can use SSHD as the magnetic disk for storing data while boosting performance without breaking the budget or adding complexity.
Virtual Desktop Infrastructure (VDI)	SSHD can be used as high performance magnetic disk for storing linked clone images, applications and data. Leverage fast read to support read ahead or pre-fetch to compliment SSD based read cache solutions. Utilize fast writes to quickly store data enabling SSD-based read or write-thru cache solutions to be more effective. Reduce impact of boot, shutdown, and virus scan or maintenance storms while providing more space capacity.

Table 1 – Example application and workload scenarios benefiting from SSHDs

Proof Points and Enterprise SSHD Validation

StorageIO has conducted various hands-on testing with the Seagate® Enterprise Turbo SSHD in our StorageIO lab environment across different real world like application workload scenarios. These include general storage I/O performance characteristics profiling (e.g. reads, writes, random, sequential or various IOP size) to understand how these devices compare to other HDD, HHDD and SSD storage devices in terms of IOPS, bandwidth and response time (latency). In addition to basic storage I/O profiling[1], the Enterprise Turbo SSHD was also used with various SQL database workloads including TPC; along with VMware server virtualization among others use case scenarios.

[1] Read more about basic Storage I/O profiling IOPs, latency and bandwidth at StorageIOblog.com http://storageioblog.com/part-ii-iops-hdd-hhdd-ssd/

Tiered Storage: Enterprise SSHD and Flash SSD – Better Together

Proof Point: Email Messaging

Email and messaging are popular applications that combine the need for cost effective performance and space capacity. This proof point looks at how SSHDs perform for Email or messaging activity handling repositories such as Exchange Databases (EDBs). For this scenario, Microsoft Jet Stress Exchange performance workloads were placed (e.g. EDB file) on each of the different devices under test with various metrics shown including activity rates and response time for reads as well as writes.

	Pro 600 SSD 120GB	Ent. Turbo SSHD 600GB	Ent. 15K V4 HDD 600GB	Savio 15K 146GB HDD	Barracuda 7.2K HDD 500GB
Avg. IOP/second	1,006.1	352.5	263.3	197.3	96.5
Avg. Reads resp. (msec)	2.6	4.7	6.7	9.4	17.8
Avg. Writes resp. (msec)	7.2	5.1	30.2	35.6	56.1
Avg. IOP (read)/second	472.2	176.1	132.4	129.6	76.7
Avg. IOP (write)/second	533.8	176.5	130.8	97.3	47.4

Table 2 - Email and messaging - Microsoft Exchange JetStress workload (single device)[2]

Table-2 above results show performance improvements with the Exchange workload simulation using the Seagate Enterprise Turbo SSHD compared to SSD and HDDs. As a Hybrid device, the SSHD demonstrate good performance and space capacity (best of both worlds) between SSD speed and HDD space capacity. Figure 1 below provides a graphical representation of the SSHD performance characteristics compared to the SSD and HDDs that were tested show read and write rates along with response times.

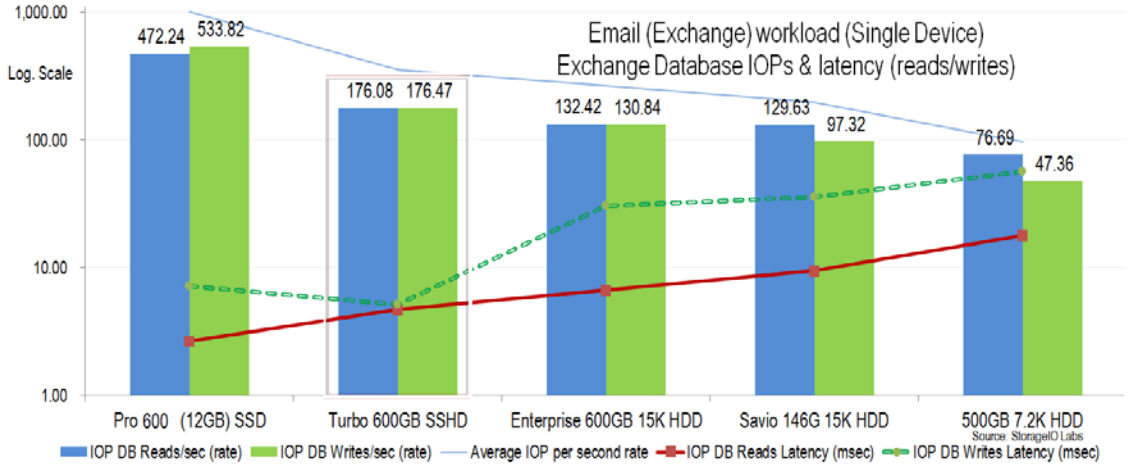

Figure 1 - Microsoft Exchange Jetstress workload (Database reads/writes and latency)

[2] Test configuration: 2.5" Seagate 600 Pro 120GB (ST120FP0021) SSD 6 Gbps SATA, 600GB 2.5" Enterprise Turbo SSHD (ST600MX) 6 Gbps SAS, 600GB 2.5" Enterprise Enhanced 15K V4 (15K RPM) HDD (ST600MP) with 6 Gbps SAS, 2.5" Savio 146GB HDD 6 Gbps SAS, 3.5" Barracuda 500GB 7.2K RPM HDD 3 Gbps SATA. Email server hosted as guest on VMware vSphere/ESXi V5.5, Microsoft Small Business Server (SBS) 2011 Service Pack 1 64 bit, 8GB DRAM, One CPU (Intel X3490 2.93 GHz) LSI 9211 6 Gbps SAS adapter, JetStress 2010 (no other active workload during test intervals). All devices being tested were Raw Device Mapped (RDM) where EDB resided. VM on a SSD based separate data store than devices being tested. Log file IOPs were handled via a separate SSD device.

Tiered Storage: Enterprise SSHD and Flash SSD – Better Together

Proof Point: Database, Data Warehouse Batch Update

SSHDs are a good fit for both transaction database activity with reads and write as well as query-based decision support systems (DSS), data warehouse and big data analytics. The following are proof points of SSHD capabilities for database activity. Two database workload profiles were tested including batch update (write-intensive) and transactional.

Activity involved running Transaction Performance Council (TPC) workloads TPC-B (batch update) and TPC-E (transaction/OLTP simulate financial trading system) against a SQL database server. Each test simulation had the SQL database object placed (MDF) on a different device with transaction log file (LDF) on a separate SSD. TPC-B for a single device results are shown for various user workloads in table 3.

Device	User Load	TPC-B TPS	Avg. Resp. Time (Sec.)	Avg. Trans. Time (Sec.)
600GB Ent. Turbo SSHD	1	52.0	0.02	0.02
	20	172.2	0.10	0.11
	50	226.7	0.20	0.22
	100	220.6	0.39	0.45
600GB Enterprise 15K HDD	1	34.8	0.03	0.03
	20	166.2	0.11	0.12
	50	140.4	0.32	0.35
	100	155.7	0.57	0.64
500GB 7.2K RPM HDD	1	16.0	0.06	0.06
	20	78.3	0.25	0.25
	50	57.3	0.85	0.87
	100	59.7	1.66	1.67
1TB 7.2K RPM HDD	1	15.8	0.06	0.06
	20	48.5	0.41	0.41
	50	44.3	1.12	1.13
	100	43.6	2.29	2.29

Table 3 - TPC-B[3] (batch update write) with 1, 20, 50 and 100 users (single device)

Note that these results are for a single device under test and additional performance could be expected using multiple drives in RAID or other system configurations. Scaling of performance should be linear with additional drives and adequate available system resources (CPU, memory, and PCIe and I/O bandwidth) to support more user workload.

[3] Test configuration: 600GB 2.5" Enterprise Turbo SSHD (ST600MX) 6 Gbps SAS, 600GB 2.5" Enterprise Enhanced 15K V4 (15K RPM) HDD (ST600MP) with 6 Gbps SAS, 500GB 3.5" 7.2K RPM HDD 3 Gbps SATA, 1TB 3.5" 7.2K RPM HDD 3 Gbps SATA. Workload generator and virtual clients ran on Windows 7 Ultimate. A SQL based database was on Windows 7 Ultimate SP1 (64 bit) 14 GB DRAM, Dual CPU (Intel x3490 2.93 GHz)), with LSI 9211 6Gbps SAS adapters with TPC-B (www.tpc.org) workloads. VM resided on separate data store from devices being tested. All devices being tested with SQL MDF were Raw Device Mapped (RDM) independent persistent with database log file (LDF) on a separate SSD device also persistent (no delayed writes). Tests were performed in StorageIO Lab facilities by StorageIO personal.

Tiered Storage: Enterprise SSHD and Flash SSD – Better Together

TPC-B (write intensive) results in figure 2 below show how TPS work being done (blue) increases from left to right (more is better) for various numbers of simulated users. Also shown on the same line for each amount of TPS work being done is the average latency in seconds (right to left) where lower is better. Results are shown from top to bottom for each group of users (100, 50, 20 and 1) for the different drives being tested (top to bottom). Note how the SSHD device does more work at a lower response time vs. traditional HDDs.

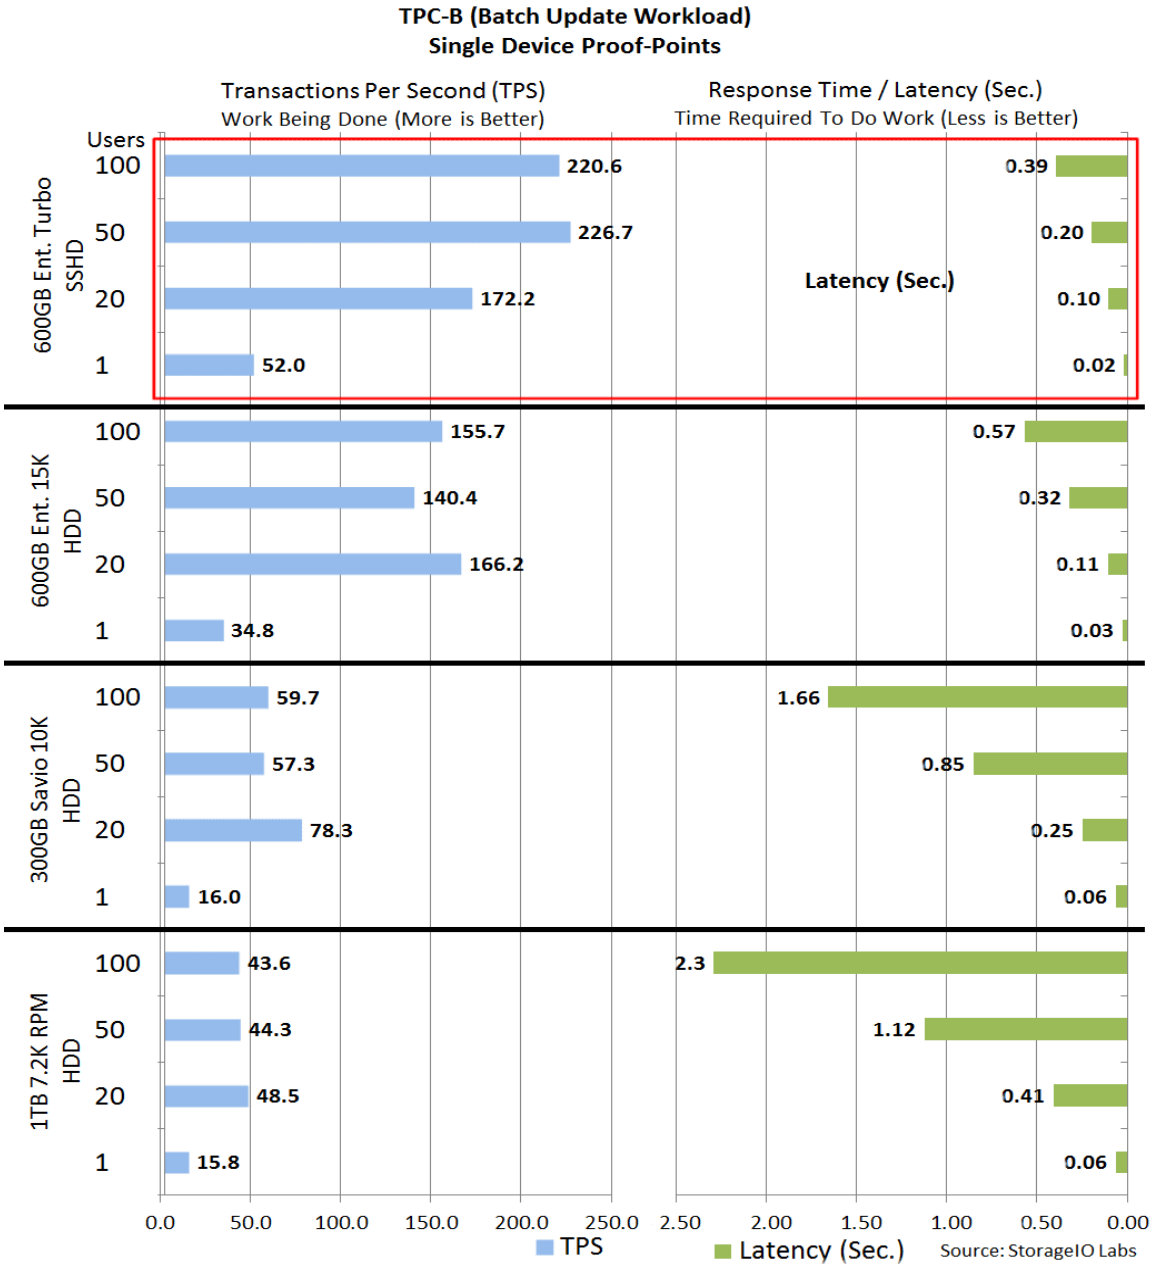

Figure 2 – TPC-B (batch update write) with 1, 20, 50 and 100 users (single device)

Tiered Storage: Enterprise SSHD and Flash SSD – Better Together

Proof Point: Database OLTP and Transactional

Table 4 shows results from TPC-E test (OLTP/transactional workload) simulating a financial trading system. TPC-E is an industry standard workload that performs a mix of reads and writes database queries.

Test simulations were performed with various numbers of users from 10, 20, 50 and 100 to determine (TPS) Transaction per Second (aka I/O rate) and response time in seconds.

Device	User Load	TPC-E TPS	Avg. Resp. Time (Sec.)	Avg. Trans. Time (Sec.)
600GB Ent. Turbo SSHD	10	3.2	1.44	3.16
	20	5.3	1.58	3.78
	50	6.2	2.76	8.09
	100	6.4	4.64	15.54
600GB Enterprise 15K HDD	10	2.2	1.72	4.45
	20	3.4	2.12	5.87
	50	3.6	4.19	13.80
	100	3.2	9.03	31.42
300GB Savio 10K HDD	10	1.8	2.07	5.48
	20	2.7	2.49	7.34
	50	3.0	5.38	16.57
	100	2.8	10.17	35.40
1TB 7.2K RPM HDD	10	1.0	3.49	10.44
	20	1.7	3.67	12.08
	50	1.9	6.79	26.37
	100	1.9	14.00	52.36

Table 4 – TPC-E[4] (OLTP transactional) with 1, 20, 50 and 100 users

[4] Test configuration: 600GB 2.5" Enterprise Turbo SSHD (ST600MX) 6 Gbps SAS, 600GB 2.5" Enterprise Enhanced 15K V4 (15K RPM) HDD (ST600MP) with 6 Gbps SAS, 300GB 2.5" Savio 10K RPM HDD 6 Gbps SAS, 1TB 3.5" 7.2K RPM HDD 6 Gbps SATA. Workload generator and virtual clients Windows 7 Ultimate. A SQL based database was on Windows 7 Ultimate SP1 (64 bit) 14 GB DRAM, Dual CPU (E8400 2.99GHz), with LSI 9211 6Gbps SAS adapters with TPC-E (www.tpc.org) workloads. VM resided on separate SSD based data store from devices being tested (e.g., where MDF resided). All devices being tested were Raw Device Mapped (RDM) independent persistent with database log file on a separate SSD device also persistent (no delayed writes). Tests were performed in StorageIO Lab facilities by StorageIO personal.

Tiered Storage: Enterprise SSHD and Flash SSD – Better Together

In figure 3 below, the TPC-E transactional results are shown for each device being tested across different user workloads. The results show how TPC-E TPS work (blue) increases from left to right (more is better) for larger numbers of users along with corresponding latency (green) that goes from right to left (less is better). The SSHD (e.g. 600GB Enterprise Turbo) is shown on the top of figure 3 with a red box around its results. Note how the SSHD as a lower latency while doing more work compared to the other traditional HDDs. This shows NAND flash and HDDs being better together providing a hybrid between SSD performance and HDD space capacity while balancing economics.

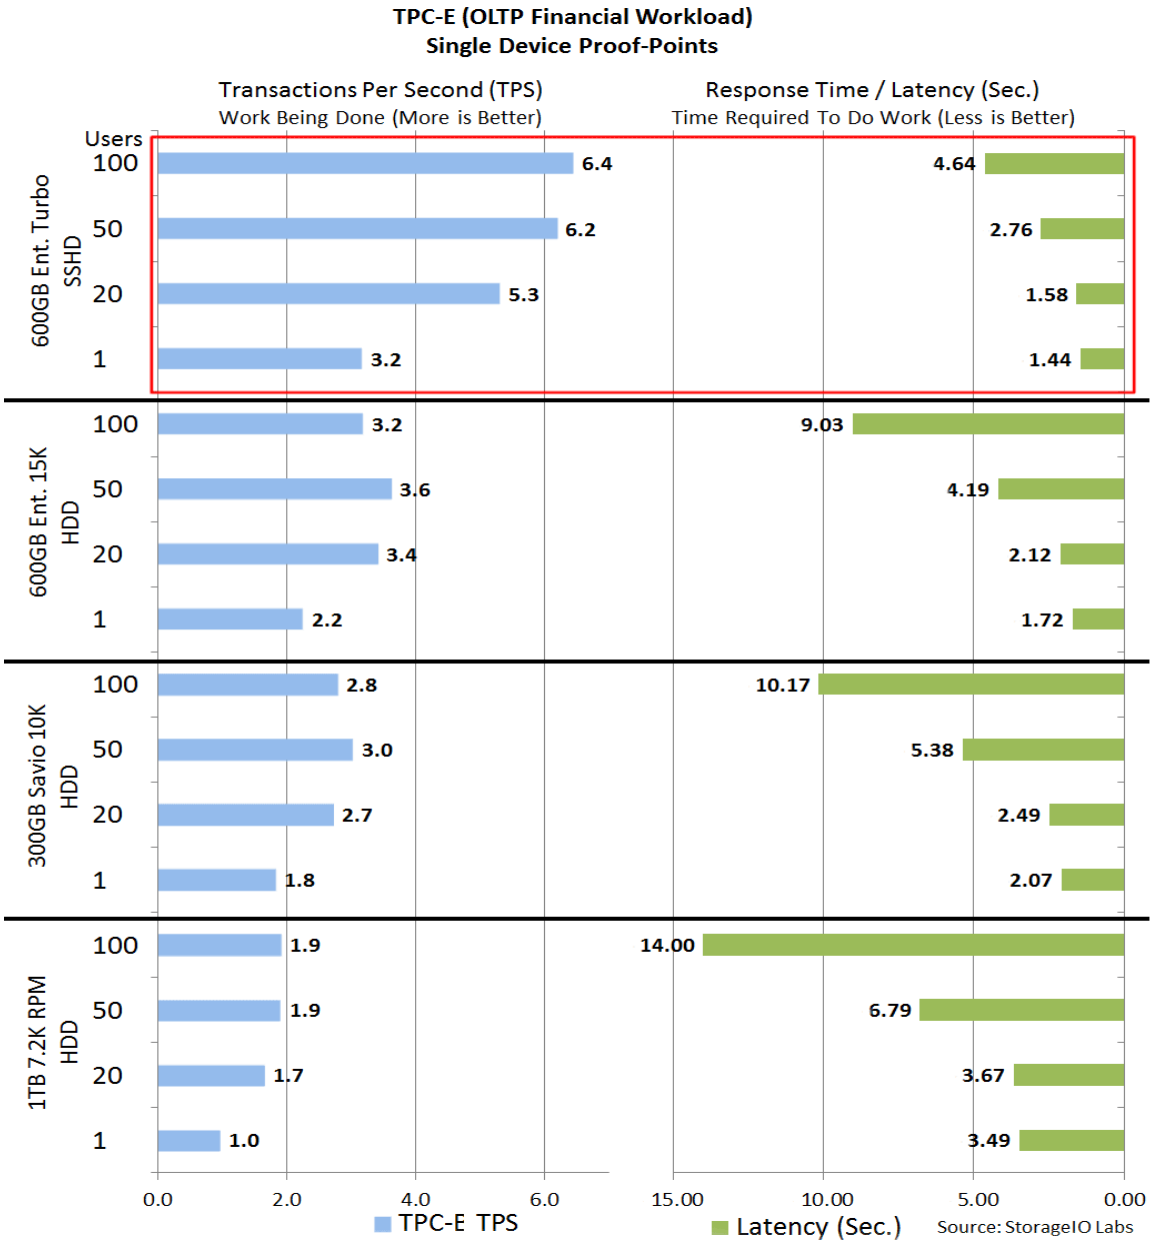

Figure 3 – TPC-E (OLTP transactional) with 10, 20, 50 and 100 users

Tiered Storage: Enterprise SSHD and Flash SSD – Better Together

Additional and Future Proof Points - Server Virtualization

Additional testing and workload proof points are being performed including VMware vSphere based among others. Watch for additional proof points involving VMware V5.5 with various workloads including I/O caching to be posted at www.storageioblog.com.

Summary

Similar to flash-based SSD technologies the question is not if, rather when, where, why and how to deploy hybrid solutions such as SSHDs. If your applications and data infrastructures environment have the need for storage I/O speed without loss of space capacity and breaking your budget, SSD enabled devices like the Seagate Enterprise Turbo 600GB SSHD are in your future.

Learn more at Seagate landing page located at:

http://www.seagate.com/internal-hard-drives/solid-state-hybrid/enterprise-turbo-sshd/

About the author

Greg Schulz is Founder and Sr. Analyst of independent IT advisory consultancy firm Server and StorageIO (StorageIO). He has worked in IT at an electrical utility, financial services and transportation firms in roles ranging from business applications development to systems management, architecture, strategy and capacity planning. Mr. Schulz is author of the Intel Recommended Reading List books "Cloud and Virtual Data Storage Networking" and "The Green and Virtual Data Center" via CRC Press and "Resilient Storage Networks" (Elsevier). He is a four-time VMware vExpert. Learn more at www.storageio.com.