Tiered Storage: Enterprise 12Gbps SAS flash SSD

Seagate® 1200 SSD (12Gbps SAS) Enterprise SAS Server & Storage SSD

Better Together - Part of Enterprise Tiered Storage Strategy

Performance, Availability, Capacity, Energy and Economic Effectiveness

By Greg Schulz

Founder and Senior Advisory Analyst
Server and StorageIO Group @StorageIO

July 9, 2014

This Industry Trends and Perspectives White Paper is compliments of

www.seagate.com

Tiered Storage: Enterprise 12Gbps SAS flash SSD

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 is not if flash Solid State Device (SSD) technologies are in your future. Instead, the questions are when, where, using what and how much to use among other related themes. Also, keep in mind that a relative small amount of flash SSD in the right place can have a positive impact on business applications without busting your budget. Thus, the importance of a tiered enterprise storage strategy includes flash SSD technology in different locations along with companion technologies. This industry trends and perspective white paper looks at tiered storage using enterprise SSD along with other technologies. In this StorageIO Industry Trends Perspective thought leadership white paper[1] we look at how enterprise class 12 Gbps Serial Attached SCSI (SAS) flash SSD drives enable performance and productivity across different application workloads.

Background and Common Data Storage Challenges

Not everything is the same in data centers or across the various applications, servers, I/O networking and storage that support them. Since there is no such thing as an information recession[2] more data needs to be cost effectively stored for longer periods. With more data being stored and accessed, not to mention that data is also getting larger, there are subsequent data center bottlenecks.

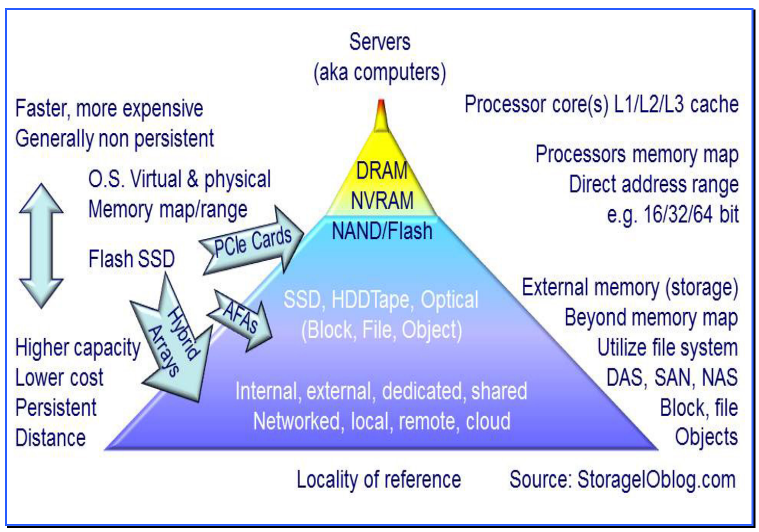

Hard Disk Drive (HDD) technology continues to increase in space capacity, resiliency including bit error rates while consuming fewer watts to do a given amount work, there are performance challenges. Key to applications performance is data locality of reference meaning how close the information to where it is being processed is. This is where data storage is used as persistent memory across different tiers or layers of technology to balance performance, availability, capacity and economic costs. This is where flash SSD comes into play in various packaging options for deployment in different locations as a storage device or as a cache.

---

[1] This is a companion to the StorageIO Industry Trends Perspective thought leadership white paper and proof points "Solid State Hybrid Drives (SSHD) - Part of an Enterprise Tiered Storage Strategy" available at http://storageio.com/whitepaper - Note that Seagate now refers to SSHD as TurboBoost™

[2] StorageIOblog - Is there an information or data recession? Are you using less storage?

Tiered Storage: Enterprise 12Gbps SAS flash SSD

Flash SSD in Different Locations that Matter

SSD technology including NAND flash is like real estate in that location matters. In addition, a relatively small amount in the right place can have a positive impact on your applications and business. As a best practice SSD should be deployed in different locations (across various storage tiers) to meet various workload and data infrastructures demand requirements. Using SSD in storage systems provides shared access benefit across multiple servers for both reads and writes.

Storage Tiers (device and mediums)

Tier 0 – Flash SSD devices (drives, cards)
Tier 1 – Flash in HDD[3]
Tier 2 – SSD cache for high capacity HDD
Tier 3 – SSD drives and cache for cloud

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, this means caching so the best IO is the one that you do not have to do while the second is the one with least impact.

Enterprise Flash SSD storage Today: Seagate Enterprise 1200 SSD

The Seagate 1200 is a 2.5" Enterprise class MLC[4] nand flash SSD with a 12Gbps SAS interface. Later in this white paper you will see various proof points of how the Seagate 1200 (400GB model ST400FM0073) SSD performs across different applications. Seagate also has 200GB and 800GB models along with versions supporting Self-Encrypting Drive (SED) along with FIPS 140-2. For data integrity these devices have a Bit Error Rate of 10^16, Annual Failure Rate (AFR) of .44%.

Seagate Enterprise 1200 SSD

As tested 400GB 12Gbps SAS

For durability and flash endurance, the devices support up to 7,300 TBytes Written (TBW) along with a 5 year limited warranty, while the 800GB version supports up to 14,600 TBW. Power consumption per Seagate spec sheets are average of 3.71 watts (400GB) while operating, 2.72 watts while in idle mode or 2.5 watts while in sleep. Obviously if powered off there would no watts consumed. Seagate spec sheets state sustained 750MB/sec data throughput StorageIO has seen similar results when using vdbench and other tools doing basic IO in a native full 12Gbps SAS environment. StorageIO has also seen random read/write (4KB) in the neighborhood of Seagate spec sheets of 100,000 IOPs (reads) and 40,000 (writes). However, as with any workload your results and mileage will vary so make sure that they are relevant to your business and application needs.

[3] For example Seagate TurboBoost™ in an Enterprise 15K HDD with integrated flash
[4] MLC = Multi-layer Cell nand flash

Tiered Storage: Enterprise 12Gbps SAS flash SSD

Proof-points: Seagate Enterprise 1200 SSD with 12Gbps SAS and PCIe Gen 3

It should not come as a surprise that a SSD device should be faster than HDD. These proof points are provided to help make decisions on when, where and how to deploy SSD and related technologies. Also included is using caching software for VMware environments that utilize the SSD drives to accelerate read performance across other devices. 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, as a high-speed repository cache and Meta data lookup.

When to use Enterprise SSD

- ✓Compliment HDD based solutions
- ✓As a storage target or cache
- ✓Replace several slower HDDs
- ✓Eliminate HDD based short stroking
- ✓Virtual, cloud and physical servers

General server operating systems	Server page and swap volumes, temporary ephemeral storage of transient data including local caches or logs using dedicated SSD drives or SSD combined with cache software.
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. Leverage cache to speed up backup processing.
HPC, Big Data, DSS, Hadoop, SAP, Data Warehouse	Support "hot-band" data caching in a cost effective manner using SSD complimenting slower higher capacity size HDDs. Persistent storage for rapid refresh of in-memory database and key-value stores. Burst buffers, import, export, "scratch" areas along with high-speed ephemeral storage.
Email, Text and Voice Messaging	Microsoft Exchange and other email journals, mailbox or object repositories can leverage faster read and write I/Os.
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. Compliment server side SSD-based read and write-thru caching.
Server Virtualization	Fast storage for data stores and virtual disks supporting VMware vSphere/ESXi, Microsoft Hyper-V, KVM, Xen and others. Compliment virtual server read cache and I/O optimization using SSD as a cache or as a target device for writes to eliminate I/O bottlenecks across different applications. In addition to working with various software cache tools SSD is also key for enabling virtual SANs (VSANs) such as VMware V5.5 Virtual SAN using SSD and HDD.
Virtual Desktop Infrastructure (VDI)	SSD can be used as high performance storage for linked clone images, applications and data. Leverage fast read to support read ahead or pre-fetch to compliment cache solutions for high capacity HDD.

Table 1 – Example application and workload scenarios benefiting from SSDs

Tiered Storage: Enterprise 12Gbps SAS flash SSD

Proof Points and Enterprise SSD Validation

StorageIO has conducted various hands-on testing with the Seagate® 1200 Enterprise 12Gbps SAS SSD in our lab environment across different real world like application scenarios. All independent hands on testing were done by StorageIO using industry standard benchmarks and workload generators. These proof point workloads include general storage I/O performance characteristics profiling (e.g. reads, writes, random, sequential or various IOP size) to understand how these devices perform. Comparisons are across various HDD, HDD with TurboBoost™, and SSD devices in terms of IOPS, bandwidth and response time (latency). In addition to basic storage I/O profiling[5], the Enterprise 1200 12Gbps SAS SSD was also used with various SQL database workloads including TPC-B and TPC-E; along with VMware server virtualization among others use case scenarios.

Note that these proof points show the performance supported by a single drive for a given workload on a particular sized processor. Your results will vary however; StorageIO believes the results shown here can be improved with normal larger configurations using multiple drivers, faster servers and other enhancements.

Testing consisted of system under test, which hosted the drives being analyzed and an application such as Microsoft SQL Server 2012 for database and Microsoft Exchange for Email. In the case of SQL, a separate server was used to drive and generate transactions, which were handled by the database server with the various drives attached. Normal configurations (e.g. defaults, no tuning) were applied to both SQL and Exchange.

Role of PCIe Gen 3 and 12Gbps SAS

Faster servers and fast storage need fast I/O interfaces. This means leveraging fast PCIe Gen 3 to move data in a given amount of time at a lower latency to remove bottlenecks.

The combination of PCIe Gen 3 and 12Gbps SAS unlocks the performance potential of fast SSD devices with fast server processors. The result getting more work done in the same amount of time, or doing the same amount of work in less time. For transactional or smaller IO sized environments, this means more IOPs at lower latency while large IO applications can do more throughput bandwidth.

All devices underwent a multi-day or longer "burn-in" or initial usage period including running various traditional storage centric benchmarks such as vdbench, iorate and iometer, along with full drive formats. Drives were formatted for Windows using defaults.

All of the servers were running as virtual machine (VM) guests on VMware 5.5 (e.g. using vSphere Essentials) each with its own dedicated physical machine (PM). Drives were Raw Device Mapped (RDM) and configured as persistent independent to VMware to avoid any cache or deferred lazy writes. For the SQL Server workloads, the VMware Para Virtual SCSI (PVSCSI) driver was used while the VMware LSI SAS driver was used for Exchange.

5 Storage I/O profiling IOPs, latency and bandwidth: http://storageioblog.com/part-ii-iops-hdd-hhdd-ssd/

Tiered Storage: Enterprise 12Gbps SAS flash SSD

Proof Point: Email Messaging

For this proof-point, Microsoft Jet Stress Exchange performance workloads were placed (e.g. Exchange Database - 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. For the Exchange testing, the EDB was placed on the device being tested while its log files were placed on a separate Seagate 400GB Enterprise 12Gbps SAS SSD.

	400GB SSD 12GSAS	Ent. 15K HDD TurboBoost™ 600GB	Ent. 15K V4 HDD 600GB	6TB 7.2K v4 HDD 12Gb SAS	3TB 7.2K HDD
Avg. Total IOP/second	2,791	334	252	174	95
Avg. Reads resp. (msec)	1.4	6.6	8.2	11.6	17.4
Avg. Writes resp. (msec)	1.3	4.7	33.9	33.0	49.2
Avg. IOP (Read)/second	1,486	181	139	89	52
Avg. IOP (Write)/second	1,305	154	113	85	42

Table 2 – Email - Microsoft Exchange JetStress workload (single device)[6]

Table-2 shows performance improvements with Exchange workload using the Seagate 1200 Enterprise SSD compared to various other devices. Note the slightly better read performance for the SSD vs. writes. Figure 1 shows average read (blue) and write (green) Exchange IOPs side by side along with a total average (combined read and write) IOP (see light blue trend line). Also shown is average read response (dark blue solid line) and write response time (dotted green line). Vertical axis is a base 5 logarithmic scale.

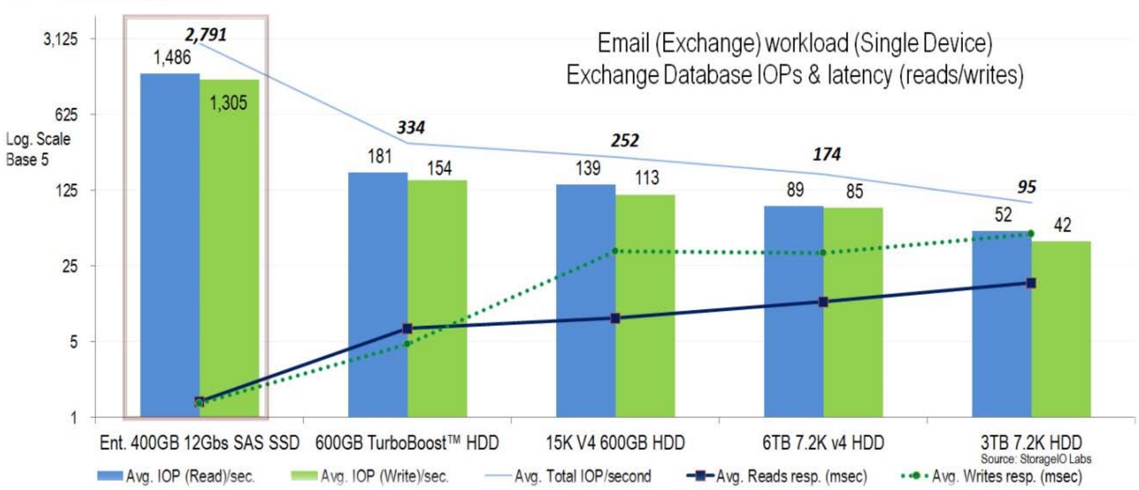

Figure 1 - Microsoft Exchange Jetstress workload (Database reads/writes and latency)

6 Test configuration: Seagate 400GB 12000 2.5" SSD (ST400FM0073) 12Gbps SAS, 600GB 2.5" Enterprise 15K with TurboBoost™ (ST600MX) 6 Gbps SAS, 600GB 2.5" Enterprise Enhanced 15K V4 (15K RPM) HDD (ST600MP) with 6 Gbps SAS, Seagate Enterprise Capacity Nearline (ST6000NM0014) 6TB 3.5" 7.2K RPM HDD 12 Gbps SAS and 3TB 7.2K SATA HDD. Email server hosted as guest on VMware vSphere/ESXi V5.5, Microsoft SBS2011 Service Pack 1 64 bit. Guest VM (VMware vSphere 5.5) was on a SSD based datastore, had a physical machine (host), with 14 GB DRAM, quad CPU (4 x 3.192GHz) Intel E3-1225 v300, with LSI 9300 series 12Gbps SAS adapters in a PCIe Gen 3 slot with Jet Stress 2010. 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 also persistent (no delayed writes). EDB was 300GB and workload ran for 8 hours.

Tiered Storage: Enterprise 12Gbps SAS flash SSD

Proof Point: Database, Data Warehouse Batch Update
SSDs 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 SSD 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 Microsoft SQL Server 2012 databases. Each test simulation had the SQL Server database (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.)
400GB Enterprise SSD	1	174	0.004
12Gbps SAS	20	1,145	0.13
	50	1,165	0.03
	100	1,205	0.07
600GB Enterprise 15K HDD	1	69	0.01
TurboBoost™	20	243	0.08
6Gbps SAS	50	287	0.17
	100	300	0.33
600GB Enterprise 15K HDD	1	65	0.13
6Gbps SAS	20	216	0.09
	50	230	0.21
	100	245	0.40
6TB 7.2K RPM HDD	1	26	0.02
12Gbps SAS	20	82	0.11
	50	111	0.20
	100	127	0.39
3TB 7.2K RPM HDD	1	31	0.30
6Gbps SATA	20	75	0.26
	50	82	0.61
	100	81	1.23

Table 3 - TPC-B[7] (batch update write) with 1, 20, 50 and 100 users (single device)

[7] Test configuration: Seagate 400GB 12000 2.5" SSD (ST400FM0073) 12Gbps SAS, 600GB 2.5" Enterprise 15K with TurboBoost™ (ST600MX) 6 Gbps SAS, 600GB 2.5" Enterprise Enhanced 15K V4 (15K RPM) HDD (ST600MP) with 6 Gbps SAS, Seagate Enterprise Capacity Nearline (ST6000NM0014) 6TB 3.5" 7.2K RPM HDD 12 Gbps SAS and 3TB Seagate 7.2K SATA HDD Workload generator and virtual clients Windows 7 Ultimate 64 bit. A SQL database was on Windows 7 guest. Guest VM (VMware vSphere 5.5) had a dedicated 14 GB DRAM, quad CPU (4 x 3.192GHz) Intel E3-1225 v300, with LSI 9300 series 12Gbps SAS adapters in a PCIe Gen 3 slot along with TPC-B (www.tpc.org) workloads. VM with guest OS along with SQL tempdb and masterdb resided on separate SSD based data store from devices being tested (e.g., where MDF (main database tables) and LDF (log file) 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) using VMware PVSCSI driver. MDF and LDF file sizes were 142GB and 26GB with scale factor of 10000, with each step running for one hour (10-minute preamble).

Tiered Storage: Enterprise 12Gbps SAS flash SSD

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 SSD device does more work at a lower response time vs. traditional HDDs.

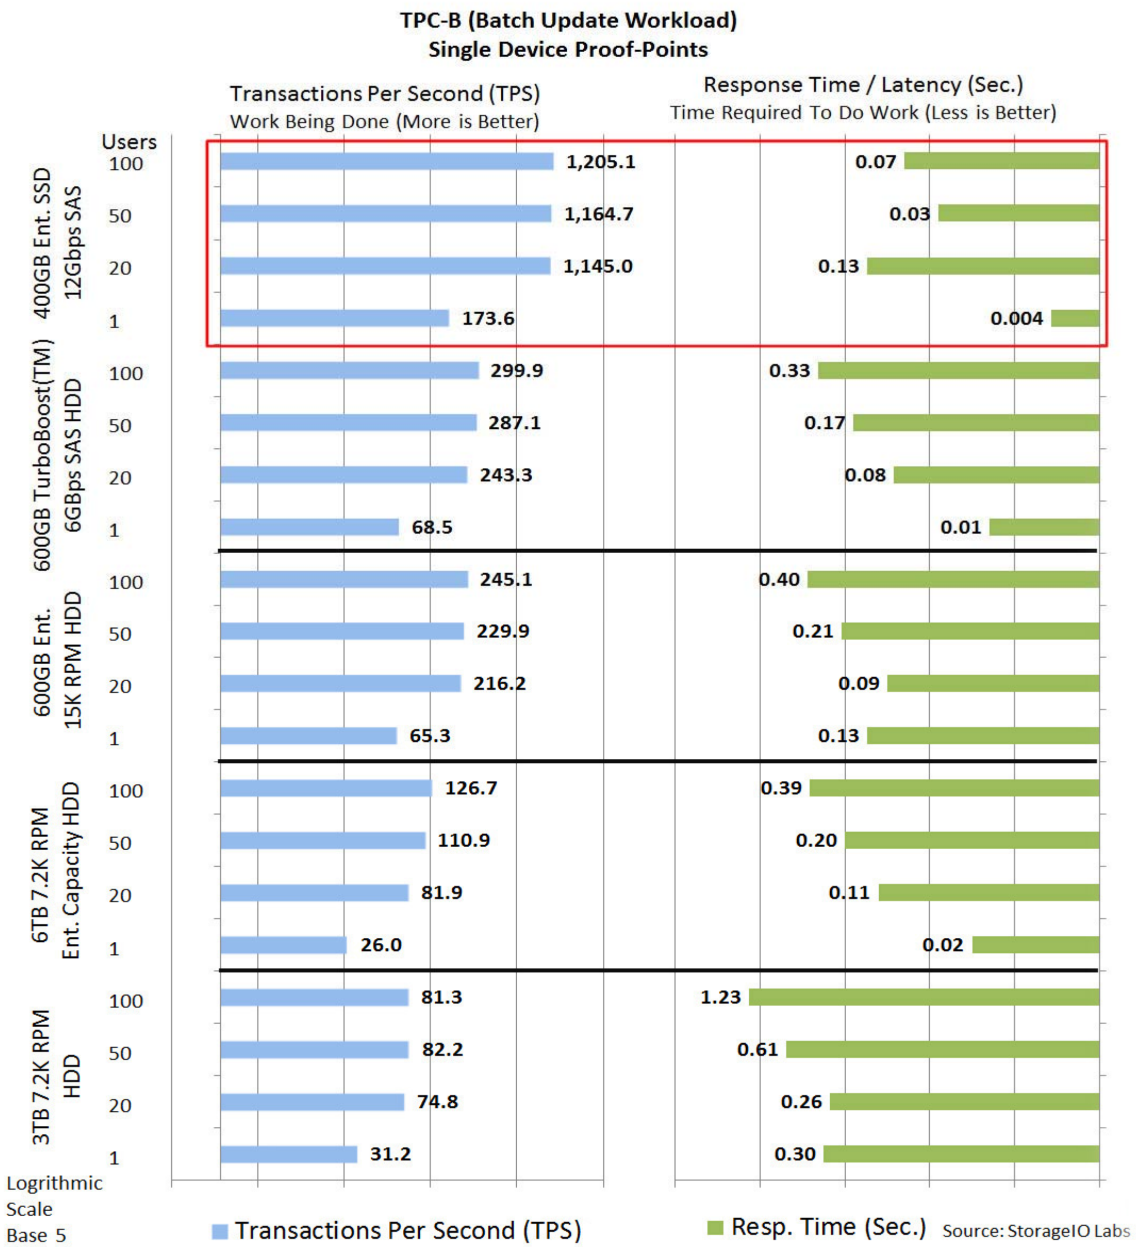

Figure 2 – TPC-B (batch update write) with 1, 20, 50 and 100 users (single device)

Tiered Storage: Enterprise 12Gbps SAS flash SSD

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.

Proof-points 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.)
Seagate Enterprise 1200	10	188.3	0.38
2.5" 400GB MLC flash SSD	20	185.4	0.23
12Gbps SAS	50	156.1	0.11
	100	123.4	0.07
Seagate Enterprise 15K HDD	10	10.2	2.58
2.5" 600GB with TurboBoost™	20	9.4	0.16
6Gbps SAS	50	9.4	0.86
	100	6.1	0.7
Seagate Enterprise 15K	10	4.2	5.73
2.5" 600GB HDD	20	4.1	3.2
6Gbps SAS	50	4.1	1.61
	100	4.0	0.1
Seagate Enterprise Capacity 6TB	10	3.7	25.57
3.5" 7.2K HDD	20	3.7	12.59
12Gbps SAS	50	3.1	6.0
	100	2.1	4.51

Table 4 – TPC-E[8] (OLTP transactional) with 10, 20, 50 and 100 users

[8] Test configuration: Seagate 400GB 12000 2.5" SSD (ST400FM0073) 12Gbps SAS, 600GB 2.5" Enterprise 15K with TurboBoost™ (ST600MX) 6 Gbps SAS, 600GB 2.5" Enterprise Enhanced 15K V4 (15K RPM) HDD (ST600MP) with 6 Gbps SAS, Seagate Enterprise Capacity Nearline (ST6000NM0014) 6TB 3.5" 7.2K RPM HDD 12 Gbps SAS. Workload generator and virtual clients Windows 7 Ultimate 64 bit. A SQL database was on Windows 7 guest. Guest VM (VMware vSphere 5.5) had a dedicated 14 GB DRAM, quad CPU (4 x 3.192GHz) Intel E3-1225 v300, with LSI 9300 series 12Gbps SAS adapters in a PCIe Gen 3 slot along with TPC-E (www.tpc.org) workloads. VM with guest OS along with SQL tempdb and masterdb resided on separate SSD based data store from devices being tested (e.g., where MDF (main database tables) and LDF (log file) 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) using VMware PVSCSI driver. MDF and LDF file sizes were 148GB and 98GB created with scale factor of 15, with each step running for one hour (10-minute preamble).

Tiered Storage: Enterprise 12Gbps SAS flash SSD

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 Seagate Enterprise 1200 SSD is shown on the top of figure 3 with a red box around its results. Note how the SSD as a lower latency while doing more work compared to the other traditional HDDs.

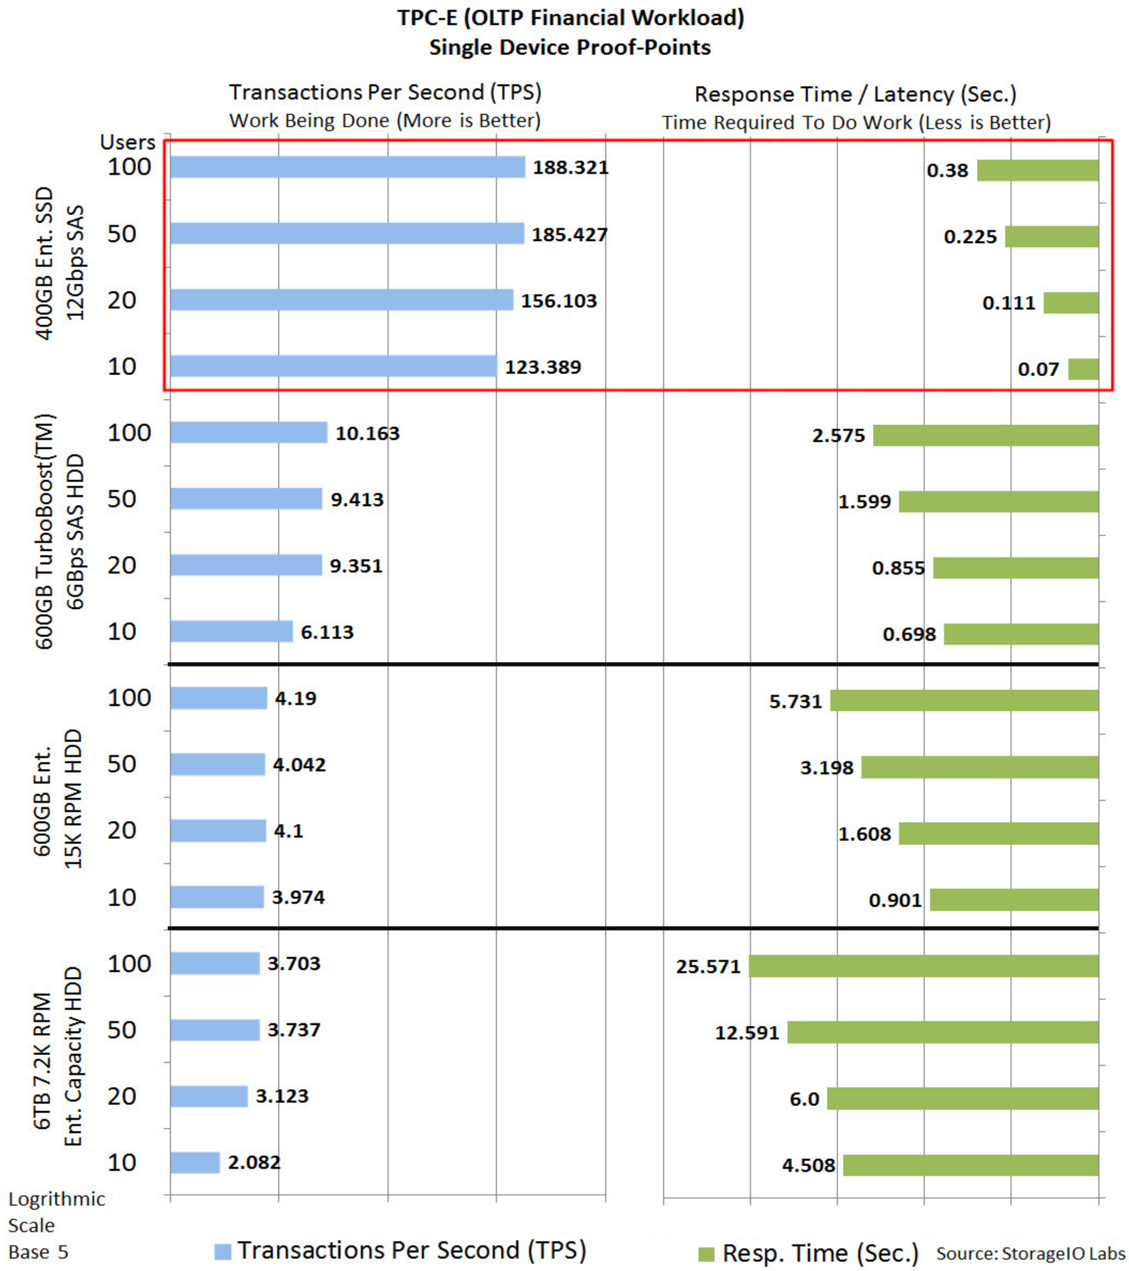

Figure 3 – TPC-E (OLTP transactional) with 10, 20, 50 and 100 users

Tiered Storage: Enterprise 12Gbps SAS flash SSD

Proof Point: What about traditional IOPs

The proof points in this white paper are from an applications focus perspective representing more of an end-to-end real-world situation. While they are not included in this white paper, StorageIO has run traditional storage component focus workloads, which can be found at **www.storageioblog.com**. These include tools such as Iometer, iorate, vdbench among others for various IO sizes, mixed, random, sequential, reads, writes along with "hot-band" [9] across different number of threads (concurrent users).

What about SSD and caching software

The previous proof-points focused on SSD as a target or storage device. In the following proof-points, the Seagate Enterprise 1200 SSD is used as a shared read cache (write-through). Using a write-through cache enables a given amount of SSD to provide a performance benefit to other local and networked storage devices.

Aggregation causes aggravation with I/O bottlenecks because of consolidation using server virtualization. Figure-4 shows non-virtualized servers with their own dedicated physical machine (PM) and I/O resources. When various servers are virtualized and hosted by a common host (physical machine), their various workloads compete with each other for I/O and other resources. In addition to competing for I/O performance resources, these different servers also tend to have diverse workloads.

Virtunet Systems VirtuCache

SSD cache proof-point workloads were run on VMware (ESXi 5.5.0 1746018) based systems (see figure-4). For the cache tool Virtunet Systems VirtuCache (write-through or read cache) was used (version 1.0-1vmw.5.5.0.1331820.64ab842.261).

The VirtuCache software installs into the VMware hypervisor. Any type of flash SSD can be used including internal SAS and SATA drives, or PCIe cards along with external shared iSCSI, SAS or FC SSD enabled storage solutions. Caching was enabled on a per VM and per device basis (local and shared block storage).

VirtuCache can be configured to use an entire SSD space capacity, or a portion can be allocated for use as another partition. This capability enables some space to be used as VMware SSD storage to place write intensive applications or workload. Learn more about Virtucache at www.virtunetsystems.com.

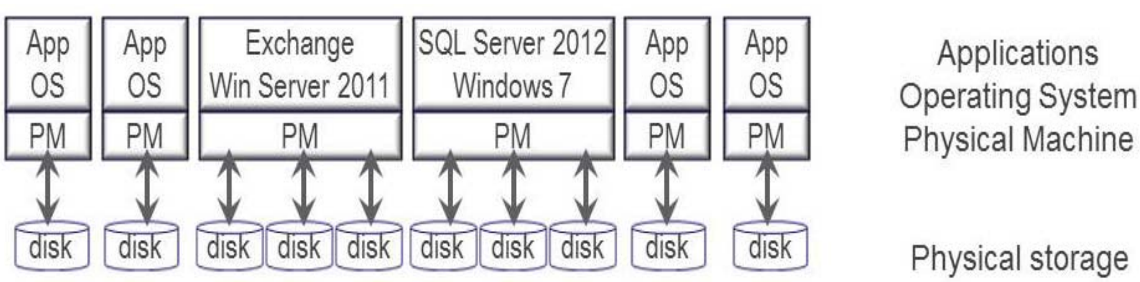

Figure-4 Non-virtualized servers with their own Physical Machines (PMs)

[9] "Hot-Band" is part of the SNIA Emerald energy effectiveness metrics for looking at sustained storage performance using tools such as vdbench. Refer to StorageIOblog.com and StorageIO.com/SSD to learn more.

Tiered Storage: Enterprise 12Gbps SAS flash SSD

The aggregation and consolidation result is a blend of random, sequential, large, small, read and write characteristics. These different storage I/O characteristics are mixed up and need to be handled by the underlying I/O capabilities of the physical machine and hypervisor. As a result, a common deployment for SSD in addition to as a target device for storing data is as a cache to reduce bottlenecks for traditional spinning HDD. Figure-5 shows aggregation causing aggravation with the result being I/O bottlenecks as various applications performance needs converge and compete with each other.

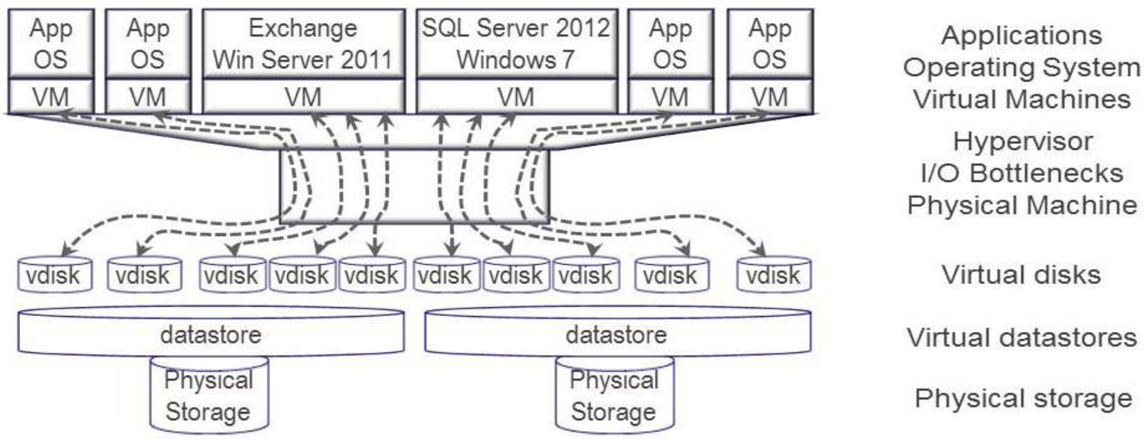

Figure-5 I/O aggravation (bottleneck) from aggregating many servers (thick dotted I/O lines)

In figure-6, a solution is shown introducing I/O caching with SSD to help mitigate or eliminate the effects of server consolation causing performance aggravations.

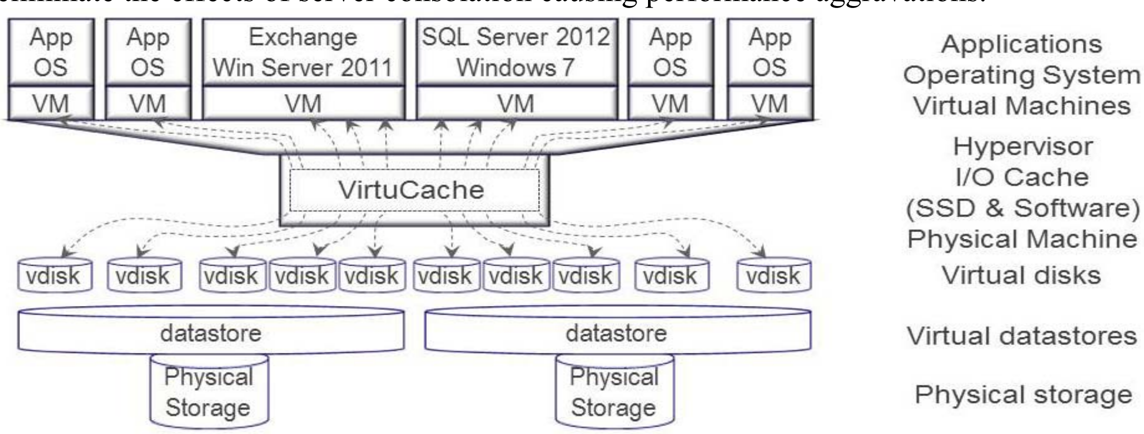

Figure-6 Using caching with SSD alleviates aggravation from aggregation (thinner I/O lines)

For these proof-points, the objective was to create an I/O bottleneck resulting from multiple VMs in a virtual server environment performing application work. In this proof-point, multiple competing VMs including a SQL database and an Exchange server shared the same underlying storage I/O infrastructure including HDDs. The 6TB (Enterprise Capacity) HDD was configured as a VMware datastore and allocated as virtual disks to the VMs. Workloads were then run concurrently to create an I/O bottleneck for both cached and non-cached results, figure-7 on next page.

Tiered Storage: Enterprise 12Gbps SAS flash SSD

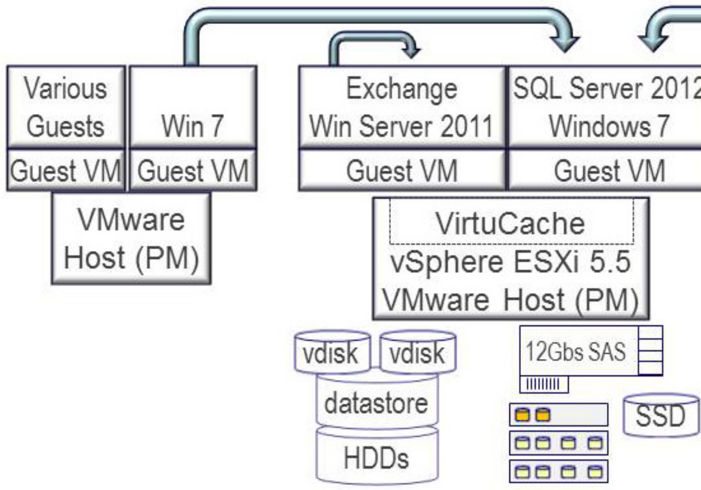

Figure 7 – Mixed workload multiple with VM (I/O bottleneck) with cache configuration[10]

Figure 8 shows two sets of proof points (cached and non-cached) with three workloads. The workloads consisted of concurrent Exchange and SQL database (TPC-B and TPC-E) running on separate virtual machine (VM) on the same physical machine host.

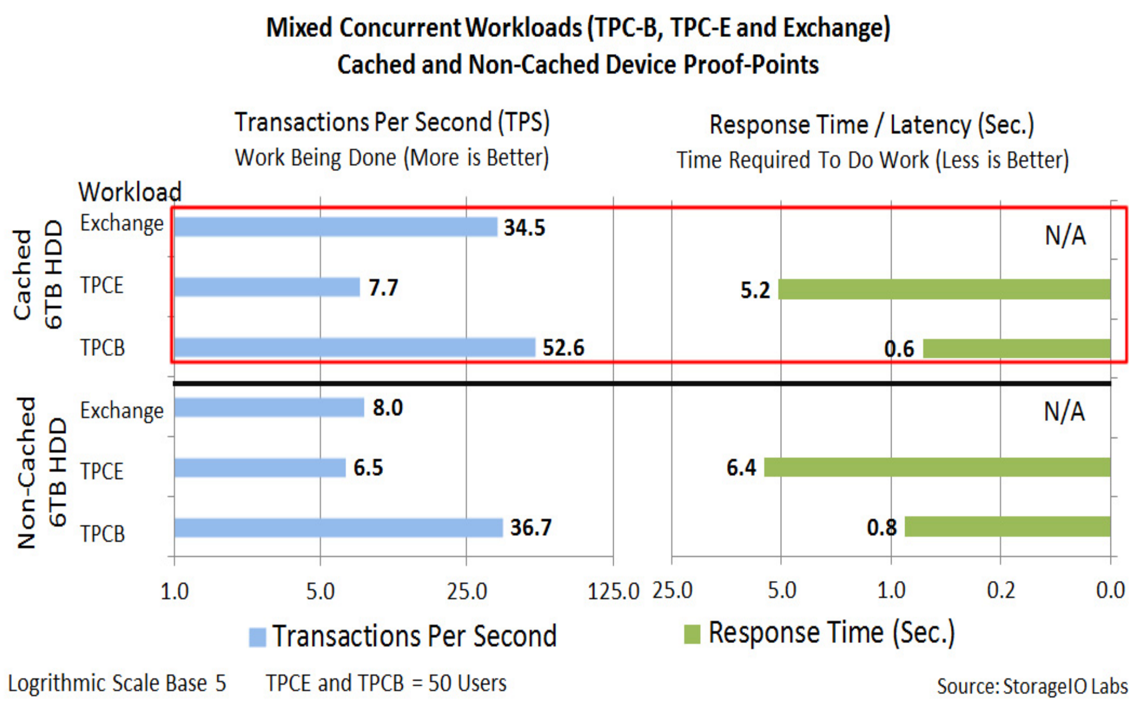

Figure 8 – Mixed workloads Cached (top) and non-cached (bottom)

[10] Note that the servers and virtual machine (VM) configurations for the mixed workload and cache test were the same for previous proof-points except for caching,, data placement, and use of VMware virtual disks on data stores vs. RDM (previous proof-points). Note that VMware virtual disks were configured as persistent independent with VMware pvscsi drivers. Cache device was a Seagate Enterprise 1200 400GB 12Gbps SAS SSD.

Tiered Storage: Enterprise 12Gbps SAS flash SSD

The cache and non-cached mixed workloads shown in figure 8 demonstrate how an SSD based read-cache can help to alleviate I/O bottlenecks. This is an example of addressing the aggravation caused by aggregation of different competing workloads that are consolidated with server virtualization.

For the workloads in figure 8, all data (database tables and logs) were placed on VMware virtual disks created from a datastore using a single 7.2K 6TB 12Gbps SAS HDD (e.g. Seagate Enterprise Capacity). The guest VM system disks which included paging, applications and other data files were virtual disks using a separate datastore mapped to a single 7.2K 1TB HDD. Each workload ran for eight hours with the TPC-B and TPC-E having 50 simulated users. For the TPC-B and TPC-E workloads, two separate servers were used to drive the transaction requests to the SQL Server 2012 database.

For the cached tests, a Seagate Enterprise 1200 400GB 12Gbps SAS SSD was used as the backing store for the cache software (Virtunet Systems Virtucache) that was installed and configured on the VMware host. During the cached tests, the physical HDD for the data files (e.g. 6TB HDD) and system volumes (1TB HDD) were read cache enabled. All caching was disabled for the non-cached workloads.

Note that this was only a read cache, which has the side benefit of off-loading those activities enabling the HDD to focus on writes, or read-ahead. Also note that the combined TPC-E, TPC-B and Exchange databases, logs and associated files represented over 600GB of data, there was also the combined space and thus cache impact of the two system volumes and their data. [11] This simple workload and configuration is representative of how SSD caching can complement high-capacity HDDs.

Additional and Future Proof Points - Server Virtualization

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

[11] Additional information about the proof-points can be found at www.storageioblog.com

Tiered Storage: Enterprise 12Gbps SAS flash SSD

Summary
Moving forward there is the notion that flash SSD will be everywhere. There is a difference between all data on flash SSD vs. having some amount of SSD involved in preserving, serving and protecting (storing) information.

Key themes include:
- Aggregation can cause aggravation which SSD can alleviate
- A relative small amount of flash SSD in the right place can go a long way
- Fast flash storage needs fast server storage I/O access hardware and software
- Locality of reference with data close to applications is a performance enabler
- Flash SSD everywhere does not mean everything has to be SSD based
- Having some amount of flash in different places is important for flash everywhere
- Different applications have various performance characteristics
- SSD as a storage device or persistent cache can speed up IOPs and bandwidth

Flash and SSD are in your future, this comes back to the questions of how much flash SSD do you need, along with where to put it, how to use it and when.

Learn more at Seagate landing page located at:
http://www.seagate.com/internal-hard-drives/solid-state-hybrid/1200-ssd/

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 five-time VMware vExpert. Learn more at www.storageio.com.