

VMware Meets the Seagate® 1200 SSD and VirtuNet Systems

Technology Paper

Boost Performance and Improve Server Utilization With 12Gb/s SAS SSDs and Server Side Caching Software

Challenge: Increase Virtualized Server Speed, Density Increase

The benefits and appeal of server virtualization are becoming increasingly well known and understood. According to VMware, 2013 marked the crossover point at which virtualized server deployments exceeded physical shipment of servers. As businesses look to the future, it's clear that virtualization is here to stay.

This growth in server utilization has resulted in greater loads on server resources and increased I/O to networked storage. Modern-day servers with multiple cores and hardware-assisted virtualization capabilities can keep pace with the demand arising from an increased number of virtual machines (VM) sessions, but when the disk I/O from each VM is sent through the hypervisor and interspersed with I/O from other VMs on the same server, the I/O pattern becomes a more concentrated, complex mix of reads and writes at varying transfer sizes. This is often referred to as the I/O blender effect. These more challenging I/O workloads create the ideal use case for SSD (solid state drive) technology.

High-end, high-performance SAN and NAS storage subsystems (and their supporting storage network infrastructure) are often used to deliver the performance needed to support these demanding workloads in virtualized server environments. Nearly all modern SAN/NAS solutions utilize some quantity of flash storage to ensure high-speed access to the virtualized server's most frequently requested data.

VMware Meets the Seagate® 1200 SSD and Virtunet Systems



While using flash storage to increase storage subsystem performance offers a compelling value proposition, this approach does have practical limitations. First, there is a limited number of SSD drives that can be added before the storage controllers are saturated and become a performance bottleneck. Second, all SAN and NAS infrastructures must communicate through I/O channels, which are limited by the speed of the basic interconnect fabric (16Gb Fibre Channel, 10Gb Ethernet, or even slower channel speeds). Even all-flash arrays, which boast of architectures optimized for high-performance business and research applications, are subject to this storage interconnect limitation.

Solution: A New Approach to Flash Implementation

By locating flash storage inside the server, it becomes accessible at the much higher speeds of the local bus, eliminating the latency of the storage network fabric and thus enabling effective and reliable access to extremely high I/O rates.

With this approach, previous I/O-bound performance constraints are raised to significantly higher thresholds. This not only delivers a measureable boost in business value by enabling much higher VM density per server, it does so without costly upgrades to the storage fabric or SAN/NAS storage subsystem. By moving significant amounts of I/O traffic off the interconnect fabric to the server, many users see improved response times from SAN/NAS storage subsystems as well.

Seagate 1200 SSD and Caching Software

As noted earlier, flash-based devices are ideal for solving I/O throughput issues, especially when the workload is read-dominated. However, because flash remains more expensive per/GB than traditional hard drives, it is most cost-effective in smaller quantities. For example, caching software can be used with small quantities of ultra-fast SSDs as a front-end cache to complement slower, capacity-optimized disk drive storage.

One example of a more sophisticated high-performance flash-based solutions is to utilize SAS-based SSDs from Seagate in concert with caching software from Virtunet Systems. Very good results can be achieved with just two 1200 SSD and Virtunet Systems Caching software. This bundled solution greatly improves application performance, dramatically increasing IOPS and significantly reducing latency, especially when compared with traditional SAN configurations or PCIe solutions like the Fusion-io with ioTurbine caching software (Figure 1).





Figure 1. Comparing Seagate 1200 SSD and Virtunet Systems, Fusion-io with ioTurbine, and Fibre Channel SAN

Server-side caching also provides tremendous benefit in terms of VM densities; that is, the number of virtual machine instances that can run concurrently on a single physical server (Figure 2).



Number of VM Instances on One Physical Server

Figure 2. VM densities of typical Fibre Channel SAN without server-side caching vs. Seagate 1200 SSD with Virtunet

Performance Boost From Accelerator Caching Software

Contributing to the performance boost delivered by the bundle's SSDs is Virtunet Systems Caching software, which addresses the relatively new approach of virtualized servers that use networked storage along with in-server flash as the caching media.

VMware Meets the Seagate® 1200 SSD and Virtunet Systems



Virtunet Systems software is arguably one of the best caching solutions on the market optimized for VMware environments. It supports read and write caching, where reads are cached on the in-server SSD. It also supports write-through caching (what we tested), in which new data is written to both the in-server solid state storage and external shared storage. While there are many reasons the Accelerator Cache software is superior to other caching solutions, a few key differentiators include:

- A hypervisor kernel-based implementation, making all caching decisions within the VMware kernel—without modifying the native VMware kernel. This is unlike competitive approaches that utilize agent software that must be loaded into each Guest VM, or others that modify the native VMware kernel.
- Sophisticated pattern-matching algorithms are used to cache only one copy of identical I/O blocks; this is especially efficient in VDI deployments, where cloned image copies of the golden image are made. Competitive solutions cache each clone instance, thereby wasting valuable solid state resources.
- Full support of the advanced features in VMware's vMotion, Storage vMotion and High Availability components.

Certainly, all server-side solid state deployments use some sort of caching solution, including the basic caching functionality that's included with any operating system and in third-party packages, but Seagate integrates the Virtunet caching software with the Seagate 1200 SSD to optimize the design for VMware environments.

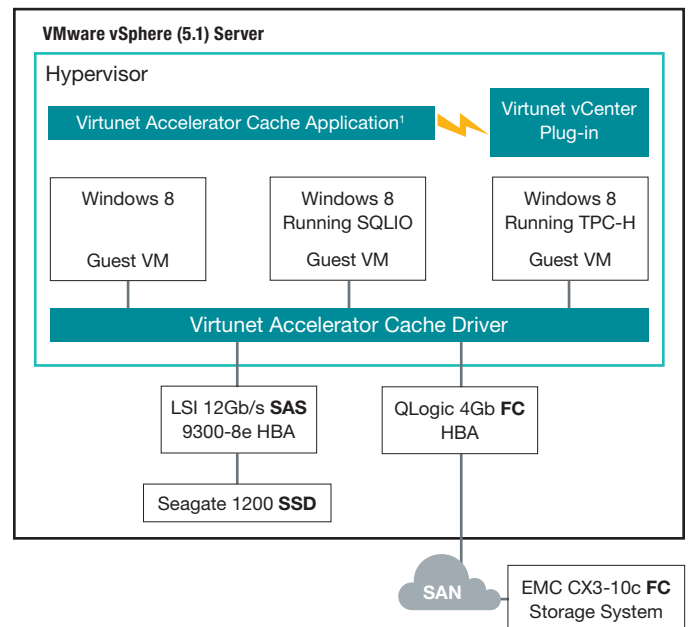
Test Methodologies

Seagate compared results of two widely-used synthetic workload benchmarks (SQLIO and TPC-H) running within a Windows Server 2008 VM instance. In a VMware environment, Seagate compared the results of its 1200 SSD and Virtunet solution against the results produced by a competing offering (Figure 3).

- SQLIO Disk Subsystem Benchmark Tool: www.microsoft.com/en-us/download/details.aspx?id=20163/
- TPC-H: www.tpc.org/tpch/

Seagate tested raw throughput and latencies for both solutions under the following conditions:

- One server with two Xeon 5620 (4 cores each) and 96GB RAM running VMware 5.1
- Connected over dual-port QLogic FC HBA (4Gb/s) to EMC CX3-10c SAN



1 One VM needed in the cluster/resource pool to host Virtunet Accelerator Cache Application
Figure 3. Testing environment

- Server in VMware cluster, version 5.1
- Accelerator cached data from LUN on CX3-10 to various Seagate 1200 SSD
- Seagate 1200 SSD, 800GB capacity
- LSI SAS 9300-8e HBA
- SQLIO tests run for Seq/Rand Read and Writes; from 2–64 threads; and from 2–128 pending queue requests; 4KB transfer sizes; 4- to 6-hour tests for each data point

Note that queue depth configuration is a primary source of easily avoidable VMware bottlenecks; further configuration details from testing include the following.

Adapter queue depth (QD) for:

- HBA—2176: 128
- Fusion-io ioDrive: 128

Disk queue depth (QD) for:

- Seagate 1200 SSD: 64
- LUN: 64
- Fusion-io ioDrive2: 32
- Six RAID10 SAS drives: 232
- One Seagate 1200 SSD drive with Series 7 RAID card, configured as volume: 116

In addition, Windows VM had 10 cores and 8GB assigned to it, with a PVSCSI driver configured for QD of 255.

VMware Meets the Seagate® 1200 SSD and Virtunet Systems



Analyzing All Solutions and All Results

Seagate compared the results of its bundling of 1200 SSD and Virtunet Software with the results of a similar configuration (Fusion-io ioDrive2 using ioTurbine caching software); it also logged the baseline performance delivered by an EMC SAN without the aid of any server-side caching. All testing configurations were running VMware 5.1 ESXi hypervisor (Figure 4). When the results were reviewed, Seagate found that their expectations were surpassed.

4K Random Read Performance With 8 Threads vs. Queue Depth per Thread




Figure 4. Performance results of all solutions under light loads

When comparing caching performance measured under extreme loads, the 1200 SSD and Virtunet configuration consistently delivered higher performance than the Fusion-io configuration (Figure 5).

4K Random Read Performance With 64 Threads vs. Queue Depth per Thread

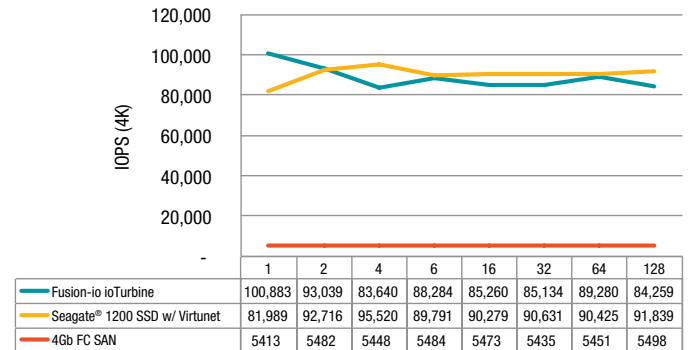



Figure 5. All solutions under extreme workloads

Operating in write-intensive configuration, performance was limited by the back-end store (in this case, a 4Gb/s FC SAN from a leading SAN vendor), which employs caching software that utilizes a write-through caching algorithm. However, the 1200 SSD and Virtunet configuration consistently outperformed the Fusion-io ioDrive2 solution by a significant margin. Viewing the test results indicates that the 1200 SSD and Virtunet configuration achieved the maximum performance possible (Figure 6)



4K Sequential Write Performance With 8 Threads vs. Queue Depth per Thread

Figure 6. Under write-intensive workloads, Seagate 1200 SSD (SAS) delivers maximum performance possible

VMware Meets the Seagate® 1200 SSD and Virtunet Systems



Conclusion

The Seagate 1200 SSD and Virtunet Systems Caching Software offers an efficient, cost-effective way to increase the performance and scalability of enterprise VMware environments at a fraction of the price of high-end enterprise SAN/NAS equipment, more expensive PCIe flash products or costly all-flash array appliances. Moreover, its simpler deployment—utilizing well-established building block technologies like HBAs and SAS drives—reduces the risk of suffering driver stability or interoperability issues (often encountered in PCIe solutions from less-mature companies). The Seagate 1200 SSD with Virtunet speeds VM response times across the board, and boosts the number of VM sessions that can be hosted on a single ESXi server—resulting in improved ROI and utilization of VMware environments.

Seagate 1200 SSD

The Seagate 1200 SSD offers the highest levels of data integrity, security, resilience and scalability for the most demanding write-intensive environments. Its industry-leading, dual-port 12Gb/s SAS interface provides twice the speed of the 6Gb/s SAS interface and optimizes investments in current SAS infrastructures—all while delivering maximum reliability and abundant capacity (up to 800GB).

Virtunet Systems

Virtunet Systems develops software products to address the performance aspects of storage for VMware, Hyper-V and Linux-based hypervisors.

The first product, VirtuCache, is a software solution that caches frequently used blocks of data from traditional disk based shared storage to in-server solid state storage, thus eliminating storage I/O bottlenecks for data-intensive workloads deployed within virtual machines.

www.seagate.com

AMERICAS Seagate Technology LLC 10200 South De Anza Boulevard, Cupertino, California 95014, United States, 408-658-1000
ASIA/PACIFIC Seagate Singapore International Headquarters Pte. Ltd. 7000 Ang Mo Kio Avenue 5, Singapore 569877, 65-6485-3888
EUROPE, MIDDLE EAST AND AFRICA Seagate Technology SAS 16-18, rue du Dôme, 92100 Boulogne-Billancourt, France, 33 1-4186 10 00

© 2014 Seagate Technology LLC. All rights reserved. Printed in USA. Seagate, Seagate Technology and the Wave logo are registered trademarks of Seagate Technology LLC in the United States and/or other countries. VMware is the trademark of VMware, Inc. All other trademarks or registered trademarks are the property of their respective owners. When referring to hard drive capacity, one gigabyte, or GB, equals one billion bytes and one terabyte, or TB, equals one trillion bytes. Your computer's operating system may use a different standard of measurement and report a lower capacity. In addition, some of the listed capacity is used for formatting and other functions, and thus will not be available for data storage. Actual data rates may vary depending on operating environment and other factors. Seagate reserves the right to change, without notice, product offerings or specifications. TP672.1-1410US, October 2014