Clipper Notes™
VENDOR-AGNOSTIC EXPLANATIONS AND ADVICE FOR THE INFORMATION TECHNOLOGY BUYER
Navigating Information Technology Horizons℠

Published Since 2007 Report #TCG2009006 February 10, 2009

A New Tier of Storage Appears — Faster, Solid-State Drives State Their Case

Analyst: David Reine

Management Summary

Whenever anyone goes shopping for a new vehicle, two of the most important factors to consider are: *Will it carry what you need to transport* and *what kind of performance will it deliver?* In addition, the price of gas also positions *fuel economy* as a significant factor in the eye, and the pocketbook, of the buyer. If you are looking to transport people, seating area becomes important. If you need to move packages, cargo space is important. In either case, gas efficiency, expressed in miles per gallon (MPG) can play an important role. The higher the MPG, the lower the operational cost. The extraordinary MPG ratings of hybrid cars, which mean lower operational costs, can become more of a factor than the typically higher acquisition price. Hybrids provide some drivers with the performance and comfort that they seek while still preserving natural resources. A hybrid has an engine that uses a traditional (you might call it *legacy*) internal combustion engine along with a new electrical power source. Using a combination of the two, the consumer can achieve an efficiency in excess of 40 MPG with operational savings in terms of gas and money, although, admittedly, with a higher acquisition cost.

In today's data center, a similar hybrid configuration is required in order to meet the high capacity demands of many Tier 2 applications, such as backup and restore, as well as meeting the Tier 1 requirements for performance and reliability, to satisfy the service levels required by mission-critical applications, while keeping configuration demands within the bounds of energy limitations. Many mission-critical, high performance applications, such as OLTP or financial analysis, place a critical requirement on performance and reliability. The number of IOPS generated by these applications can far exceed the capability of a single HDD, or even hundreds of HDDs. In these cases, the data center must provision many more drives than the size of the data would warrant, distributing the data over many drives in order to provide more read heads and I/O pipes to transmit transactions to the server network. For some applications, performance rules the roost. If you can not respond to the customer in a timely fashion, that prospect will go elsewhere. Today, with the Internet, it is easier to find responsive vendors than ever before. Unfortunately, performance is not the only factor; energy and floor space *are* issues today and must be considered when deploying any hybrid IT infrastructure.

Today, the IT staff building a SAN has a new, innovative technology to work with in the form of Solid State Disks (SSDs). In addition to the traditional Tier-1 FC and Tier-2 high-capacity SATA HDDs, IT can deploy Tier-0 SSDs with significantly higher levels of performance, while at the same time, significantly lower power consumption. SSDs are not right for all environments due to a higher acquisition cost, but if performance is an issue, you may be able to replace tens, or even hundreds, of expensive FC drives with only a few SSDs, improving response time and lowering energy costs. To learn more about Solid State Disks, please read on.

IN THIS ISSUE

The Clipper Group, Inc. - Technology Acquisition Consultants ♦ Strategic Advisors
888 Worcester Street ♦ Suite 140 ♦ Wellesley, Massachusetts 02482 ♦ U.S.A. ♦ 781-235-0085 ♦ 781-235-5454 FAX
Visit Clipper at *www.clipper.com* ♦ Send comments to ***editor@clipper.com***

Data Center Challenges

Challenges are growing in the data centers of every enterprise around the world. They center on each data center's ability to resolve the three most important pain points facing a staff coping with an expanding storage requirement during an economic downturn.

- The total cost of ownership (TCO) of the IT infrastructure,
- The reliability of the infrastructure, and
- Its performance.

Total Cost of Ownership

The IT staff must be concerned about the TCO of all IT components, not just acquisition costs. The costs of energy, maintenance, administration, and floor space may contribute as much, if not more, to the ROI equation as acquisition costs. With the global economy in a downward spiral, IT budgets are frozen, if not reduced; CIOs are implementing cost-reduction plans wherever possible. Recently, data centers have been engaged in specific programs to reduce costs by consolidating and virtualizing their server environment to reduce the number of servers in the IT infrastructure, while at the same time making them more efficient, with the deployment of multi-socket, multi-core, multi-threaded microprocessors, and continuing to meet performance SLAs. The major driving forces behind the need to control TCO have been CPU inefficiencies, along with the cost and availability of energy to both drive the compute environment and cool the data center. No matter what changes are made, however, reliability and performance can not suffer.

With the availability of the latest x86 processors, such as *Xeon* (*Nehalem*) from Intel and the newest *Opteron* (*Shanghai*) from AMD, the data center now has the capability to achieve greater processor efficiency from the compute engine through consolidation and virtualization, driving an increasing number of mission-critical application operations, placing ever greater demands on an overtaxed storage area network (SAN). Unfortunately, the mission-critical storage environment has not been able to keep pace with the innovation in processor technology.

Reliability

The reliability of any IT infrastructure is a key factor for any enterprise in its ability to communicate with customers, vendors, and other partners. The MTBF for both the server network and the storage area network (SAN) help to determine total system availability. In addition, an efficient and cost-effective backup and recovery system will enable the enterprise to get back up and running in the event of a failure. It is important to remember that the cost of many backup/recovery applications is predicated on the total amount of storage under the purview of the backup application.

Storage Performance

Over the past 20 years, the data center has seen the rotational speed of disk media advance from 7200 RPM, to 10K and 15K RPM, with few prospects for improvements in this area in the foreseeable future, due to the physical limitations of spinning media. Seek time has been reduced from ~9ms to ~4ms, with latency time being halved from 4ms down to 2-3ms, resulting in an improvement in access time from ~13 ms down to 6-7ms, setting a lower limit on read response time. Unfortunately, these physical improvements only result in a typical hard disk drive (HDD) performance level of 150-200 IOPS[1] per drive, with a maximum of 300-400 IOPS under the best of conditions.

At the same time, the capacity of an enterprise-quality HDD has grown to 450GB for a 15K FC HDD and 1 TB for a 7200 RPM SATA HDD. These higher *capacities* indicate that a mission-critical application could be deployed with fewer drives. Unfortunately, while the performance level may be satisfactory for single-socket, single-core *infrastructure servers*, it is definitely an issue in many of today's high-performance environments that require IOPS levels measured in the tens of thousands, and even higher. This requires the data center to over-provision[2] the SAN with hundreds of HDDs to meet the SLA, even though the application may not need that kind of capacity. These superfluous HDDs will also burn up a significant amount of unneeded energy and floor space.

The processor industry has experienced a paradigm shift from single-socket, single-core CPUs to quad-processor blades, with quad-core architecture, providing the IT staff with 16 multi-threaded cores in "single server" packaging, with ever-increasing clock speeds. Storage technology, on the other hand, has been fairly stagnant with its improvements, at least until now.

[1] Input/outputs per second.

[2] In order to meet required IOPS levels, the IT staff may have to distribute mission-critical data over multiple HDDs in order to put a sufficient number of IOPS in play. This could result in HDDs with, say, only 10% utilization, or less.

Because of the difference in performance between processors and HDDs, the data center may require a significant increase in the number of HDDs required in high-demand SANs to provide an acceptable access time and a sufficient level of IOPS to achieve the required SLA for *system-wide* performance. All of the efficiencies gained through innovation in processor technology will be wasted if storage cannot perform I/Os fast enough. Currently, this required level of performance is only available in solid-state drives (SSDs), which can replace spinning media with memory chips.

SSDs are not a new technology. They have been available for over 30 years in the form of RAM Disks. The original models, developed by StorageTek in 1978, were very fast, but, unfortunately, extremely expensive, as both the memory and batteries were susceptible to frequent and costly failures. It has only been in the past few years that the technology has evolved to flash memory, with a commensurate reduction in acquisition cost along with improved performance and reliability, along with reduced energy consumption.

What is a Solid State Drive?

A solid-state drive (SSD) is a data storage device that uses non-volatile random-access memory, typically a hybrid with a combination of both flash memory and DRAM that can be electrically erased and rewritten, to store persistent data with a conventional disk interface at very high performance levels. This enables an SSD to emulate an HDD, easily and transparently replacing it in most applications. With lower power requirements[3], SSDs run cooler and achieve greater reliability, although with a shorter life expectancy.

The original usage of the term *solid-state* refers to the use of semiconductor devices to distinguish solid-state electronics from rotating electromechanical devices. With no moving parts, solid-state drives are less fragile than hard disks and also are silent, as there are no mechanical delays. They usually enjoy very low access time and latency, with latencies at least 10 times faster than the fastest HDDs. In fact, when considering random read performance, an SDD can deliver the same throughput as 30 or more HDDs while consuming significantly less power, with a typical SSD using less than half the energy of an HDD. Depending upon the application, this may enable the IT staff to reduce – signifycantly – the number of spinning drives required for the SAN.

The first flash-based SSDs appeared over a decade ago, however, as a result of the high demand for flash memory for a variety of hand-held devices, and recent innovations in flash memory, SSDs have gotten better, faster, and less expensive, although they remain more expensive than HDDs in terms of acquisition cost, on a one-to-one basis. Flash SSDs are available in two technologies: *single-level cell (SLC)* and *multi-level cell (MLC)*. SLC Flash is very fast and more expensive, but with greater reliability and lower capacity. MLC Flash has higher capacity, but is slower and less expensive. SLC Flash provides a longer life expectancy with 100K writes as compared to 10K writes for MLC, making it a more likely choice for the enterprise data center. When we consider performance as *the* critical factor, the difference in cost for SSD becomes significantly less for high-demand processing.

With an access time measured in microseconds (typically, 20-120), an SSD can be 250 or more times faster than an HDD with an access time measured in milliseconds (typically, 3-5). This enables the IT staff to make better use of their expensive x86 server network. Instead of wait-states, the server can be engaged in increasing performance via higher IOPS, increasing transaction volumes, thus improving the bottom-line for the enterprise. With a read performance level upwards of 50,000 IOPS for an SSD[4], the data center can achieve a better relative TCO then with HDDs rated at only 300-400 IOPS (in read mode). A high-performance (FC) HDD typically carries a cost of $1-$2 per gigabyte while an SSD has a premium price of anywhere from $15-$100 per gigabyte. This metric changes significantly, however, when we look at $/IOPS, especially in random read intensive environments. While an HDD has an estimated cost of $1-$2 per IOP, a fully utilized SDD has a cost of only $0.20 per IOP in an I/O bound application. If you are forced to deploy hundreds of HDDs to satisfy SLA levels, the TCO advantages of replacing them with only a few SSDs is obvious, especially when you factor in the storage management charges, which often are based upon total storage capacity, for over-provisioned HDDs. When you include the issue of the reliability of SSDs as compared to their electro-mechanical brethren, the data center can achieve further cost advantages.

[3] SSDs consume 1/10 the power of HDDs.

[4] Write performance is in the area of 17K IOPS.

Admittedly, there has been a concern about SSDs in the past, concerning the number of times any given block can be erased and rewritten (currently about 100,000 write cycles for SLC flash). However, innovations in flash technology have resulted in new "wear-leveling" algorithms enabling the application to use the SSDs' entire capacity evenly, to avoid failures, and increase drive life. The MTBF of an enterprise-class SSD should now exceed that of an HDD.

SSDs can be protected in the same fashion as HDDs with RAID technology, either striping with parity or by mirroring, although mirroring does carry a significant cost penalty (as it does with HDDs).

Selection of the right SSD for the enterprise storage environment involves consideration of the very same factors involved in the selection of an HDD: speed, capacity, interface, and price – these must all be right for the application if the IT staff expects to meet TCO and SLA considerations. The right choice for an HDD in a laptop might not be the same choice for an enterprise server. The same is true for SSDs. In addition, the IT staff must also factor in the design of the SSD controller and the specific host bus interface required to achieve satisfactory performance, as these elements will have the same impact on an SSD as on an HDD, although the impact of the SSD controller can have a much greater impact on performance.

Application Considerations

Applications that generate a high random read workload, such as OLTP[5] where caching techniques do not apply, will push conventional HDDs to their limit and require spreading the data across many lightly-loaded HDDs, are an excellent candidate for SSDs. This environment is where SSDs can shine, achieving a 30:1 ratio in terms of throughput, compared to an HDD. Flash technology does not experience the same advantage in applications that involve high levels of write commands (due to the ability of HDDs to write large blocks of data without having to reposition the disk heads).

Environments that can take advantage of the performance levels of SSD include high-performance computing (HPC) applications such as financial modeling, seismic data analysis, computer-aided engineering, and video media server for video on demand, with a requirement for upwards of 20GBps of read/write bandwidth.

The IT staff must consider the TCO of the available alternatives. This includes acquisition and deployment costs as well as maintenance, administration, energy, and floor space. What is the ROI of a decision to deploy SSDs? Will the expected gains in performance and reduction in energy and space offset the obvious increase in acquisition expense? The IT staff can rightfully expect a longer warranty, and therefore, lower maintenance costs for a semiconductor-based technology than one based on spinning media. A longer life-cycle will result, as justifications for the replacement of HDDs are often based upon increased maintenance costs, forcing the data center to upgrade to newer devices with higher capacity, even though capacity may not be the issue.

Conclusion

For some enterprises speed isn't just the most important factor, it is the only factor! Being green is simply a bonus – but one that will help to sell the solution at the CxO level.

SSDs are attaining a maturity level and a price point that make them very attractive as a drop-in replacement for HDDs for the mission-critical enterprise application with very high OLTP rates. They satisfy the most stringent requirements for reduced access times and lower power consumption. While every HDD in a server can be replaced by an SSD, not every application requires it. From a performance standpoint, SSDs will supplant the fastest HDDs in terms of ultra-high Tier-0 performance needs.

From a TCO standpoint, fewer drives mean lower cost, in many dimensions. The fact that SSDs will help to reduce enterprise storage management costs is a significant plus that will help to improve the ROI of SSDs.

If your data center has a requirement for extremely high transaction rates, in the order of 50,000 IOPS or higher, your enterprise needs to look at SSDs as the vehicle to improve performance and reduce total cost of ownership.

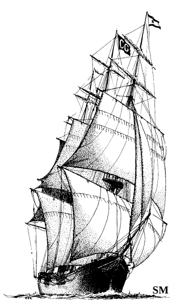

[5] OnLine Transaction Processing.

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

***David Reine* is Director, Enterprise Systems for The Clipper Group.** Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. He joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- ***Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 123. (Please dial "123" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

Disclosure

Regarding the Information in this Issue