White Paper:

New Equations for Measuring the Value of High Performance Storage

by
Tomas Havrda, Managing Partner
Chas Chesler, Director of Sales
Solid Access Technologies, LLC

April 21, 2008

Solid Access Technologies, LLC
27 Green St., Newburyport, MA 01950 USA
Telephone: 978-463-0642 Fax: 978-463-7480
www.solidaccess.com

New Equations for Measuring the Value of High Performance Storage

Contents

Introduction	1
Making the IOPS Comparison	1
How SSDs Deliver Performance	2
Non-Proprietary Hardware Provides Lower Cost, Better Reliability	3
Some Real-World Examples	4
Conclusion	6
Company Background: Solid Access Technologies, LLC	7

New Equations for Measuring the Value of High Performance Storage

Introduction

It's widely known that solid state disks (SSDs) can increase I/O operations per second (IOPS) by orders of magnitude in transaction-intensive enterprise applications. However, even managers who know this sometimes balk at implementing SSD technology because of its presumably high cost. This paper explains how the ultra-fast category of SSDs compares favorably to hard disk drive (HDD) and RAID technology in both performance and value.

Making the IOPS Comparison

Effectively managing today's storage infrastructure means balancing required capacity, desired performance, available budget for acquisition, and on-going operation and maintenance. Historically, the approach to improving performance has been an increase in HDD capacity: adding enough drives to distribute data among them. Unfortunately, a single HDD adds storage capacity but not very much performance improvement. This leads to high levels of unused storage capacity, or over-provisioning of drives.

The Solid Access USSD 200 solid state disk has achieved random read/write benchmarks of 95,000 IOPS using a single Fibre Channel link and over 70,000 IOPS using SAS. By comparison, a good HDD performs at about 200 IOPS (the SSD is 300 times faster). This means a single Solid Access USSD 200 solid state disk is equivalent to 300 hard disk drives when the yardstick is performance instead of storage capacity. What's more, the SSD requires only 5 percent of the power and cooling and occupies only one 2U rack compartment.

To make useful value comparisons such as these, organizations must begin to use additional measurements to determine the value of various storage device types. Ratios such as IOPS/$, IOPS/watt, IOPS/kW-hour and IOPS/U (U = each 1U of rack space) will enter the lexicon as appropriate measures of a device's suitability for a task.

At the same time, the use of existing database utilities and emerging software to identify storage-related performance issues will cause organizations to embrace the concept of Tier 0 storage. This is running the most heavily used files with the highest performance requirements on a suitably architected device, with the balance of the storage infrastructure operating in the well known online, near-line and archival fashion.

How SSDs Deliver Performance

Ultra-fast solid state disks are SDRAM-based peripheral devices, comprised of an array of semiconductor chips. They function as hard disk emulators, but unlike mechanical disks, ultra-fast SSDs have no moving parts in the data path and can process data at RAM speed. (Ultra-fast SSDs are different from Flash-based SSDs, which are not suitable for intensive random read/write operations because of slow write speeds and durability issues.)

SSDs eliminate seek and rotational latency times associated with HDDs during read and write requests and access data in microseconds (compared to milliseconds for hard disks). This increases IOPS by hundreds or even thousands of times, depending on the application. SSDs also make it possible for systems to exploit the incredible speed of today's CPUs. This speed is severely under-utilized when systems limited by relatively slow HDDs experience bottlenecks in transaction-intensive enterprise applications.

I/O acceleration is achieved in applications by off-loading I/O-demanding files ("hot files", typically less than 5 percent of the content) onto an ultra-fast SSD for processing at RAM speed and using mechanical disks (or RAIDs) to process the remaining "cold files." This instantly improves the efficiency of the application servers by recovering CPU cycles formerly lost in I/O wait loops.

While file caching provides some improvement for HDD performance in demanding applications, it does not raise HDD performance to SSD levels. File caching and mathematical algorithms do an amazing job predicting the next files to be read or written, but even the best products succeed only 70 percent of the time (a 70-percent "hit rate"). This results in a significant performance drop 30 percent of the time and inconsistent I/O performance. Ultra-fast SSD technology stores 100 percent of the data in memory, which is equivalent to a 100 percent hit rate for HDDs. Meanwhile, SSDs deliver solidly predictable performance at the speed of silicon!

Ultra-fast SSDs have been recognized for many years as a proven technical solution to IOPS imbalance and I/O bottleneck problems. They also provide a multitude of benefits throughout the system network. As peripheral devices, properly implemented ultra-fast SSDs do not disrupt existing IT infrastructure. They complement RAIDs (file cache vs. block cache) and allow desirable processing load shifts by using CPUs more efficiently and reducing the hits on wear-and-tear-prone hard disks.

Ultra-fast SSDs can dramatically increase the life span of legacy systems, extending their return-on-investment period. Additionally, multiple users can be served concurrently with no peak-load performance decrease. Furthermore, multiple servers using different I/O interfaces can be used with a single ultra-fast SSD device. As standards-based products, SSDs from Solid Access Technologies also provide investment protection when an organization replaces its computing environment with products from different vendors or moves the SSDs to other systems. As plug and play devices, Solid Access SSDs are independent of their computing environment.

Non-Proprietary Hardware Provides Lower Cost, Better Reliability

The SSD industry is one of the last to rely heavily on proprietary hardware components, even while almost every other hardware platform moves from proprietary to open architecture. Unlike most of the industry, however, Solid Access Technologies created an ultra-fast SSD using only standard, unmodified building blocks – best-of-breed, proven hardware components.

Solid Access SSDs are firmware-driven and based on innovative, patent-pending technology. They merge the disk drive emulation of ultra-fast SSDs with the rapid advancement and superior economics of server ecosystem building blocks. The components in the Solid Access USSD 200 have proven reliability, already deployed by the thousands in the world's top data centers.

The company's core product – the Universal Solid State Disk 200 (USSD 200) – delivers the full performance of enterprise-level solid state disks without any compromise in performance or reliability. This disruptive technology delivers the first affordable means of conquering I/O bottlenecks in transaction-intensive enterprise applications, and it allows complete transfer of large scale databases to memory-based storage, to run at almost RAM speed.

The new breed of I/O acceleration devices from Solid Access represents the future of enterprise-class SSD technology which will both close the server-storage "performance gap" and be affordable in mainstream applications. Unlike Flash SSD, the performance of ultra-fast SSD is consistently fast performing both sequential reads and writes of data as well as the much more challenging random reads and writes.

Available with Fibre Channel, Serial Attached SCSI (SAS) and SCSI interfaces, the USSD 200 is most suitable for direct attached storage and SAN storage environments. iSCSI-based Networked Attached Storage (NAS) currently adds too much network latency to justify the investment in ultra-fast SSD technology, but upon market adoption of 10 Gb/second IP-based networks, replacing current 1 Gb/second technology, ultra-fast SSD will penetrate this market also.

USSD 200 units of 16, 32 and 64 GB use 2 AMD Opteron processors per unit; the 128 GB (available now) and 256 GB (available June, 2008) units use 4 Opteron processors as well as key components from trusted vendors such as LSI, Super Micro and world class memory suppliers.

Some Real-World Examples

Here are three examples of the value of the Solid Access USSD 200 I/O acceleration appliance:

1. Solution to the over-provisioning problem.
2. Ultra-fast SSD storage beats traditional HDDs on significant total cost of ownership (TCO) measures.
3. The "two win" application performance improvement strategy.

Solving the Over-Provisioning Problem

Improving I/O bottlenecks historically meant over-provisioning: adding HDD/RAID devices to gain performance (IOPS) even if the capacity of the individual drives was not even close to full utilization. It is common for HDDs to be used at only 20 to 30 percent of capacity. Each drive adds only approximately 200 IOPS, and an application requiring high transaction throughput might need up to 100,000 IOPS.

In addition, time consuming and costly performance tuning is often employed to ensure adequate I/O throughput. With the Solid Access USSD 200, the designated "hot files" can be moved to the "hot storage device," typically with no additional work other than simply integrating the plug and play unit into the storage infrastructure. This frees unused storage capacity while also providing noticeable performance improvements. Customer benchmarks have consistently shown re-directed files to run several hundred times faster!

Ultra-Fast Storage Beats Traditional HDDs on Total Cost of Ownership Measures

Because the moving parts in HDDs make them failure-prone and their associated storage infrastructure requires active management, on-going operating expenses for HDDs are much higher than those of ultra-fast SSDs. This means all aspects of the total cost of ownership (TCO) equation favor the use of SSDs. As the table below illustrates, ultra-fast SSDs are less expensive when comparing IOPS/dollar, power and cooling costs, and space requirements. In addition, SSDs require virtually zero maintenance while HDDs have on-going costs associated with replacing failed drives in the disk array and performing other maintenance. Thus, for high transactional throughput requirements, if it is cheaper to acquire, operate, house and maintain the required storage, forward thinking organizations should investigate the potential benefits of SSDs.

TCO Comparison: HDD vs. Universal Solid State Disk (USSD)

	Units	IOPS/ Unit	IOPS	Price	Power/Unit (Watts)	Cooling/Unit (Watts)	Total Watts	Space
HDD	300	200	60,000	$200K	15	7.5	6,750	50U
USSD	1	60,000	60,000	$50K	218	109	327	2U

"Two Win" Application Performance Improvement Strategy

If new servers are introduced into sites with significant I/O bandwidth issues, the added server capacity will be underutilized, meaning wasted customer investments and applications that run slower than desired or expected. In some cases, organizations cannot buy the new computing capabilities, even though they would like to, because of the inherent and growing gap between increasing CPU performance and much, much slower HDD improvements (IOPS, not platter rotational speeds). Significant adoption of dual- and quad-core processors and the new 2 quad processors are only making this problem worse.

Solid Access has experienced this situation and can report that introduction of the USSD 200 brings much needed I/O acceleration as reported to us by customers, allowing them to add desired server capacity, having solved the I/O bottleneck issue. Both storage performance and server performance increase, leading to faster database read/write and shorter user response times, a "two win" solution.

Conclusion

The future of storage will continue to be a critical topic in IT as a "perfect storm" of events is brewing:

- Continually increasing requirements for both storage capacity and performance
- Spiraling energy costs for power and cooling to operate a growing volume of storage devices
- Organizational desires for space consolidation for both business and ecological reasons.

Any technology that can help reconcile these opposed forces should interest forward-thinking managers. Firmware-driven ultra-fast SSD technology using non-proprietary hardware components should be included in an overall strategy to architect a well performing, energy efficient storage infrastructure while optimizing space utilization and reducing on-going maintenance efforts.

Company Background: Solid Access Technologies, LLC

Solid Access Technologies, LLC, was founded in 2002 as a privately held Massachusetts corporation for the purpose of translating the new and innovative solid state disk (SSD) technology it had developed into a line of enterprise SSD I/O accelerators.

Solid Access Technologies' development activity over the past 6 years has been fueled by the recognition that to achieve a balanced system, the electro-mechanical deficiencies of spindle-based hard disk drive (HDD) storage would have to be overcome. To accomplish this, the most critical I/O demanding segments of commercial applications would have to be placed on a medium with electronic access speed. The company's goal was to develop technology that would result in an unprecedented level of SSD device production efficiency and user cost savings, thus motivating mainstream users to place their entire databases or their most performance intensive files on silicon-based storage.

The SSD industry was one of the last "closed" areas that relied on proprietary hardware components. This existed in the face of almost every computer IT hardware platform moving from proprietary to open architecture. Solid Access predicted the same trend would emerge in the ultra-fast SSD segment to gain the highest degree of efficiency possible. Solid Access wanted to create a new generation of SSD devices using 100-percent standard, unmodified, off-the-shelf hardware components that are proven best-of-breed.

The primary difference between the Solid Access Technologies SSDs and its competitors' products is a novel design architecture that achieves efficiency on multiple levels and breaks the cost barrier that had previously restricted SSD applications. By procuring proven industry standard components on a just-in-time basis and manufacturing to order according to strict specifications, the company's business model involves minimal real estate, manufacturing, overhead, and inventory costs. This translates into a fully functional, reliable, compact, plug-and-play SSD with unsurpassed performance that is now affordable and readily deployed in mainstream IT applications.

By embracing a 100-percent "buy vs. build" philosophy since the company's inception, and instead focusing on software-driven innovation, Solid Access is well positioned to establish leadership in this category. With the inherent capability to pre-engineer instead of re-engineer, Solid Access can be first to market with important industry advancements.

The CPU to HDD performance gap shown in the graph below has become more critical over time because of the ever-growing size of databases and demand for greater I/O throughput in real-time applications. As Intel's Knut Grimsrut explained, "...while capacities of hard drives have improved, latency continues to be a performance bottleneck. As CPU performance has improved some 30 times over the past decade, hard disk latency has improved only 1.3 times in the same period. It's so grossly out of whack that for many workloads the disk drive ends up determining what the effective system performance is."

Being able to deliver the benefits of SSDs to mainstream markets necessitated a fundamental shift in SSD technology that would result in a transparent and convincing rationale for adoption of SSD solutions.

Solid Access has customers and partners in North America, Europe and Asia in the private sector, as well as the U.S. Federal and European Union Governments.

Note: The Solid Access technology is SDRAM-based SSD (also referred to as ultra-fast SSD), not Flash SSD. Flash SSD is useful for "Read Only" fast application start up situations but not suitable for random read/write scenarios due to write speeds that can actually be slower than traditional HDDs. Another factor limiting the utility of Flash SSD is its reduced durability during write data operations, which the SDRAM technology used by Solid Access does not suffer from.

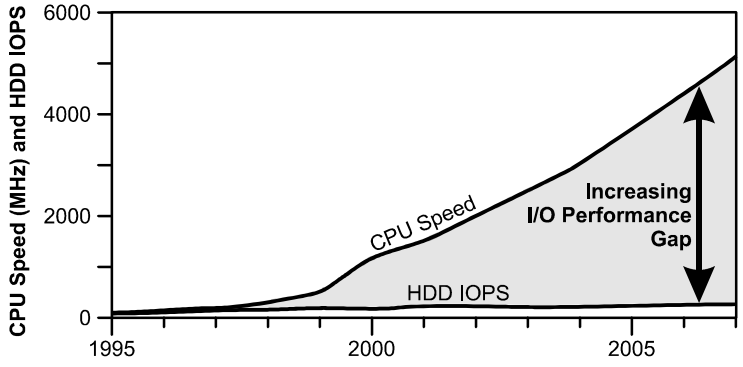

SSDs address the increasing gap between CPU and HDD performance