Solid State Storage Initiative

Solid State Storage 101
An introduction to Solid State Storage

January 2009

SSSI Members and Authors:
Neal Ekker, Texas Memory Systems
Tom Coughlin, Coughlin Associates
Jim Handy, Objective-Analysis

The Solid State Storage Initiative

Recognizing the dramatically escalating interest in solid state storage, and thus an increased need for accurate technology information, education, and standards development, the Storage Networking Industry Association (SNIA) recently formed a new body, the Solid State Storage Initiative (SSSI). Many educational articles, presentations, and tutorials are being developed and made available to the public by the SSSI to fulfill its mission of fostering the growth and success of solid state storage in both commercial and consumer environments. In 2008 some standardization efforts related to Flash SSS in other standards bodies began, and the SNIA SSSI formed the SSS Technical Working Group (TWG) to complement these other standards activities and provide system level standards related to SSS, capitalizing on SNIA's expertise in storage systems.

No matter how the future of solid state storage plays out, the SNIA and its SSS Initiative and TWG will play an important role. The ultimate goal is to further the interests of solid state storage through education, standards development, and overall promotion of appropriate SSS deployment. The group will be a resource to users and the solid state storage community in defining, testing, and refining the rapidly developing solid state storage market. For more information about the SSSI, visit www.snia.org/sssi.

Overview

Solid state data storage is gaining significant acceptance today. Storage based on Ram Access Memory (RAM) and Flash chips instead of mechanical hard disk drives is earning much greater attention by meeting the market requirements for reliability, performance, and cost more effectively than ever before. Additionally, the surge in consumer electronics sales has increased awareness of Flash in particular and solid state storage (SSS) solutions in general. This technology, though available for decades, is finding new ways of providing cost effective solutions in a wide range of commercial and governmental information systems. These varied environments, often referred to collectively as enterprise applications, typically require higher performance, reliability, and capacity than demanded in consumer products such as music and video players, mobile phones, PCs, and laptops.

Because enterprise-grade solid state storage accesses data directly from RAM or Flash chips, it can achieve input and output data rates far greater than conventional, magnetic storage devices such as hard disk drives (HDD). SSS products are available in a number of form factors (shapes and sizes) and solution types implemented at different points in the data path. Today, vendors offer SSS products in the size and shape of add-in cards that plug into server motherboards, HDD form factor modules inserted into existing servers and storage arrays, and stand-alone systems intended to be rack mounted alongside other storage media in a data center. In the past, enterprise SSS solutions have been mostly RAM-based, with batteries and back-up HDDs to ensure data persistence. More recently, NAND Flash-based SSS has been introduced into the marketplace as manufacturers have learned to make it reliable and fast enough to meet enterprise needs.

What Is Solid State Storage?

Simply stated, solid state storage is data storage made from silicon chips instead of spinning metal platters or streaming tape. Figure 1 shows a hard disk drive and a representative type of SSS commonly referred to as a solid state drive. The images illustrate the differences in the components that make up each device. At the heart of every solid state storage implementation lies a simple fact: CPUs can process data much faster than HDD storage can supply it. Mechanical spinning disks have dominated the enterprise storage market for decades. HDD storage capacities have increased over 200,000 times since the 1950s with continuous declines in hard disk drive prices over that time, making HDD-based storage increasingly more affordable. Unfortunately, HDD performance, measured either by its response times (latency) or by its inputs/outputs per second (IOPS), is another story. In the past 30 years the speed of processors has increased exponentially, while the basic read and write response time of hard disk drives has only increased marginally, resulting in a significant gap between HDD and processor access times.

Figure 1. Images of HDD and Representative SSS

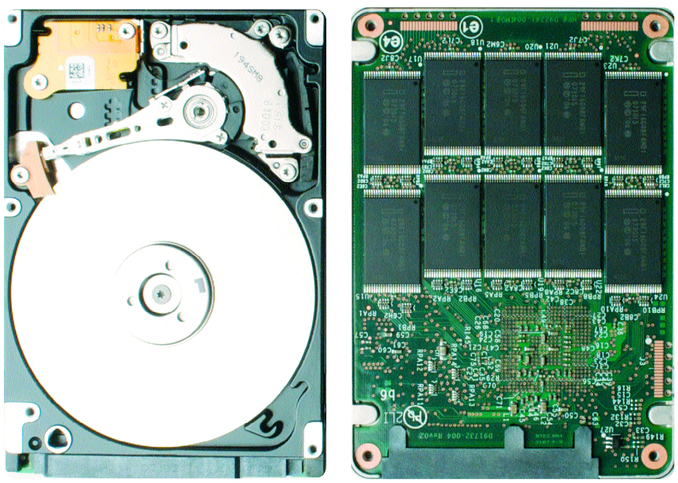

Photo provided by Intel Corp 2008

Increasing Speed Increases Profits

Until now, enterprise solid state storage solutions have been employed primarily to accelerate mission critical enterprise applications. Due to significant improvements in memory densities and falling prices, SSS is much more affordable now per unit of capacity and can be utilized in a wider array of applications. For example, if an equities trading firm can increase profits by processing trades faster, they deploy SSS, thus reducing processing times. If a telecommunications company can handle more calls by speeding up the responses from their massive databases, they deploy solid state storage and increase call volume. If video editing and streaming dramatically accelerates when served faster, the data files are moved to SSS. The list of solid state storage applications is almost endless – all types of high performance computing, countless military applications, data warehousing, online games and gaming – in any environment where traditional HDD storage cannot respond fast enough, that's where SSS is and will increasingly be deployed.

Solid state storage has the following characteristics:

- ***Lowest possible access times:*** SSS offers access times 100 to 1000 times faster than mechanical drives.
- ***High bandwidth:*** Solid state storage can achieve multiple gigabytes (GB) per second of random data throughput.
- ***High IOPS:*** SSS offers extraordinarily high random input/output (I/O) performance because of its low access times and high bandwidth.
- ***Low price for performance:*** SSS provides the best possible price/performance of all storage devices.
- ***High reliability:*** SSS offers the same levels of data integrity and endurance as other semiconductor devices.

Other advantages include more consistent I/O response time as well as predictable wear and lifespan.

No More Waiting

Servers today can complete millions of operations during a hard disk's typical 5 milliseconds of data access time. As the performance gap between processors and hard disk drives increases, the requirements for supporting a pure HDD storage system become wasteful and expensive. Following are a list of issues associated with the performance gap between processors and hard disk drives:

- ***Wasted server infrastructure:*** Caused by businesses responding to performance problems by buying more processors and more servers
- ***Wasted storage infrastructure:*** Caused by businesses responding to performance problems by buying large arrays of dramatically underutilized HDDs to improve bandwidth
- ***Long application response times:*** As systems become I/O bound
- ***Decreased user productivity and customer satisfaction:*** Caused by application response time delays
- ***Long running batch jobs:*** Caused by traffic bottlenecks due to low performing hard disk drives
- ***Decreased CPU utilization for I/O intensive applications***
- ***Decreased profitability:*** Increased facilities costs and ever increasing power and cooling requirements for HDD-based storage increase the operational costs of facilities.

SSS offers an excellent solution for I/O bottlenecks, particularly those bottlenecks caused by the high access times of traditional disk-based storage systems. Figure 2 shows one approach to increasing the performance of Flash SSS devices by striping data across multiple chips that are accessed in parallel. Replacing HDD storage with solid state storage offers numerous benefits, including:

- ***Improved Server Efficiency:*** When slow, conventional storage holds back the potential of expensive processors and servers, efficiency is reduced and money is wasted. Conversely, introducing fast SSS enables those servers to be fully utilized, resulting in maximized ROI.
- ***Fewer Servers:*** The drive towards server consolidation means squeezing every drop of performance out of the remaining servers. When SSS can improve server efficiency, then that efficiency increase can lead to server consolidation without performance loss.
- ***More Concurrent Users:*** When SSS is installed it typically relieves I/O bottlenecks. When processors can run closer to maximum throughput, more processing power becomes available for other applications and tasks. In the case of transaction-based applications, adding SSS can increase the number of concurrent users receiving data at higher speeds. Conventional thinking suggests that adding concurrent users requires more servers, but with SSS the number of concurrent users scales by improving server efficiency.
- ***Lower Power Consumption:*** Adding SSS to the storage mix can cut storage system power and cooling requirements significantly.
- ***Less Data Center Space:*** Reductions in computational and storage infrastructure free up badly needed floor space.
- ***Faster Response Times:*** SSS read speeds are significantly faster than HDD storage. Write speeds can also be faster with greater parallel paths to Flash memory for write data.

Figure 2.

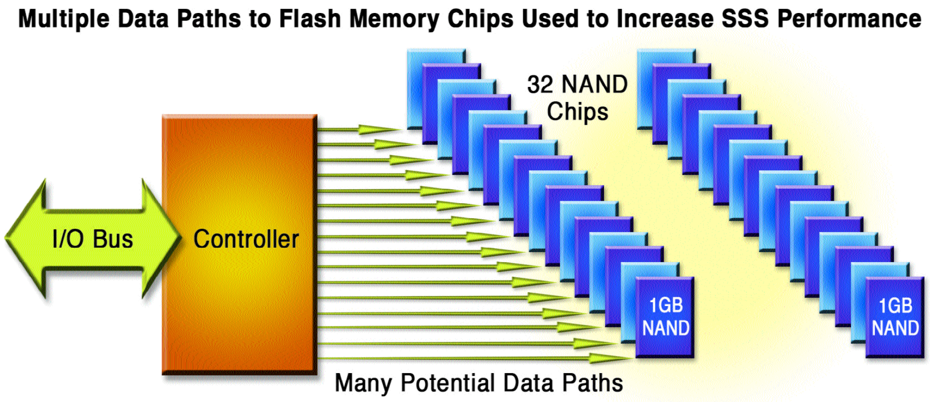

Tiered Storage

In many application workloads, users understood that some data would be accessed frequently, such as database index tables, while other data – the records of a closed account, for example – might need to be archived or rarely accessed. To respond to this need, storage "tiers" developed, with database managers increasing application responsiveness and lowering costs by moving frequently accessed data onto faster, more expensive media and less frequently accessed data onto slower, less expensive media. Tiered storage developed for the best of reasons: it helps IT managers increase application performance and lower overall data storage costs. In many enterprises it makes sense to store most of the data on much slower and less expensive HDD arrays because the majority of the dataset is accessed infrequently and thus does not affect application performance. However, there is a smaller portion of the data that is accessed frequently, that does affect system performance, and therefore deserves to be moved to a higher level in the memory hierarchy, either DRAM memory or very fast solid state storage; doing so improves overall system performance and increases the productivity of mission critical applications, all while actually driving down costs. Figure 3 shows multiple storage tiers used to provide a balance of performance and capacity for a Network Attached Storage (NAS) system.

Figure 3.

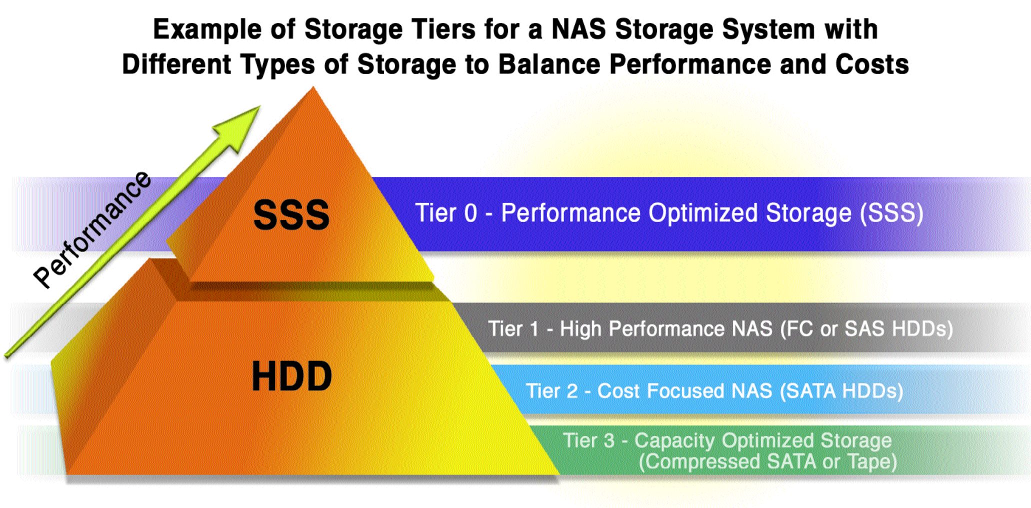

The benefits of tiered storage architectures ripple up and down. Use of SSS to accelerate transaction processing allows servers to return to doing what they do best, processing data. With the use of SSS in the storage hierarchy hard disks are no longer stressed by performance requirements. They can also return to what they do best, store a lot of data relatively inexpensively.

SSS vs HDD

Some enterprise storage market observers have predicted that solid state storage will eventually and perhaps inevitably replace hard disk storage, but the evidence suggests otherwise. Instead, most storage analysts see enterprise SSS and hard disk storage as complementary, not antagonistic. HDD storage is not going away any time soon, and its sales will increase, because the overall volume of data is increasing dramatically. But almost certainly the role of hard disk drives will evolve to accommodate the increased presence of SSS in enterprise storage solutions. HDDs will be used where high capacity is the main requirement. SSS will be used where high performance is the main requirement.

Common Solid State Storage Solutions

Though solid state storage has to date been limited to Flash and RAM-based products, the approaches SSS manufacturers have taken in terms of where these products fit in the data path have proven much more varied. In fact, careful observers of SSS technology and market trends see these particular differences evolving into the major differentiators between products. At this point, several SSS solution types have emerged. First, a number of companies have started manufacturing hard disk form factor Flash drives which are integrated by system-level vendors via OEM arrangements. Some of these system-level vendors are designing the Flash disks into their existing storage arrays. On a very different tack, other vendors have begun offering in-server solid state storage where Flash chips are mounted on printed circuit boards and integrated into the data stream through bus (PCI-Express) connections. Finally, several manufacturers have continued to build rack-mounted stand alone RAM and Flash-based SSS systems.

Storage Array Solutions

It makes sense that storage vendors with successful enterprise HDD storage array product suites would take a close look at incorporating Flash disks into their existing HDD array architectures. In fact some of the early Flash solid state storage product announcements were for solutions of this type. Flash disks incorporated into existing storage array architectures can take advantage of infrastructure features that vendors have been adding to these products for years. These vendors are also betting that customers will value a solution where multiple storage tiers are offered within one system, including hard disks and Flash disks, while maintaining and further optimizing large RAM caches in the controllers for increased performance.

Storage array-based solutions often employ solid state drives (SSD) that are solid state storage solutions with form factors and electrical interfaces compatible with common HDD bays. These products allow for a more evolutionary replacement of HDDs with higher-performance solid state storage. SSDs put semiconductor memory behind a memory controller and interface electronics that allow the Flash storage device to emulate the commands and operation of a hard disk drive. Putting semiconductor memory into such a device allows that device to be easily substituted for a hard disk drive in stand-alone, in-server, or array environments. Such storage devices can also be substituted for small form factor hard disk drives in environments such as server blades for local storage.

Figure 4.

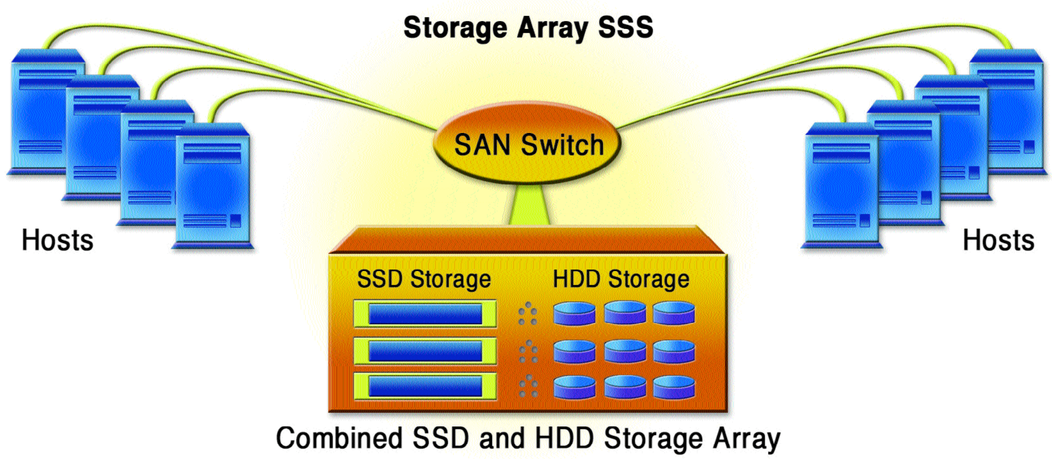

In-Server Solutions

The decline of Flash prices opened additional horizons for Flash-based solid state storage products, and engineers began to examine where SSS might fit best in the data path between CPU and storage. One line of reasoning suggests that the closer fast storage resides to processors, the better the performance. Based on this rationale, in-server SSS solutions began to evolve.

One option to bringing fast SSS storage close to the processor is to offer Flash-based storage that connects directly into the server motherboard. Another option is to have it plug directly into PCI (Peripheral Component Interconnect bus) express slots. The reasoning is that legacy disk controllers and protocols become a bottleneck when the storage media increases performance by one or more orders of magnitude. Having solid state storage that fits in disk drive slots is convenient when you have no other interface to the system (like in laptops). But if such a disk controller is not required, storing persistent data in SSS not encumbered by legacy controllers and protocols can leverage its direct access to the CPU and memory when directly coupled via high performance, low latency interfaces. In this case, the Flash media is integrated with its own tailored controllers. This type of solid state storage offers the advantage of bringing the solid state storage as close as possible to the CPU and minimizing the latency added by disk and RAID controllers, network protocols, switches, and the additional hardware and software involved in both NAS and Storage Area Network (SAN) topologies.

Figure 5.

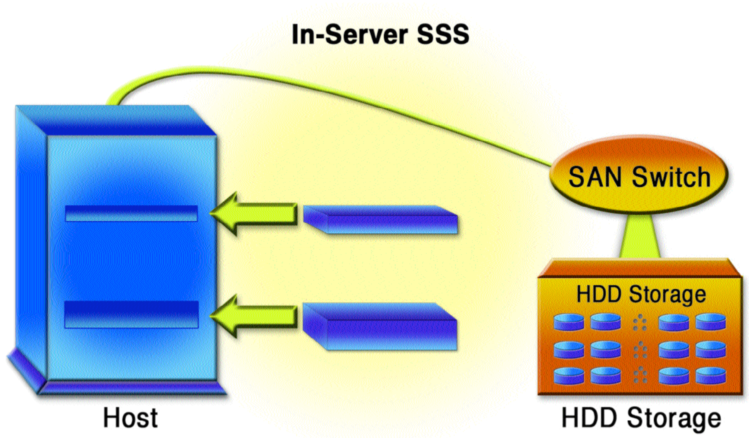

Stand-alone Solution

Flash SSS in enterprise storage arrays is a relatively new concept; in contrast, RAM-based solid state storage systems as stand-alone solutions are not. Several companies have been offering these types of products for years. The RAM-based systems, as previously mentioned, involve RAM chips, batteries, and connectivity and management controllers all engineered into "one box." These units are usually rack-mounted and can be attached directly to servers or connected as either NAS or as components of a SAN. As with most other SSS solutions, when deployed they appear simply as another disk drive to the operating system; this is the genesis of the term "solid state disk."

Stand-alone Flash SSS solutions have also been available on the market for many years. Due to Flash's greater ruggedness and higher cost in comparison to other storage, until recently these products have mostly exploited the extreme environment market niches such as embedded military and reconnaissance applications. In these environments, on the battlefield or high in the sky, Flash systems have performed very well, offering a ruggedness and performance that mechanical hard disks could not.

Figure 6.

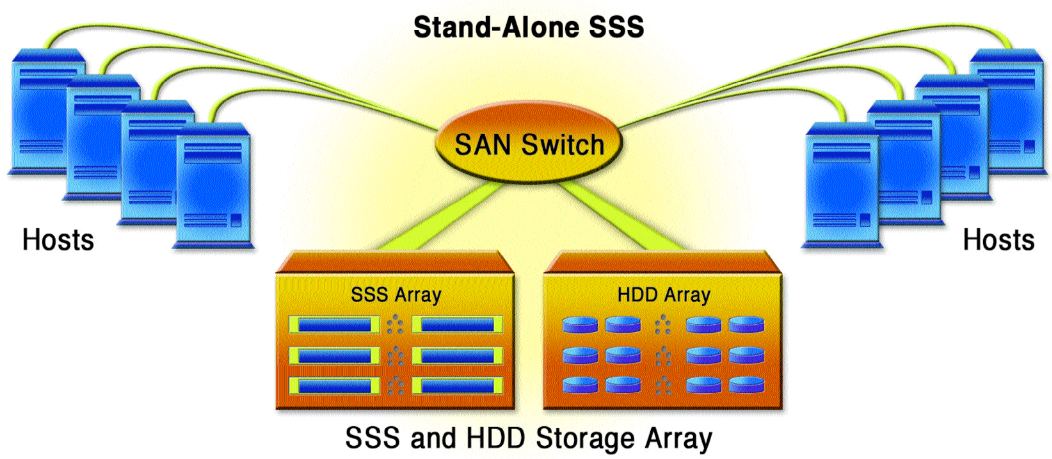

Solid State Storage: Cost Effective and Reliable

With all the advantages that solid state storage solutions bring to the table, one might imagine that enterprises would be buying and deploying truckloads of SSS. To some extent that has happened; many of the world's largest financial institutions and telecommunications companies have already implemented various SSS solutions. But many more commercial firms and governmental entities that could truly benefit from SSS technologies have remained reluctant to add it to their storage systems. Two apparent problems with SSS are most often cited by potential users as reasons for not yet deploying it: cost and reliability.

In regards to cost, it is true that SSS solutions still cost more per unit of storage capacity than "raw" HDD storage. It is also true that cost-capacity has been one of the primary drivers in the storage market for decades. But, use of this one metric to decide the cost effectiveness of a storage solution is misleading and frequently wrong. Accurate and meaningful calculations of storage costs at the enterprise level must necessarily include more variables – one of the most important being performance.

Storage Costs Include Performance

Most IT managers still make a simple but fundamental mistake when they calculate their data storage requirements: they consider only storage capacity, not performance. But for even moderately performing applications the number of HDD spindles needed to meet the required storage capacity is generally much lower than the number of spindles required to meet the performance objectives. In some cases, companies will use smaller, faster HDDs when they are trying to improve performance. Drives with smaller storage capacities tend to offer higher spindle speeds, higher performance interfaces, and better access density (a metric that considers how many IOPS are achieved per storage capacity). But this type of drive does not have the best cost-to-capacity ratio.

The difficulty that enterprises run into is that it can take a large number of small capacity, power-hungry, and relatively expensive drives to meet the performance requirements of an application. Companies that want to configure their storage arrays to give them the best performance often have to buy a significant number of drives, definitely more drives than they need for capacity.

Along with adding disks, the most popular way to improve the performance of an application is to add more servers or deploy servers with more CPUs or cores. By adding servers, many enterprises believe that they can improve the performance of their applications. Up to a certain point, adding servers can improve the performance of an application, but storage is a bottleneck for many applications. Most data center servers are actually running well below optimum processor utilization in spite of server virtualization techniques. In fact, virtualization may change storage access patterns in ways that are detrimental to the efficiency of traditional storage. Quite often enterprises will add more servers than they need in order to attempt to improve the performance of an application, with little resultant benefit.

Instead, enterprises need a solution that delivers high performance, high capacity, and low cost, preferably while reducing power consumption. As noted earlier, mixing SSS with slower but high density HDD storage can result in a perfect blend of high IOPS and low cost per capacity. If a company has an application that requires the purchase of a large number of HDDs for performance, they can move the files or database logs/journals that demand the highest performance off the inefficient and slow HDDs to solid state storage. Also, by implementing an SSS solution and moving the most frequently accessed files or logs onto it, enterprises can gain greater performance from their servers because the processors are no longer waiting on I/O from relatively slow storage. By getting more performance from each processor, fewer servers are needed to meet application performance requirements. And by using fewer servers, enterprises can reduce overall hardware and maintenance costs, plus reduce the amount of software licensing they must purchase. For example, in a SSS study done by a database software company, the elapsed time for a complete batch processing run improved by up to 35x, I/O wait times improved by up to 56x, and IOPS increased by 37x. Clearly, the cost advantage of SSS cannot be measured based on cost-capacity alone.

The ability to scale both performance and capacity are big advantages that solid state storage paired with HDD arrays can offer. In most cases, if an enterprise wants to increase the performance of their storage architecture they simply add additional SSS capacity. If more storage space is needed, they can add more shelves of HDDs to the array. But the best part of using the combination of SSS and HDD storage is that it can often be less expensive to purchase, contrary to the impressions most potential buyers might have about SSS solutions. For large configurations, the purchase price for high performance enterprise HDD storage arrays can easily run into the millions of dollars, while the purchase price of an equivalent SSS/HDD system may be only a fraction of this. By using both technologies, more performance and more storage capacity scalability can be purchased for less money.

Space and Power Savings

Initial purchase price isn't all that goes into total cost of ownership. Data center space is becoming more and more valuable every day. SSS products normally take up much less space than equivalent HDD solutions. When used in conjunction with appropriate HDD storage, the mix can reduce storage rack and floor space by more than half. Along with saving rack space, the power savings produced by using both technologies is very significant. Consider that Flash-based SSS can perform around 1,000 IOPS per Watt, depending on how it's deployed and configured, while hard disks perform in the range of 5-15 IOPS per Watt. This doesn't include the power and cooling efficiency gains for the system as a whole.

The cost – both the initial purchase price and on-going costs of ownership – of these mixed storage solutions, where SSS stores the frequently accessed data and HDDs handle the rest, has already dropped well below the true cost of equivalent enterprise storage based only on HDDs in many cases. The problem for SSS isn't its cost; it's in finding clear ways to demonstrate to potential customers that solid state storage already is extremely cost effective as an element of the storage hierarchy.

Reliability

Reliability may be viewed simply as a product performing as specified. In the case of solid state storage, data integrity and the endurance of certain types of memory chips are often included under the topic of reliability. RAM-based SSS rarely suffers from reliability issues. In fact, RAM chips often have much higher reliability rates than any mechanical storage. But, until recently, the same could not be said for Flash-based SSS. Though NAND Flash chips have many valuable attributes that make it attractive for use in solid state storage, the tendency to wear out with excessive erase/write cycles and corrupt data when accessing adjacent memory cells has until recently hindered its adoption into mission critical enterprise environments. But Flash offers so many benefits that solid state storage engineers have been motivated to overcome chip reliability issues.

Making NAND Flash-based solid state storage solutions sufficiently robust for enterprise use has required innovative and advanced controller breakthroughs. One of the most effective solutions to the Flash chip endurance challenge is called wear leveling. In its simplest form, wear leveling ensures that writes are evenly distributed amongst all blocks within the Flash chip, thus greatly increasing the overall endurance and reliability of the SSS. Storage vendors have developed specialized Flash controllers programmed with wear leveling algorithms that distribute data writes evenly over all the blocks in each Flash chip and across all the Flash chips on the board.

Wear leveling can be dynamic or static. Dynamic wear leveling works on data blocks that are being written to. Dynamic wear leveling addresses wear to Flash memory cells by redirecting new writes to different physical blocks. This avoids premature wear out of actively used blocks. If a data block is not being written to, it will not be wear leveled by a dynamic wear leveling algorithm. In the case of static wear leveling, the memory controller wear-levels all data blocks, including those that are not being written to. This is done in the background, with no impact on the host system. Different triggers can be used to initiate static wear leveling.

Flash SSS solutions have also added more Flash memory to their products than what is reported as usable to spread the writes over even more chips. By employing these wear leveling techniques at the board level, vendors are producing Flash SSS products designed so that even under maximum write loads it isn't possible to wear out the Flash for many years.

In addition to wear leveling, at the circuit board level technologies such as error correction codes (ECC) have also been widely implemented. And at the system or whole product level, additional reliability features are being incorporated such as chip RAID technology (one example of SSS RAID is known as ChipKill). All of these techniques and engineering ingenuity have allowed Flash SSS to function reliably in the most demanding environments.

Many technologies developed to monitor and report on storage device status in HDDs can also be incorporated into SSS devices. Self Monitoring and Reporting Technology (SMART) that monitors and logs important characteristics of HDDs such as the number of hard and soft errors, internal temperatures, etc., can be modified to provide useful information about the state of SSS. The self monitoring data logs created by a SMART device can be used by storage systems and system administrators to detect problems before they lead to failure so that field replaceable units can be swapped out before they fail. This mechanism can easily be used to replace an SSS device before it fails due to reaching its write endurance limit. Said differently, the capacity and performance of NAND Flash SSS products tend to degrade slowly and predictably, allowing for scheduled maintenance with replacement, or simply repurposing.

What is Next?

Solid state storage has a bright future. The technology is improving rapidly. Price per unit of storage continues to fall. The need for higher performance storage is predicted to increase over the next five years at a 70%* annual growth rate, according to some analysts. We expect that in the next few years there will be the successful infiltration of solid state storage solutions into "mainstream" enterprise applications and environments. Will SSS remain limited to accelerating only applications that need relatively high performance, or will a combination of lower component prices and more demanding software allow SSS implementations to bleed down into less high performance environments? The early indications are that SSS will indeed become mainstream.

Solid state storage will have significant impacts on server, data center, and storage markets in the next five years:

1. ***Server market:*** Every major server vendor will have an SSS offering. Initially, these offerings for rack-mount servers will be either bus-attached or disk drive form factor. In addition, blade solutions using SSS devices are likely to multiply as blade enclosures evolve.
2. ***Data center market:*** The future is here and hundreds of companies are already benefiting from SSS technology to accelerate applications. At the high end, data center HDD storage is already being augmented and in some cases displaced by SSS. Over the next five years this trend will expand significantly. SSS will also enable significant power and space savings, in addition to server consolidation, helping IT managers achieve their "Green IT" goals.
3. ***Storage market:*** Though it's unlikely that mechanical hard disks will ever disappear, it is certain that enterprise storage solutions will more and more often include solid state storage.

We expect that HDDs will continue to be used where high capacity is the main requirement, and SSS used where higher performance is the main requirement.

* IDC Market Analysis
Worldwide Solid State Drive 2007-2012 Forecast and Analysis: Entering the No Spin Zone
Jeff Janukowicz, John Ridning, and David Reinsel May 2008

Storage Networking Industry Association
425 Market Street, Suite 1020 • San Francisco, CA 94105 • Phone: 415-402-0006 • Mobile: 415-420-2145 • Fax: 415-402-0009 • www.snia.org