Does Serial Attached SCSI make sense for Solid State Storage?

An in-depth analysis

Esther Spanjer
Director of Technical Marketing
Adtron Corporation

TABLE OF CONTENTS

1 EXECUTIVE SUMMARY

Serial Attached SCSI (SAS) has been gaining popularity in the past few years, and is expected to become the dominant enterprise storage interface by 2010. To better understand SAS's gaining market share over competing storage interfaces, such as FC, SATA, and SCSI, it is important to understand the key reasons behind this growth. SAS delivers the following benefits to any storage system:

- Better performance
- Configuration flexibility
- Scalable to over 16,000 devices in a single domain
- Superior Device Reliability and System Availability

It needs to be put in perspective that SAS was designed with the purpose to overcome the inherent performance and reliability problems of hard disk drives, while also providing additional benefits of scalability and flexibility over legacy storage interfaces. When replacing a hard disk drive with solid state storage, these benefits are diminished.

SAS will still provide superior flexibility and scalability in combination with solid state storage, but improvements in performance and device reliability are much less obvious. SSD performance is not reliant upon the size of platters and the speed at which they rotate. A SAS interface in a SSD may or may not improve performance, as the performance of an SSD is more dependent on the internal architecture and controller than on the storage interface. The same reasoning can be used when reviewing device reliability. Systems using HDDs clearly benefit from SAS in terms of device reliability and system availability; however, since solid state storage is already far more reliable than HDD technology, SAS-interface SSDs result in a marginal system availability improvement over using SSDs with a non-SAS interface.

If the reason for adopting SAS is based only upon achieving improved reliability and performance, then SSDs already provide this and the interface choice becomes less relevant. Instead, the decision process should focus on the choice of SSD versus HDD and the related Total Cost of Storage Ownership.

2 INTRODUCTION

Dramatic advances in processor speed (2000x faster), RAM size (1000x greater) and RAM speed (2000x faster) have combined to accelerate system performance to levels unthinkable 20 years ago. Yet one aspect of system capability has conspicuously lagged behind: hard disk drive performance. Tremendous advances in aerial density have yielded exponential growth in drive capacity (3600x greater), but parallel transfer rates have achieved only modest gains (32x)[1].

With the continuing explosive growth in transactional enterprise data, significant performance challenges are surfacing for any storage solution and parallel interfaces are reaching the limits of their capabilities. With increased pressure on IT budgets, the need for higher performance and more reliable storage is more visible today than it has ever been.

In the light of the above, two relevant trends have happened in the storage market in recent years:

- As parallel storage technologies such as SCSI and ATA approach their performance limits, the storage industry began a transition to next-generation serial technologies that meet the bandwidth requirements of today's enterprise data centers. One of the crucial serial storage interfaces emerging is Serial Attached SCSI (SAS). Designed as the next step in the evolution of SCSI technology, SAS meets the performance and availability demanded by I/O-intensive, mission-critical applications.
- The price of solid state disks have plummeted over the past decade with an average of 20-30% per year, resulting in price/performance that is very competitive with enterprise-class HDDs (see Adtron's "Total Cost of Storage Ownership" White Paper)

With solid state storage reaching price points that make them suitable to be the online storage component in many transactional enterprise applications, the question can be asked "why not have a solid state drive with a SAS interface"? This would solve the storage bottleneck once and for all, both in performance, as well as reliability.

This white paper discusses the advantages of having a SAS interface for a hard disk drive versus the advantages of a SAS interface for a solid state drive.

Figure 1: Serial Attached SCSI in HDD and SSD

[1] SCSI Inflection Point: The New Era of Serial Attached SCSI, Seagate TP-528 Technology Paper

3 THE RISING POPULARITY OF SAS

SAS has been gaining popularity in the past few years because it alleviates some of the problems inherent to HDDs. Data from both the SCSI Trade Association (STA)[2] and Gartner[3] show that an estimated 4.15 million SAS hard disk drives shipped in 2006, with a projected 210% CAGR through 2011.

According to IDC, SAS accounted for only 5% of total enterprise-level disk drive shipments in 2006, but is expected to account for about 33% of the market by the end of 2009, as shown in Figure 2 below.

Figure 2: IDC - Interface Market Share

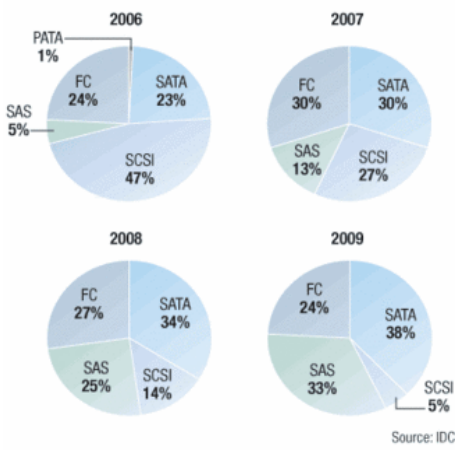

To better understand SAS's rising popularity and market share gains over competing storage interfaces, such as FC, SATA, and SCSI, it is important to understand the key reasons behind this growth. In comparison to legacy parallel SCSI and SATA, SAS delivers the following benefits:

- Better performance
- Configuration Flexibility
- Scalable to over 16,000 devices in a single domain
- Superior Device Reliability and System Availability

SAS leverages the manageability and more than 20-year history of parallel SCSI, which today represents the majority of the enterprise disk market. However, SAS takes SCSI to the next level to meet tomorrow's demands for mission-critical, heavy-I/O applications, such as databases, web servers and transaction processing. These applications demand the highest performance with around-the-clock availability.

3.1 Higher Bandwidth

Serial Attached SCSI significantly improves the interface bandwidth by utilizing full-duplex, point-to-point architecture. With a transfer rate of 3.0 Gbit/sec and a clear roadmap to 12.0 Gbit/sec, as shown in Figure 3 below, Serial Attached SCSI provides the means to reach far beyond the performance capabilities of the fastest parallel SCSI interface to date (320MB/sec).

Full-duplex operation can double the effective throughput by enabling simultaneous signal transfers in both directions. The duplex feature reduces a major cause of latency and allows aggregation of several ports as “wide” links so data can travel at a higher transfer rate.

[2] SAS Hard Drive Shipments Grow Exponentially in 2006, SCSI Trade Organization Press Release June 25, 2007
[3] Gartner, Inc., "Dataquest Alert: Serial Attached SCSI to Dominate in 2007", April 2007

Figure 3: SAS Roadmap

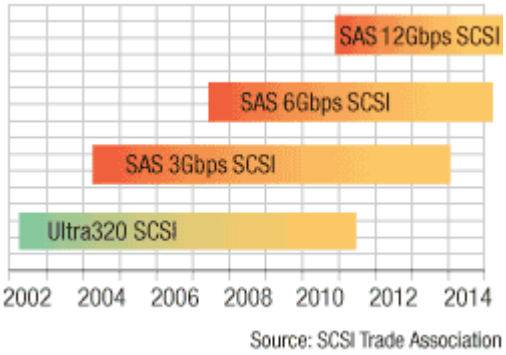

3.2 Greater Flexibility

Probably the best known benefit of SAS for external storage is its configuration flexibility – the freedom to select the most appropriate disk drive for a given task. While SAS hard disk drives will be used for mission-critical enterprise applications where I/O performance and reliability are crucial, SATA drives will most likely be used for lighter-duty tasks where cost-per-gigabyte (vs. price/performance) is crucial. The ability to take a SATA drive and plug it into a SAS backplane without any modifications or changes is a tremendous benefit of SAS. A SATA host on the other hand can only receive SATA drives, and cannot accommodate SAS drives in the same architecture.

Figure 4 below shows the configuration flexibility of a SAS backplane, allowing for both SATA and SAS drives to be deployed.

Figure 4: SAS Backplane provides storage flexibility[4]

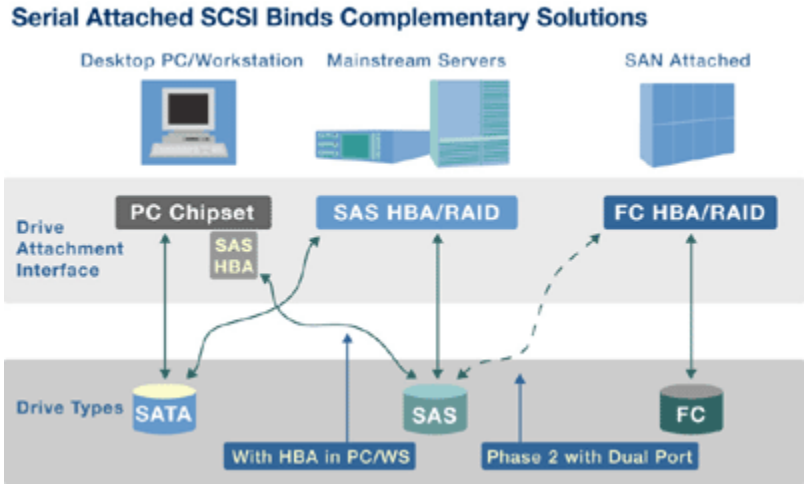

3.3 Improved scalability

Scalability was never among SATA or parallel SCSI's strengths, and a key goal of Serial Attached SCSI was to vastly improve the ease with which drives could be added to increase capacity and throughput. Unlike SATA, SAS can be connected to multiple hard drives through expanders, allowing a single SAS domain to contain up to 16,384 devices without performance degradation.

As the volume of enterprise data continues to expand, any viable long-term storage strategy must include simple, cost effective scalability as an essential component. SAS's ability to scale beyond the known

[4] Source: Seagate

boundaries of SCSI of 15 devices to envisaged 4,032 devices per port promises an unprecedented path of future capacity and performance growth.

3.4 Improved device reliability and system availability

The SAS interface is more reliable, has improved signal integrity, and hence a higher level of data availability, due to the following characteristics as compared with parallel SCSI and SATA:

- The dual-ported interface ensures high availability - in the event one SAS host controller fails, the extra data port maintains uninterrupted communication with a second controller.
- Data Integrity Field helps improve data integrity all the way from the device to the host controller
- The serial point-to-point architecture of SAS technology eliminates crosstalk and signal skews (added benefit for SATA too)
- Hot plug capability further ensures uninterrupted data availability by permitting drive swapping without having to shut the system down

4 SOLID STATE DRIVE VS. HARD DISK DRIVE

SAS was clearly designed with the purpose overcoming the inherent problems of hard disk drives related to performance and reliability/availability. When trying to understand whether SAS can benefit solid state storage, it is important to understand how these two parameters compare for both HDD and SSD.

Note: The added benefits of scalability and flexibility of SAS are not related to the storage media used, and would be the same for HDD vs. SSD.

4.1 HDD vs. SSD Performance

The performance requirements for a storage device in an enterprise environment depends where the storage devices sits in the architectural hierarchy. The need for high I/O performance and high reliability/availability is the most prevalent for online storage, where it is of utmost importance that the storage device delivers content in a fast and reliable manner. Today SCSI and FC are the dominant interfaces for online and nearline storage, and it is forecasted that these will be replaced with SAS in the years to come. HDD manufacturers position their 10K-15K RPM drives for this segment of the enterprise market.

Off-line storage is typically characterized by the need for high sustained performance and high capacity drives. These drives are not in the critical path of content delivery and their reliability/availability requirements are lower – cost/GB is the primary metric here. In the past, the ATA interface was the preferred choice, but it has been replaced rapidly over the last several years with SATA drives. HDD manufacturers position their 7,200 RPM drives for this segment of the market.

Until recently, solid state drives did not have any market share in the enterprise storage market due to their relatively high cost. With the year over year decline in flash cost, solid state drives have started to become a viable enterprise storage alternative. Due to its lack of rotating media, the access and seek time of solid state storage is superior to that of HDDs, and SSDs that have recently entered the market are matching up equally well in terms of sustained read/write performance.

Table 1 below shows a performance comparison between solid state drives and the two types of hard disk drives discussed above. It is clear that solid state drives could replace either type of hard disk drive, but the higher cost/GB for solid state drives makes it a more viable alternative for online/nearline storage today.

Table 1: Performance Comparison - SSD vs. HDD

Performance	Solid State Drive	Enterprise-class HDD 15K RPM[1]	Enterprise-class HDD 7,200 RPM[2]
Sustained Performance	Up to 100MB/sec Read Up to 80MB/sec Write	93-108MB/sec Read 57-80MB/sec Write	78.5MB/sec Read 44.3MB/sec Write
Access Time	< 0.3ms Read/Write	5.4-5.8ms Read 5.9-6.3ms Write	12.3ms Read 13.5ms Write
IOPS	3,500 – 20,000	165-411	76-93

[1] Parameters based on Seagate 147GB SAS Cheetah 15K.4 (ST3146854SS) and Seagate Savio 15K.1 73GB SAS (ST973451SS), tested by www.StorageReview.com under various workloads.
[2] Parameters based on Seagate Baracuda ES 750GB SATA (ST3750640NS), tested by www.StorageReview.com under various workloads.

4.2 HDD vs. SSD Reliability and Availability

The serial point-to-point architecture and dual port interface of SAS technology results in improved signal integrity and hence a higher level of system availability.

To address the enterprise market need for high reliability drives, hard disk manufacturers rate the MTBF numbers of their enterprise-class HDDs around 1Million hours. This number needs to be interpreted correctly though, since HDD manufacturers follow a different statistical model than solid state drive manufacturers. The MTBF for a HDD is the mean distribution of the HDD service life, where every disk within an environment must be replaced within its defined service life (3 to 5 years). The MTBF number is the sum of service life periods before the probability of a disk failure accumulates to 70%. This means that to achieve 1Million hours MTBF (~114 years), a HDD with a 5 year service life must be replaced 23 times.

Solid state drives have inherently higher reliability, since no platters are spinning or read/write heads moving. The MTBF number is far less influenced by duty cycle, power-on-hours, and environmental factors such as temperature and shock and vibration. Instead, MTBF numbers are calculated upon statistical failure rates of the internal components.

Table 2 below shows a comparison of some environmental and reliability parameters of the Adtron Flashpak SSD with three enterprise-class HDDs. As can be seen, the solid state drive shows much better environmental parameters than the enterprise class HDDs.

Table 2: Reliability Comparison - SSD vs. HDD

Reliability	Adtron Flashpak SSD	Enterprise-class HDD 15K RPM[1]	Enterprise-class HDD 10K RPM[2]	Enterprise-class HDD 7,200 RPM[3]
Temperature	-40 to 85C	5 to 55C	5 to 55C	5 to 55C
Shock	50G (11 ms)	30G (2ms)	65G (2ms)	65G (2ms)
Vibration	16.4G (10-2000Hz)	1.2G (5-500 Hz)-	3G (10-300Hz)	0.75G (20-300Hz)
MTBF	1.6M hrs	N/A	1.2M hrs	1.2M hrs

[1] Parameters based on Hitachi Ultrastar 15K300 (HUS153014VLS300)
[2] Parameters based on Western Digital Raptor 74GB (WD1500ADFD).
[3] Parameters based on Western Digital RE2 750GB SATA (WD7500AYYS).

As mentioned, comparing MTBF numbers should be done carefully, since different calculation methods are used by the different manufacturers. Instead, statistical annual failure rates are a far better indication of the reliability of HDDs and recent studies[5] by Google and Carnegie Mellon University show HDD failure rates of up to 8.6% in controlled environments.

[5] Failure Trends in a Large Disk Drive Population, Google Inc., 2007; Disk Failures in the Real World: What Does an MTTF of 1,000,000 Hours Mean to You? Carnegie Mellon University, 2007

4.3 SAS in Solid State Disks

Due to its point-to-point architecture and dual-port interface, SAS is the most reliable interface available in the industry today. HDD manufacturers that have been positioning their high reliability and high performance drives (10K-15K RPM) for the enterprise market are now providing these drives with a SAS interface. The performance specifications can be achieved with SATA too, but the high reliability is achieved by using the SAS protocol.

Can the same analogy be given for solid state drives? Would a solid state drive with SAS interface provide higher reliability than a SATA-based SSD? The answer is yes, but only marginally. Solid state drives are inherently so much more reliable than HDDs that the increased reliability due to the SAS interface will not significantly influence the overall system reliability. Instead, replacing a HDD with a SSD in any system will boost overall reliability, regardless of the chosen storage interface.

What about performance? Is a SAS-based SSD capable of achieving higher performance than a SATA-based SSD? The perception that SAS stands for high performance and high reliability is based upon the fact that HDD manufacturers design enterprise hard disks specifically for heavy workloads and large power-on hours. High performance in SSDs is related to internal architecture and controller technology only, and is not related to interface protocol that is used.

Table 3 below shows an overview of where SAS will be beneficial for both SSD and HDDs.

Table 3: SAS Benefits - SSD vs. HDD

SAS Benefit	Hard Disk	Solid State Disk
Better performance	√	No
Greater flexibility	√	√
Improved scalability	√	√
Improved device reliability	√	No
Improved system availability	√	√

5 CONCLUSION

As IT budgets continue to undergo intense scrutiny, the pressure to rationalize enterprise storage systems has never been greater. The ultimate goal of any IT manager has always been to employ the best, most cost-effective storage solutions for the applications at hand. The greater the flexibility in choosing and deploying those solutions, the more likely optimal performance and cost-effectiveness will be achieved.

With SAS backplanes providing the flexibility to deploy both SATA and SAS drives, customers used to have the choice of deploying cost-effective SATA HDD drives for low-cost bulk storage or higher performance SAS HDD drives for mission-critical applications. With the price points of solid state drives approaching levels that make them a viable option to replace HDDs, customers can now also consider solid state drives as the choice for online and nearline storage.

Deploying solid state drives as the online storage device will increase the overall I/O performance, as well as reliability. The choice for SAS or SATA is irrelevant when using SSDs, and the decision process should focus instead on the choice on SSD vs. HDD and the related Total Cost of Storage Ownership. Please refer to Adtron's "Total Cost of Storage Ownership" White Paper for more in-depth comparison on HDD vs. SSD.

About the Author

Esther Spanjer is Director of Technical Marketing at Adtron Corporation. With more than 10 years of experience with flash-based solutions, Ms. Spanjer has gained valuable insight into the use of this rapidly evolving technology in a wide range of embedded applications in the military, aerospace, communications and industrial markets. She joined Adtron to help evangelize the use of flash technology for new types of applications in both the traditional markets and the emerging enterprise market. As users become more familiar with the full range of benefits of flash technology, she believes that flash-based storage will be adopted in a broader range of what have typically been thought of as hard-disk applications.

Ms. Spanjer received a B.Sc. degree in Electronic Engineering from the Technical University Amsterdam (Netherlands) in 1991. She can be reached at espanjer@adtron.com.

About Adtron

Founded in 1985, Adtron is the leading designer and global supplier of high performance and high capacity solid state flash disk drives. Adtron Flashpak SSDs integrate seamlessly into defense/ aerospace, industrial automation, medical, transportation, telecom and enterprise applications. Based on Adtron's advanced ArrayPro™ performance engine, Adtron solid state flash disk drives deliver superior sustained read and write rates and are designed to reliably meet stringent environmental requirements. The Adtron Quality Management System is ISO 9001:2000 certified. Adtron is headquartered in Phoenix, Arizona with channels in all global markets.

Learn more about Adtron at http://www.adtron.com

ADTRON
Smart Storage, Smart People.®

SAS-R0-113007