Serial Attached SCSI

A White Paper Prepared for the SCSI Trade Association
November 2004

ENTERPRISE MANAGEMENT
ASSOCIATES

EXECUTIVE SUMMARY

Since its development in the early-1980's, SCSI has seen steady and dramatic increases in speed to move more information, starting with support for small disk drives at 1.5 megatransfers per second, to moving data at 320 megabytes per second between multi-terabyte storage arrays. Now the industry is gearing up for Serial Attached SCSI (SAS) – a new, faster version of SCSI designed to meet the demands of enterprise IT for the coming decade.

SAS builds on parallel SCSI by providing higher performance, improving data availability, and simplifying system design. Most importantly, however, because the SAS interface supports both SAS disk drives for data-intensive applications, and Serial ATA (SATA) drives for low-cost bulk storage of reference data, IT managers can now incorporate greater price and performance flexibility in their storage strategies. As a result, managers can now make storage decisions that better align with business requirements, balancing their storage investment between compatible technologies and assigning storage devices to data based on the information's business value.

20 YEARS OF SCSI DEVELOPMENT

The specification for the SCSI device interface has been evolving since the early 1980's. During that time, SCSI has moved well beyond its original function – to provide connectivity to storage devices attached to high performance workstations and smaller servers – so that today it has become the interface of choice for high performance servers in commercial, engineering, and scientific environments. In fact, the SCSI interface to disk, tape, and other storage devices has become so prevalent in the enterprise that younger IT managers are often surprised to learn that "SCSI" is an acronym for Small Computer System Interface. SCSI may have been designed for "small" systems once, but it now provides connectivity to the largest devices in the computer room, and it is clear that the technology has outgrown its original name.

The increasing demands of high performance and commercial computing have meant a constant requirement on SCSI designers to provide faster speeds and greater design flexibility. The technology developed from a 1.5-megatransfer, eight-bit-wide interface, to supporting 320 megatransfers per second at 16-bits. It moved from only supporting internal drives, to providing connectivity to external storage in RAID and JBOD devices through differential connections that maintained the signal integrity over greater lengths of cable. At the present time, SCSI devices have become so ubiquitous in the enterprise, that current estimates show that SCSI represents about 80% of the enterprise drive market.

Until now, all communication over the SCSI bus has been over a parallel connection, with all data (and the command overhead that controls the manipulation of the data) traveling over the cable at the same time, in parallel. Now, however, parallel SCSI is being replaced with a serial interface.

TABLE OF CONTENTS

WHY A NEXT STEP?

SAS takes advantage of existing SCSI development, in that it makes full use of the SCSI command set that has been developed and proven over the years. Because of its serial interconnect, it is also a major departure from the older physical technology. The application layer and command set remain the same, however, so users can make the transition easily.

The move from parallel to serial connections became necessary when SCSI development hit performance obstacles upon achieving transfer rates above 320 megabytes/second (Ultra320 SCSI). The requirements for implementing parallel data transfer, with all signals departing and arriving in parallel over longer distances and at greater speeds, became prohibitive. Use of a serial bus eliminates this timing requirement for signals, and enables many additional efficiencies.

WHAT YOU GET WITH SERIAL ATTACHED SCSI

The primary features of SAS are:

Point-to-point connectivity – SAS is a point-to-point architecture with each device connecting directly to a SCSI port (rather than having a number of devices sharing a common bus, as is the case with both parallel SCSI and Fibre Channel). Devices can be connected to expanders. This "unshared" connectivity improves data throughput, as well as data reliability, (point-to-point is more reliable than loops) and improves the ability to locate and repair failures.

Higher performance, increased efficiency – The current data transfer rate is 3.0 gigabits/second, with speeds of 6.0 gigabits/second planned for the next phase of development. Full duplex data transfer means reads and writes execute simultaneously, reducing a major source of latency. Wide links enable aggregate data transfers of 24 gigabits/second. Additionally, improved signal arbitration means that all devices have equal access to the host.

Connectivity and scalability – A single SAS host adapter can connect to hundreds of drives with throughput of 3.0 Gigabits per second (full duplex) in each link (4.8 Gigabytes in bursts on an eight-port adapter) when used with SAS expanders, providing significantly more build out capability than is available in current IT installations. When used with expanders, SAS supports up to 16,384 physical links.

Easier administration – All SAS devices are assigned unique worldwide names (SAS Addresses), making it easier to identify initiator ports, initiator devices, target ports, target devices, and expander devices. There is no setting of switches or IDs; all discovery and configuration is automatic.

Flexibility – SAS devices and less expensive SATA devices are plug compatible. The SAS backplane accommodates both types of device, allowing IT buyers to use both high performance SAS drives and low-cost SATA drives in the same box and in any sort of mix-and-match configuration. As with traditional SCSI implementations, drives are hot swappable, so that arrays can be reconfigured on the fly. Investment in disk drives can thus be more flexible than was previously possible in terms of matching performance and cost structure to business requirements.

Simpler cabling – With SAS, the large SCSI cables of old, with multiple large connectors, have given way to thin cables and small connectors; connected point-to-point to the SAS ports. Thinner cables are easier to manage, and enhance air circulation to help devices run cooler.

Improved reliability – Simpler system design, simplified cabling, simple point-to-point connections, elimination of cross-talk and signal skew (inherent problems with signals sent in parallel), elimination of any single point of failure through dual porting, and improved signal integrity mean more reliable systems and higher data availability.

	SAS	Parallel SCSI	SATA	Fibre Channel
Performance (present)	3 Gigabits/second	3 Gigabits/second	1.5 Gigabits/ second	4 Gigabits/second
Performance (planned)	6 Gigabits/second	-	3 Gigabits/second	8 Gigabits/second
Performance (wide links)	24 Gigabits/second	-	-	-
Scalability	16,000 devices per domain	15 devices per bus	1 device	127 devices per loop; many more when used with fabric switches
Dual ports	Yes	No	No	Yes
Connector	7 pin	68 pin	7 pin	Copper or fiber optic
Cabling	8 meters (per hop)	12 meters	1 meter	15 meters
Interface Compatibility	SAS and SATA	Parallel SCSI only	SATA only	Fibre Channel only

Table 1: Comparing enterprise hard disk drive interface technologies

SAS AND SATA IN THE SAME SYSTEM

While SAS was being developed, vendors of the less expensive ATA drive technology were upgrading their offerings to provide optional features, which may add cost to SATA drives. Serial ATA (SATA) shares many characteristics with SAS, and provides an inexpensive alternative for situations where investment in SAS drives is not warranted. SCSI drives have higher performance, (typically 10K or 15K RPM) for support of mission-critical enterprise data in applications where availability and reliability are requirements. SATA drives have lower performance, (typically 7200 RPM) lower cost, and are generally suited to less demanding storage environments. The slower rotational speed, higher capacity, and cost of SATA drives are suitable for data storage that is accessed less often.

Major similarities between the two technologies are:

- Both types of drive plug into the SAS backplane
- The drives are interchangeable within an SAS drive bay module
- Both are long-proven technologies, with worldwide acceptance

Major differences between the two technologies are:

- SATA devices will be less expensive
- SATA devices use the ATA command set
- SAS drives have dual porting capability, faster spindle speeds, and lower latencies
- While both types of drive plug into the SAS backplane, a SATA backplane cannot accommodate SAS drives
- SAS drives are tested against much more rigid specifications than are SATA drives, and have a significantly longer MTBF and duty cycle
- SAS drives are faster, and offer several features not available on SATA, including variable sector sizes, LED indicators, dual ports and data integrity
- SAS offers multi-initiator software support with a 20-year legacy

When choosing between SAS and SATA, follow this rule: Choose SATA when cost is the most important issue; choose SAS whenever data availability, reliability and performance count most.

Independent of the drive being used, the SAS infrastructure (device controllers, expanders, connectors, backplanes, cables) allows a system to be provisioned for either category of drive, permitting common infrastructure components to be integrated across a wide range of storage requirements within the enterprise.

MANAGING THE INFORMATION LIFECYCLE

Explosive data growth, budgetary pressures, and the time constraints associated with IT maintenance windows are causing a revolution in storage management techniques. Increasingly, these three countervailing forces are causing IT managers to come to a better understanding about the value of their data, and to make hard decisions about the resources they assign to protect their company's most vital asset.

Information Lifecycle Management (ILM) seeks to manage data according to its value within a company.

In theory, with ILM, high-value data is assigned to the best arrays and receives the best management services, while data of lesser value gets proportionately less investment in terms of the infrastructure and services that protect it. Among the many challenges associated with this approach to management is the fact that the value of data changes over time, (this week's sales reports are more useful than last week's, and much more valuable than last year's data) and that the data needs to be migrated from one class of device to another when its value changes.

Making sure that only high value data resides on high value storage arrays, and ensuring that such data is properly serviced in terms of backup and recovery applications, is an important role played by storage resource management (SRM) software. Software vendors are also providing improved software to help in migrating data from one tier of storage device to another in a timely (and in the best cases, an automated) fashion. But what is a storage tier?

Storage tiers will vary from site to site, but will typically follow a pattern similar to this: high-end storage (SCSI and Fibre Channel RAID devices) for a company's most valuable data; midrange storage (again, SCSI and Fibre Channel, but with slower devices and perhaps configured as JBODs rather than as RAID subsystems) for data that is of lesser value; still lesser-valued data assigned to SATA-based devices; and older data not typically referenced, and therefore of minimal value, archived to tape. Clearly an investment in SAS systems – and in SAS drives that are interchangeable with low-priced SATA disks – will allow a site to make disk investments that map directly to the site's data protection requirements, allowing site managers to reconfigure storage assets by replacing one type of drive with another, or by adding additional storage on an as-needed basis.

WHY VENDORS HAVE EMBRACED SERIAL ATTACHED SCSI

Vendors like SAS for the simplest of reasons: it enables them to produce faster equipment more quickly and at less cost than do other alternatives. This efficiency is achieved primarily because of the following:

Lower power requirement – Each SAS port requires only 0.3 watts* of power, approximately one quarter the requirement for Ultra320 SCSI.

Smaller cables and connectors – SAS cables are thin, round cables, like SATA cables, as contrasted with ribbon cables that parallel SCSI requires. SAS connectors are much smaller than SCSI connectors, all of which allows more room and better air flow and cooling.

The SAS backplane – The SAS backplane accepts both SCSI and SATA connections, giving vendors a "two-for-one" advantage, reducing development costs, and allowing them to produce storage devices that offer customers maximum flexibility in designing for their particular storage needs.

Simpler, less expensive system design – Serial data transfer means that backplane designers no longer have to worry about timing issues associated with signal routing. Because SAS devices are dual ported, it is easy to design systems with no single point of failure.

Faster design cycles – Because of all of the above.

Familiar technology, known protocols, proven devices – This is the same SCSI command set that vendors' engineering staffs have been working with since the 1980's. SCSI disk drives presently hold more enterprise data than any other online device. The same proven management tools will work with SAS.

Better Customer Value – Vendors can deliver high value products to their customers by building modular arrays. IT managers can then easily build pay-as-you-grow systems that are expanded with additional modules and populated with additional drives, (either SAS or SATA) only when such investments are really needed. Systems may be reprovisioned to reflect changing requirements within the enterprise.

* 0.3 watts is typical; however, it is dependent on the vendor's implementation.

EMA's View: What Users Can Expect, and When

The analysts at Enterprise Management Associates have been looking at storage and storage management for several years, and it is quite evident that, insofar as disk storage is concerned, the prospects for aligning storage with both technology requirements and business demands have improved significantly during the past 12 months. The two serial technologies, SATA for low-end (or near line) devices, and SAS for industrial strength and high performance, offer a variety of excellent alternatives with a common connector both for direct-attach applications, and for use in arrays. The common connector is important, because when users purchase SAS systems, they can continue to use them economically even when their data protection needs have changed.

A SAS system is the preferred approach to aggregate SATA drives for low-end, near-line storage applications even when SAS drives are not immediately installed. It's a better, easier way to integrate a large number of SATA disk drives. At a later time, if storage requirements change, SAS drives can be added to the system, or replace the SATA drives entirely.

The differences between the two classes of devices are clear-cut; buyers should have little difficulty deciding which type of device is most appropriate for their needs. The most common scenario, in EMA's opinion, will find SATA disks on desktops and in devices whose value lies in good performance at very low cost, such as virtual tape devices and a wide variety of single application appliances. The SCSI protocol, now in the form of SAS disk drives, will continue to dominate in all parts of the enterprise where performance, reliability, and flexibility of design for primary storage are the most important considerations. Serial Attached SCSI will continue to represent the lion's share of enterprise storage well into the future.

Neither SAS nor SATA should be viewed as a competitive threat to Fibre Channel, which will also continue to play a significant role in SAN and campus-wide storage. It is likely, however, that SAS and SATA devices will begin to appear with increasing frequency in the SAN as buyers begin to take advantage of what the technologies have to offer. By early 2005, deliveries of SAS devices will have begun. From that point on, look for SANs of the future to have a mix of SAS, SATA, and Fibre Channel devices, providing tiers of storage to which managers can assign and migrate data according to the data's current value to the company.

We expect future SANs to contain a mix of:

- Integrated SAS Storage – High performance SAS drives and low-cost/GB SATA drives integrated in a SAS system that utilizes a Fibre Channel bridge to connect to the SAN.
- Integrated Fibre Channel Storage – High-performance Fibre Channel drives and low-cost/GB Fibre Channel drives integrated in a Fibre Channel system that seamlessly connects to the SAN.
- Discrete SATA Storage – Low-cost/GB SATA drives in a SATA box that utilizes a Fibre Channel bridge to connect to the SAN.

All of this, when coupled with efficient storage resource management software, holds the promise of creating IT environments that will be far more efficient than were their counterparts of just a few years ago. Faster throughput and improved data integrity from the SAS devices, economical application of SATA devices, and Fibre Channel used for SAN connectivity will all contribute to an improved, efficient and increasingly cost-effective storage environment.

IT management needs to keep in mind that the SAS infrastructure has to be designed for the rotation vibration changes that SATA drives present. These

vibrations are not an issue at lower vibration transfer rates; however, if the SATA drives are running at high rates, the likelihood of MTBF failure and the opportunity for misreads due to rotational vibration increases. For those managers who have learned to rely heavily on SCSI technology over the last 20 years or so, it will also be comforting to know that, as parallel SCSI approaches its waning years, (it will probably last five to eight years) the move to Serial Attached SCSI will be an easy transition for enterprise IT organizations to make.

SCSI technology has a long history in enterprise storage, and SAS should be viewed as a next-step in SCSI's evolutionary development. As such, it offers enterprise IT managers proven technology with no surprises. As long as increasing data demands and continued budget restraints remain the major influences on all IT buying, it is hard to beat the combination of SAS and SATA on the floor of the enterprise data room.

About SCSI Trade Association

The SCSI Trade Association was established in 1995 to provide a focal point for members to communicate the benefits of SCSI to the industry. STA promotes the understanding and use of SCSI technology and influences the evolution of SCSI standards to meet future industry needs. The Association has an eight-member Board of Directors, which oversees the Marketing Communications and Technology Committees and all STA activities. For more information, please visit the STA web site at http://www.scsita.org, send an email to info@scsita.org or call the STA office at (415) 561-6273.

Corporate Headquarters
2585 Central Avenue, Suite 100
Boulder, CO 80301
Phone: 303.543.9500
Fax: 303.543.7687
Email: info@enterprisemanagement.com
Web: www.enterprisemanagement.com