Seagate
We turn on ideas

From: Global Product Marketing
Date: January 2004 • Number: TP-519

Technology Paper

Bridging the Enterprise Storage Continuum With Serial Attached SCSI

Introduction

The advent of Serial Attached SCSI devices reflects a growing demand for more targeted solutions in enterprise storage, to enable both higher performance and greater cost-effectiveness. Tightening IT budgets certainly play a central role in this movement, but it's also driven by a fundamental change in the nature of enterprise data management. Gone are the days when there were only two types of data storage: online and offline. A new concept, near-line or secondary storage, has become increasingly popular because it recognizes that data has a lifecycle, during which its need for accessibility and security will vary. Rather than a simple online/offline dichotomy, enterprise storage now comprises a continuum along which data can easily migrate as its value and relevance change.

From the beginning Serial Attached SCSI was designed to strengthen that continuum, both by retaining its parallel predecessor's reliability while setting new standards in performance and scalability, and by introducing an unprecedented degree of compatibility and flexibility to enterprise storage solutions. The emergence of Serial ATA (SATA) drives accompanied a profound advance in areal density, to the extent that IT managers began to deploy SATA for near-line storage of non-critical data and backups/restores, where SATA drives offer vastly faster access (at appropriate capacities) compared to tape. The authors of the Serial Attached SCSI standards well understood the significance of SATA's new role in the enterprise, and the synergies (both fiscal and physical) that would result if SAS and SATA drives could share a single storage subsystem.

THE ENTERPRISE STORAGE CONTINUUM

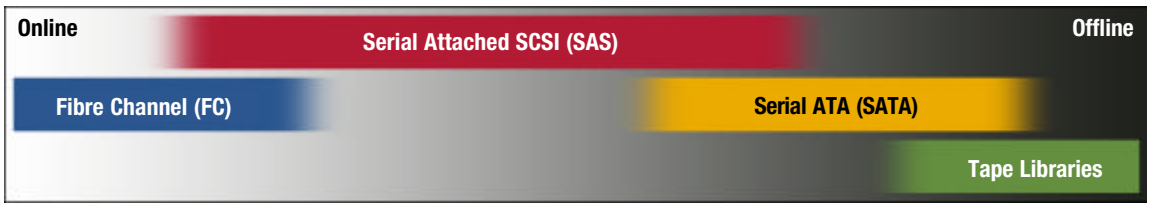

Background: Parallel SCSI vs. Parallel ATA

To fully appreciate the significance of Serial Attached SCSI and Serial ATA compatibility, it's instructive to first review why parallel SCSI and parallel ATA were not designed to work together, and how evolving storage requirements have compelled a new era of flexibility and connectivity.

Over two decades ago, when parallel SCSI and parallel ATA standards were written, compatibility was simply not an issue. The roles of the two interfaces were seen in black and white terms, with no expectation they'd ever be called upon to interoperate in common applications. At one end of the spectrum was SCSI, clearly intended for enterprise duty; at the other was ATA, optimized for use in desktop computers. The former required versatility, control and customization, while the latter demanded simplicity and low cost.

To achieve these disparate goals, the two interfaces employed markedly different (and utterly incompatible) technologies. Parallel SCSI utilizes a bus-structured cable supporting up to 16 ports daisy-chained on a single bus, with an addressing scheme that allows a host controller to select and communicate with any one of the daisy-chained ports. Cable lengths of up to 12 meters are possible, enabling the interface to accommodate both internal (within server enclosure) and external (JBOD or cabinet) drives. The SCSI protocol supports a rich command set suitable for complex enterprise disc management operations such as custom applications, advanced error recovery and performance optimization.

By contrast, parallel ATA's cost-sensitive desktop mission dictated a point-to-point interface between two ports that neither supports nor requires an addressing scheme (a rudimentary master/slave arrangement suffices). It handles internal drives only, as its short cable length (1 meter) is incompatible with external storage. ATA's protocol is register-based, thus simplistic and inexpensive, as is its command set.

New Challenges, New Solutions

Over the next two decades the enterprise storage landscape changed dramatically, with the rapid growth of the Internet spurring explosive demand for online/transactional storage solutions that delivered data with high availability and reliability. The traditional Direct Attached Storage (DAS) model was augmented by Storage Area Network (SAN) and Network Attached Storage (NAS) environments. SCSI received repeated speed boosts, but its parallel architecture and shared bus limited both its performance and scalability potential.

Fibre Channel arrived with an advanced serial architecture to achieve stunning throughput and scalability, soon superceding parallel SCSI as the premium storage solution for mission-critical data in large enterprises. For mid-size enterprises Fibre Channel's infrastructure cost was daunting, hence many continued with SCSI for mission-critical use while investigating ways to move their vast quantities of non-critical, infrequently accessed data from online SCSI subsystems onto less costly near-line alternatives. As areal density continued its rapid growth, inexpensive servers with internal ATA drives became viable near-line storage solutions, though as noted earlier such configurations are severely limited in flexibility and scalability.

Completing the Storage Continuum

It was the debut of low-cost, high-capacity Serial ATA drives that signaled greater flexibility and performance in near-line storage. SATA offers the requisite high capacity at low cost, and does so without the scalability constraints of its parallel ancestor. Now the enterprise storage continuum had strong solutions at the high end (Fibre Channel) and low end (SATA), but the midrange was conspicuously lagging behind. Parallel SCSI was reaching the practical limits of its development potential—clearly a new solution was required.

Enter Serial Attached SCSI, which leverages the proven strengths of parallel SCSI (rock-solid reliability, a rich and mature command set) while employing a new serial architecture to achieve blazing throughput (3.0 Gbits/sec) and remarkable scalability (up to 16,384 devices through the use of expanders). Because development of Serial Attached SCSI began after the Serial ATA 1.0 standard had been completed, the storage industry wisely chose to incorporate compatibility with SATA as a key element of the SAS standard. The enterprise storage continuum was now complete, offering a comprehensive range of solutions that address storage needs at every level (see chart on next page).

Solution	Application	Notes
Fibre Channel	Online, high availability, random reads	Suitable for mission-critical data in large enterprises [for example, Storage Area Network (SAN)], supports up to 16 million addresses and maximum cabling distance of ten kilometers, relatively costly
Serial Attached SCSI (SAS)	Online, high availability, random reads	Suitable for mission-critical data in mid-size enterprises, high performance coupled with exceptional scalability on local level, more affordable than FC, compatible with SATA
Serial ATA (SATA)	Near-line, low availability, sequential reads	High capacity at low cost, not suitable for mission-critical data, compatible with SAS
Tape Libraries	Offline	High capacity, ability to move data offsite for archival and security

Compatibility Step by Step

Before delving into the many benefits SAS/SATA compatibility offers, a brief review of how it was achieved is in order. To ensure full interoperability with SATA, Serial Attached SCSI needed to address a variety of issues:

PROTOCOLS

Serial Attached SCSI employs three different protocols to transport information over its serial interface: Serial SCSI Protocol (SSP), SCSI Management Protocol (SMP) and Serial ATA Tunneling Protocol (STP). SSP communicates with SAS devices and existing SCSI software, while SMP manages SAS's point-to-point topology. The third protocol, STP, enables SAS controllers to identify and communicate with Serial ATA devices. Thus when a SATA drive is added to a SAS backplane with an edge expander, an STP connection is immediately opened to enable SATA frames to pass through the connection to the drive.

THE THREE SAS PROTOCOLS: SSP, SMP AND STP

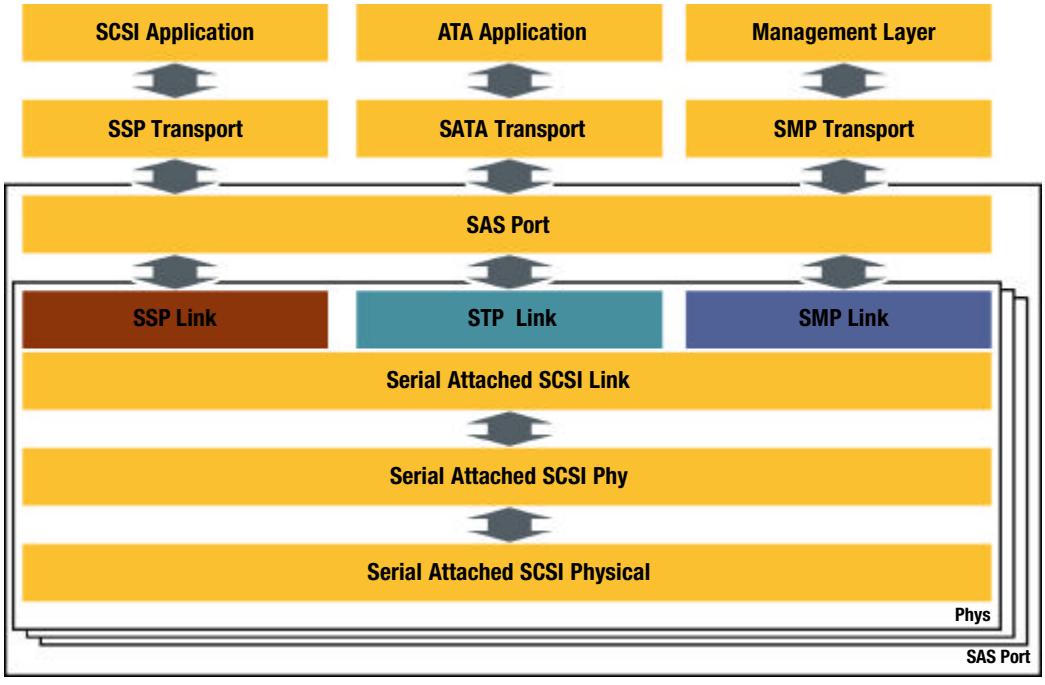

CABLES AND CONNECTORS

SAS cables and connectors are fully compatible with SATA devices. Note that while SATA and SAS connectors share the same form factor, they are not identical. Unlike SATA's single-port connector, which utilizes seven pins (four signal and three ground), the SAS connector fills in the keyway on the SATA connector and adds seven more pins on the opposite side to carry the additional signals in SAS's dual-port interface. The modified keyway thus enables SATA-to-SAS compatibility while ensuring a SAS device cannot be plugged in to a SATA port.

Note: As is the case with SATA connectors, SAS connectors are designed for hot plugging and blind mating.

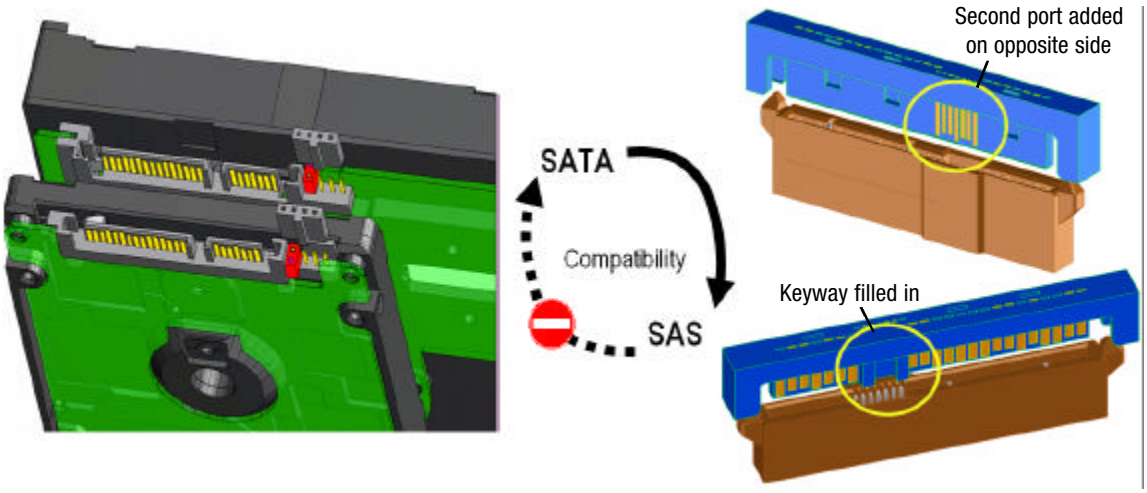

BACKPLANES AND HOST BUS ADAPTERS

SAS backplanes and host bus adapters (HBA) offer full compatibility with SATA (single port) devices, while still incorporating the second data port that gives SAS its fail-over capability. In the event one SAS host controller fails, the additional data port ensures uninterrupted communication with a second controller. In addition, these two ports can be combined into a single "wide port" for higher throughput.

Reflecting its enterprise-class roots, SAS employs higher transmit (Tx) and receive (Rx) differential voltages in order to drive signals over backplanes and long (up to ten meters) cables. As a desktop solution, SATA need only drive short (one meter) cables and thus can use lower differential voltages. To prevent damage to SATA drives when connected to a SAS port, device identification as noted in the STP protocol is used to ensure the proper voltage is delivered to each device.

FORM FACTOR

SAS gives IT users the freedom to choose the right drive for their applications. For denser computing environments in which raw capacity takes a back seat to higher throughput (IOPS/U), 2.5-inch SAS drives will play an increasingly prominent role. Choosing 2.5-inch SAS drives for transactional applications such as database storage for ERP and CRM software does not prevent IT users from achieving compatibility with 3.5-inch SATA drives that perform periodic backup and restore functions. For example, servers and storage subsystems of varying sizes (1U, 2U, 4U, etc.) stacked on top of one another in a cabinet can transfer data interchangeably between SAS and SATA drives. Depending on the cabinet's configuration, Serial Tunneling Protocol (STP) could be accomplished at an HBA or expander, thereby freeing the need for SATA and SAS drives to share the same backplane.

Of course it goes without saying that Serial Attached SCSI drives will be available in the industry-standard 3.5-inch form factor for storage systems utilizing a common backplane. A single storage subsystem will now be able to house a low-cost 7200-RPM SATA drive in the same enclosure as the preferred online enterprise solution, the industry-standard 15K-RPM SCSI drive.

The Benefits of SAS/SATA Compatibility

As IT budgets continue to undergo intense scrutiny, the pressure to maximize storage efficiency has never been greater. But in truth the ultimate goal of any enterprise 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 optimum performance and cost-effectiveness will be achieved.

The concept of Information Lifecycle Management (ILM) has helped to clarify and codify the need to allocate storage resources efficiently, based on the time value of data. Migrating data from high availability solutions like SAS to low-cost SATA drives will play a key role in successfully implementing ILM, and the compatibility of those two solutions will streamline the process.

Because SAS and SATA drives can share a common backplane and be housed in a single enclosure, a single SAS-based subsystem is capable of handling the full gamut of enterprise storage duties, from high-availability online storage to data migration, disc-to-disc backup and near-line storage. Consolidating drives into fewer, more rationalized subsystems can save a substantial amount of both money and space, as well as reducing IT storage maintenance costs.

Growing firms can purchase SAS cables, backplanes and enclosures for initial use with SATA drives only, secure in the knowledge that such equipment will not become obsolete when their storage needs expand and enterprise-class drives are needed. SAS drives can simply be plugged in to their existing storage subsystems.

Finally, SAS cable/connector compatibility with SATA yields benefits for VARs and system integrators (fewer SKUs to stock), as well as customers (fewer SKUs to purchase and qualify, fewer cable types for IT managers to inventory). In addition, SAS's thin cabling allows easier routing and improved airflow/cooling within enclosures, while its compact connectors are ideally suited to the eagerly-anticipated 2.5-inch form factor SAS drives.

Common Interface, Different Solutions

To fully benefit from the remarkable flexibility SAS/SATA compatibility offers, it's important to remember the fundamental differences between these two serial-based storage solutions. SAS was designed and engineered for high availability, enterprise-class data storage, where random-read performance, reliability and data integrity are absolutely paramount. SATA is a desktop-class solution that offers high capacity at low cost and impressive throughput on sequential reads (for example, backup and restore operations). The key is to understand that SAS and SATA complement each other, and that their compatibility in no way suggests that they are interchangeable.

It's easy to identify cases where moving data from SAS to SATA would be beneficial. For example, using SAS drives for bulk storage of non-critical, low-availability data is a wasteful allocation of resources, and SAS backplanes will enable IT managers to seamlessly substitute inexpensive SATA drives for such near-line storage applications. That's an effective, entirely appropriate use of SATA drives in an enterprise environment.

However, problems arise when IT managers lose sight of the fact that SAS and SATA are intended to fulfill radically divergent needs. Pressured to reduce costs and flushed with enthusiasm for SATA's apparent price/performance advantage, some IT managers have begun deploying SATA drives in enterprise roles for which they are wholly unsuited.

Make no mistake, though their serial interfaces make them appear quite similar, SAS and SATA drives are vastly different. SAS drives are engineered for rugged enterprise duty, and every component (drive motor, spindle, actuator, magnetic recording heads, control and servo processors, firmware, etc.) is specifically designed and manufactured for that rigorous use. SATA drives are impressive examples of clever engineering to achieve solid performance at low cost, but they are in no way capable of meeting the performance and reliability standards demanded by high-availability enterprise storage solutions.

Device	Duty Cycle	MTBF (Typical)	Seek Time (Typical)
SAS (15K RPM)	24 hrs/day, 7 days/week	1,200,000 hours	3.6 msec
SATA (7200 RPM)	8 hrs/day, 5 days/week	600,000 hours	9.5 msec

Conclusion

Enterprise storage has entered a new era in which targeted solutions are now available to address a broad variety of needs, from the most rigorous high-availability applications to cost-sensitive bulk data storage. Serial Attached SCSI will form the central core of this new enterprise storage continuum, and its seamless compatibility with SATA further enhances SAS's appeal as a powerful, flexible enterprise storage solution.

AMERICAS **Seagate Technology LLC** 920 Disc Drive, Scotts Valley, California 95066, United States, 831-438-6550
ASIA/PACIFIC **Seagate Technology International Ltd.** 151 Lorong Chuan, New Tech Park #06-01, Singapore 556741, 65-6488-7498
EUROPE, MIDDLE EAST AND AFRICA **Seagate Technology SAS** 130-136, rue de Silly, 92773 Boulogne-Billancourt Cedex, France, 33 1-41 86 10 00