

SAS and SATA Team Up for the Enterprise

Executive Summary

The enterprise is currently experiencing unprecedented need for storage while at the same time being constrained by flat IT budgets.

Serial data transmission, which offers performance capable of scaling with ever-growing enterprise storage needs, is rapidly making inroads into disk drive interconnect technology, promising to eventually replace existing parallel architectures.

Legacy parallel Advanced Technology Attachment (ATA) drives, long a commodity desktop storage item, are being replaced by serial ATA (SATA) drives, which raise the ceiling on performance and do away with bulky cables while maintaining a commodity pricing structure. Newer SATA drives that are specifically designed with enterprise storage features offer improved reliability that rivals traditional Fibre Channel (FC) and Small Computer System Interface (SCSI) drives.

Complementing these advancements in SATA are the equally impressive performance gains being made with Serial Attached SCSI (SAS) drives. Small form-factor (2.5") SAS drives are being offered with greater I/O throughput, reliability improvements, and the ability to scale far beyond what was possible with traditional parallel SCSI drives.

A robust SAS infrastructure, which supports both the new SAS and SATA drives without any bridging, enables new options for delivering diverse storage solutions. For the first time in the history of enterprise computing, IT managers are now able to mix storage drives within the same infrastructure, targeting different types of storage to different application needs. This provides an exceptional opportunity for OEMs and systems integrators to leverage a common infrastructure in servicing a broad range of customer needs.

LSI Logic Corporation

Harry Mason

Director Industry Marketing
Storage Standard Products

Intel Corporation

Scott Peiffer

Product Line Manager
Storage Platforms
Storage Server Marketing

Western Digital Corporation

Hubbert Smith

Director
Enterprise Marketing

Enterprise Storage Expanding Rapidly

A number of factors are pushing rapid growth in enterprise storage and creating demand for technology companies to respond to many different customer needs. Demands for faster, cheaper, better, and more storage fall into well-defined areas:

- **More storage capacity.** Demand for enterprise storage is coming from new areas such as rich media, Enterprise Resource Planning (ERP) applications, and growing regulatory requirements, as well as an increasing emphasis on backup and data recovery.
- **Low-cost storage.** While the demand for capacity is increasing, IT budgets aren't budging. Traditional Storage Area Network (SAN) infrastructures in many instances are too expensive to meet the growing demand and will need to be replaced by new lower-cost alternatives that offer comparable reliability and performance.
- **Scalable storage.** More than ever, IT organizations are being tasked to develop an infrastructure that can flex and grow with the business. In many instances, IT managers must make hard tradeoffs as to whether they want a storage infrastructure that can grow based on low cost or high performance.
- **Higher performance storage.** New enterprise requirements such as ERP, e-Commerce, data warehousing, and enterprise grid computing require higher performance. There is also a need to back up data without affecting performance.

Enterprise Storage Solutions Today

The storage industry provides three hard disk drive options with varying levels of cost and performance: They are Fibre Channel, SCSI/SAS and SATA (see Table 1).

Each storage technology offers different advantages suited to different usage models. Unfortunately, they are not interchangeable within an infrastructure without the use of expensive bridging technology. This tends to lock enterprises into a particular technology and limit scalability.

FC and SCSI systems offer high-performance options

such as multi-host failover support that are inherently designed for high-availability mission-critical applications. Both FC and SCSI use common command structures based on the 20-plus-year SCSI legacy of enterprise-proven middleware. These drives are critical in environments where transactions-per-second is important and have a rugged construction that maintains low-vibration characteristics under extreme workload environments. Hundreds of these drives can work cooperatively in an infrastructure without impacting drive performance, data availability or data reliability. Moreover, the legacy middleware applications that support these drives are enterprise-proven and offer a low-risk approach to storage.

As enterprise data grows a variety of new storage applications arise, users demand lower cost alternatives, especially in situations where data may be less frequently accessed. Applications that require infrequently accessed storage, such as archiving

Table 1: Enterprise Storage Solutions

Features	Fibre Channel	SCSI/SAS	SATA	SATA for enterprise
Performance (Speed)	2 Gbps (now) 4 Gbps (new) 8 Gbps (future)	Ultra320 (now) 3 Gbps (new) 6 Gbps (future)	1.5 Gbps (now) 3.0 Gbps (new)	3.0 Gbps (new)
Performance (RPM)	10K -15K	10K - 15K	7.2K-10K	7.2K- 10K
Typical Large Capacity	146GB	146GB	400GB	400GB
Typical MTBF (Hour)	1.4 M	1.4 M	600K	1.2 M
\$/GB	\$3.50 - \$3.00 ¹	\$3.40 - \$2.60 ¹	\$1.00 - \$0.75 ¹	\$1.00 - \$0.80 ¹
Usage Rating	High duty cycle, server and networked storage	High duty cycle, server and networked storage		High duty cycle, server and networked storage
Form Factor	3.5 & 2.5	3.5 & 2.5	3.5 & mobile SATA ²	3.5

1. Prices are approximate for small quantity (<10K purchases), and may vary based upon performance, reliability and form factors.

2. Pricing for small form factor drives (2.5") including mobile SATA drives generally have a 40%+ \$/GB premium.

files, backing up systems, and referencing data, can benefit from lower cost storage technology.

SATA drives incorporate improved reliability features over their parallel predecessors. New improvements, such as extended disk drive platform reliability, error recovery, queuing, and jumper elimination, have been responsible for SATA moving into the enterprise world where it has found acceptance as Web server and Redundant Array of Inexpensive Disks (RAID) storage. Vendors have eliminated the bulky flat-ribbon cables and incorporated hot-swapping, a requirement of many enterprise environments, into the SATA plug. SATA drives support these application environments by increasing the performance ceiling established by parallel ATA drives but retaining the commodity pricing established by desktop volumes.

The Evolution to Serial Technology

Serial Data Transmission will Eventually Replace Parallel.

Engineering challenges with parallel transmission, including crosstalk, ground bounce, ringing, and clock skew, grow harder to overcome as data transmission speeds increase, to the point where parallel technology has reached its upper bound. As disk drive platter speeds and data density increase, the interface between the drive electronics and the controller card must handle ever faster data transfers. The only way to meet this demand is with serial data transmission.

Serial disk drive technology is already pervasive as the interconnect technology for SAN and Network Attached Storage (NAS) environments, where it is required to meet the need for distance and storage consolidation. However, this usage is only about 20 percent of the disk drive interconnects. Now the market is discovering that point-to-point serial architectures are required to improve system reliability, reduce cabling congestion (where cables are necessary), and to enable the infrastructure to scale with the insatiable demand for accessing more data quickly.

Parallel technologies comprise a large portion of the installed storage base in today's enterprise, and it's likely to remain that way for some time. Most drive suppliers expect to continue supporting parallel SCSI and ATA drives for enterprise customers for many years to come. However, these implementations are not without problems, especially as storage systems expand far beyond their original intent. Some implementations are already experiencing problems with the limits of parallel data transmission and these are becoming increasingly difficult to manage as data transmission speeds increase. Other implementations have had problems related to scalability, as demands on their

data storage environment have increased. As a result, many enterprises are now finding that they need to migrate to serial interface technology sooner rather than later in order to keep up with the mounting business pressures that are demanding more and more from their storage infrastructure.

Serial Choices

Fibre Channel drives have always been serial. Fibre Channel systems operating at 4 Gbps are expected to begin shipping later this year, and the roadmap to 8 Gbps has been ratified by the Fibre Channel Industry Association (FCIA).

Fibre Channel drives are typically found in external storage subsystems that have been connected to the host through a Fibre Channel SAN. While based on the same standard, the interfaces can take on different characteristics depending upon whether its primary use is as a SAN-attach or drive-attach. The connections to the host are primarily Host Bus Adapters (HBAs) that communicate through FC switches before reaching the external storage subsystems. While the infrastructure can be costly in some instances, users are able to share data more freely and more efficiently between a large number of distributed hosts. Also, the dual-port nature of this interface provides a failover capability that other interfaces like SAS are just beginning to exploit.

SAS drives will be available in mid-2005. The emergence of 2.5" SAS form factors enables new classes of high-performance, highly available RAID systems in racks as small as a 1U. Faster interface rates for the drives and better scalability than parallel SCSI offers users the capability of incrementally expanding storage capability. It also offers better investment protection than the parallel interfaces of today.

Market data forecasts for 2006 predict that installation of SAS drives will surge and become a significant percentage of overall SCSI shipments. SAS drives will find their way into traditional SCSI markets, primarily standard high-volume servers, as well as external storage enclosures that can be cascaded with the use of SAS expanders. Unlike Fibre Channel, the SAS infrastructure is not considered a SAN, but the fabric can be used to greatly expand near-box storage far beyond traditional parallel SCSI applications. In some cases, SAS is also finding its way into remote storage subsystems, especially in NAS-based implementations.

Enterprise class SATA drives are available now. Advancements in SATA drives oriented specifically to the enterprise have resulted in improved drive reliability, as confirmed in server and video security applications requiring 24x7 duty cycles. Active vibration correction technologies have solved the problem created when vibrations from one drive impact the tracking of an adjacent drive. Also, the new generations of SATA 3.0 Gbps drives with Native Command Queuing (NCQ) allow an onboard processor to reorder I/O reads and writes. This minimizes head movement and the amount of time the head spends waiting for the platter to spin around to the proper sector. While NCQ is not as powerful as the queuing found on FC or future SAS drives, the market is accepting NCQ as adequate for many mainstream server tasks, with the exception of heavily loaded, highly random database workloads typically performed by larger servers.

These additional drive improvements have greatly enhanced data reliability and availability and will expand the acceptance of SATA drives in mission-critical enterprise applications. Although the enterprise version of these drives may be priced at a premium (25-40 percent higher) relative to their commoditized desktop brethren, in most cases they offer a lower priced alternative in enterprise environments. The ability to run with existing SATA software means that no new device drivers are required. All mapping of operating system I/O calls from legacy software drivers to the new serial interface is done on the controller card.

SATA in a SAS Infrastructure – Enterprise Solutions for Tomorrow

SAS Infrastructure Supports Both SAS and SATA Drives

The SAS infrastructure supports both SAS and SATA drives. The SAS interface was designed from the outset as a superset of the SATA interface. SATA drives can plug into a SAS connector, though the opposite is not true. No bridging is required. There are no jumpers to select and no software to change. The SAS protocol provides a discovery capability that allows the host controller to sense whether the drive attached to the system is a SAS or a SATA drive and will adjust accordingly.

Flexible Scalability

For the first time in the evolution of storage technology there will be the opportunity to establish storage infrastructure in the enterprise that is capable of growing in two directions – affordable yet reliable capacity or high performance. SAS expanders allow for systems with either SAS or SATA drives to scale greatly, and the dual-port nature of the SAS infrastructure provides a failover structure that works independently of the drive type being used.


Native SATA has an inherent one-to-one association with the host, which does not lend itself to a multi-host configuration. The SATA tunneling mechanism provided with the SAS protocol allows a SATA device to operate without knowing the host from which a request came. The SATA one-to-one association is managed at the end-point that connects to the SATA drive – another advantage of the point-to-point nature of both SAS and SATA.

The value of stored data changes over time within an enterprise. Concepts like Information Life-cycle Management (ILM) address the dynamic nature of data. Information that one day is critical for stocking the shelves with the right items in the right location at the right time, may now only be useful for business analysis, historical references or regulatory compliance.

The SAS infrastructure delivers compelling value to this challenging enterprise environment by allowing storage enclosures to be re-provisioned within the enterprise as the demands upon these data systems change. The high-performance ERP applications of today may use SAS drives initially to address I/O performance requirements. As data accumulates and newer and faster storage subsystems are acquired to satisfy increasing performance demands, SAS enclosures can be re-provisioned with SATA drives to expand the available storage needed for regulatory compliance.

SAS + SATA Configuration Examples

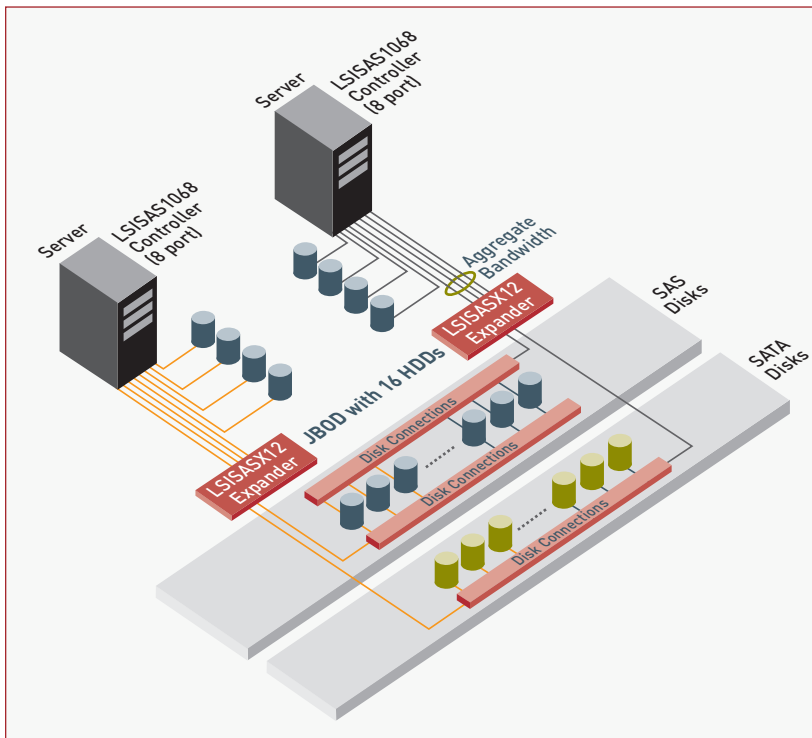
Figure 1 – SAS System Expansion Capabilities



SAS System Expansion Capabilities

SAS has the ability to address upwards of 16,000 devices per port through low-cost expansion devices called expanders. Simple expanders are referred to as edge expanders and allow up to 128 devices to be addressed without the cost of complex routing tables. Edge expanders are often cascaded to provide the required number of ports to support a large numbers of drives contained within an enclosure or rack-mounted shelf. More sophisticated implementations require fanout expanders that route traffic between the various edge expanders and provide for much greater expansion capabilities.

Figure 2 – SAS Performance and Availability



SAS Performance and Availability

Because SAS allows ports to be aggregated, high-bandwidth connections between the host and storage enclosures give SAS the bandwidth to support large numbers of drives. The 4-port JBOD (Just a Bunch Of Disks/Drives – as distinct from RAID) connection shown in Figure 2 may be implemented using standard cabling techniques and provides an aggregate bandwidth to the host of 24 Gbps. The dual-port nature of SAS also provides high-availability failover capabilities, providing access to all storage devices in the case of a failed host. Either SAS or SATA drives may be attached directly to host controller devices or addressed through expanders.

Conclusion

For the first time, IT buyers can implement a single storage infrastructure that offers both flexibility and scalability. Its ability to uniquely serve the high-performance needs of the enterprise with SAS drives and also the high-capacity needs with enterprise-class SATA storage allows it to scale flexibly as enterprise needs change. This flexibility, combined with efficiencies provided by combinations of SAS and SATA drives that have common components (cables, connectors, enclosures and backplanes), will allow the enterprise to meet new demand without substantially growing existing IT budgets.

Contact Information



Harry Mason
LSI Logic Corporation
hmason@lsil.com



Scott Peiffer
Intel Corporation
scott.t.peiffer@intel.com



Hubbert Smith
Western Digital Corporation
hubbert.smith@wdc.com

www.lsilogic.com

www.intel.com

www.wdc.com

Copyright ©2005 Intel Corporation. Intel and the Intel logo are trademarks or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others. 0505/MJR/HBD/PDF