An Introduction to Serial Attached SCSI (SAS)

ULTAMUS™ RAID

Overland STORAGE®

Simply Protected

Save

How SAS Evolved

Serial Attached SCSI (SAS) is the next generation of business-class storage technology, the successor to UltraSCSI (Small Computer System Interface) and the top contender to the dominance of Fibre Channel in high performance storage. SAS enables a new generation of storage products that exploit its innovative scalability, bandwidth and performance architectures. This document overviews the architecture, the technology, host adapters, disk drives and storage appliances.

SAS serializes the transmission of data, requiring far fewer physical connections and eliminating the need for the familiar SCSI bus. The current SAS specification defines a 300 MB/sec fully-bidirectional bandwidth per port, as well as exchanging the shared bus for point-to-point connections between devices. Eliminating the shared bus is a significant benefit as it removes performance-robbing bus contention and the potential for bus saturation. Finally, SAS employs the familiar and well-known SCSI protocol, ensuring ease of implementation and application compatibility.

The SCSI specification defines hardware and software standards for command and data transfers between a server and a storage device. For over two decades, SCSI has been a mainstay in the enterprise, evidence that the technology is relied upon by many users. Its market dominance and the large number of developers who design products to meet the specification has created economies of scale that make SCSI one of the most affordable interconnect technologies available today. Over the years SCSI has evolved, adding increased functionality, performance and reliability with each release. Yet, the current generation of SCSI, Ultra320, represents the end of the road for parallel disk interconnect technologies.

The SAS Architecture

SAS uses a serial, point-to-point interface in which disk drives are linked directly to RAID controllers, enabling systems with many drives to scale in bandwidth far beyond the capability of SCSI. To complement this bandwidth scalability, SAS expands the maximum number of devices to 16,384, well beyond the maximum number (16) that can exist in a parallel SCSI domain.

Up to Four Times (4x) the Performance Bandwidth

SAS offers a direct link from the RAID or JBOD controller to disk drives. In contrast, SCSI drives use a parallel interface in which drives share a common bus and compete for a portion of the available bandwidth.

SAS offers the ability to aggregate ports to form "wide ports" capable of supporting higher throughput. ULTAMUS RAID products employ SAS connectors (IN and OUT) in which 4 SAS ports are aggregated to yield a "wide port" with a bandwidth of 1.2 Gb/sec (1,200 MB/sec) (2.4 Gb/sec or 2,400 MB/sec in full duplex mode). This produces a total bandwidth that ranges from nearly four-times (4X) to eight times (8X) that of a SCSI bus (320 MB/sec).

These bandwidth advantages taken together give SAS an edge over SCSI, delivering more headroom for scaling application performance.

The SAS Expander—the Key to SAS Performance and Scalability

There are two types of SAS devices: end devices and expander devices. End devices include SAS initiators, SAS targets, and SATA devices. Expanders include port expanders (also referred to as "edge expanders) and fanout expanders.

SAS initiators can connect to a total of 128 SAS targets, SAS expanders, or SATA devices. SAS expanders can be interlinked using "narrow links" or "wide links." SAS targets (like disk drives) can connect to SAS initiators or SAS expanders using a single (narrow) link or connection. SAS expanders connect to each other using wide links.

SAS port expanders are non-blocking switch-like devices that are used to connect SAS (or SATA) initiators to SAS (or SATA) targets as shown in Figure 1. These expanders vary in the number of SAS ports provided, enclosure services supported and the number of SAS devices permitted to be attached to the expander.

SAS links use "10b/8b" encoding to transmit data over the physical link, delivering a powerful error correction capability to ensure data integrity.

SAS fanout expanders are starting to be seen in the market. Fanout expanders connect to edge (port) expanders to create SAS domains that are much larger than those that can be built from edge expanders alone. Fanout expanders enable new and exciting devices like SAS switches that can be used to build SAS Storage Area Networks (SANs) not unlike Fibre Channel SANs.

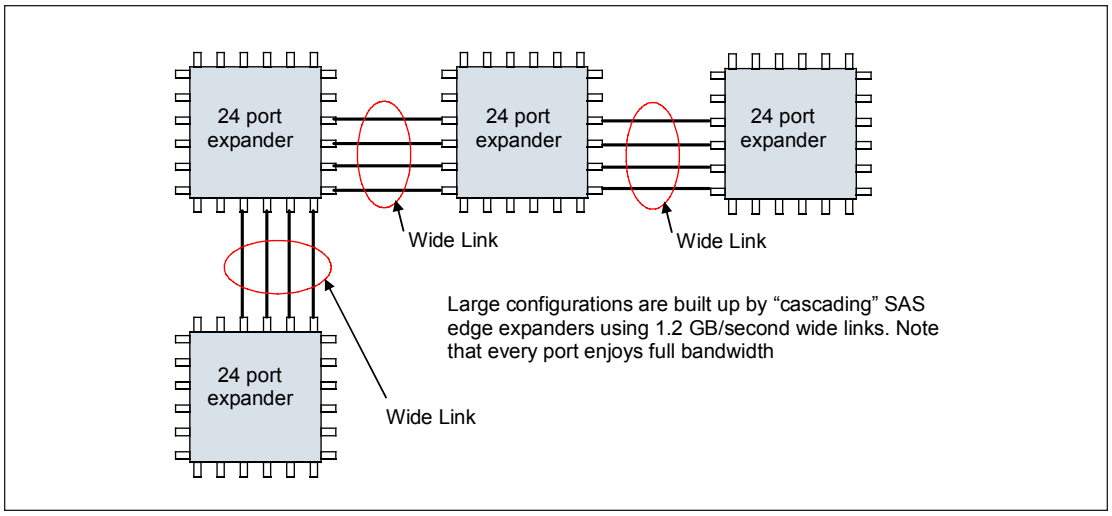

Figure 1. SAS port expanders interconnected using wide links

SAS port expanders can have "inside the enclosure" and "outside the enclosure" connections to other SAS expanders and to disk drives. ULTAMUS RAID products use inside the enclosure SAS expanders connected to a RAID controller using 4x wide links and disk drives. The RAID controller (or active/active controller pair) uses outside the enclosure SAS expanders to connect to capacity expansion chassis equipped with a SAS port expander and disk drives as shown in Figure 2.

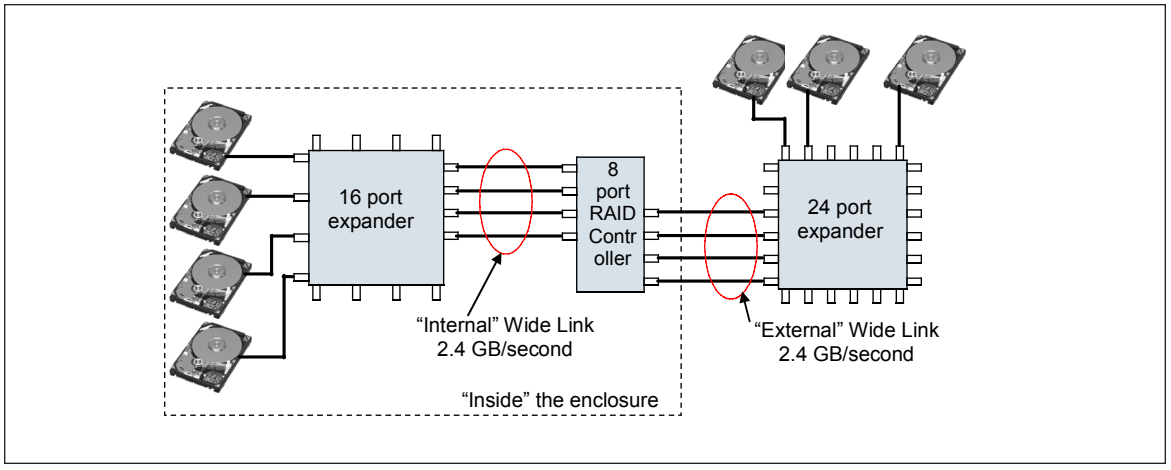

Figure 2. RAID controller interconnects to "inside the enclosure" and "outside the enclosure" SAS port expanders and disk drives

Comparing Storage Interfaces

SAS joins SATA II as the state-of-the-art in storage interfaces. SAS joins Fibre Channel as the premier storage interface technology. Figure 3 below describes the differences among the popular storage interfaces.

	ATA/IDE	SATA II	SCSI	SAS	FC
Number of devices	2	1	16	16,384	16 million
Maximum distance	0.5 meters	1 meter	3 - 25 meters	10 meters	10 km
Cable type	Copper	Copper	Copper	Copper	Copper or fiber optic
Interface type	Parallel	Serial	Parallel	Serial	Serial
Disk drive form factor	3.5 in	3.5 in	3.5 in	2.5/3.5 in	3.5 in
Transfer speeds (MB/second)	Up to 133	300	Up to 320	300/600	100, 200, 400
Random I/O rates (IOPS)	Up to ~50	Up to ~130	Up to ~140	Up to ~210	Up to ~210
Drive rotational speed	Up to 7,200	Up to 10,000	10,000/15,000	10,000/15,000	10,000/15,000
Fault isolation architecture	N/A	Point-to-point	Shared bus	Point-to-point	Looped
Device compatibility	ATA/IDE only	SATA only	SCSI only	SAS & SATA	FC only
High reliability features	N/A	Port selector	N/A	Dual active porting; multi-initiator	Dual active porting; multi-initiator

Figure 3. Popular storage interfaces

Figure 4 illustrates the differences in connectivity offered by the popular storage interfaces.

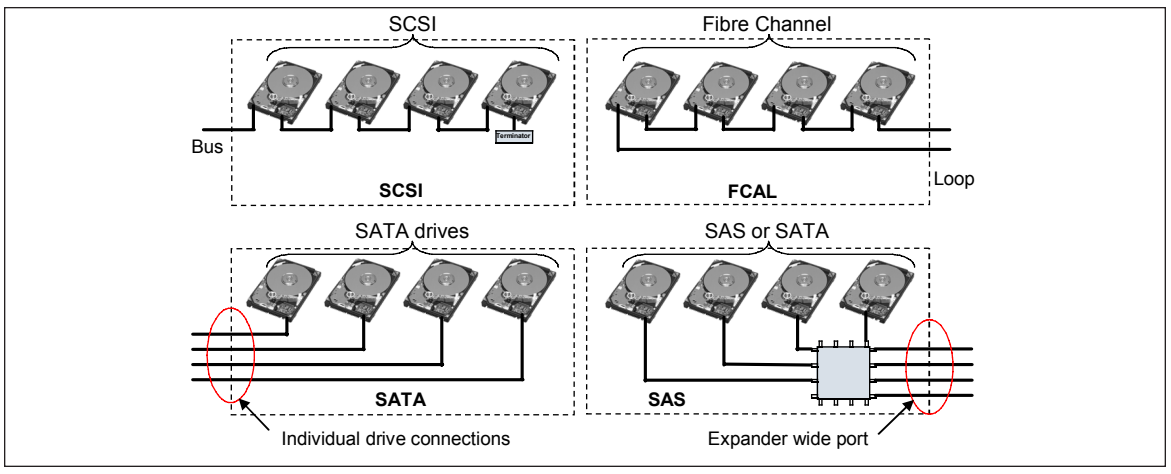

Figure 4. Differences among storage interfaces

SAS and SATA II Disk Drives

To further address performance, SAS drives use 10,000 - 15,000 RPM rotational speeds, which are well-suited for handling random database type transactions such as email, databases, data warehousing and OLTP (Online Transaction Processing).

SAS enables a direct, unshared 3 Gb/sec link from each drive to the RAID controller as opposed to SCSI that delivers a maximum bandwidth of 320 MB/sec that must be shared across every drive on the same bus.

There are currently two SAS disk drive form factors: the 3.5-inch form factor and the emerging 2.5-inch form factor. 3.5-inch form factor drives offer significant advantages over the emerging 2.5-inch disk drives, including nearly 30% higher performance (a total of five 2.5-inch drives are needed to match the performance of four 3.5-inch 15,000 RPM disk drives).

A 3.5-inch SAS disk drive can store as much as 300GB. A 2.5-inch drive can store up to 73GB.

SATA II disk drives are well-suited for less performance intensive applications that require large storage capacities at the lowest cost per GB. Current SATA II disk capacities range up to 500 GB each with 750 GB drives nearing production (end of 2006).

Figure 5 describes disk drive technology suitability for typical applications.

	High-End Applications (OLTP, HA, random write & retrieval)	**Mid-Range Applications (OLTP, HA, random write & retrieval)**	**Low-End Applications (Nearline, disk-based backup, sequential reads & writes)**	**Available Drive Capacities**	**Typical Applications & Characteristics**
Fibre Channel (FC)	YES (Suitable for mission-critical data in large enterprises, e.g. SAN)	NO (Relatively costly unless FC infrastructure is already in place)	NO (Too costly)	18 - 300 GB	Transaction data (DBs) • Maximum IOPS for OLTP • Calculation intensive files • Small block random reads/writes • Reservation systems • Billing systems
Serial Attached SCSI (SAS)	YES (Suitable for mission-critical data in large enterprises, e.g. SAN)	YES (Suitable for mission-critical data in SMB and mid-range enterprises)	NO (Not as cost-effective as SATA for high capacity applications)	36 - 300 GB	Transaction data (DBs) • Maximum IOPS for OLTP • Calculation intensive files • Small block random reads/writes
Serial ATA (SATA)	NO (Not suitable for OLTP/database applications)	YES (Suitable for mission-critical data in SMB and mid-range enterprises with RAID 6)	YES (Highly cost-effective for high capacity applications)	80 - 750 GB (1 TB soon)	Throughput data • High MB/sec, large data intensive files • Reference data • Fixed content • Archival data for secondary/nearline applications • Image storage • Backup to disk

Figure 5. Disk drive suitability for typical applications.

Combining SAS and SATA Disk Drives in Tiered Storage Applications

The SAS architecture enables combining the high I/O performance of SAS drives with the low cost per GB of SATA II drives in tiered storage applications. SAS and SATA disk drives can both reside in a storage enclosure with separate RAID sets built from SAS drives and SATA drives.

High IOPS performance applications can store their data on the RAID sets built using SAS drives and then transfer the data to RAID sets built using slower, less costly SATA II drives as the performance requirement decreases.

Application-specific tiered storage is another capability enabled by the judicious merging of SAS and SATA II technologies. For example, Exchange data may be maintained on RAID sets with the highest performance available while archived data may be stored on lower performance RAID sets.

Technology Summary

- SAS is a standardized, next generation I/O interface that provides an evolutionary growth path from the existing SCSI bus
- SAS is a new storage architecture, not just a drive interface or an adapter (HBA, RAID controller, etc.) interface
- SAS takes advantage of the latest advances in high-speed serial transmission and switching
- SAS provides fully bi-directional data transfer rates of 300 MB/sec with 600 MB/sec on the horizon —to each drive!
- SAS drives are dual-ported for failover
- SAS drives support >2 TB LUNs and NCQ
- SAS cables are thinner than SCSI cables
- SAS connectors are less bulky than the familiar SCSI connectors
- SAS allows the use of longer cables than SCSI
- SAS crosstalk is less likely because there are fewer conductors in the cables
- Hardware for serial interfaces is less costly than the hardware for equivalent parallel interfaces
- SAS architectures are based on non-blocking switch technology
- All connections use switched point-to-point links
- SAS switches are called "Expanders"
- SAS supports up to a maximum of 16,384 devices
- SAS supports full duplex operation
- SAS and Serial ATA (SATA) drives can be mixed and matched on the same RAID controller and mid-plane (although this is not advised in the same RAID set)
- "4-Lane Wide Links" (2.4 GB/sec) are used to connect expanders
- ULTAMUS RAID products are built using the SAS architecture.

Business Summary

ULTAMUS RAID arms storage managers with the tools they need to implement storage strategies and build IT infrastructures that align with both today's and tomorrow's business and data management needs.

ULTAMUS RAID harnesses today's most advanced storage technologies to create a SMB/SME RAID solution that:

- Delivers a rich feature set
- Offers differentiated technology at a commodity price point
- Provides investment protection for storage managers
- Creates a new class of end-to-end performance in the entry-space market
- Provides a complete tiered storage strategy (ULTAMUS RAID, REO™, NEO™, ARCvault™)

Technology	Benefits	Overland Leadership
RAID 6 Hardware Accelerated	• Higher Availability	• Better for mission-critical application availability. • Overland's hardware accelerated RAID 6 allows storage to remain online in the event of two drives failing simultaneously.
SAS Drives	• Higher Performance • Lower Costs • Investment Protection • Higher Availability	• Ideal for email, web serving, database and other performance-intensive applications • provides the performance and reliability of Fibre Channel drives at 30% lower cost • SAS and SATA II drives can be combined in the same RAID chassis to support applications with vastly different capacity and performance requirements.
Enterprise SATA II Drives	• Higher Availability • Higher Performance • Better Price/TB • Investment Protection	• Better performance for D2D backup and media • Built for 7 x 24 operating environments • 1.4 million hour MTBF at 100% duty cycle—more than double the reliability of desktop SATA II drives • Directed Offline Scan monitors every write and performance diagnostic tests when drives are idle, enhancing data availability • Perpendicular Recording reduces the number of moving parts, for higher availability
SAS expansion & scalability	• Grows with your application requirements	• Effortlessly scales up to 60 drives in 12-drive/2U increments • No performance penalty
4Gb Fibre Channel	• Higher Performance • Investment Protection	• Better performance for all applications. • 4Gb FC improves IOP intensive applications with reduced latency and improves streaming applications with higher data rates
Snapshots	• Disaster Recovery • Data Protection	• Snapshot technology offers the best value in data protetion by creating low-impact copies of data that support point-in-time recovery
Pricing	• Low Acquisition Cost	• Excellent price/TB for RAID storage • Integrated RAID ASIC technology reduces controller costs and an optimized supply chain reduces platform costs.

Save

Worldwide Headquarters
4820 Overland Avenue
San Diego, CA 92123 USA
TEL 1-800-729-8725
TEL 1-858-571-5555
FAX 1-858-571-3664
EMAIL sales@overlandstorage.com

United Kingdom (EMEA Office)
Overland House
Ashville Way
Wokingham, Berkshire
RG41 2PL England
TEL +44 (0) 118-9898000
FAX +44 (0) 118-9891897
EMAIL europe@overlandstorage.com

France Office
Overland Storage
126 rue Gallieni
92643 Boulogne Cedex France
TEL +33 (0) 1 55 19 23 93
FAX +33 (0) 1 55 19 25 02
EMAIL europe@overlandstorage.com

Germany Office
Humboldtstr. 12
85609 Dornach Germany
TEL +49-89-94490-214
FAX +49-89-94490-414
EMAIL europe@overlandstorage.com

Asia Pacific Office
Level 44, Suntec Tower 3
8 Temasek Boulevard
Singapore, 038988
TEL +65.6866.3847
FAX +65.6866.3838
EMAIL asia@overlandstorage.com

www.overlandstorage.com

Overland STORAGE®
Simply Protected