adaptec

Serial Attached SCSI: Meeting the Growing Needs of Enterprise Storage

Introduction

Today's predominant server storage interface standards (SCSI for enterprise environments and ATA for cost-sensitive desktop applications) rely on parallel transmission of data streams for device-level attachment. However, parallel technology's signal skew and crosstalk, signal termination restrictions, cable and connector reflections, and device addressability, stand as barriers to throughput performance as servers are pushed to meet advancing system and application capabilities and requirements.

SAS and SATA

Serial technology – SAS (Serial Attached SCSI) and SATA (Serial ATA) – was introduced to overcome these barriers, delivering greater speed, reliability and scalability. While SATA is designed for desktops, making it a good choice in storage environments requiring configuration simplicity or optimal cost/capacity, SAS delivers the high performance, scalability and reliability required for bandwidth-hungry mainstream servers and enterprise storage. SAS lends itself to the high-frequency, immediate random data access required for transactional data applications such as online purchases and bank transactions.

Only SAS combines the proven reliability and functionality of SCSI with the performance and design power of serial technology by delivering the follow features and capabilities:

Performance – First-generation performance of up to 3Gb/s (300 megabytes/second), wide ports for aggregated bandwidth, full duplex, port aggregation, advanced command queuing and the rich command feature set of SCSI

Scalability – Broad address range to physical devices, long cables with small connectors, and connectivity to external storage systems

Reliability and Availability – Point-to-point connections, multi-initiator capability through expanders for simultaneous access, dual active port support, and redundant paths to targets

Flexibility – Physical and software compatibility with SAS and SATA drives, and backward compatibility with SCSI software and middleware

Performance

First-generation SAS will deliver throughput of 3Gb/s and succeeding generations up to 12Gb/s to keep pace with technology and application advances. In addition, SAS provides multiple point-to-point connections that enable fault-tolerant designs.

The full-duplex, point-to-point nature of SAS enables simultaneously active connections among multiple initiators and high-performance SAS targets. Devices can transfer data in both directions at once to effectively double the useable bandwidth of the link rate. Narrow ports allow for a single serial link, while wide ports support multiple links, allowing the aggregation of eight SAS or SATA targets to increase total available bandwidth to 24Gb/s (Figure 1).

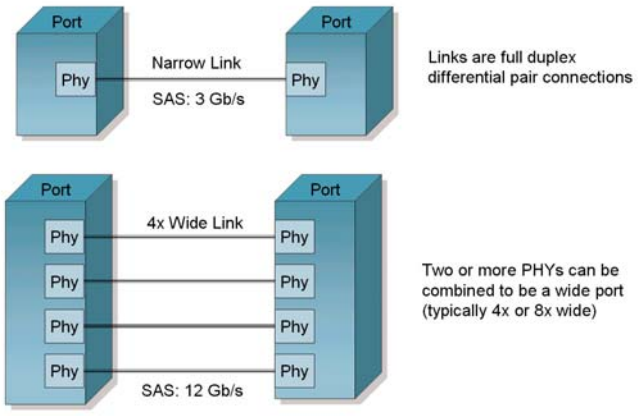

Figure 1: Wide SAS Ports

By combining multiple PHYs together, Wide Ports are created that can support the significant bandwidth requirements of large SAS topologies.

Like SCSI, SAS includes advanced command queuing with 256 queue levels, providing unique intelligent data handling features such as head of queue and out of order queuing. These queuing features are critical to enterprise applications, allowing a system to reorder and reprioritize commands within the interface.

Large high-performance topologies are also made more practical by advances in cabling for SAS. SAS replaces the wide 68-pin ribbon cables for signals and separate power cable with a single thin 4-wire SAS cable of up to 8 meters long, a compact design that improves airflow inside enclosures and simplifies hot-plug connections.

Flexibility

A key advantage of SAS is that its backplane design and protocol interface allows the use of both SAS and SATA drives in the same system. Though each drive type is typically used in different applications, most enterprise users have needs for both drive kinds. The ability to mix and match these drives is a powerful benefit for designers and users.

SATA drives will be primarily designed for cost-effective bulk storage. To achieve economies of scale, SATA drives feature lower spindle speeds (typically 7,200 rpm), lower mean-time-between-failure rates, and lower cost. Consequently, they tend to be applied where transaction rates are low and data availability is not critical.

SAS drives, on the other hand, are built for high-performance, high-availability use. SAS drives will operate at higher spindle speeds (10,000 to 15,000 rpm) with compensation for rotational vibration to assure data integrity, and are built for higher reliability. SAS drives will be used in environments where data volumes are high and data availability is essential.

Since SATA connector signals are a subset of SAS signals, SATA devices are fully compatible with SAS controllers—the SATA Tunneled Protocol (STP) included within SAS passes through SATA commands to SATA drives. And since the SAS connector itself is designed as a single uniform backplane, designing a system with both drive types is simple. This compatibility reduces the cost and complexity of storage designs, which increases the variety of design options.

This compatibility also allows system builders to design hybrid storage systems using common connectors and cabling. Installing or upgrading either SATA or SAS drives in the same system is simply a matter of replacing one drive type with the other—the SAS backplane connectors receive both SAS and SATA devices. However, since SATA backplanes connect only to SATA devices, backplanes should use SAS connectors to provide the greatest system design flexibility.

SAS's backward compatibility with previous generation SCSI software and middleware also makes it easy to incorporate legacy components—hosts and drives—into evolving SAS topologies, eliminating new training or integration costs and the need for modifications to legacy software.

Scalability

The scalability of parallel buses is limited because they share connection paths, and adding more buses with multiple initiators does little to extend this limited sharing ability. SAS uses expander hardware as a switch to simplify configuration of large external storage systems that can be easily scaled with minimal latency while preserving bandwidth for increased workloads. This expander hardware enables highly flexible storage topologies of up to 16,256 mixed SAS and SATA drives. SAS expander hardware functions as a switch to simplify configuration of large systems that can be scaled with minimal latency while preserving bandwidth for increased workloads.

A fan out expander, for example, can be connected up to 128 devices, including initiators, SAS and SATA drives, Edge expanders and other Fan Out expanders in either narrow or wide formats. These additional Fan Out and Edge expanders can in turn be linked to other hosts and drives, providing additional connection nodes. The SCSI Management Protocol (SMP) within SAS manages the point-to-point connections in the topology (Figure 2).

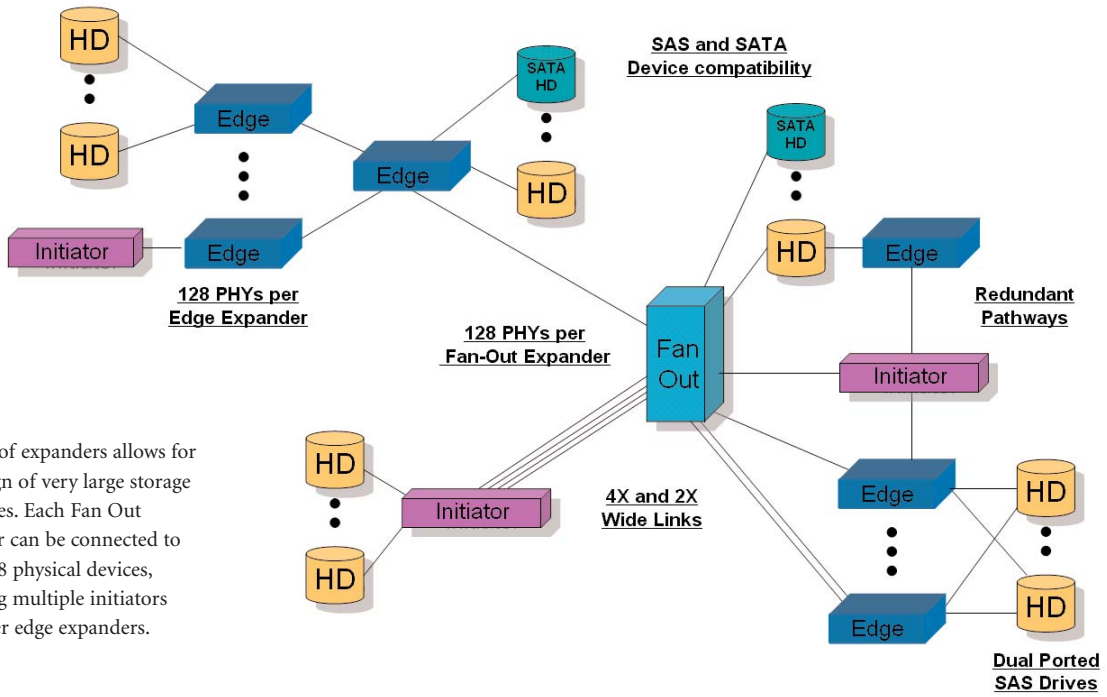

The use of expanders allows for the design of very large storage topologies. Each Fan Out expander can be connected to up to 128 physical devices, including multiple initiators and other edge expanders.

Figure 2: SAS Topology with Expanders

Reliability

Multiple initiators have long been used in enterprise computing to provide disk drive access to multiple hosts, host bus adapters or both and ensure continuous data access in case one fails. However, using multiple initiators in parallel technology configurations leaves single points of failure that can block access to a device. With SAS, devices featuring dual porting can be used to build high-availability systems with no single points of failure.

Another way to increase fault-tolerance with SAS is by using expanders to connect multiple devices to multiple initiators, which maintain concurrent operation to many devices. Commands can be sent down one link and data returned on another in a separate connection to increase fault tolerance.

The SAS expander feature, in combination with dual-port SAS drives and SATA drives with 2-port adapters, makes it easy to design redundant systems for maximum fault-tolerance. The availability of 2.5-inch dual-ported SAS drives in addition to the standard 3.5-inch drive enables fully fault-tolerant designs in new technically challenging, higher compute density applications. This highly scalable and reliable connection scheme enables enterprise-level topologies that support multi-node clustering for automatic failover availability or load balancing—vital for mission critical applications.

Conclusion

The advantages of SAS are compelling. Before the advent of serial technology, system designers had to carefully balance the trade-offs between cost-per-gigabyte of storage and cost-per-I/O with drive reliability, scalability, and availability—even as server performance and application requirements grew. Plus, parallel interfaces forced users to implement separate systems for each drive type and bear the costs of additional servers, drives, and support.

Now, as a uniform interface platform for both SAS and SATA emerges, users will be able to quickly and easily create storage systems where all these elements can co-exist without additional system and support costs, and the different drive types can be hot-plugged in and out as the enterprise's storage needs change.

Serial Attached SCSI represents both the next generation in the ongoing evolution of SCSI performance, and a significant new advance in the architecture of I/O interfaces. When it enters the market in 2004, SAS will begin to change the storage landscape by enabling the easy custom design of highly flexible and fault-tolerant storage topologies.
For more information, please visit;
http://www.adaptec.com
http://www.scsita.org

adaptec®

P/N 666572-011 Printed in USA 3/04