Serial Attached SCSI technology

technology brief

Introduction

The parallel SCSI bus has been a workhorse in the data storage industry for over 20 years. The success of SCSI as an I/O interface can be attributed to its performance, intelligence, and backward compatibility. However, after years of steady performance improvements, parallel SCSI has reached technological hurdles that may be too difficult to leap. To get to the next level of performance, the industry is turning to Serial Attached SCSI (SAS).

This technology brief begins with a description of parallel SCSI technology and the technological challenges it faces to get to the next level of performance beyond Ultra320 SCSI. Subsequent sections describe SAS technology and performance expectations, SAS and Serial ATA (SATA) interoperability, and what makes SAS the ideal solution for mission-critical enterprise storage applications.

Parallel SCSI technology

The SCSI specification was developed to provide a common interface that could be used across all peripheral platforms and system applications. The SCSI interface addresses a wider range of applications, such as Redundant Array of Independent Disks (RAID) storage, and has a broader command set than the parallel ATA interface. The SCSI system contains the SCSI controller (initiator), the SCSI bus (cable or backplane), and one or more target devices. The SCSI controller may be built into the motherboard or housed on a SCSI host bus adapter (HBA) card in a PCI or PCI-X slot. Both configurations are shown in Figure 1.

Figure 1. SCSI components

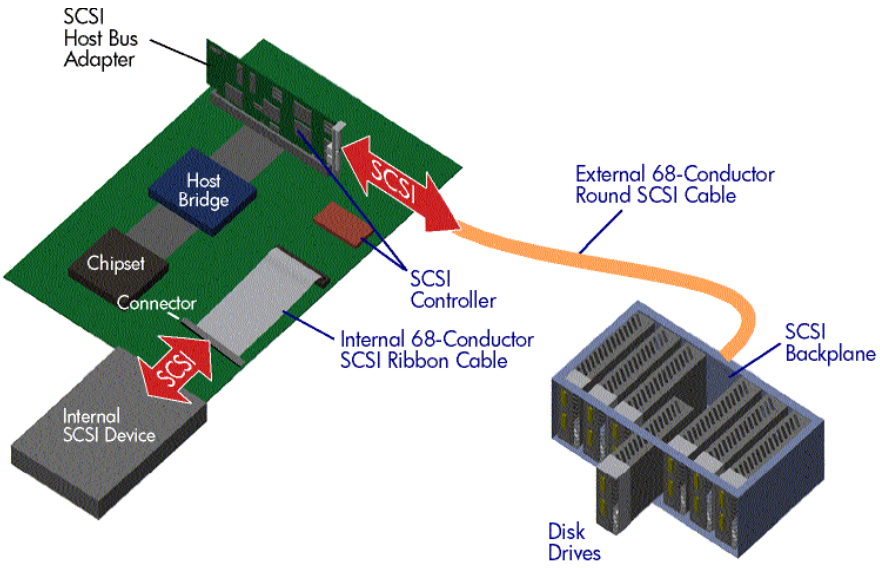

SCSI cables can connect up to 16 devices, including the SCSI controller. SCSI cables consist of 34 twisted pairs of multi-stranded flexible copper wires for a total of 68 conductors. SCSI devices inside the server are connected to the SCSI controller by a 68-pin ribbon cable. The ribbon cable has a connector at each end and one or more connectors along its length. External SCSI devices are attached to the SCSI HBA by a round 68-pin cable. Two sets of terminators, one at each end of the SCSI bus, prevent signal reflections within the cables.

Since 1981, there have been seven generations of the SCSI protocol. Each new generation has doubled the performance of the previous one (Figure 2). SCSI performance has ranged from an 8-bit, single-ended interface transferring data up to 4 MB/s (SCSI-1) to the latest 16-bit, low-voltage differential interface transferring data at 320 MB/s per channel (Ultra320 SCSI).

Figure 2. Bandwidths of seven generations of SCSI

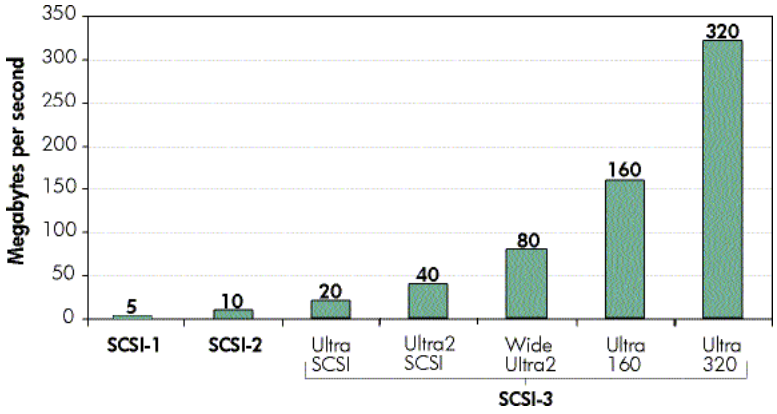

The development of Ultra320 SCSI was very challenging because it went beyond simply doubling the clock frequency of Ultra160 SCSI. In fact, the higher performance of Ultra320 SCSI was made possible by the implementation of several performance enhancements, including:

- Read and write data streaming
- Quick arbitration and selection
- Flow control

Ultra320 SCSI also introduced two new technologies—pre-compensation and training—to minimize the negative impact of signal skew (slight signal delays from one wire trace to the next) and attenuation on signal integrity. Because the development of Ultra320 SCSI presented significant electrical engineering challenges, it is generally believed that Ultra640 SCSI cannot be deployed reliably without using new and expensive technologies. Therefore, parallel SCSI architecture is being replaced with a serial I/O architecture that has the performance headroom to meet enterprise storage needs for years to come.

Serial Attached SCSI technology

Serial Attached SCSI is a point-to-point architecture in which all storage devices connect directly to a SAS port rather than sharing a common bus as traditional SCSI devices do. Point-to-point links increase data throughput and improve the ability to locate and fix disk failures. More importantly, the SAS architecture solves the clock skew and signal degradation problems of parallel SCSI at higher signaling rates. SAS inherits its command set from parallel SCSI, frame formats from Fibre Channel, and physical characteristics from Serial ATA.

HP (Compaq), IBM, LSI Logic, Maxtor, and Seagate founded the Serial Attached SCSI Working Group in 2001. SAS is being developed in an open standards (ANSI T10) environment to ensure that the technology is an industry standard.

The SAS and SATA technologies have several common features, including low-voltage differential signaling, 8b/10b encoding, and full duplex communication.[1] The Serial Attached SCSI standards committee designed the SAS infrastructure to be compatible with SATA drives, allowing the coexistence of both storage technologies in the same system and opening the door to SATA scalability. Because the SAS architecture features a proven SCSI command set, advanced command queuing, and advanced verification/error correction, SAS is the ideal solution for mission-critical enterprise storage applications.

Performance

The speed of the first-generation SAS link is 3.0 gigabits/second (Gb/s). The speed of the second-generation SAS link will be 6.0 Gb/s. SAS links are full duplex; they send and receive information simultaneously, thereby reducing a major source of latency. The SAS interface allows for combining multiple links to create 2x, 3x, or 4x connections for scalable bandwidth.

SAS/SATA interoperability

The SAS architecture enables system designs that deploy both SAS and SATA devices, a breakthrough for enterprise customers. This capability provides a broad range of storage solutions (Figure 3) that give IT managers the flexibility to choose storage devices based on reliability, performance, and cost.

SAS supports three protocols to handle communications with various devices: Serial Management Protocol (SMP), Serial SCSI Protocol (SSP), and SATA Tunneling Protocol (STP). SMP is used to manage the point-to-point topology of expanders and enclosure services. SSP allows the controller to communicate with SAS devices and existing SCSI software. STP, on the other hand, allows SAS controllers to communicate with SATA devices through expanders. After power up or restart, the SAS domain goes through an initialization sequence to determine the proper protocol to communicate with the various types of devices.

Figure 3. The SAS architecture gives IT managers the flexibility to chose solutions.

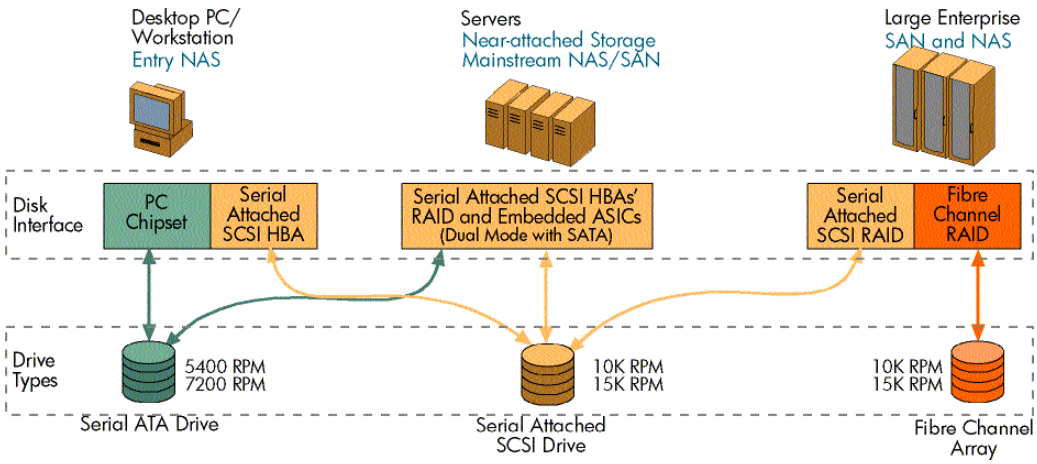

[1] Please refer to the technology brief "Serial ATA technology", document number TC050303TB, at http://www.hp.com/servers/technology for more information about these features.

Connectors and cabling

SAS uses a 7-pin connector (four signal lines and three ground lines) and a small-diameter cable up to 8 meters long. SAS and SATA devices will share the same physical device connector, except for an extension within the notch on the SAS connector (Figure 4). This extension will allow a SAS connector to accept SATA device connections, but it will not allow a SATA connector to accept SAS device connections. The thinner SAS and SATA cables improve both airflow and routing inside the server chassis.

Figure 4. SAS and SATA device connectors and SAS backplane connector

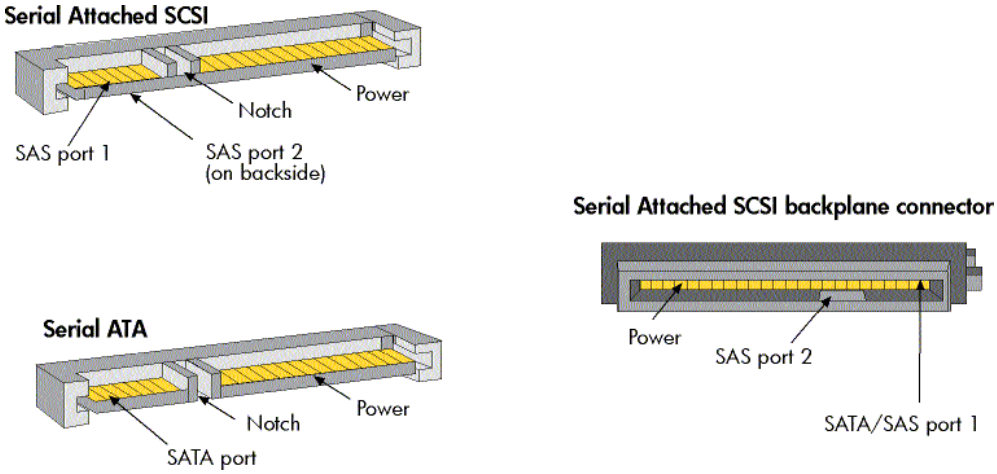

Expanders

SAS RAID controllers can be deployed as either internal device links or external links (Figure 5). All SAS devices can have one or more ports, with each port configured as narrow (single link) or wide (multiple links).

Low-cost, high-speed switches called expanders enable SAS controllers to connect to a greater number of devices than parallel SCSI. There are two types of expanders: edge expanders and fan-out expanders. Edge expanders can be grouped into edge expander sets, with each set connecting up to 128 devices (SAS and/or SATA drives, or expanders). Only one fan-out expander is allowed, which may connect up to 128 edge expander sets (or drives).

Figure 5. Serial attached SCSI components

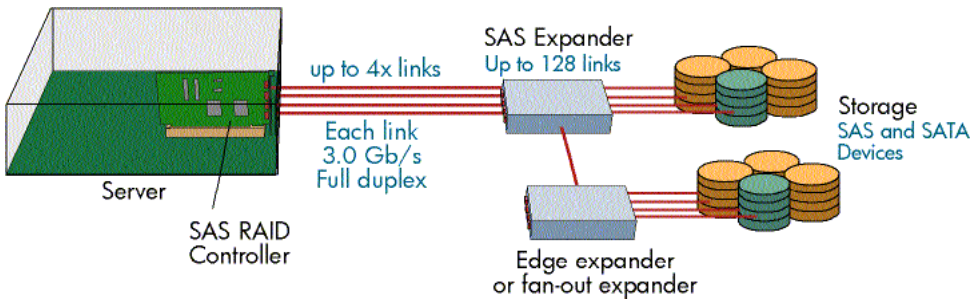

Administration

Unlike parallel SCSI devices, SAS devices do not need to be assigned an ID. All SAS devices are automatically assigned unique worldwide names (SAS addresses), making it easier to identify initiator devices, target devices, and expanders.

Greater scalability

Serial Attached SCSI enables highly scalable topologies—internal, external, or a combination of both—to give manufacturers and customers the flexibility to design and deploy a range of solutions. The Serial ATA Tunneling Protocol (STP) enables SAS HBAs to communicate with SATA devices through expanders and, therefore, is key to SATA scalability in the SAS domain.

Internal

Figure 6 shows a topology that can be used for internal RAID systems. The SAS RAID controller can support from two to eight internal HDDs, which can be either SAS or SATA devices. The small SAS connector will allow connections to 3.5-inch or 2.5-inch internal HDDs, enabling redundant storage configurations in dense server form factors like blade servers.

Figure 6. Topology for internal storage applications

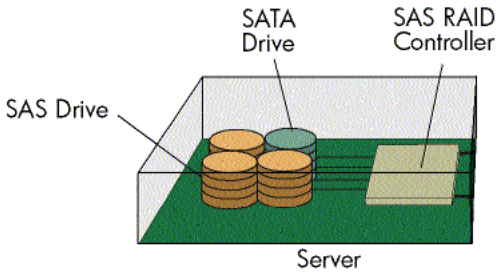

Internal/external with JBODs

Figure 7 shows a topology that can be used for internal RAID systems. The SAS interface will enable customers to plug either SAS or SATA drives into one backplane. This will allow the use of enterprise devices and desktop devices in the same server or networked storage subsystem.

Figure 7. Topology for internal and external storage applications

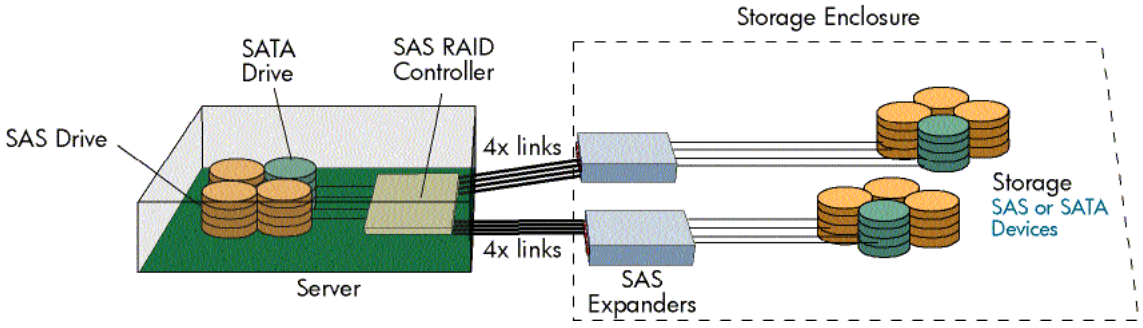

Multi-node clusters

The highly scalable architecture enables topologies that support multi-node clustering for high availability (failover) or load balancing (Figure 8).

Figure 8. Topology for multi-node cluster applications

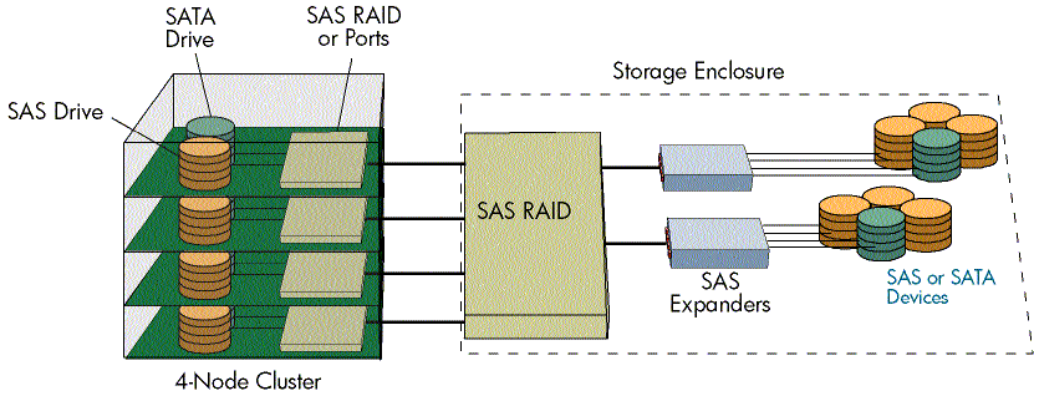

SAS migration

SAS compatibility with legacy parallel SCSI software will provide a migration path for existing SCSI users and allow SAS to replace SCSI as the dominant interface for DAS, NAS, and SAN systems. SAS is not a fabric topology, so it is intended to complement, not replace, Fibre Channel in the SAN market. SAS compatibility with SATA will enable the deployment of multi-function systems that can provide high performance and high availability as well as low-cost-per-megabyte solutions.
The current SAS data transfer rate is 3.0 Gb/s, with a rate of 6.0 Gb/s planned for the next phase of development. At 6.0 Gb/s, wide links (2x, 3x, and 4x) will enable aggregate data transfers up to 24 Gb/s.

Figure 9. Planned SAS performance enhancements

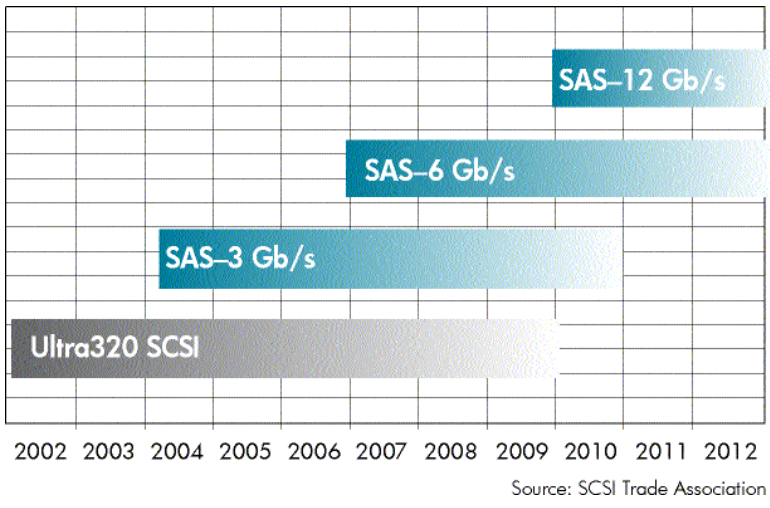

Conclusion

The benefits of SATA and SAS interoperability will extend from the desktop to the data center. These benefits include investment protection in SCSI software, higher HDD performance, smaller form factors, and greater device addressability. Because the SATA and SAS architectures use the same physical device connector, customers have the flexibility to choose solutions that use both SAS and SATA devices. This flexibility is crucial for the adaptive enterprise.

For more information

For additional information, refer to the resources detailed below.

Resource description	Web address
SCSI Trade Association website	http://www.scsita.org
Serial ATA technology	http://www.hp.com/servers/technology

TC050304TB, 03/2005

invent