Serial ATA technology

Technology brief, 4th edition

Introduction

Serial ATA (SATA) has evolved beyond a serial replacement for parallel ATA in desktop computers. SATA is now an alternate, lower cost solution for non mission-critical enterprise storage applications. SATA technology can provide adequate scalability and hot plug capability needed for server and network storage devices at the entry to mid-range levels.

The Serial ATA Revision 3.0 and 3.1 specifications increase the maximum data transfer rate from 3 gigabits per second (Gb/s) to 6 Gb/s and add new features. The higher speed provides more headroom for solid-state drives (SSDs) and for using multiple disk drives through expanders. New features such as native command queuing (NCQ) and NCQ management improve performance. NCQ lets a drive rearrange the order of requests from the host to maximize throughput, and NCQ management lets the host manage and process outstanding NCQ commands.

This technology brief describes key technologies in the SATA specification, its implications for SATA storage devices, the interoperability of SATA and Serial Attached SCSI (SAS) devices, and examples of internal and external SATA topologies. Table 1 in the Appendix contains terms related to SATA technology.

SATA devices and interoperability with SAS devices

SATA devices include initiators (controllers or hosts) and targets (hard drives, optical disk drives, and solid-state drives). The initiator attaches to one or more targets or expanders. You can increase the number of SATA targets attached to an enclosure by adding a SATA-compatible SAS expander (low-cost high-speed switch). The SATA Tunneling Protocol (STP) enables SAS controllers to communicate with SATA devices, allowing you to deploy both SAS and SATA drives.

Figure 1: This is an example of SATA/SAS connectivity.

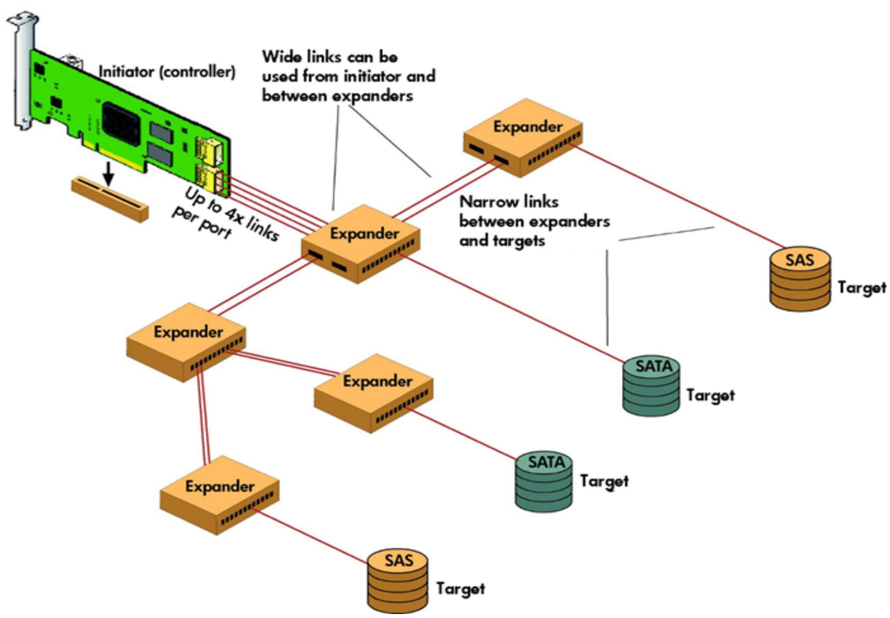

Initiators

An HP initiator is an embedded SATA or SAS controller or a PCIe SAS controller. Our embedded controllers only support internal devices. We have Host Bus Adapter (HBA) and Smart Array options that support internal devices, external devices, or both. We support SATA optical drives only on an embedded SATA controller.

Expanders

Expanders connect SAS initiators with SAS and SATA targets. They receive commands and data in one port and route them to another port based on the address of the target. Expanders allow multiple drives to connect to a single initiator.

Targets and their recommended uses

SATA targets can be hard disk drives (HDD) or solid state drives (SSD), each with a single port, or optical disk drives. We offer SATA HDD and SSD in 2.5-inch small form factor (SFF) and 3.5-inch large form factor (LFF), in both hot plug and non-hot plug versions. Optical disk drives are available in half-height and slimline (9.5mm) form factors.

We categorize server disk drives in three levels based on performance, reliability, and capacity: Entry, Midline, and Enterprise. We offer SATA drives in Entry and Midline levels.

We categorize Enterprise SSDs as Value, Mainstream, and Performance. Today's HP SATA SSDs are available as Enterprise Mainstream storage devices.

Entry SATA disk drives

Entry-level SATA drives offer the lowest unit cost for server solutions in non mission-critical, low workload environments. This includes boot drives and large-volume, low-cost implementations such as file-and-print and video. Entry-level drives also provide economical storage for static web pages, Domain Name Systems (DNS), or firewalls.

Midline SATA disk drives

Economical HP SATA Midline drives offer reliability and performance for high capacity implementations that are not mission-critical, for example external storage, backup, and archives.

Enterprise mainstream SATA SSDs

SSDs achieve high random-read performance by eliminating the seek time and rotational latency of traditional disk drives. HP SSDs are also as reliable as current HP Enterprise disk drives. More important for some applications, our server SSDs deliver this level of reliability under conditions that are unsuitable for traditional disk drives, including environments exposed to greater shock and vibration.

Our SATA SSDs use multi-level cell (MLC) NAND flash technology and are suitable for environments with constrained workloads and less than 100% duty cycle. The NAND flash memory used in SSDs has a limited life cycle for writes. Advanced techniques such as wear leveling, the TRIM command (dependant on controller, driver, and OS support), and capacity over-provisioning help overcome this limit.

SATA SSDs are best suited for read-intensive, high input/output operations per second (IOPS) environments (where writes are 30% or less of the total IO load) and for applications that do not require large capacity RAID configurations.

SATA technology

SATA architecture uses point-to-point connections that allow each drive to communicate with a host without waiting. SATA uses low voltage differential (LVD) signaling to address the electrical interference and signal integrity issues of parallel ATA. The LVD signaling scheme uses two pairs of data lines to transmit low-voltage signals between the host port and the SATA device. The voltage potential between each wire pair represents a data bit. SATA technology lets engineers shrink connectors, cabling, and device form factors while lowering power consumption.

Serial data transmission

Serial communication requires a serializer/deserializer (SerDes) to convert parallel data into a serial bit stream and vice versa (Figure 2). A SerDes contains these components:

- A parallel digital interface
- First-In-First-Out (FIFO) caches
- 8 bit/10 bit (8b/10b) encoder and decoder
- A serializer
- A deserializer

Figure 2: The SerDes core integrates 8b/10b coding and decoding logic.

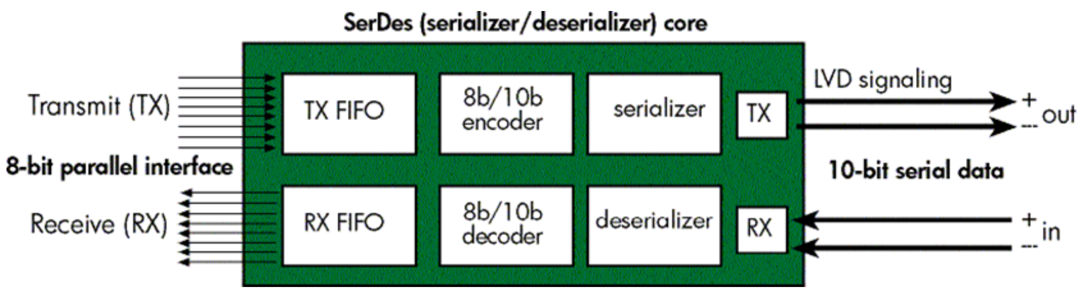

The 8b/10b encoder converts each 8-bit data byte to a 10-bit transmission character. That allows encoding the clocking information into the data stream. It adds about 20% embedded overhead to the data stream, but it eliminates clock skew problems.

SATA devices transmit signals in a single stream across the SATA interface in packets called Frame Information Structures (FIS). SATA uses a half-duplex scheme that transmits data in one direction at a time. One wire pair transmits the FIS, and the other wire pair transmits feedback from the receiving device. Figure 3 shows as serial actions the transmission of the original signal on one pair and the return of status information on the other pair.

Figure 3: Half-duplex LVD signaling transmits data in one direction at a time.

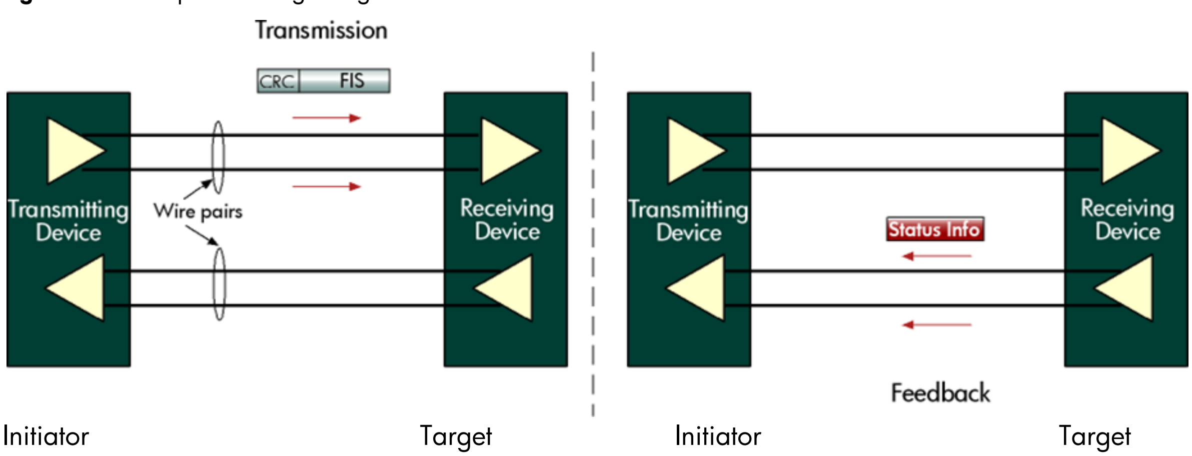

Each FIS includes a Cyclic Redundancy Check (CRC) that can detect single-bit or double-bit errors. The SATA drive reports CRC errors to the host. The host then re-transmits all packet types except data packets. SATA does not support data packet retry because data packets can be very large (up to 8 KB) and the SATA interface would have to buffer them for re-transmission.

Signal integrity

Serial architectures encode (embed) the clock signals in the data stream, eliminating the parallel bus skew problem of aligning data and clock signals. Serial architectures reduce electrical noise because they have fewer data lines to switch simultaneously and the low voltage reduces the effects of capacitance, inductance, and noise. As a result, designers can increase serial signaling rates well beyond rates that are possible with a parallel bus.

Connectors and cabling

SATA uses an L-shaped 7-pin connector (four signal lines and three ground lines) and a small diameter cable up to 1m in length (Figure 4). The thin serial cable minimizes airflow resistance inside an enclosure and eases cable routing. SATA also has a 15-pin, single-row power connector. The power connector provides optional hot-plug capability, so you can swap out a drive without powering down the server. Servers with backplanes provide data and power through the backplane connector.

Figure 4: SATA connections use a 7-pin data cable and 15-pin power cable.

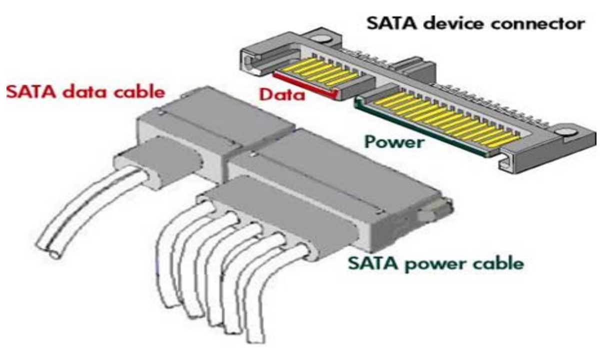

SATA and SAS devices use a similar connector. The SATA device connector has a notch between data and power pins (Figure 5), and the SAS device connector has a bridge between the data and power pins for a secondary data port. The SATA host connector requires this notch, which prevents it from accepting SAS devices. SAS host connectors can accept both SAS and SATA devices.

Figure 5: The SATA device connector is similar to the SAS device connector.

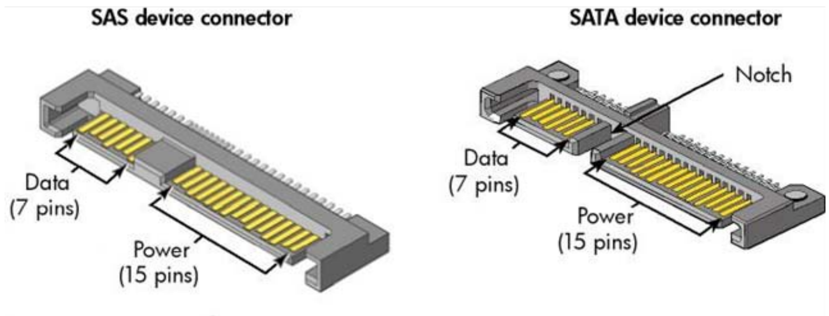

You can connect SATA 6 Gb/s devices using the same cables and connectors as SATA 1.5 Gb/s and SATA 3 Gb/s devices. To maintain signal integrity of a SATA 6 Gb/s device requires high-quality cabling and connectors.

SATA topologies

You can use SATA and SAS drives in the same storage enclosure. SAS architectures can support internal and external configurations. With this broad range of storage solutions, you can choose storage devices based on reliability, performance, and cost.

Internal

Figure 6 is an example of a system incorporating internal SATA drives. Each drive has a point-to-point connection to the controller. This can be a single drive, a JBOD, or a RAID configuration.

Figure 6: Internal drives connect directly to a controller.

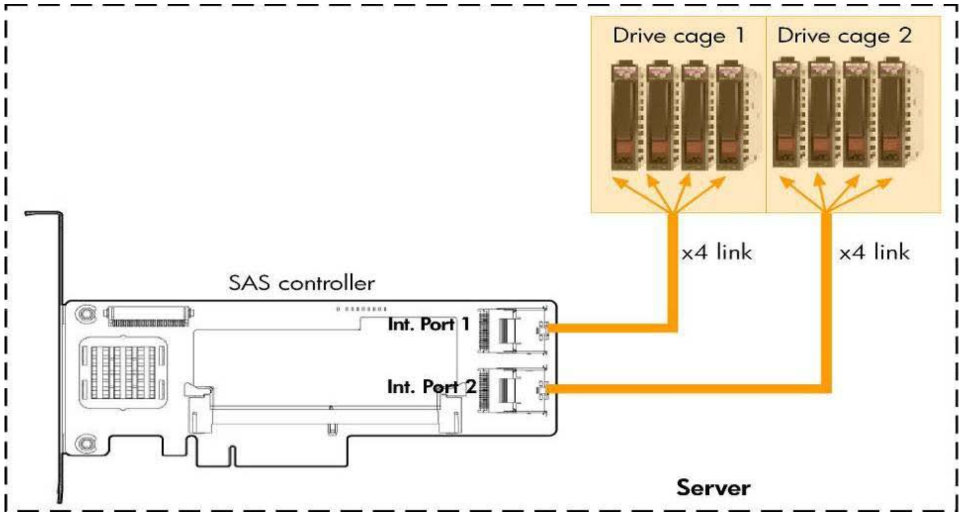

External

Figure 7 shows a topology for connecting a 4-lane-wide port on a controller to an external storage enclosure. This storage enclosure example contains an internal 36-port expander that supports cascading an additional enclosure in a 1+1 configuration.

Figure 7: This external port topology can support cascading to an additional enclosure.

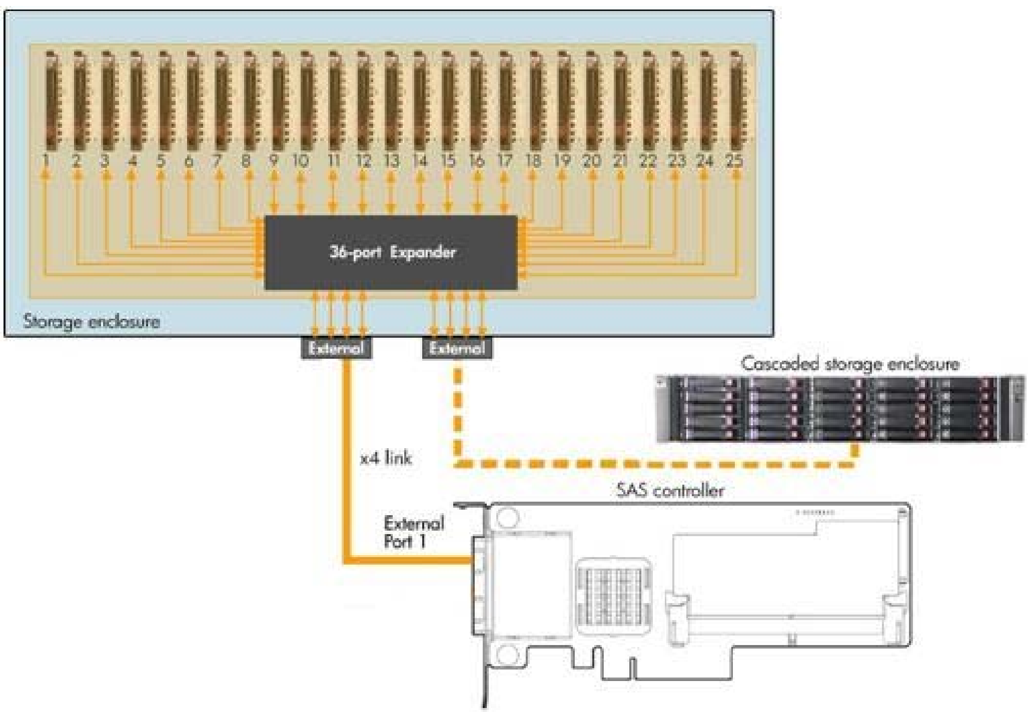

Dual-path SATA implementations use a single-domain method of providing tolerance to cable failure. Dual paths can prevent a single point of failure in complex enterprise configurations such as cascaded JBODs. This configuration connects a controller to the IO module at each end of a set of cascaded JBODs. Each JBOD contains a single expander, as shown in Figure 8. This configuration requires half as many IO modules and expanders as other redundant storage configurations, so it is more economical.

HP dual-path enclosures support both SAS and SATA configurations. Dual-path SATA implementations do not provide the full redundancy of a dual-domain SAS solution.

Figure 8: This is an example of a dual-path configuration for cascaded JBODs. The "4" notation indicates a 4-lane data path.

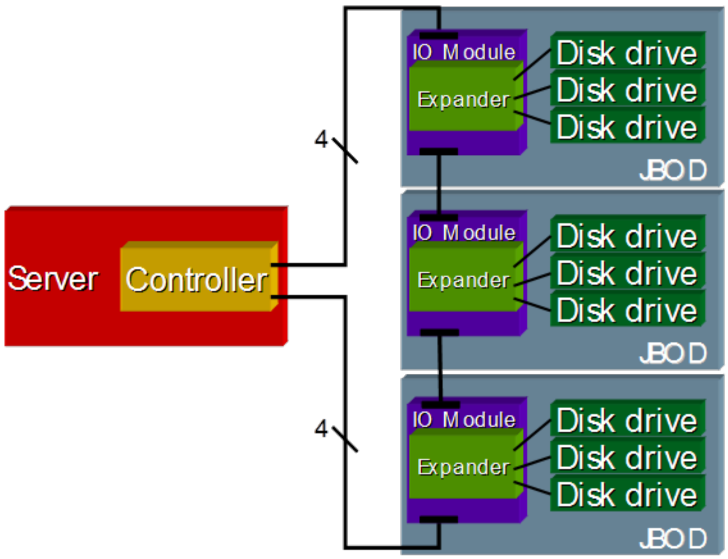

Conclusion

With SATA, we offer you affordable configuration options for a broad range of storage applications. For best performance and reliability, limit the use of SATA hard drives to servers that perform low-workloads (<40%) and to multi-drive storage configurations such as JBOD and RAID. Limit the use of SATA SSDs to servers performing high read/low write (70% read/30% write) random data operations. We do not recommend using SATA disk drives in applications that are not fault-tolerant, require high IOPS, or are mission-critical. The SAS infrastructure gives you the flexibility to install SAS drives, SATA drives, or both in the same enclosure.

Appendix: SATA terminology

Table 1: SATA terminology

Term	Definition
Entry drives	SATA drives that provide a basic level of reliability and performance for environments that are not mission critical.
Expander	A SAS device that functions as a switch to attach one or more initiators to one or more targets
Frame Information Structure	A FIS is a packet that may contain ATA register contents, data, control information, or other information.
Host	The initiator of communication with a SATA device
Logical drive	A partitioned space on physical devices
Midline devices	SATA drives that provide larger capacity and greater reliability than Entry-drives, making them better suited for multi-drive configurations.
PHY	The mechanism that contains a transceiver, which electrically connect to a physical link
Physical link	Two differential signal pairs, one pair in each direction, that connect two physical PHYs
SATA Tunneled Protocol (STP)	A protocol used to communicate with SATA devices in a SAS topology.
Target	A disk drive or solid-state drive

For more information

Visit the URLs listed below if you need additional information.

Resource description	Web address
SATA IO Web site	www.sata-io.org
"Serial Attached SCSI technologies and architectures" technology brief	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01613420/c01613420.pdf
HP ProLiant drives and data storage	http://h18004.www1.hp.com/products/servers/proliantstorage/drives-enclosures/index.html
"Server drive technology" technology brief	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01071496/c01071496.pdf
"Solid state storage technology for ProLiant servers" technology brief	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01580706/c01580706.pdf
FAQs about the SATA rev. 3.0 specification	http://www.sata-io.org/documents/SATA-Revision-3.0-FAQ-FINAL.pdf

Send comments about this paper to TechCom@HP.com

Follow us on Twitter: http://twitter.com/ISSGeekatHP

TC1108815, October 2011