



Why SATA Express?

By Paul Wassenberg
SATA-IO Board Member
Marvell Semiconductor





Why SATA Express?

The SATA Express specification developed by SATA-IO (www.sata-io.org) defines host and device connectors that support SATA or PCIe. In this document, references to PCIe devices are as defined in the SATA Express specification. We will get into more details on SATA Express later, but we'll first address why SATA-IO has chosen this development path.

Background

SATA-IO has nearly 200 members, including all the major HDD and SSD makers and quite a few of the smaller SSD companies. So as a group, we have a good idea of what's coming up for SSDs. We knew that some client SSDs would soon need an interface faster than the 6Gb/s we get from SATA. [Client storage refers to single user environments, primarily PCs].

The main factors driving the decision to pursue PCIe were:

- 1) With the next speed increase, the SATA infrastructure has to change in any case
- 2) Some client SSDs would soon require an interface faster than 6Gb/s
- 3) The portion of client SSDs that will require greater than 6Gb/s is fairly small
- 4) 6Gb/s SATA will be more than adequate for HDDs for the foreseeable future
- 5) Last, but not least, the new interface must be low cost

12Gb/s SATA would seem to be the logical next step, and the T10 (SAS) committee has done a lot of the work on 12Gb/s already. From this work, we know that the transition from 6Gb/s to 12Gb/s is not simple. SAS 3.0 (12Gb/s) requires transmitter equalization, which adds a great deal of complexity to the interface controller and the PHY. In silicon, complexity equates to more die area, which means higher cost. Also, the protocol needs to change to support transmitter training, and that turns out to be fairly significant. Additionally, many of the backplanes and cables that worked fine at 6Gb/s won't reliably carry data at 12Gb/s. Thus, as stated in factor #1, a lot of the SATA infrastructure would have to change in order to support 12Gb/s.

Even with the background work done on 12Gb/s SAS, it would not have been possible to have a 12Gb/s SATA spec out in time (factor #2).

Factors #3 & #4 tell us that the new interface will be fairly low volume for the first few years at least. 6Gb/s SATA will continue to serve the majority of the client storage market. So 12Gb/s SATA would not be able to take advantage of the economies of scale for quite a while yet, and would therefore be relatively expensive, which conflicts with factor #5.

Considering factor #1, we looked at other options to meet market needs and decided that PCI Express was a good fit.





PCIe has been shipping for years and is a mature technology. PCIe 3.0 (8Gb/s or 1GB/s) provides the needed bump up in speed with a single lane. By comparison, SATA at 6Gb/s equates to 0.6GB/s. Also, PCIe has been increasingly used as a storage interface because of the ability to scale up by simply adding additional PCIe lanes.

SATA Express standardizes the connection between the host and a HDD-type form factor, most typically the 2.5-inch size. Including a 2nd PCIe lane in the SATA Express specification provides the ability to support up to 16Gb/s (or 2GB/s), which gives us plenty of headroom for even faster client SSDs. Next generation PCIe 4.0 will double the bandwidth to 16Gb/s (or 2GB/s) per lane, so SATA Express has a growth path.


Hopefully, that clarifies the thinking that went into deciding to pursue SATA Express.

Now let's get into more detail on SATA Express.

What is SATA Express?

SATA Express is pure PCIe. There is no SATA link or transport layer, so there's no translation overhead – users will see the full performance of PCIe. Perhaps a good way to think about SATA Express is as the standardization of PCIe as an interface for a client storage device.

The benefit of standards is that they provide interoperability and interchangeability, meaning that devices from different manufacturers work together and can be interchanged with no loss of functionality. In this way, standards drive volume. To achieve these goals, SATA Express needs standard connectors and common operating system drivers.





SATA-IO defined SATA Express host and device connectors. Both connectors are slightly modified standard SATA connectors and are mechanically compatible with today's SATA connector. This plug-compatibility is important, as it enables SATA and PCIe to co-exist. The new host connector supports up to two SATA ports or up to two PCIe lanes. There is a separate signal, driven by the drive that tells the host if the device is SATA or PCIe, so the host knows what "language" to speak. Thus the motherboard can have a single connector that supports a current SATA drive or a PCIe drive.

A PCIe device can also mate with the SFF-8639 multifunction host connector (<ftp://ftp.seagate.com/sff/SFF-8639.PDF>), which is the primary means for PCIe to connect to enterprise environments.

Although beyond the scope of the SATA Express specification, developers must decide on the operating system driver interface. AHCI (www.intel.com/content/www/us/en/io/serial-ata/ahci.html) is the driver interface used for SATA, and is built into Windows, Linux, and pretty much any O/S that supports SATA. A PCIe device with AHCI is compatible with SATA software environments. However, AHCI was designed for hard drives and does not always take full advantage of the lower latencies of SSDs.

On the other hand, NVMe Express (www.nvmeexpress.org) was designed specifically for PCIe SSDs and provides optimal performance. NVMe Express is supported in Windows, starting with version 8.1; drivers for other operating systems are also available.

SATA Express is included in SATA specification revision 3.2. Meanwhile, SATA is being enhanced with additional features, and will continue to be the standard interface for the vast majority of PC storage for the foreseeable future.

For more information on SATA Express, visit <http://www.sata-io.org/sata-express>.

