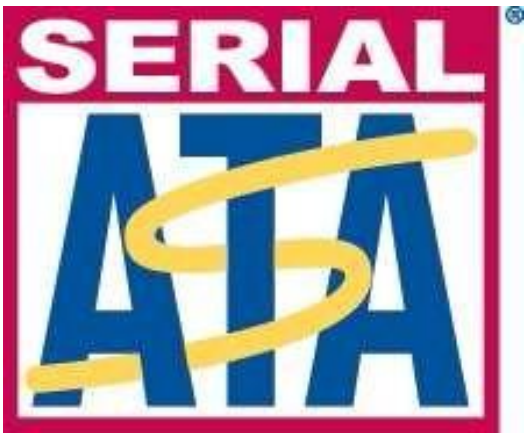

Why SATA Express?

By Paul Wassenberg
SATA-IO Board Member
Marvell Semiconductor

Why SATA Express?

SATA-IO (https://www.sata-io.org/) is developing SATA Express, which will utilize PCI Express as the physical interface, and will have an option to support SATA software compatibility. We will get into more details on SATA Express later, but we'll first address why this development path was chosen by SATA-IO.

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 know that some client SSDs will soon need an interface faster than the 6Gb/s we get from SATA today. [Client storage refers to single user environments, PCs primarily]

The main factors driving the decision to pursue SATA Express were:
1) With the next speed increase, the SATA infrastructure has to change in any case
2) Some client SSDs will require more than 6Gb/s within 2012
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, SATA & SATA Express must be low cost, per the SATA-IO charter

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 for products developed in 2012 (factor #2).

Factors #3 & #4 tell us that SATA Express will be fairly low volume for the next 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 the market needs, and decided that PCI Express was a good fit.

PCIe has been shipping for years and is a mature technology. PCIe Gen 3 (8Gb/s) provides the bump up in speed that we need with a single lane. Also, PCIe has been increasingly used as a storage interface, because of the ability to scale up by simply adding more PCIe lanes. To date, most PCIe SSDs have been in a standard PCIe card form factor. 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 SATA Express provides the ability to support up to 16Gb/s, which gives us plenty of headroom for even faster client SSDs. PCIe Gen 4, which is anticipated to come out in 3-4 years, will double the bandwidth to 16Gb/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 in an HDD-type form factor.

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 is defining host and device connectors for SATA Express. 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 SATA Express to co-exist. The new host connector supports up to two SATA ports or up to two PCIe lanes. The host connector pins for the SATA ports are multiplexed with the PCIe pins. There is a separate signal, driven by the drive, that tells the host if the device is SATA or SATA Express (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 SATA Express drive.

Although beyond the scope of the SATA Express specification, developers must decide on the operating system driver interface. AHCI (http://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 SATA Express device with AHCI would be compatible with SATA software environments. On the other hand, NVM Express (http://www.nvmexpress.org/) was designed specifically for PCIe SSDs, and will provide optimal performance. Being relatively new, NVM Express is not widely used yet, but drivers for Windows, Linux, and VMWare do exist.

SATA Express will be available later in 2012. Meanwhile, SATA will live on, 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.