


Computer Storage: Form factors, protocols and interfaces



I was chatting recently with some of our newest members of staff about various computer storage and transfer technologies, new and old, such as IDE, NVMe and UASP. There were a lot of abbreviations in that conversation. There are three in this paragraph!

This paper is all about the form factors, interfaces and protocols used by some common data storage media, and the abbreviations, acronyms and initialisms that come with them. It is impossible to describe the current crop of storage technologies properly without discussing some of their predecessors first. It is also necessary to consider that a standard such as SATA is made of separate physical, electrical and logical interfaces which work together. Let's begin with IDE.

IDE, ATA and PATA

IDE stands for Integrated Drive Electronics, which indicates that the drive electronics (or controller) is built in to the drive itself, a departure from previous generations of hard drives (excluding Hardcards). It was developed by Compaq, Western Digital and Control Data Corporation for Compaq's IBM-compatible computers in 1986.

The Personal Computer/AT was IBM's second generation PC (AT stood for Advanced Technology). Many of the technologies introduced in the AT in 1984, such as the ISA (Industry Standard Architecture) bus, persisted for many years afterwards. IDE was based on the ISA bus and was later standardised as ATA (AT Attachment). Once SATA (Serial ATA) came along, IDE/ATA came to be known as Parallel ATA or PATA.

The following is a diagram of an IDE ribbon cable (source: Wikipedia):



The ATA standard has been updated several times and the supported transfer speed has increased to 133 MB/s. One of the many improvements made to the original ATA standard was ATAPI (ATA Packet Interface), which enabled us to connect peripherals such as CD-ROM drives and Zip drives to our computers. ATAPI uses the ATA electrical and physical interfaces to transport the SCSI commands used by those devices.

SCSI, SAS etc.

SCSI (Small Computer System Interface) is older than IDE, having been around in some form since 1978 and standardised in 1986. This interface was not designed specifically for hard drives only and is still used by various peripherals such as tape drives. A separate adapter card is required usually as most motherboards do not have SCSI support on-board. This used to be the case with IDE as well in the early days, but consumer motherboards soon added support for that interface. SCSI support on the other hand tended to be reserved to high-end workstation motherboards. But there were a few manufacturers such as Amiga, Apple and Sun who adopted SCSI as a storage interface on their computers.

The SCSI standard defines the transfer protocol and physical interfaces. Various SCSI physical interfaces exist now including the original SCSI Parallel Interface or SPI as well as newer connections such as SAS (Serial Attached SCSI), FCP (Fibre Channel Protocol), iSCSI (Internet Small Computer Systems Interface) and UASP (USB Attached SCSI Protocol).

USB, UMS and UASP

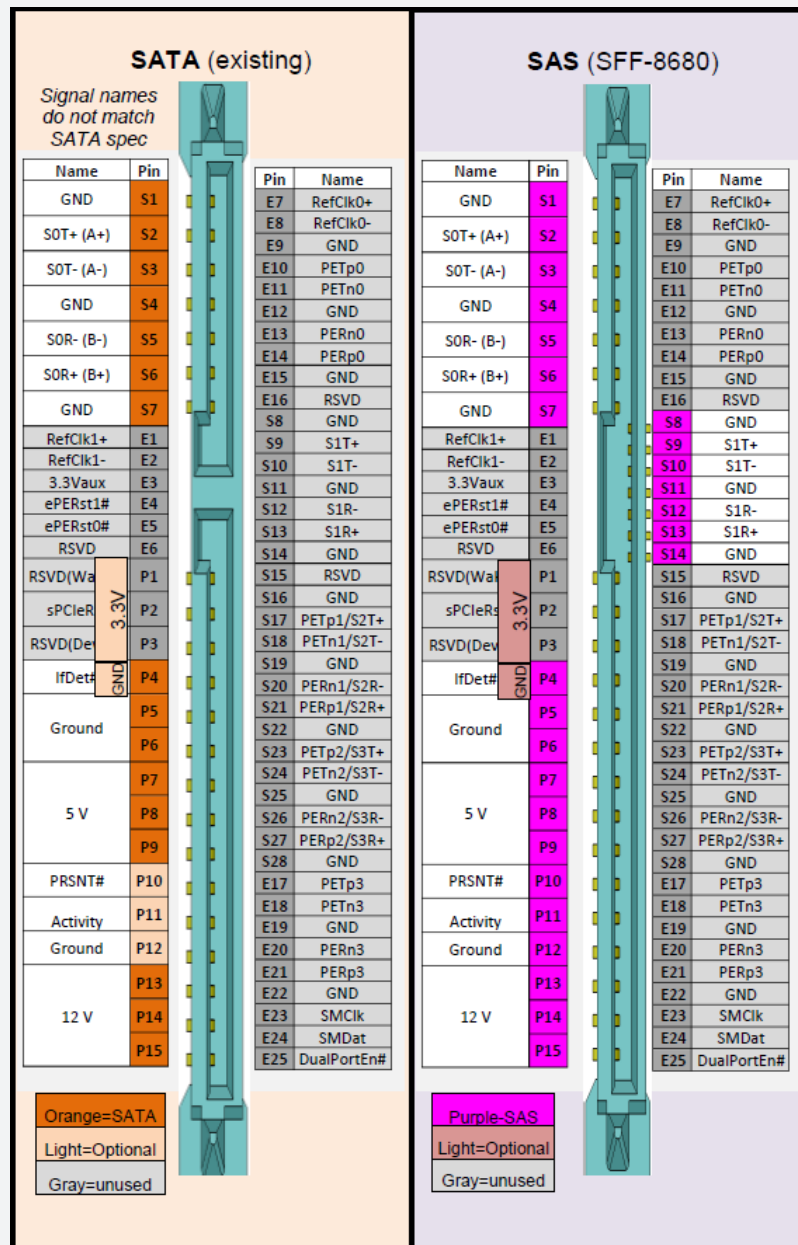
Let's take a break from internal hard drives and consider USB (Universal Serial Bus) Mass Storage (UMS) devices. Peripherals such as USB flash drives, most external hard drives, as well as some cameras and phones can use this feature of USB to appear to a computer as a regular storage device. I say most drives and some cameras and phones because alternatives such as NFS (Network File System), SMB (Server Message Block), AFP (Apple File Protocol) and MTP (Media Transfer Protocol) are also prevalent. There was a time, not more than a few years ago, when it was possible to connect a typical mobile phone to a computer as a mass storage device. The portion of the phone's non-volatile memory which was allocated for media storage would usually be presented as a FAT partition and a forensic image could be created with little effort. In contrast, a modern Android phone makes use of MTP. No part of its internal storage medium is exposed as a block device via UMS and therefore it isn't possible to create a forensic image in the same way.

If you're like us and utilise hard drive docking stations in your workflow, it may be worthwhile to check if they support UASP, also known as UAS (USB Attached SCSI). UASP combines the USB protocol with the command set of SCSI, resulting in faster transfer speeds than USB mass storage mode. To receive the benefits, UASP has to be supported by the hardware on both sides of the USB connection as well as the operating system (Linux 3.15, Windows 8 or OS X 10.8). StarTech claims that read and write performances can be up to 70% and 40% faster respectively, with a significant reduction in CPU usage.

SATA and AHCI

Back to drives. SATA introduced several improvements over PATA, including smaller cables, higher transfer speeds and hot-swapping, but it uses the same command set as ATA. It is accompanied by a new storage controller interface called AHCI (Advanced Host Controller Interface), which was developed by Intel. The AHCI controller is a PCI interface which moves data between a drive and system memory. If AHCI is not available, or if it is disabled on the motherboard, SATA controllers usually permit PATA emulation (IDE mode) as well, which disables most of the advantages of SATA (e.g. hot-swapping).

The following diagram shows the SATA connector and the (SCSI-based) SAS connector side by side (source: SSD Form Factor Working Group):



PCI and other buses

Another deviation is necessary to expand an abbreviation introduced in the previous section, PCI, and to introduce its successor PCIe.

Buses are communication pathways which facilitate the transfer of data and signals between various computer components. In a broad sense, this could include communication between components on a motherboard, between computers and peripherals, and between computers. On a motherboard, an expansion bus allows the use of expansion cards to introduce additional hardware such as graphics cards and network adapters.

The bus used on the original IBM PC came to be known as ISA, thanks to Compaq and other manufacturers of PC clones. This was succeeded by EISA (Extended Industry Standard Architecture), VLB (VESA Local Bus), PCI (Peripheral Component Interconnect), AGP (Accelerated Graphics Port) and most recently PCIe (PCI Express). There have been several derivatives of ISA, including ATA, PCMCIA (Personal Computer Memory Card International Association) and CompactFlash. Similarly, PCIe has spawned ExpressCard, SATA Express and M.2.

NVMe

NVMe (Non-Volatile Memory Express) is a new interface specification which was first published in 2013. Unlike previous technologies such as AHCI, NVMe was designed specifically for storage media based on non-volatile memory chips (e.g. flash memory). NVMe based storage devices are connected directly to the PCIe bus. Whereas AHCI (SATA) and SAS (SCSI) have a single command queue of 32 and 254 commands respectively, NVMe can have up to 65535 queues with 65536 commands per queue. The need for a SATA/SAS controller chip for data transfer is eliminated. These and other improvements result in NVMe offering significantly better performance than the old rotating magnetic media era interfaces.

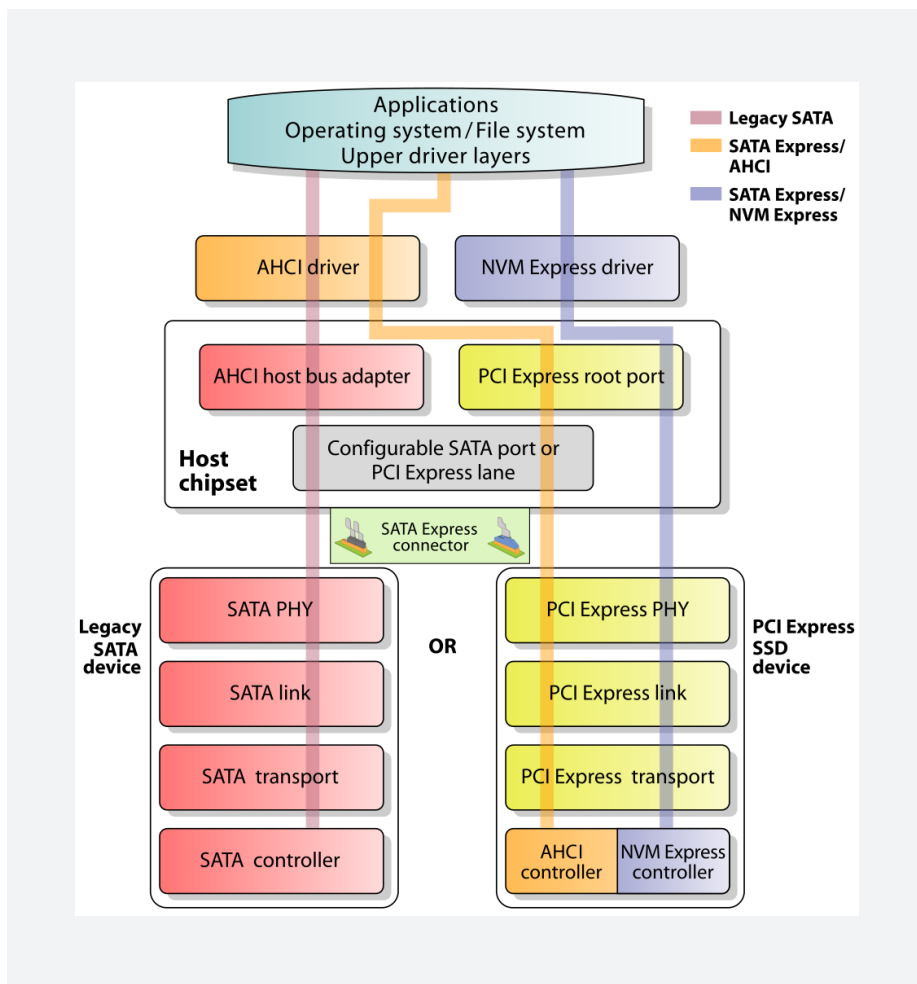
The first enterprise NVMe drives were released in Q4 2013, with consumer devices following in 2014. NVMe is rapidly becoming the de facto interface standard for SSDs in all environments.

Let's consider some PCIe drive form factors and connectors now. But beware, not all of them use NVMe. Some SATA drives use the newer connectors too!

SATA Express

The theoretical maximum transfer speed of a SATA 3.0 compliant storage device is 6 Gb/s (750 MB/s). When various overheads are taken into account, a SATA SSD (Solid State Drive) could achieve up to 4.8 Gb/s (600 MB/s). As SSD technologies improved, the speeds offered by the SATA interface became inadequate. One response to this problem was the introduction of SATA Express, as part of SATA 3.2. The SATA Express connector on a motherboard combines two SATA ports and a 2 lane PCIe port into the same connector. The PCIe connection can offer speeds close to 16 Gb/s (2 GB/s).

As illustrated in the following diagram (source: Wikipedia), SATA Express is compatible with plain old SATA, PCIe with AHCI and PCIe with NVMe. However, other connectors such as M.2 and U.2 are proving to be more popular.



mPCIe and mSATA


SATA Express introduced a new interface, but the shape of the drive remained the same. Unlike a Hard Disk Drive (HDD), which uses heavy rotating magnetic platters to store data, an SSD will have memory chips inside that requires a lot less space. Early SSDs used the same physical and electrical interfaces as HDDs, say the 2.5 inch form factor with a normal SATA connector. Since there was no good reason for SSDs to be so big, manufacturers started using smaller form factors such as mSATA.

The mSATA (mini SATA) connector is physically identical to the PCIe Mini Card form factor (mPCIe) which is commonly used by devices such as Wi-Fi adapter cards in laptops. But mSATA and mPCIe are not electrically compatible. To further complicate matters, there are mPCIe SSDs out there which use proprietary interfaces. So if you need to image one of those, make sure you have the correct adapter first.



M.2

M.2 (formerly NGFF) is an even smaller form factor introduced to supersede mSATA. At the physical level, the M.2 interface can support up to 4 PCIe lanes, SATA 3 and USB 3, depending on the host and the device. The M.2 interface standard permits both AHCI and NVMe based SSDs (not at the same time). A M.2 drive with 4 PCIe 3 lanes can offer transfer speeds close to 32 Gb/s (4 GB/s).



As the diagram above illustrates (source: Wikimedia Commons), M.2 cards can use one of three different connectors: B key, M key or B+M key. Cards with a B key can support SATA, 2-channel PCIe or one of many other interfaces. M key is suited to SATA or 4-channel PCIe. Five different card lengths are defined: 2230, 2242, 2260, 2280 and 22110. 22 signifies the width and the remaining digits indicate the length, both in millimetres.



The Western Digital SSD (green) in the previous photograph uses SATA/AHCI. The Intel SSD (blue) on the other hand uses NVMe. The NVMe SSD cannot be connected to a computer with a regular SATA adapter or write-blocker. Despite the fact that both cards have the same B+M key M.2 physical interface, their electrical interfaces and command sets are completely different.

That Intel SSD is interesting for another reason too. The storage components on it are not flash memory chips. Instead it uses 3D XPoint (pronounced cross-point), a technology which is designed to be faster and more resilient than flash memory. Intel is using 3D XPoint to offer both conventional storage solutions (Optane SSDs) as well as caching devices for legacy SATA based drives (Optane memory).

A forensic image of the SATA drive may be created by attaching it to any hardware write blocker which supports SATA; the only additional hardware required would be a M.2 to SATA adapter. To create an image of the other drive one would need a hardware write blocker which supports NVMe and has a suitable physical interface or adapter. Fortunately, the likes of WiebeTech and Tableau do support NVMe and M.2.

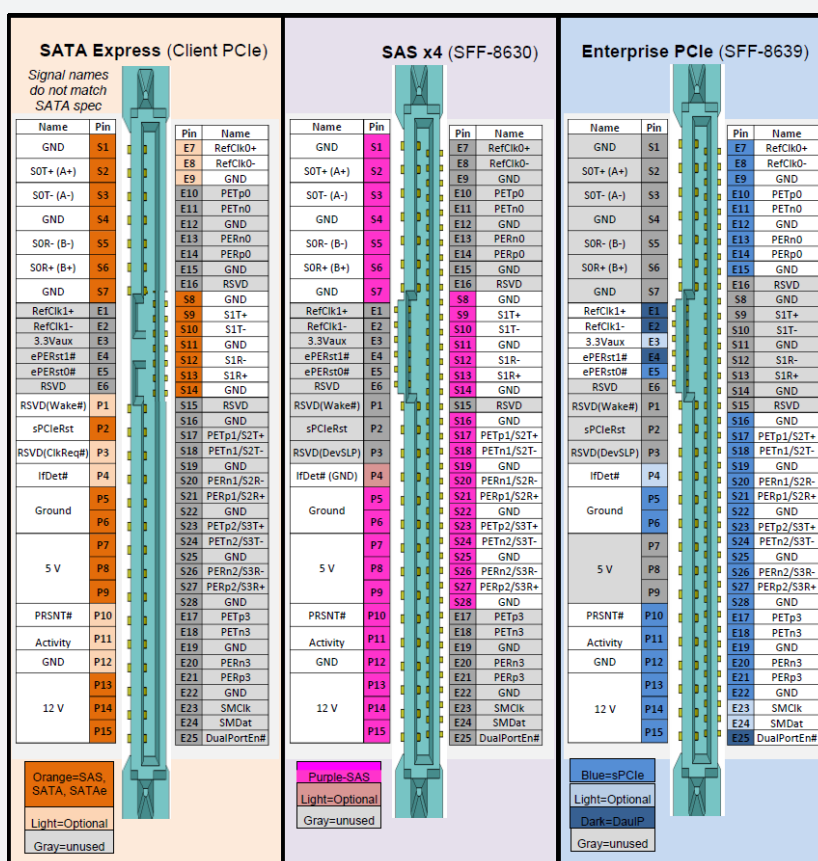
I have to make a point about Apple at this point. As is its wont, Apple has used several proprietary SSD form factors in recent devices and they can look deceptively similar to M.2 drives. Devices introduced from 2010 to 2012 used SATA, whereas newer ones use PCIe. Once again, an appropriate adapter will be necessary to read these SSDs outside their native environments.

PCIe AIC and U.2 SSDs

In addition to the newer PCIe connectors such as M.2, SSDs are also available in the traditional PCI Express AIC (Add In Card) form factor. Most of these SSDs are based on the NVMe specification, but some early devices used SATA or even proprietary interfaces requiring special drivers to operate. The high performance offered by these devices still come at a relatively high price which puts them beyond the reach of most home consumers.

Another connector designed for enterprises and enthusiasts is U.2 (formerly SFF-8639). This connector is based on the SATA Express drive connector, which in turn is an extension of the SAS connector. It can use up to 4 PCIe lanes, but is backwards-compatible with SAS and SATA. Nevertheless, the term U.2 is being used by manufacturers to describe NVMe SSDs which carry this connector and are housed in 2.5 inch style drive enclosures. The Mini-SAS HD (SFF-8643) connector is usually used to connect U.2 drives to motherboards.

The following diagram shows the SATA Express, SAS x4 and U.2 (SFF-8639) connectors (source: SSD Form Factor Working Group):



More form factors, more connectors...

I shall end by mentioning just a few more new storage technologies.

- Intel has a 325.35 mm long SSD form factor they call the ruler (based on EDSFF). 32 of those could fit in a 1U server.
- Samsung has a shorter version (NGSFF) and has recently launched an 8 TB NVMe SSD. It has 12 GB of RAM!
- SD Express is going to bring PCIe and NVMe to SD cards.
- CFexpress (CompactFlash + PCIe/NVMe) is already a thing.
- When hardware write-blocking isn't required, a USB enclosure is a convenient way to read a SATA M.2 SSD. This wasn't possible with NVMe drives until recently, but USB to PCIe bridge chips are available now (e.g. JMS583 and ASM2362). As a result NVMe to USB Type-C enclosures can be purchased for around £30. I haven't tested one yet though.
- Spinning rust drives aren't going away any time soon. Some of them have lasers (HAMR) or microwaves (MAMR) inside them now.

Arun Prasannan, Principal Analyst (Research & Development)

July 2018