11 Myths About M.2 and U.2 SSDs

Electronic Design
Doug Rollins and Jason Echols
Thu, 2016-09-08 12:36

The [SATA-IO](#) M.2 and the [SSD Form Factor Working Group](#) U.2 storage form factors combine multiple interfaces in their standards, including PCI Express and support for NVMe. They are relatively new form factors, and often times, questions arise regarding their capabilities and use. This article will help clear up any lingering confusion.

1. M.2 is the only PCI Express (PCIe) storage form factor.

The high throughput and low latency of PCIe is causing many new deployments to focus on PCIe SSDs to get the most out of their investment. While M.2 is a great form factor for fast SSDs, your application may not work with M.2 for some of the reasons outlined here. The good news is that you have two other options to look at. The first is native PCIe SSDs that plug right into the existing PCIe x4 slots on your server or client device. The other is U.2, which gives you more options for deployment, including dual porting for high availability.

Related

[3D NAND Flash Powers SSD and NVMe Drives](#)

[3D XPoint: A New Revolution in Memory?](#)

[3D XPoint Technology Claims to Find Sweet Spot Between NAND and DRAM](#)

2. M.2 doesn't support hot swap, so it has no place in enterprise platform designs.

Although it's true that the M.2 connector was not designed to support hot swap (the ability to remove a device from a system that's powered on without electrically damaging either the device or the host), M.2 is still a great option in many enterprise platforms. The compact M.2 form factor is seeing growing adoption in blade and hyperscale platforms as well as in embedded designs.
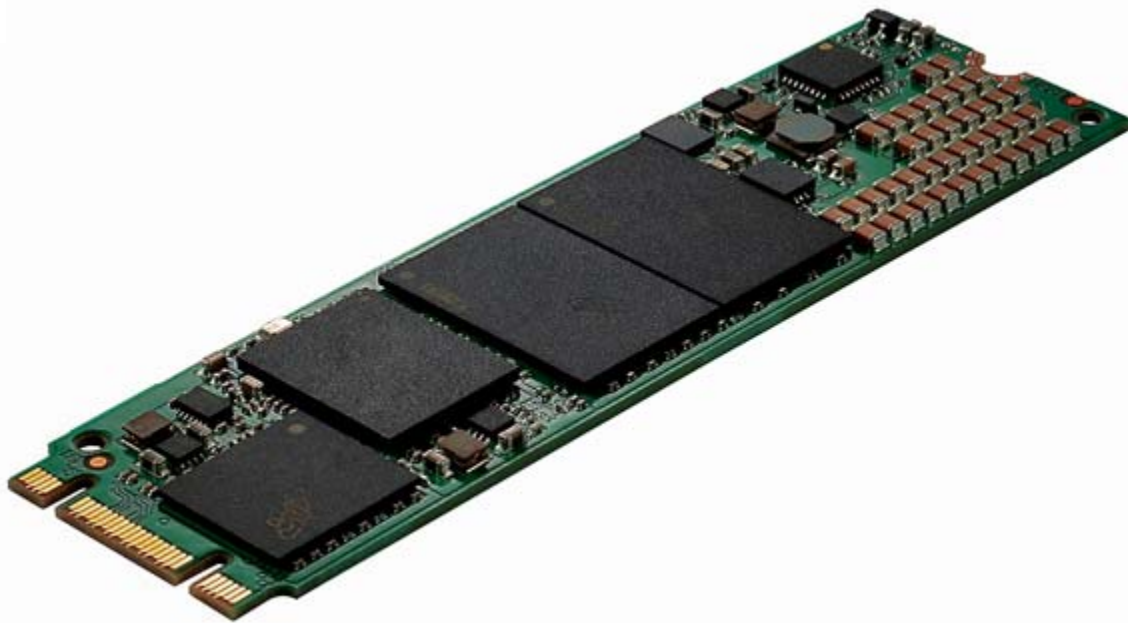
3. M.2 is physically too small to support useful storage capacities.

The M.2 form factor can be physically quite small, and both current technologies like 3D NAND and those yet to come like 3D XPoint Technology are fostering the development of M.2 SSDs with terabyte-class capacities. A single, embedded M.2 socket can support an SSD large enough for a full-sized, full-featured operating system and data.

4. M.2 is the same as mSATA.

M.2 is extremely versatile and supports PCI Express (up to x4), SATA 3.0, and USB 3.0. M.2 supports different keying configurations to distinguish among interface types. M.2 is extremely versatile and supports PCI Express (up to x4), SATA 3.0, and USB 3.0. It supports different keying configurations to distinguish among interface types. Many current personal-computing platforms support automatic device type detection, allowing for use of either SATA 3.0 or PCI Express SSDs in the same physical slot. M.2 also overcomes the bandwidth limitations of mSATA (mSATA supported up to SATA 3.0Gb/s, whereas M.2 was developed to allow SATA 6.0Gb/s and PCIe speeds).

5. All M.2 drives are the same physical size.

The M.2 standard supports 13 different keying options (note that some are reserved for future designs), four device widths (12, 16, 22, and 30 mm) and eight device lengths (16, 26, 30, 38, 42, 60, 80, and 110 mm). In addition, eight device thicknesses are supported.

6. M.2 can't dissipate enough heat to be a useful interface for high-performance systems.

M.2 doesn't require passive heat sinks as part of the connector specification, but they're not prevented if the M.2 devices meet the specification requirements. Due to low-power operation, self-monitoring and adjusting capabilities, and workload-focused design inherent in many M.2 designs, heat dissipation is less of a concern.

7. PCI Express doesn't support hot plug.

The U.2 connector (originally called SFF8639) design leveraged existing hard-disk drive mechanical specifications as well as PCI Express specifications. The connector was designed from the outset to support hot plug (additional host, target, and operating-system considerations may need to be understood for end-to-end hot-plug support and behavior as might be commonly expected). See www.ssdformfactor.org for a copy of the specification.

8. PCI Express doesn't support dual-port redundancy.

The U.2 form factor is available in two formats: a four-lane, single-port ("server") configuration in which all four PCI Express lanes are connected to a single path, and second configuration in which a pair of two-lane dual-port connections are used for a "storage" configuration. The dual-port storage configuration is optional (see section 5 of "Enterprise SSD Form Factor Version 1.0a" on www.ssdformactor.org for implementation details).

9. Nobody supports U.2 SSDs or chassis designs.

While they used to be rare, currently several enterprise SSD and server platform suppliers support the U.2 form factor. Multiple models are supplied by large-scale multi-national companies, OEM providers, and smaller, custom manufacturers. Chassis and systems come in 1U and 2U rack mounts with multiple configurations and processor types.
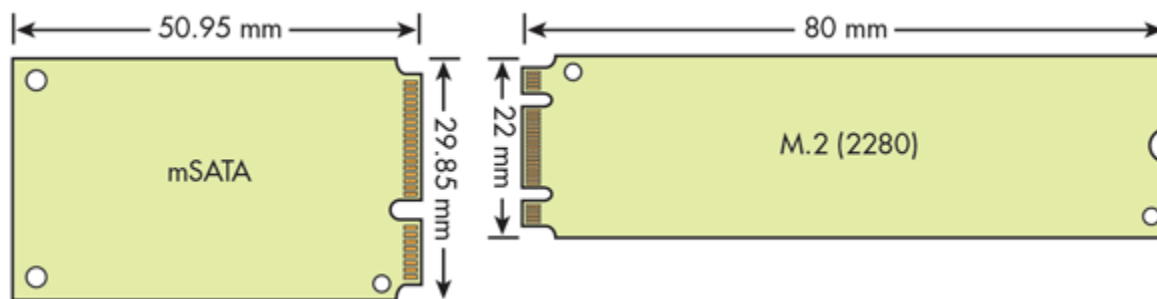
10. U.2 is unique and will see limited use in enterprise platforms.

U.2 supports SATA, SAS, and PCI Express on the same connector. For SAS and SATA, U.2 uses the same pin locations as legacy SAS and SATA for pin-wise backwards compatibility. U.2 adds the required PCI Express pins to the existing SAS/SATA layout by leveraging the space between the data and power connectors—an area that used to be left empty. By integrating all of these signaling pins into U.2, the drive and chassis designer have far more flexibility without forgoing historical device compatibility.

The U.2 PCI Express form factor can't support enough bandwidth to be a relevant enterprise erface.

While U.2 is limited to four lanes, its bandwidth per lane has kept pace with conventional PCI Express, with U.2 moving from Gen2 to Gen3 to realize double the total interface bandwidth with the newer generation to roughly 3.9GB/s. As the PCI Express standard continues to evolve, U.2 is expected to maintain the same performance trajectory.

Source URL: http://electronicdesign.com/memory/11-myths-about-m2-and-u2-ssds