




White Paper:

M.2 SSDs: Aligned for Speed

Comparing SSD form factors, interfaces, and software support



SAMSUNG
BUSINESS

M.2 SSDS: ALIGNED FOR SPEED

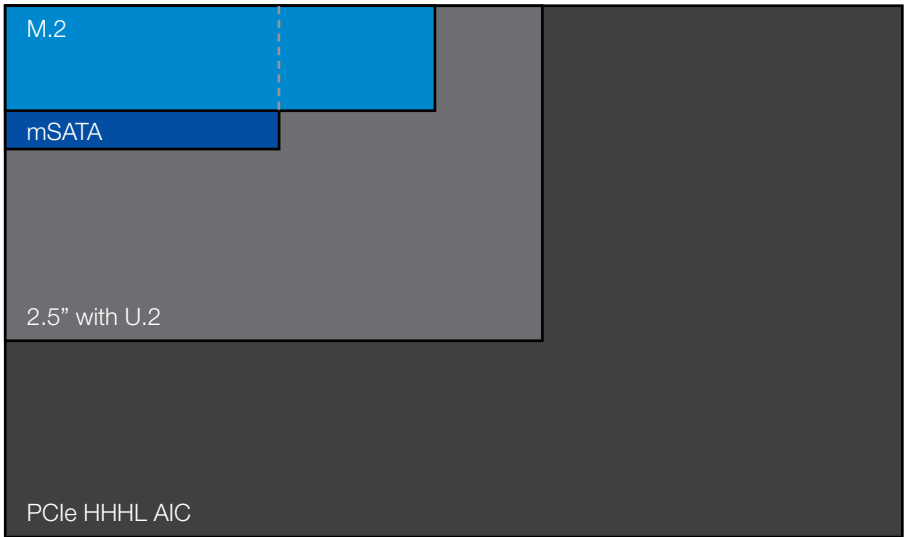
A flurry of new standards activity has opened a wide range of choices for solid-state drives (SSDs). Client devices are adopting new approaches quickly to gain advantages in density and performance. Evaluating the technology choices in a sea of acronyms – including AHCI, HHHL AIC, M.2, mSATA, NVMe, U.2, and V-NAND – means exploring three issues: form factors, interfaces, and software support. Aligning the latest technology onboard, M.2 SSDs free users from limits of earlier generations.

FOUR FORM FACTORS IN PLAY

mSATA was developed specifically for SSD modules, borrowing its 51x30mm low-profile outline from the PCI Express (PCIe) Mini Card form factor popular in laptops. For compatibility with more systems, mSATA replaced the PCIe interface with SATA 3.0, but users have already run into performance limits. Three other form factors for SSDs have turned to PCIe to deliver performance users want.

2.5” drives with the new U.2 (formerly SFF-8639) connector are well suited for enterprise drive farms with PCIe, SATA, or SAS controllers. In hot-swappable, front-removable applications, 2.5” drives with U.2 connectors work well. Client devices can cable PCIe slots to a 2.5” U.2-enabled drive with an adapter.

PCIe Half-Height, Half-Length Add-In Card (HHHL AIC) provides PCIe interface speeds on a relatively large 167x111mm card. These are simple to install in a PCIe slot, but resting perpendicular to a backplane or



Form factor	L (mm)	W (mm)	H (mm)	Interfaces
M.2	16-110	12-30	2.2-4.2	PCIe, SATA, USB
mSATA	51	30	3.8	SATA
2.5" with U.2	100	70	9-19	PCIe, SATA, USB
PCIe HHHL AIC	167	111	19	PCIe

riser card the overall system size ends up more suited for servers than for clients.

All three approaches – mSATA, 2.5” drives with U.2 connectors, and PCIe HHHL AIC – are space-inefficient for today’s client devices. Thin, compact platforms such as tablets, 2-in-1s, game consoles, and NUC-like machines present a new challenge. Succeeding mSATA, which achieves density but lacks performance, is a new

specification released at the end of 2013: M.2.


M.2 modules provide more choice of sizes and interfaces, at typically less than 1/10 the weight of a 2.5” hard drive. Optimized and larger layouts give M.2 more capacity than mSATA modules. Available PCIe, SATA, USB, and other interfaces support many functions – including Wi-Fi, Bluetooth, 4G LTE modems, and SSDs, making the M.2 form factor versatile and increasingly popular.

MORE CHOICES WITH M.2 MODULES

Standard M.2 module widths are 12, 16, 22, and 30mm, while standard lengths are 16, 26, 30, 38, 42, 60, 80, and 110mm. These digits begin an M.2 type designation. For instance, “2280-S3-M” indicates a 22x80mm module, the most popular size for SSDs, with single sided components and M keying. Each module has a 75-position edge connector, with a cutout on the opposite edge for securing via a screw. A module installs parallel to the host board to maintain its low profile.¹

BEFORE INSTALLATION, CHECK THE KEYING

Keying helps ensure electronic compatibility for an “M.2 slot”, which may be one of several varieties. Notches on the M.2 edge connector match keys in the M.2 receptacle, allowing installation of only compatible modules. These profiles carry key IDs (reserved profiles not shown):



Key ID	Notched pins	Interfaces
A	8-15	2x PCIe x1, USB 2.0, I2C, DisplayPort x4
B	12-19	PCIe x2, SATA, USB 3.0, USB 2.0, I2C, Audio
E	24-31	2x PCIe x1, USB 2.0, I2C, SDIO
M	59-66	PCIe x4, SATA

For an M.2 module to fit in their device, users should check the device slot size and keying. Most M.2 SSDs fall into either B or M profiles, with the primary difference being two or four lanes of PCIe.

SHIFTING FROM SATA TO PCIE

SATA (Serial Advanced Technology Attachment) interfaces dominated client storage for a generation. PC chipsets integrated SATA controllers, 2.5" and smaller hard drives became ubiquitous, and point-to-point cabling made for easy and inexpensive connections. In enterprises, SAS (Serial Attached SCSI) interfaces

provided more performance from dedicated storage controllers and enhanced cabling.

When SSDs debuted, they emulated hard drives at first, borrowing the 2.5" form factor and SATA interfacing. mSATA introduced smaller, lighter SSD modules, but stayed with SATA

for compatibility. As flash memory and storage controller technology improved, many SSDs have hit the SATA throughput ceiling in the latest revision of the specification. Actual performance depends on many implementation factors, but theoretical bandwidth limits per direction are easy to compare.

Specification	SATA			SAS			PCIe		
Revision	1.0	2.0	3.0	1.0	2.0	3.0	Gen 1	Gen 2	Gen 3
Bandwidth (Gb/sec)	1.5	3.0	6.0	3.0	6.0	12.0	2.0	4.0	7.877
Bandwidth per direction with PCIe multi-lane links (Gb/sec) ²						x2	4.0	8.0	15.75
						x4	8.0	16.0	31.5

At first glance, SAS looks like it might be attractive for SSDs, especially with support for multiple initiators and longer cable lengths (10m compared to 1m for SATA). However, economics do not favor SAS for client devices. Typically, client PCs need a dedicated SAS host bridge adapter, or servers use a high-end chipset with an integrated SAS controller.

Faster is better. After intensive debate, the SATA community decided further speed enhancements were extremely challenging. In smaller systems with affordable chipsets and single or dual drives, SAS tends toward

cost prohibitive. An easier path to speed is to embrace an interface that offers better performance, scales better, and comes integrated in every new PC chipset – PCIe.

PCIe has been supported in PC chipsets for some time, used primarily for add-in graphics adapters, networking interfaces, and RAID controllers. As PCIe-enabled SSDs appeared, Intel (who also supplies SSDs) began focusing on PCIe storage support in the Intel Series 9 chipset. AMD and others quickly followed suit. With chipset support and the M.2 specification in place,

PC motherboard manufacturers are incorporating M.2 SSD slots directly into products, as are client device makers seeking the benefits of space efficiency. This drives volume adoption, reducing costs for all M.2 users.

The latest storage specification, SATA Express, redefines a combination of SATA and PCIe interfacing, allowing legacy devices to move forward and new devices to capture PCIe throughput. SATA Express also incorporates both legacy protocols and a new, faster driver scheme for advanced operating systems.

NVME BECOMES THE LOGICAL CHOICE

Once the hardware limitations of SATA are overcome, the next barrier in performance is software. The first wave of SATA-based SSDs, again seeking maximum compatibility with legacy hard drives, used the Advanced Host Controller Interface (AHCI). This standardized protocols so that hard drives and SSDs were interchangeable, but compromised performance.

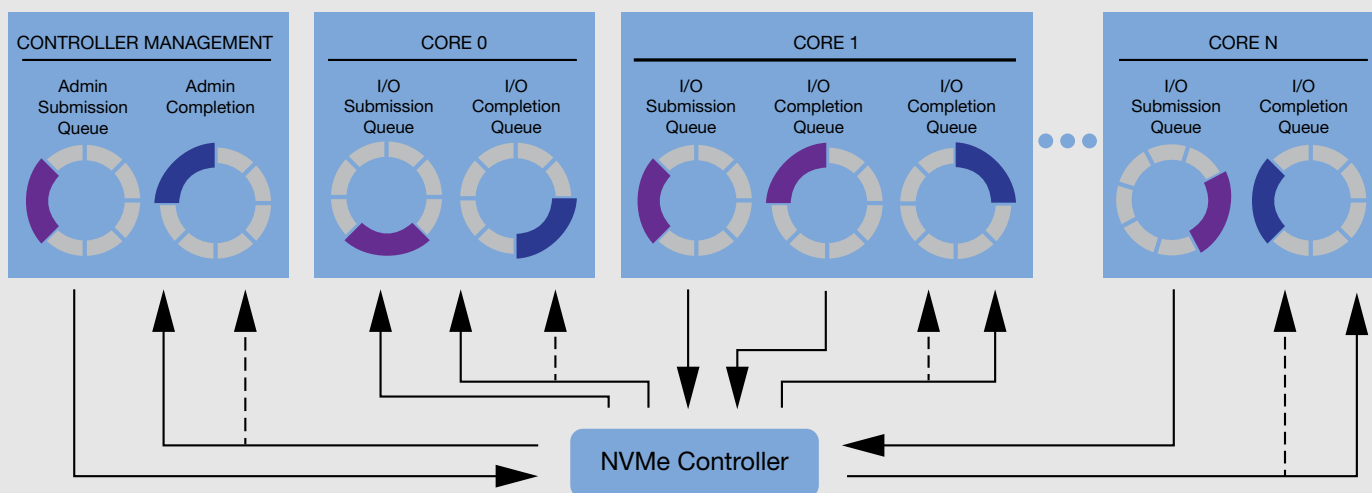
First-generation PCIe-based SSDs often used proprietary protocols, enabling them to achieve better performance than using AHCI, but requiring device-specific drivers and potentially new BIOS. The standards community responded

with a new logical interface specification, NVM Express or NVMe, and the definition of SATA Express allowing AHCI and NVMe to co-exist in the same system.³

NVMe was designed from the ground up to provide a consistent, high performance logical interface for SSDs over PCIe. It provides features such as enhanced queueing, interrupt steering, and multiple threading. Instead of the single command queue with 32 commands in AHCI, NVMe has 64K queues with 64K commands each. One interrupt in AHCI is supplanted by 2048 in NVMe, and the locking synchronization scheme has been eliminated.⁴

Both SSD vendors and software firms have rapidly embraced the NVMe logical interface. Microsoft Windows 8.1 added native support for NVMe, carried forward into Windows 10. NVMe support is also available in Mac OS X Yosemite, and the latest versions of Linux and Google Chrome OS. With advanced, multi-threading operating systems and multicore processors, NVMe support is essential to getting the most out of new SSDs with PCIe interfacing and V-NAND flash.

MORE CORES, MORE QUEUES, MORE COMPLETED REQUESTS



Adapted from NVMe Express, Inc

M.2 SSDS ROLL OUT V-NAND

To reach the performance potential of PCIe and NVMe, attention is now turning to SSDs with faster flash memory. In previous SSD implementations, V-NAND flash delivered improved write endurance and reduced transactional latency, but real-world sequential read figures top out at 560MB/sec on the obligatory SATA interface.

Tests on the first Samsung M.2 SSDs with PCIe Gen 3 x4 and NVMe have shattered the

read limit users experienced with SATA SSDs, hitting up to 2500MB/sec. The first M.2 SSD with V-NAND flash memory, PCIe Gen 3 x4, and NVMe is ready for consumers: the Samsung 950 PRO M.2 SSD. Adding V-NAND promises to improve latency, reduce power consumption, and extend write endurance.



SAMSUNG ENTERPRISE SSD PORTFOLIO

950 Pro Series Client PC SSDs

- Designed for high-end PCs and workstations
- Form-factors: M.2
- NVMe

850 EVO Series Consumer SSDs

- Optimized performance for everyday business
- Form-factors: M.2, mSATA or 2.5"
- AHCI

850 Pro Series Enterprise SSDs

- Designed for light enterprise and Client PCs
- Form-factors: 2.5"
- AHCI

PM863 Series Data Center SSDs

- Designed for read-intensive applications
- Form-factors: 2.5"
- AHCI

SM863 Series Data Center SSDs

- Designed for write-intensive applications
- Form-factors: 2.5"
- AHCI

¹"All About M.2 SSDs", webcast presentation, Solid State Storage Initiative (SNIA), <http://snia.org/sites/default/files/SSSI%20M%20%20Webcast%20Slides%20v10fml.pdf>

²"How does the PCIe 3.0 8GT/s 'double' the PCIe 2.0 5GT/s bit rate?", Frequently Asked Questions, PCI-SIG, https://pcisig.com/faq?field_category_value%5B%5D=pci_express_3.0&keys=bandwidth

³"AHCI and NVMe as Interfaces for SATA Express Devices", Dave Landsman, The Serial ATA International Organization, https://www.sata-io.org/sites/default/files/images/NVMe_and_AHCI_as_SATA_Express_Interface_Options_Overview_final.pdf

⁴"A Comparison of NVMe and AHCI", Don Walker, The Serial ATA International Organization, https://www.sata-io.org/sites/default/files/documents/NVMe_and_AHCI_long.pdf

Learn more: samsung.com/enterprisesd | 1-800-SAMSUNG | 1-866-SAM4BIZ

Follow us: [youtube.com/samsungbizusa](https://www.youtube.com/samsungbizusa) | [@SamsungBizUSA](https://twitter.com/SamsungBizUSA)

© 2015 Samsung Electronics America, Inc. All rights reserved. Samsung is a registered trademark of Samsung Electronics Co., Ltd. All products, logos and brand names are trademarks or registered trademarks of their respective companies. This white paper is for informational purposes only. Samsung makes no warranties, express or implied, in this white paper.

WHP-SSD-M2INTERFACE-SEP15J

SAMSUNG
BUSINESS