KIOXIA

Frequently Asked Questions

Key Questions to Ask About EDSFF E3 Form Factors

Today's server and storage system designs are more use case-specific and require data storage that can easily adapt to optimizations in performance or capacity (or a hybrid of both). With advancements in CPU, memory, accelerator, networking and IoT technologies, coupled with faster communication protocols and more robust interfaces, the server storage architecture is being pushed to its limits. Storage drives based on 2.5-inch[1] and M.2 formats are impacted and cannot keep technological pace with today's server and storage demands and future high-speed interfaces, such as those based on PCIe® 5.0 and 6.0 technologies.

The 2.5-inch drive form factor has served the industry for almost 30 years. Several years ago, the industry realized that this drive format was nearing the end of its useful effectiveness, and a new set of storage drive form factors would be needed to not only address 2.5-inch limitations, but M.2 as well.

What is EDSFF?

With limitations and challenges associated with 2.5-inch and M.2 drive formats, the Small Form Factor (SFF) working group within the Storage Networking Industry Association (SNIA) architected a new set of form factors for present and future data center architectural requirements. From this effort, the Enterprise and Datacenter Standard Form Factor (EDSFF) E3 family of form factors was created.

What is the E3 Family of Form Factors?

The E3 family of form factors consists of four types: (1) E3 Short Thin or E3.S; (2) E3 Short Thick or E3.S 2T; (3) E3 Long Thin or E3.L; and (4) E3 Long Thick or E3.L 2T as described below:

E3 Short Thin (E3.S)	Targeted to NVMe® SSDs with x4 PCIe link widths though it can mechanically fit an x16 card edge. It supports power profiles up to 25W and positioned to be a primary form factor for mainstream NVMe server storage subsystems as it can be used across a wide variety of platforms including modular and short depth chassis. The thickness designator of '1T' is not required.
E3 Short Thick (E3.S 2T)	Targeted to higher performance NVMe SSDs, Compute Express Link (CXL), Storage Class Memory (SCM), computational storage or front I/O implementations. It supports x4, x8 or x16 PCIe link widths and power profiles up to 40W.
E3 Long Thin (E3.L)	Targeted to be a primary form factor for storage subsystems and server platforms requiring maximum capacity for each 'U' configuration that utilize deeper chassis, and for high-capacity NVMe SSDs or SCM devices with support for x4, x8 or x16 PCIe link widths and power profiles up to 40W.
E3 Long Thick (E3.L 2T)	Targeted to Field-Programmable Gate Arrays (FPGAs) or accelerators, with support for x4, x8 or x16 PCIe link widths and power profiles up to 70W.

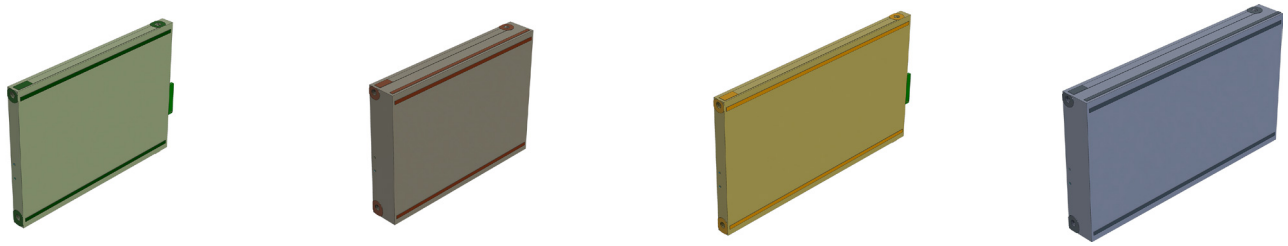

(Images provided by KIOXIA)

E3.S	E3.S 2T	E3.L	E3.L 2T
Height: 76mm Length: 112.75mm Width: 7.5mm	Height: 76mm Length: 112.75mm Width: 16.8mm	Height: 76mm Length: 142.2mm Width: 7.5mm	Height: 76mm Length: 142.2mm Width: 16.8mm

Are specifications available that define the E3 family of form factors?

E3 form factors are defined by the following SNIA SFF specifications:

Specification	Description
SNIA-SFF-TA-1008 Rev. 2.0	Enterprise and Datacenter Device Form Factor
SNIA-SFF-TA-1002 Rev. 1.3	Protocol Agnostic Multi-Lane High Speed Connector
SNIA-SFF-TA-1009 Rev. 3.0	Enterprise and Datacenter SSD Pin and Signal Specification (EDSFF)
SNIA-SFF-TA-1023 Rev. 0.8	Thermal Characterization Specification for EDSFF E3 Devices

These specifications are publicly available at https://www.snia.org/technology-communities/sff/specifications.

What challenges are associated with 2.5-inch form factors?

The 2.5-inch form factor originated with hard disk drives and is not optimal for flash memory packaging or optimized for flash memory channels. As performance scales to exercise all of the flash memory and activate the dies, power increases on flash memory and the PCIe interface. Since the 2.5-inch SSD format 'caps out' at 25W, there's a performance limit as to what can be achieved. Additionally, the connector used with 2.5-inch form factors are not designed to compensate for extended signal integrity challenges from high-speed interfaces such as PCIe 5.0.

How does E3 form factors address 2.5-inch form factors challenges?

E3 form factors can scale both power requirements and PCIe speed increases, and allow for wider PCIe link widths to enable full-throttle Input/Output Operations Per Second (IOPS) performance per terabyte[2] (TB) of capacity. E3 supported SSD packaging features a new airflow design that delivers more efficient heat dissipation and system cooling within the data center. Additionally, E3 supported SSDs are better prepared for the signal integrity challenges associated with next-generation high-frequency interfaces, and where the SFF-8639 connector (used on 2-5-inch drives) may not suffice.

What challenges are associated with M.2 form factors?

The M.2 form factor is a more recent design than 2.5-inch and brings SSD storage to the size of a stick of gum. They are deployed in servers and client notebooks/laptops as primary storage and boot drives. M.2 SSDs are slim and offer clean, uncluttered server and PC interiors that occupy less real estate and require only a few cables versus 2.5-inch drives. However, there are also challenges and limitations of the form factor that need to be addressed.

Performance on an M.2 SSD is limited at 8.25W, especially with higher capacities such as 4 TB and beyond. When driven by the PCIe 4.0 interface, M.2 SSDs cannot achieve full PCIe 4.0 16 gigatransfers per second (GT/s) x4 bandwidth with an 8.25W power envelope. Additionally, these drives have thermal challenges when operating in extreme temperature environments. If they fail, M.2 SSD replacement requires an entire server to be powered down. As M.2 drives are packaged in a 22mm width, flash memory placements are not optimal, which in turn limits the ability to increase drive capacities. M.2 SSDs also do not support any type of 'presence detect' capability, so if an M.2 drive is powered off, the system will not know if the drive is active or not.

How does E3 form factors address M.2 form factor challenges?

E3 form factors are designed to achieve full PCIe 4.0 (16 GT/s x4) bandwidth and 'soon-to-be-available' PCIe 5.0 (32 GT/s x4) bandwidth, and even PCIe 6.0 (64 GT/s x4) bandwidth, up to a 25W power envelope and much higher than 8.25W. Thermal throttling mechanisms can detect if an SSD has reached an operating temperature beyond its specification and will throttle performance to reduce the heat generated by flash memory die activation.

Addressing physical serviceability, E3 form factors enable supported SSDs to be hot-swapped from the front of the server, saving time, system down time and ease of maintenance. E3 form factors enable more space in the SSD than M.2 for additional flash memory chips that can enable higher capacity SSDs with more capacity. E3 form factors also include a 'presence detect' capability so if the device is powered off, it can be detected in the system and removed if required.

Which server/storage vendors support E3 form factors?

Dell EMC® and HPE® are among server vendors developing E3-based solutions. KIOXIA announced its CD7 PCIe data center NVMe SSD series in support of E3 form factors. Initial E3-based development and demonstration systems are in development with early full functional systems shown below:

EDSFF E3.S prototype system with KIOXIA CD7 Series PCIe data center NVMe E3.S SSDs.
(Sources: Dell Technologies and KIOXIA)

What are the key benefits of E3.S form factors vs. 2.5-inch and M.2 drives?

E3 form factors provide a number of key benefits to both system architects/designers and end-users alike:

Higher Performance / Higher Power Budget vs. 2.5-inch and M.2 Drives
More than doubles the power budget vs. M.2 devices enabling E3.S SSDs to saturate PCIe Gen4 performance

Standardized Thermal Solutions
Improves interoperability across vendors and platforms while providing the flexibility to select the right balance of cooling and storage density via E3.S enclosures

Improved Physical Serviceability
Improves drive serviceability with hot-plug support that no longer requires an entire server to be taken down in order to replace a single SSD

Better NAND Flash Memory Orientation
Wider PCB design enables optimized orientation of the NAND flash memory packages and provide more headroom for higher capacity drives

Supported by Leading Server and Storage OEMs
Dell EMC and HPE, leading authors of the Open Compute Project (OCP) NVMe Cloud SSD specification[3], are using E3.S designs on new and upcoming PCIe 5.0 platforms is leading the way in industry-wide support and adoption

When will E3-enabled products be market-available?

Upon development, E3-enabled products will be subjected to extensive validation and interoperability testing before market-availability. Early E3-enabled system-level solutions may sample in 2022, with ramp up expected in late 2022 and early 2023.

Notes:

[1] 2.5-inch indicates the form factor of the SSD and not its physical size.

[2] Definition of capacity - KIOXIA Corporation defines a megabyte (MB) as 1,000,000 bytes, a gigabyte (GB) as 1,000,000,000 bytes and a terabyte (TB) as 1,000,000,000,000 bytes. A computer operating system, however, reports storage capacity using powers of 2 for the definition of 1Gbit = 2^{30} bits = 1,073,741,824 bits, 1GB = 2^{30} bytes = 1,073,741,824 bytes and 1TB = 2^{40} bytes = 1,099,511,627,776 bytes and therefore shows less storage capacity. Available storage capacity (including examples of various media files) will vary based on file size, formatting, settings, software and operating system, and/or pre-installed software applications, or media content. Actual formatted capacity may vary.

[3] The Open Compute Project standardized on EDSFF E1.S form factors and authored the "OCP NVMe Cloud SSD Specification" for storage vendors to develop E1.S SSD standardized designs.