KIOXIA

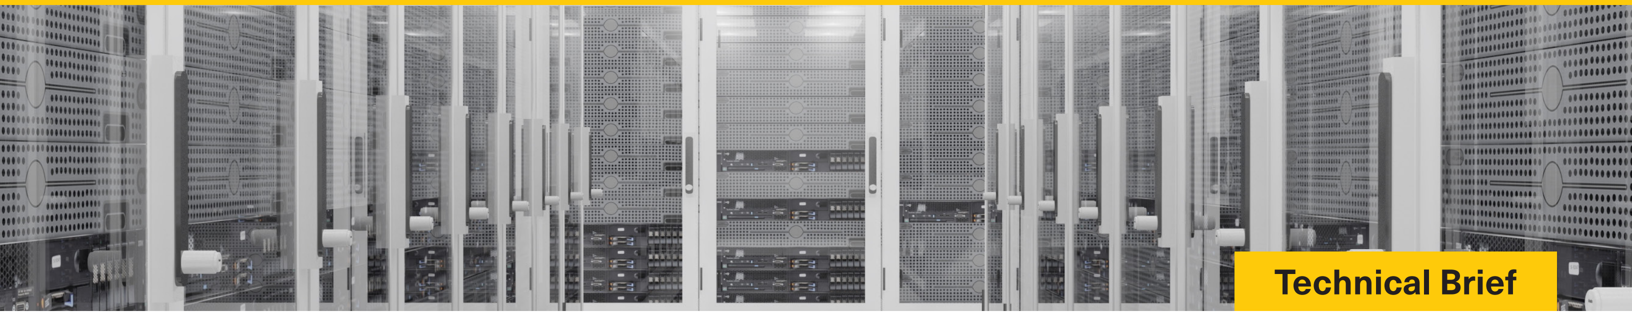

Technical Brief

EDSFF E3 Form Factor Overview

Next-Generation Server Storage Form Factors Address Future Data Center Requirements

Solid State Drives (SSDs) used in server storage subsystems have been primarily designed in 2.5-inch[1] and M.2 form factors. The 2.5-inch drive form factor has served the industry for almost 30 years while the first M.2 SSDs became market-available in 2013. With the advancements in CPU, memory, accelerator, networking and IoT technologies, the server storage subsystem must also scale to higher performance interface speeds (PCIe® 5.0/6.0 interfaces and beyond), and new protocols (Non-Volatile Memory Express (NVMe®) specification, Compute Express Link (CXL), etc.). The current connector and profile used in 2.5-inch and M.2 devices do not scale to higher speed interfaces and thermal resistance is needed for future generations of server storage subsystems.

The Small Form Factor (SFF) working group within the Storage Networking Industry Association (SNIA) architected a new set of form factors that not only address 2.5-inch and M.2 drive limitations, but future data center system requirements as well. From this effort, the Enterprise and Datacenter Standard Form Factor (EDSFF) working group was formed and the E3 family of form factors was created.

Next-Generation Form Factor Requirements

Given 2.5-inch and M.2 SSD limitations, the EDSFF working group considered many different aspects for creating a next generation form factor that will serve the industry for many years. These requirements included:

Signal Integrity (SI)	Support for next generation high frequency interfaces. At a minimum the connector system must support PCIe 5.0 and PCIe 6.0 interfaces.
Multiple Device Types	Support for multiple device types such as NVMe SSDs, CXL and Storage Class Memory (SCM) devices, computational storage devices, low end accelerators, and front-facing I/O devices such as Network Interface Controllers (NICs).
Link Width	Support for multiple host connection link widths. Different device types require different link widths such as PCIe x4, PCIe x8 and PCIe x16 connections.
Form Factor Size	Support for different size requirements to work optimally in both 1U and 2U platforms and to be large enough to accommodate multiple device types and high performance flash memory controllers.
Power Envelopes	Support for appropriate power envelopes in the future that scale to higher power devices. For NVMe SSDs, 25W is required to saturate a PCIe 4.0 (16 gigatransfers per second (GT/s)) link. For PCIe 5.0, 30W is the expected max power. For low end accelerators, 70W max power is sufficient.
Thermal Environments	A casing design that is optimized for heat dissipation of the controller and flash memory inside an SSD to enable more efficient heat transfer and system cooling than legacy 2.5-inch and M.2 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 the PCIe 5.0 interface.

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. 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 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.

E3 Form Factor Family

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:

<table>
<tr><th colspan="2">E3 Short Thin (E3.S)</th></tr>
<tr><td>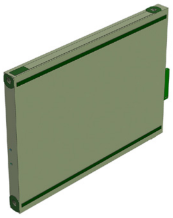</td><td>Targeted to NVMe SSDs with x4 PCIe link widths though it can mechanically fit an x16 card edge. It supports power profiles up to 25W and positioned to be a primary form factor for mainstream NVMe server storage subsystems as it can be used across a wide variety of platforms including modular and short depth chassis. The thickness designator of '1T' is not required.
Height: 76mm | Length: 112.75mm | Width: 7.5mm</td></tr>
<tr><th colspan="2">E3 Short Thick (E3.S 2T)</th></tr>
<tr><td></td><td>Targeted to higher performance NVMe SSDs, CXL and SCM devices, computational storage devices, and front I/O implementations. It supports x4, x8 or x16 PCIe link widths and power profiles up to 40W.
Height: 76mm | Length: 112.75mm | Width: 16.8mm</td></tr>
<tr><th colspan="2">E3 Long Thin (E3.L)</th></tr>
<tr><td>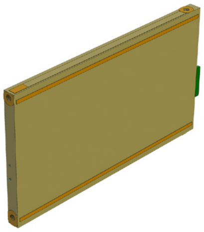</td><td>Targeted to be a primary form factor for storage subsystems and server platforms requiring maximum capacity for each 'U' configuration that utilize deeper chassis, and for high-capacity NVMe SSDs or SCM devices with support for x4, x8 or x16 PCIe link widths and power profiles up to 40W. The thickness designator of '1T' is not required.
Height: 76mm | Length: 142.2mm | Width: 7.5mm</td></tr>
<tr><th colspan="2">E3 Long Thick (E3.L 2T)</th></tr>
<tr><td></td><td>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.
Height: 76mm | Length: 142.2mm | Width: 16.8mm</td></tr>
</table>

(Images provided by KIOXIA)

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.

E3 System Design Capabilities

The E3 family of form factors enable supported SSDs to make efficient use of the flash memory chips implemented for storage density, and efficient use of the number of SSDs deployed in server and storage systems as follows:

Server Platform	Device Pitch	E3 SSDs Supported *(7.5mm wide)*	E3 2T SSDs Supported *(16.8mm wide)*
1U	9.3mm	20	10
2U	9.3mm	44 – 46*	22 - 23*

*depends on the 2U chassis mechanical structure

An E3.S 2T SSD is designed to enable more flash memory chips when compared to today's U.2 form factor SSD (at 2.5-inch, 15mm Z-height). E3.L variants provide roughly the same amount of flash memory chips as well as ~2x the capacity of today's 2.5-inch drives.

To accommodate today and tomorrow's rising system power, the airflow through the system needs to increase. E3 form factor support enable increased airflow through a system by depopulating supported SSDs from the center of the chassis. This enables increases in overall system airflow while a large number of SSDs can still be utilized to increase storage subsystem performance (Figures 1 and 2).

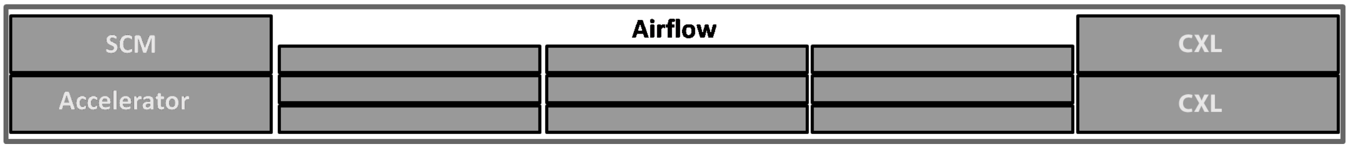

Figure 1 is an example of a 1U system supporting up to 4 alternate device types with 9 SSD slots

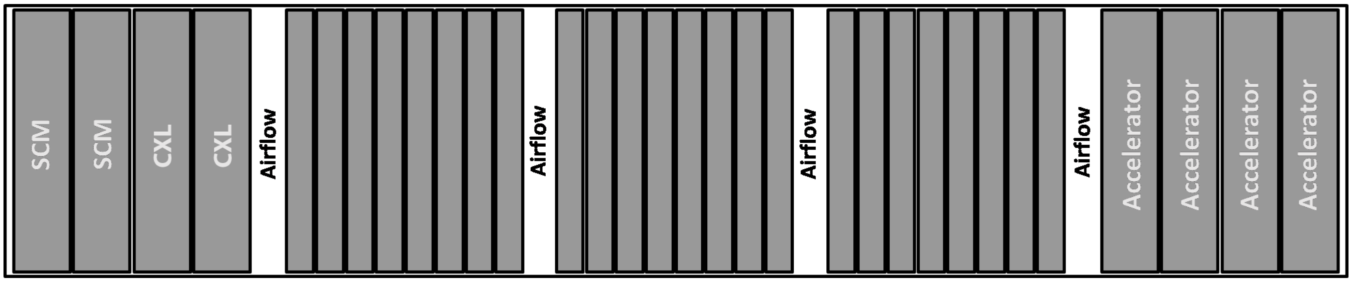

Figure 2 is an example of a 2U system supporting up to 8 alternate device types with 24 SSD slots

The server storage subsystem itself provides yet another challenge for platform architects as traditionally the front of a server has been dedicated to storage devices. With new E3 system designs, future server architectures will share the front server bulkhead space with a multitude of device types including more powerful NVMe SSDs, CXL and SCM devices, computational storage devices, accelerators and front-facing I/O devices such (Figures 1 and 2).

E3 supports multiple mechanical sizes, host link widths and power profiles with a family of interchangeable form factors for next-generation system designs and support for multiple use cases. Its mechanical architecture supports both carrier and non-carrier based implementations. Using a wrap-around carrier, system designers can create a single drive bay mechanics set that enables E3 (thin) and E3 2T (thick) device interchangeability. The ability to optimize around either density, host bandwidth, system power or device type makes the E3 family of form factors the next-generation choice for platform architects and system designers.

E3 Form Factor Benefits

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

Presence Detect Capability
E3 supported SSDs include a presence detect capability so if the E3 SSD is powered off, it can be detected in the system and removed if required. If the SSD needs to be removed, E3 form factors enable supported SSDs to be hot-swapped from the front of the server

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[2], are using E3.S designs on new and upcoming PCIe 5.0 platforms is leading the way in industry-wide support and adoption

E3 Form Factor Industry Support

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 (Figure 3):

Figure 3: EDSFF E3.S prototype system with [KIOXIA CD7 Series PCIe data center NVMe E3.S SSDs](#).
(Source: Dell Technologies and KIOXIA)

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] 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.

TRADEMARKS:

DISCLAIMERS: