KIOXIA

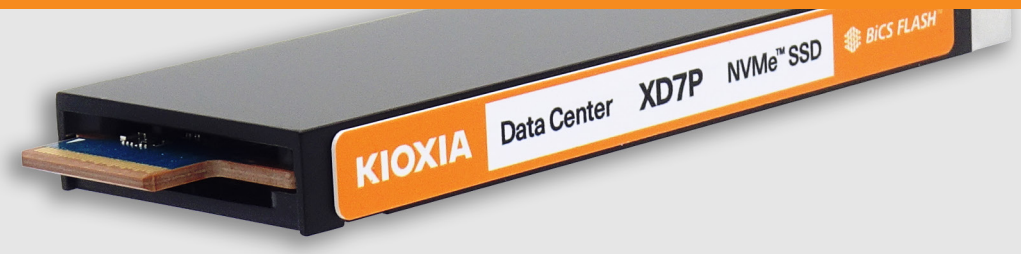

Top 5 Reasons

E1.S Form Factor: The Future of Data Centers

E1.S form factor is ushering in a new era of SSD design. It's crafted exclusively for the abundance of data centers worldwide, not just for hyperscalers, and its potential is limitless. With a streamlined total cost of ownership (TCO), high performance, tailored power, adjustable thermal management, and a purpose-built feature set, E1.S form factor may become the heartbeat of every flash-centric data center in the future.

The main building block of cloud computing is the hyperscale data center - an infrastructure built on a scale many orders of magnitude greater than a standard data center. Hyperscale architectures deliver extremely efficient, scalable and cost-effective platforms that meet today's big data needs. However, many benefits of this compelling value proposition are limited to a small number of cloud providers able to operate at this hyperscale level.

How can standard data centers benefit from the hyperscale value proposition?

For the larger demographic of standard data centers who would like to emulate the hyperscale value proposition, the OCP specification is enabling similar benefits and economies of scale as the large hyperscalers possess. It takes an open source, industry collaborative approach to the data center landscape, addressing storage challenges as previously mentioned.

One of the primary areas of OCP focus is on storage device form factors that, in many cases, could address these storage challenges mentioned above within a data center landscape. OCP standardized on the Enterprise and Datacenter Form Factor (EDSFF) E1.S and E3 form factors developed by the Small Form Factor Technical Affiliate (SFF-TA) working group as part of the Storage Networking Industry Association (SNIA).

EDSFF specifications allow for greater power envelopes, enabling higher performance using the PCI® transport and NVMe™ command set, compared to legacy form factors. From an interface perspective, PCIe 4.0 provides a much larger pipe for data to move through versus the SATA interface. While many infrastructures still utilize SATA-based storage, hyperscalers have deployed NVMe/PCIe 4.0 storage to ensure efficient, cost-optimized solutions.

KIOXIA Corporation provides next-generation PCIe 4.0 data center NVMe SSDs with its XD Series, which is designed to the OCP NVMe Cloud SSD specification and available in E1.S form factors. There are a number of key reasons why these new EDSFF E1.S SSDs can enable a data center server architecture to adapt to changes, but the top five supported by KIOXIA include:

1. **Highest Performance Scaling**
   More than doubles the power budget versus legacy devices, enabling E1.S SSDs to saturate PCIe 4.0 (16 gigatransfers per second (GT/s) x4) performance
2. **Flexible Thermal Solutions**
   Improves interoperability across vendors and platforms while providing the flexibility to select the right balance of cooling and storage density through different E1.S heatsink options
3. **Improved Physical Serviceability**
   Vastly improves serviceability with hot plug support and does not require an entire server to be taken down in order to replace a single SSD
4. **Enables High Density Solutions by Better Accommodating NAND Flash Memory Packages**
   Wider PCB design enables optimized orientation of the NAND flash memory packages, providing headroom for higher capacity drives when compared to legacy devices
5. **Supported by Leading Hyperscalers**
   Meta™ and Microsoft®, leading authors of the OCP NVMe Cloud SSD specification, are using E1.S designs on new and upcoming platforms that has resulted in industry-wide support and adoption

1. Highest Performance Scaling

KIOXIA XD7P Series PCIe 4.0 data center NVMe SSDs are representative of a new category of OCP cloud-optimized SSD designs and are based on the EDSFF E1.S form factor. The XD7P Series product specifications[1] include:

XD7P Series NVMe/PCIe 4.0 Data Center SSD Specifications:

Form Factor	E1.S (9.5/15mm)		
Flash Memory Type	BiCS FLASH™ 3D flash memory		
Interface Specification	PCIe 4.0 (16 GT/s x4), NVMe 2.0		
User Capacities *(in gigabytes[2])*	1,920 GB	3,840 GB	7,680 GB
Performance			
Sequential Read: 128 kibibytes[3] (KiB), Queue Depth (QD) =32	7,200 MB/s*	7,200 MB/s	7,200 MB/s
Sequential Write: 128 KiB, QD =32	3,100 MB/s	4,800 MB/s	4,800 MB/s
Random Read: 4 KiB, QD =256	1,500,000 IOPS	1,540,000 IOPS	1,540,000 IOPS
Random Write: 4 KiB, QD =128	95,000 IOPS	180, 000 IOPS	200,000 IOPS
Power			
Supply Voltage	12V		
Active Power Consumption	up to 20 watts (typical)		
Endurance *(per 5 years)*	1 Drive Write Per Day[4] (DWPD)		
MTTF Reliability	2.0 Million Power-On Hours (MPOH)		
Operating Temperature	0 to 70° C		
Security Option(s)	Self-Encrypting Drive: TCG-Opal 2.0		

**MB/s = megabytes per second;*

Utilizing a purpose-built controller with the latest PCIe 4.0 interface and 3,840 GB capacity, the KIOXIA XD7P Series offers 100% greater improvement in random write performance, 75% greater improvement in random read performance, 104% greater improvement in sequential write performance and 10% greater improvement in sequential read performance over the previous generation (XD6 Series SSDs). The sequential read bandwidth nearly saturates the PCIe 4.0 bus, transferring data at a speed up to 7,200 MB/s (Figure 1).

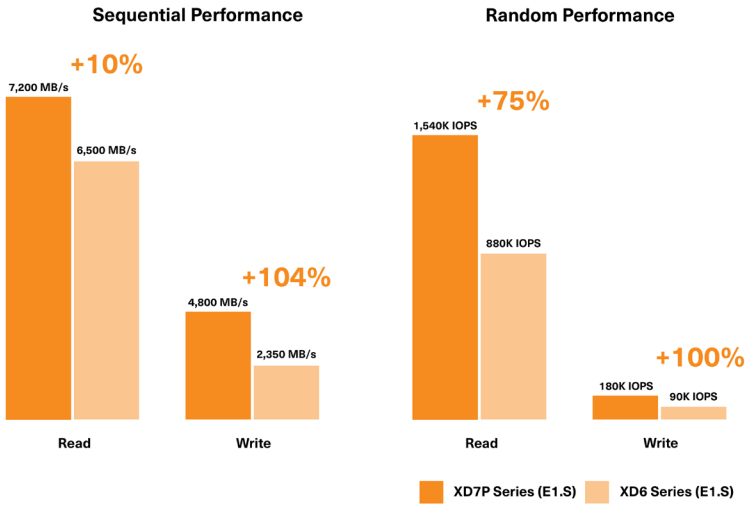

Figure 1: KIOXIA XD7P Series SSD (E1.S) vs. KIOXIA XD6 Series SSD (E1.S) performance comparison

In addition to raw performance, power efficiency is a key metric for data center administrators who care about TCO. The XD7P Series offers PCIe Gen4x4 bandwidth saturation performance while maintaining a less than 20-watt active power draw during these sequential read operations. Additional OCP power states are supported enabling a host server to set the right balance of power to performance for each application.

Depending on the heat sink, capacity and performance requirements, an E1.S SSD can achieve full Gen4x4 bandwidth and even Gen5x4 bandwidth with support for up to a 25W power envelope.

2. Flexible Thermal Solutions

OCP provides the ecosystem with a standardized, open source specification for data center hardware, enabling similar thermal benefits and economies of scale used by the world's largest hyperscalers.

E1.S SSD Type	Width	Length	Thickness
E1.S (9.5mm)	33.75mm	118.75mm	9.5mm
E1.S (15mm)	33.75mm	118.75mm	15mm
E1.S (25mm)	33.75mm	118.75mm	25mm

The KIOXIA E1.S SSDs[5] include:

XD7P Series – 9.5mm[6]

XD7P Series – 15mm[6]

XD7P Series – 25mm[6]

For OCP cloud-specified SSD designs based on the E1.S form factor, the XD7P Series supports a number of thermal management features that monitor drive health and data integrity. It is available in 9.5mm, 15mm and 25mm thick heatsink sizes as defined in the EDSFF E1.S standard, while its performance profile and power envelope are identical across all XD7P Series variants.

Additionally, a thermal throttling mechanism has been implemented to protect XD7P Series drives if a thermal event is detected as defined in the product specification. The thermal throttling is engaged and a '**WCTEMP**' Asynchronous Event Notification (AEN) is issued to the host when the operating range is reached. If the temperature continues to rise, as detected by thermal sensors and reported via the Self-Monitoring, Analysis and Reporting Technology (SMART) health monitoring system, the XD7P Series SSD will throttle its performance, reducing heat that was generated by NAND flash memory die activation, with the goal of reducing the temperature. If the temperature continues to rise in the event of a system fan failure, accidental fire or catastrophic hot conditions, a CCTEMP AEN is issued to the host and the XD7P Series SSD will perform a shutdown to preserve the data that was written to flash memory. After the thermal shutdown has been issued, the SSD will remain in power off mode until power is turned back on and the system cycles.

3. Improved Physical Serviceability

The EDSFF E1.S specification allows for hot swap (also known as hot plug support) which is featured within XD7P Series SSDs. E1.S hot plug support directly resolves physical serviceability concerns associated with previous generation form factors such as M.2/U.2. For example, to replace the older SSDs, normally an entire server would need to be powered down in order to replace a single drive - a process that is time-consuming and potentially challenging to work around without impacting serviceability.

With E1.S SSDs, the drive swap is simple and can occur on the fly. As data center administrators manage an overabundance of systems, the cliché that 'time is money' holds true. The ability to improve physical serviceability is critical for saving time and improving TCO.

4. Enables High Density Solutions by Better Accommodating NAND Flash Memory Packages

Packaging on legacy form factor SSDs include a 22mm wide printed circuit board (PCB) that constrains NAND flash memory placements and limits higher drive capacities. As defined by the EDSFF E1.S specification, the PCB width dimensions have increased to 33.75mm. This wider width provides more space on the PCB for additional flash memory chips, which in turn enables higher capacity SSDs and provides more capacity per allowable space.

When compared to the highest capacity previous form factor SSD generation available[7] at 3.84 TB, E1.S SSDs can support up to 15.36 TB, or a 4x increase in storage capacity.

Supported Capacities: E1.S SSDs

960 GB, 1.92 TB, 3.84 TB, 7.68 TB, 15.36 TB

5. Supported by Leading Hyperscalers

Leading authors of the OCP NVMe Cloud SSD specification, Meta and Microsoft are using E1.S designs on new and upcoming platforms with scalability and versatility relating to capacity, thermal considerations, performance, hot-plug capabilities, etc. These large volume deployments from these key hyperscalers are receiving industry-wide adoption that positions E1.S as a preferred form factor for future storage projects, particularly 1U deployments.

Jason Adrian, Senior Director of Azure® Platform Architecture, Microsoft and OCP Storage Chair

"Microsoft and the OCP storage workgroup demonstrated how an open collaboration across the industry could align hyperscalers, system designers and SSD vendors around next-generation storage form factors. The EDSFF E1.S form factor is the future of flash storage in hyperscale data centers, including Azure platforms. SSDs designed to the OCP NVMe Cloud SSD specification, such as the KIOXIA XD7P Series data center SSDs, will power the next-generation of EDSFF E1.S based servers."

Additional XD7P Series SSD information is available here.

Notes:

[1] Product testing was conducted in a lab environment by KIOXIA Corporation. Tested content are believed to be current and accurate as of the date that the document was published, but is subject to change without prior notice. Read/write sequential and random performance results may vary depending on the host device, read and write conditions, and file 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] KiB: A kibibyte (KiB) means 2^{10}, or 1,024 bytes, a mebibyte (MiB) means 2^{20}, or 1,048,576 bytes, and a gibibyte (GiB) means 2^{30}, or 1,073,741,824 bytes.

[4] Drive Write(s) per Day: One full drive write per day means the drive can be written and re-written to full capacity once a day, every day, under the specified workload for the specified lifetime. Actual results may vary due to system configuration, usage, and other factors.

[5] Each XD7P Series E1.S SSD product image may represent a design model.

[6] The 'millimeter (mm)' size for each XD7P Series E1.S SSD indicates the form factor of the SSD and not its physical size.

[7] Based on the capacity of a leading M.2 SSD with information publicly available at the time of this publication.

Trademarks:

Meta is a trademark of Meta Platforms Inc. Microsoft and Azure are registered trademarks of Microsoft Corporation in the United States and/or other countries. NVMe is a registered or unregistered trademark of NVM Express, Inc. in the United States and other countries. PCIe is a registered trademark of PCI-SIG. All other trademarks or registered trademarks are the property of their respective owners. All other company names, product names and service names may be trademarks or registered trademarks of their respective companies

Disclaimers: