
# Elevate Performance, Scalability and Serviceability of Edge Server Applications with KIOXIA XD7P E1.S Data Center SSDs

Data centers provide centrally-shared access to applications and data using compute, storage and network infrastructures. Edge servers are a key component of today's data centers and reside at the periphery of the network, strategically closer to the devices (and users) they serve. They process and store data locally, reducing communication with a centralized server or cloud infrastructure, which in turn enables low latency, fast response times to queries and higher application/device performance. Localized edge server processing is valuable for real-time data analysis or in use cases where low-latency interactions are required.

Some edge servers utilize M.2 SSDs to store data, mostly due to their compact size, which takes up less space and uses much less power than 2.5-inch[1] SSDs. For edge servers with limited real estate, M.2 SSDs are a good match. However, they lack the performance and capacity capabilities to handle data-intensive applications such as from the Deep Neural Network (DNN) enabling machine learning (ML) for artificial intelligence (AI). Content delivery networks (CDNs) and autonomous vehicles represent other applications that benefit greatly from frequent data exchange at the edge. With challenges associated with M.2 SSDs, another approach to edge computing and storage is required.

This tech brief presents the benefits of an Enterprise and Datacenter Standard Form Factor (EDSFF) E1.S design for edge servers and includes additional benefits in edge storage when KIOXIA XD7P Series E1.S data center SSDs are deployed. Traditionally, M.2 SSDs were chosen as the preferred storage device for edge servers due to their compact sizes and low cost, but due to the evolution of edge applications that require higher performances and more capacity, a large number of E1.S servers have been deployed already, and others are anticipated to launch over the coming year. This tech brief presents an E1.S-based edge system as an alternative to a legacy edge system based on M.2.

## Challenges Associated with M.2-based Edge Servers

M.2-based edge servers are prevalent in use cases where space is limited, power consumption needs to be low and client-grade performance is acceptable. For data-intensive applications, the following represents some of the challenges that these servers face in key use cases:

**[1] Write Latency Performance of Large Datasets**

Edge servers perform data analysis, ML tasks and AI algorithms locally, without having to transmit data to a central server or cloud storage. For these use cases, fast response time is crucial in order to gain the quick insights from large datasets. From internal tests conducted by KIOXIA America, Inc. in a lab environment, M.2 SSD write latency (lower result is better), at a 4 kilobyte[2] (KB) block size, was 0.025 milliseconds (ms) at a queue depth (QD) of 1, 6 ms at a QD of 8, and 13 ms at a QD of 16, per the results below:

| | QD 16 | | | QD 8 | | | QD 1 | | |
|---|---|---|---|---|---|---|---|---|---|
| **Metric** *(in ms)* | **M.2 SSD** | **E1.S SSD** | **E1.S Gain** | **M.2 SSD** | **E1.S SSD** | **E1.S Gain** | **M.2 SSD** | **E1.S SSD** | **E1.S Gain** |
| **Write Latency** | 13 | 0.53 | 24.5x | 6 | 0.159 | 37.7x | 0.025 | 0.006 | 4.1x |
| **Read Latency** | 0.605 | 0.223 | 2.7x | 0.601 | 0.171 | 3.5x | 0.183 | 0.068 | 2.6x |

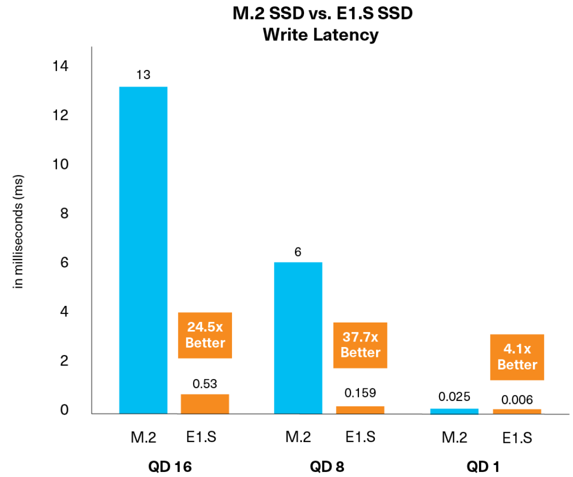


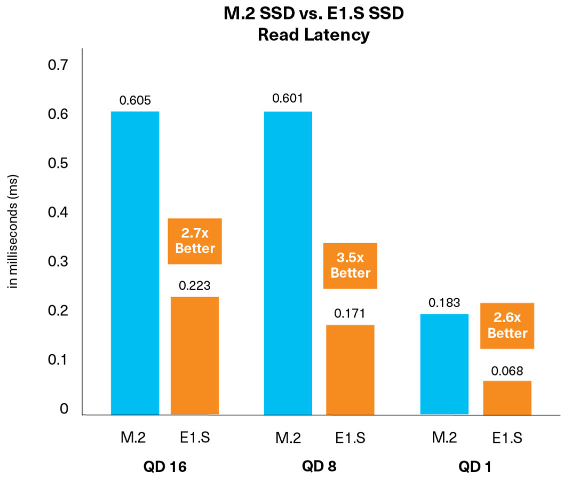


When tested with an E1.S SSD in the same hardware/software configuration at a queue depth of 16, write latency improved by 24.5x. Under a similar scenario with the queue depth changed from 16 to 8, the write latency of the E1.S SSD was 0.159 ms, compared with the M.2 SSD at 6 ms, demonstrating a 37.7x improvement.

In an ML application, edge servers must be able to work quickly and efficiently in order to process and train on the vast amounts of data needed to create an accurate model. This requires compute resources that can handle the various steps of the ML process, such as preprocessing, training and inference. Although compute resources are becoming increasingly faster at performing these steps, they cannot start their work until the data is preprocessed and fetched from the underlying storage system. As such, the underlying storage system must be able to respond to multiple queries very quickly, and with low latency. If the storage system lags behind, expensive compute resources may become underutilized. In order to keep the many cores utilized at all times, the underlying storage system must deliver high sustained throughput so that the compute resources can do training on different pieces of data in parallel. As demonstrated by the latency comparison, the E1.S SSD delivered significant improvements in write latency.

*NOTE: The hardware and software configuration for the latency comparison included the following:*

| Test Server Information | |
|---|---|
| CPU Model | Intel® Xeon® Gold 6444Y |
| Rack Size | 1U |
| No. of Sockets | 2 |
| No. of Cores | 16 |
| Test Software | Flexible I/O[3] |

| SSD Information | | |
|---|---|---|
| SSD Model | Vendor A* | KIOXIA XD7P Series |
| Form Factor | M.2 | E1.S |
| Interface | PCIe® 4.0 | PCIe 4.0 |
| No. of SSDs | 1 | 1 |
| SSD Capacity | 1,024 gigabytes[2] (GB) | 1,920 GB |

**The M.2 SSD used for testing represented one of the highest performing drives in the client SSD market. This module size was 2280 (22 millimeter (mm) wide by 80 mm long), which is widely used in legacy edge servers.*

**[2] Limited Capacities**

The current mainstream capacity of a single-sided M.2 2280[4] SSD is 4 terabytes[2] (TB) or 4,096 GB. Its 22 mm width constrains flash memory placements and limits high drive capacities. With the growing size of media and web content, the need for higher capacity storage devices in edge servers is a requirement for CDNs. The primary purpose of a CDN is to reduce latency, or reduce the delay in communication created by a network's design. Without larger capacities to store this growing content, a group of geographically distributed edge servers will be unable to speed up the delivery of content and bring it closer to where users are located.

**[3] Inefficient Cooling and Thermal Optimization**

Due to unstable environmental conditions[5] that edge server deployments need to endure, more external cooling mechanisms are typically implemented in M.2-based server chassis to provide continuous operations. Without built-in heat sinks, additional heat sinks need to be installed in the M.2-based server chassis as required. There are other measures to reduce ambient temperature[5], such as higher fan speed[6] or bigger vent holes[7], that depend on the environmental requirements of the use case, and have limitations not optimized for today's deployments.

**[4] No Hot-Swap**
An M.2 SSD cannot be hot-swapped from an M.2-based edge server. The lack of hot-swap makes faulty M.2 SSD replacement and field maintenance difficult. To replace an M.2 SSD, the edge server would typically need to be powered down - a time-consuming process that requires removing the system from the rack, removing the faulty M.2 SSD, inserting a new M.2 SSD and reconfiguring the assembly, all of which causes downtime to services.

*NOTE: It is possible to hot-swap M2 SSDs by placing them in a carrier but additional layers of components would still need to be maintained, such as the adaptor board, standalone heat sinks and the carrier itself.*

## An Alternative Paradigm: E1.S Edge Servers

E1.S is a flexible, power-efficient building block for hyperscale and enterprise compute nodes and storage. The basic M.2 form factor was popular in hyperscale data centers due to its low cost, compact size and low power consumption – but has challenges as noted. E1.S form factors solve many M.2 limitations while also maintaining a smaller design with a multitude of heat sink choices. At 33.75 mm wide and 118.75 mm long, an E1.S SSD enclosure is wider and longer than a basic M.2 SSD (22 mm wide and 110 mm long) and can accommodate more flash memory chips that fit vertically in a 1U chassis.

By incorporating an E1.S design in edge servers, the following improvements can be obtained over M.2 platforms:

**[1] High Capacities**
E1.S SSDs support capacities up to 15.36 TB or about 4x the capacity of M.2 SSDs (single-sided). This level of supported capacity is ideal for CDNs that require larger capacities to store growing web and media content as well as speeding up delivery of the content.

**[2] Efficient Cooling and Thermal Optimization**
Each E1.S SSD includes built-in heat sink options that enable selection of thermal profiles that are well suited to cool the edge server and are a good mechanism to balance power versus cooling capabilities in specific applications. A thermal throttling mechanism is also included that detects if the E1.S SSD has reached an operating temperature beyond its specification. If this occurs, E1.S SSD performance will be throttled back and heat generated by flash memory die activation will be reduced. The E1.S packaging is also designed for better system airflow when compared to M.2 SSD packaging.

**[3] Hot-Swap Serviceability**
An E1.S SSD can be hot-swapped from an E1.S-based edge server enabling time-saving and efficient physical serviceability that reduces total cost of ownership. It's an easy field maintenance with plug and play capabilities.

## Introducing KIOXIA XD7P Series E1.S SSDs

Many data-intensive applications require that the underlying storage system be able to respond to multiple queries very quickly, and with low latency. Otherwise, compute resources may become underutilized. The KIOXIA XD7P Series is the latest E1.S generation of PCIe® 4.0 data center NVMe™ SSDs representative of a new category of Open Compute Project® (OCP) cloud-optimized SSDs and based on the EDSFF E1.S form factor. The series includes 9.5 mm and 15 mm E1.S form factor thickness and supports up to 7.68 TB capacities at 1 Drive Write Per Day[8] (DWPD). The XD7P Series performance specifications[9] include:

*KIOXIA XD7P Series SSD[10]: E1.S 15 mm*

| SPECIFICATION | | Units | 1,920 GB | 3,840 GB | 7,680 GB |
|---|---|---|---|---|---|
| **Sequential Read** | *(128 KiB[11]; QD=32; ~16-20 W)* | MB/s | 7,200 | 7,200 | 7,200 |
| **Sequential Write** | *(128 KiB; QD=32; ~16-20 W)* | MB/s | 3,100 | 4,800 | 4,800 |
| **Random Read** | *(4 KiB; QD=256; ~16-20 W)* | IOPS | 1,500,000 | 1,650,000 | 1,550,000 |
| **Random Write** | *(4 KiB; QD=32; ~16-20 W)* | IOPS | 95,000 | 180,000 | 200,000 |

## KIOXIA XD7P Series Features That Enhance Edge Servers

In addition to the benefits of E1.S designs, KIOXIA XD7P Series E1.S SSD performance can enhance edge server operations the following ways:

**Reduced Latency**

The KIOXIA XD7P Series E1.S SSD, when paired with an edge server, demonstrated better overall response times, especially relating to write latency. By narrowing the performance gap between the network card and the E1.S SSD, AI and Internet of Things (IoT) services are enhanced as decisions and complicated issues can be made and dealt with immediately. The image to the right represents examples of AI applications enabled by edge servers that require instantaneous responses to queries.

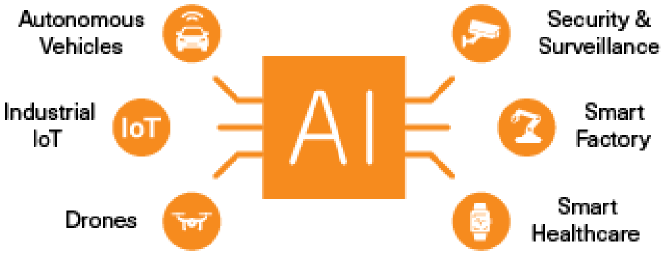


**High Resiliency**

The KIOXIA XD7P Series E1.S SSD hot-swap capability ensures seamless, uninterrupted service for applications such as critical control systems and healthcare. Drive replacements do not require a system shut down as is common for M.2 SSDs.

*High resiliency medical imaging application*
*Image courtesy of Gorodenkoff/Shutterstock.com*

*Hot-swap capabilities with KIOXIA XD7P E1.S SSDs*

**Data Security**

As edge servers can be prone to theft in remote areas, KIOXIA XD7P Series E1.S SSDs are available with a Self-Encrypting Drive[12] (SED) security option, and includes support for the Trusted Computing Group®[13] (TCG) Opal V2 specification. These SSDs are also compliant with the OCP 2.0 Security Project.

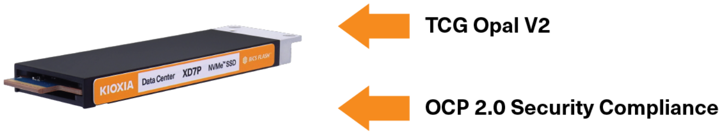


*KIOXIA XD7P Series SSD[10]: E1.S 9.5 mm*

## Summary

As the need for large datasets continues to grow, edge servers equipped with M.2 SSDs face increasing challenges and lack the performance and capacity capabilities to handle data-intensive applications. With inefficient cooling or thermal optimization, and no hot-swap capabilities, these SSDs may not be a good match for large datasets that continue to grow.

In its place, the EDSFF E1.S form factor is well-suited for various enterprise storage applications with a focus on latency performance, especially in use cases where response times are crucial. As evident from the latency comparison, the tested E1.S SSD delivered better read and write latency when compared with the tested M.2 SSD. For write latency at a queue depth of 16, the E1.S SSD was 24.5x better, and at a queue depth of 8, it was 37.7x better.

When deploying storage devices in edge servers, it is essential to consider the operating environment so that easy maintenance can be performed and service can always be available with minimum interruption. E1.S SSDs feature hot-swapping to ensure seamless, uninterrupted service in which drive replacements do not require a system shut down as is common for M.2 SSDs.

Future edge server applications call for higher standards in performance, cooling and thermal optimization, along with higher capacity and better reliability. E1.S SSDs stand out when compared with traditional M.2 SSDs in these areas. KIOXIA is ready to fulfill the needs of edge servers in the future with a full product line-up of E1.S SSDs.

Additional information on E1.S form factors is available here.

Additional information on KIOXIA XD7P Series E1.S data center SSDs is available here.

**NOTES:**

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

[2] Definition of capacity: KIOXIA Corporation defines a kilobyte (KB) as 1,000 bytes, a megabyte (MB) as 1,000,000 bytes, a gigabyte (GB) as 1,000,000,000 bytes, a terabyte (TB) as 1,000,000,000,000 bytes and a petabyte (PB) as 1,000,000,000,000,000 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, 1TB = $2^{40}$ bytes = 1,099,511,627,776 bytes and 1PB = $2^{40}$ bytes = 1,125,899,906,842,624 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] Flexible I/O (FIO) is a free and open source disk I/O tool used both for benchmark and stress/hardware verification. The software displays a variety of I/O performance results, including five corner testing that measures sequential read/write performance, random read/write performance and average latency.

[4] It is possible to achieve an up to 8 TB capacity in an M.2 2280 SSD double-sided configuration, however, this configuration is not common in edge server designs due to signal integrity issues and thermal challenges.

[5] Typical SSD temperature specification is approximately 70 to 85 degrees Celsius. A reasonable ambient temperature of 25 degrees Celsius or below is needed to sustain the drive temperature from exceeding 70 to 85 degrees Celsius.

[6] Higher fan speed can result in more power consumption and noise.

[7] Bigger vent holes have a higher risk to dust and other contaminations. They require the installation of additional filters to prevent dusts and contaminating material from entering the chassis. Additionally, vent hole space competes with other connector outlets resulting in less flexibility to implement all of the necessary connectors, especially in smaller edge servers.

[8] DWPD: 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, for the specified lifetime. Actual results may vary due to system configuration, usage, and other factors.

[9] KIOXIA XD7P Series SSD product specifications provided by KIOXIA Corporation and is accurate as of this publication date.

[10] The product image shown is a representation of the design model and not an accurate product depiction.

[11] KiB: a kibibyte (KiB) means $2^{10}$, or 1,024 bytes.

[12] SED (Self-Encrypting Drive) encrypts/decrypts data written to and retrieved from an SSD via a password-protected alphanumeric key, continuously encrypting and decrypting the data. Optional security compliant drives are not available in all countries due to export and local regulations.

[13] The Trusted Computing Group (TCG) is a not-for-profit international standards organization that applies hardware-based encryption to solid state drives.

**TRADEMARKS:**

Intel and Xeon are registered trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. NVMe is a registered or unregistered trademark of NVM Express, Inc. in the United States and other countries. Open Compute Project is a registered trademark of the Open Compute Foundation. PCIe is a registered trademark of PCI-SIG. Trusted Computing Group is a registered trademark of Trusted Computing Group. All other company names, product names and service names may be trademarks of third-party companies.

**DISCLAIMERS:**