KIOXIA

Application Brief

KIOXIA XD7P Series PCIe® 4.0 EDSFF E1.S SSDs Accelerate Database Storage Performance and Improve Data Center Density Challenges

Featuring KIOXIA XD7P Series PCIe 4.0 Data Center E1.S SSDs deployed in a Supermicro® SSG-121E-NES24R E1.S Server, running Apache Cassandra® Database Software

Introduction

Data centers deploy many applications that generate large volumes of data. This abundance of data has become more difficult to manage and quickly access, and less room is available in systems for storage media. Data intensive applications demand increased storage requirements especially relating to performance and capacity. Database applications must be accessed in real-time to deliver the high Quality of Service (QoS) experiences that users expect. As such, today's data centers require high performing, dense, thermally efficient storage capabilities that address the storage challenges associated with legacy form factors.

Hence, with the recent availability of Enterprise and Datacenter Standard Form Factor[1] (EDSFF) SSDs, storage performance and total capacity per server can increase when compared with legacy hard drive interfaces and protocols. Servers with EDSFF 1U short (E1.S) slots, deployed with E1.S SSDs, deliver fast data throughput and input/output operations per second (IOPS) performance, low latency, high density and thermally optimized storage capabilities.

This application brief presents E1.S server/storage performance relating to database throughput and latency as tested in a lab environment. It also presents CPU utilization where high performance does not rely solely on system memory. The E1.S test set-up includes a Supermicro SSG-121E-NES24R server running a NoSQL[2] Apache Cassandra database cluster driven by synthetic tests. There are five workload tests in the comparison that includes three 3.84 terabyte[3] KIOXIA XD7P Series E1.S SSDs and three SSDs from Vendor A with the same capacities. Appendix A covers the hardware and software configuration – Appendix B covers the configuration set-up and test procedures.

Apache Cassandra is a highly regarded NoSQL database that delivers fast write performance. The brief includes test results from write-based insert operations. The KIOXIA XD7P Series E1.S SSD and the Vendor A E1.S SSD comparison featured synthetic benchmarks from Yahoo!™ Cloud Serving Benchmark (YCSB) software. The test database created on the Apache Cassandra cluster consisted of three hundred million records/operations (and regarded as the 100% insert[4]).

The tests measured: (1) Database Throughput; (2) Average Insert Latency including at the 99th and 95th percentiles; and (3) CPU Utilization. KIOXIA XD7P Series E1.S SSDs featured the PCIe 4.0 interface and NVMe™ 2.0 protocol, and tested at PCIe 4.0 x4 speed.

Test Results Snapshot

The KIOXIA XD7P Series PCIe 4.0 E1.S SSDs deliver exceptional NoSQL database performance when compared with PCIe 4.0 E1.S SSDs from Vendor A as depicted by the test results:

Database Throughput:
23% Higher

Average Insert Latency:
14% Lower

95th Percentile Insert Latency:
12% Lower

99th Percentile Insert Latency:
4% Lower

CPU Utilization:
3% Better Usage

Test results show that the KIOXIA XD7P Series E1.S SSDs demonstrate higher throughput, lower latency and better CPU usage of NoSQL database workloads than the Vendor A E1.S SSDs in a Supermicro SSG-121E-NES24R E1.S server environment.

The test results presented include a brief description of each workload test, a **graphical depiction** of the test results and a test analysis.

TEST RESULTS

Workload 1: Database Throughput (operations per second)

This workload test represents the number of operations per second that a system can complete on average and includes three hundred million records/operations inserted into the database. The test results below are in operations per second. The highest result represents the fastest throughput.

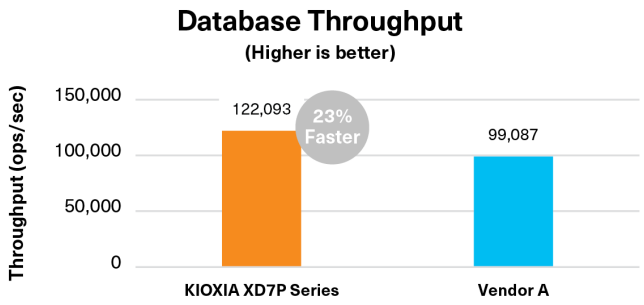

Test Analysis:

The KIOXIA XD7P Series SSDs demonstrated 23% faster times to complete database operations versus Vendor A SSDs. These results also indicate that the number of incoming operations is much higher than the achievable database throughput. If this occurs, the server can get overloaded creating longer wait times per operation that can negatively affect application performance and the user experience.

Workload 2: Average Insert Latency

This workload test represents the time it takes to perform a database insert operation. It includes the average time it takes for the YCSB workload generator to issue the insert operation and for the operation to complete successfully. The lowest result represents the best latency. The test results are in milliseconds (ms).

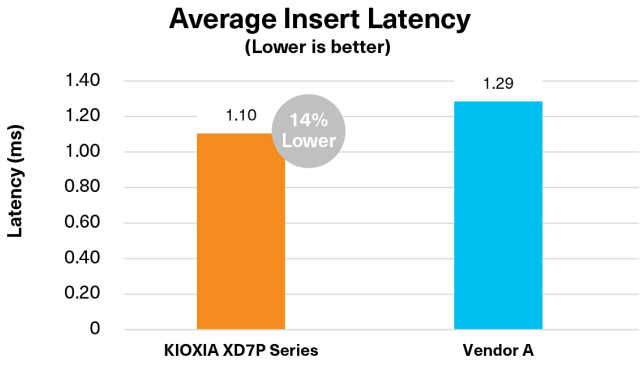

Test Analysis:

The KIOXIA XD7P Series SSDs demonstrated 14% lower latency to complete database insert operations versus Vendor A SSDs. Low latency positively affects database performance and application response times that can translate into better user experiences.

Workload 3: Average Insert Latency at the 95th Percentile

This workload test represents the time it takes to perform an insert database operation at the 95th percentile, or the point at which 95% of the database workload traffic is experiencing latency that is less than the 95% value. An alternative way to look at this metric is that 5% of the database traffic is experiencing values that are out of range. The lowest result represents the best latency. The test results are in milliseconds.

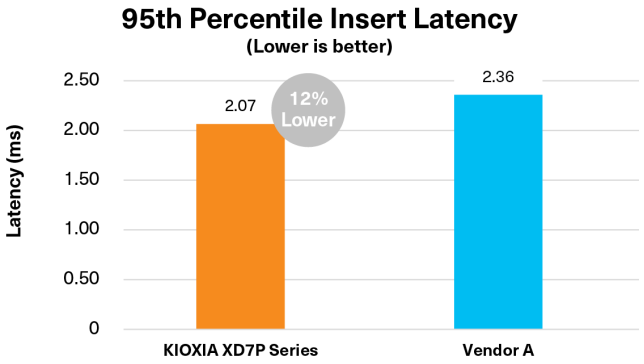

Test Analysis:

The KIOXIA XD7P Series SSDs demonstrated 12% lower latency to complete database insert operations at the 95th percentile versus Vendor A SSDs. Low latency positively affects database performance and application response times that can translate into better user experiences.

Workload 4: Average Insert Latency at the 99th Percentile

This workload test represents the time it takes to perform an insert database operation at the 99th percentile, or the point at which 99% of the database workload traffic is experiencing latency that is less than the 99% value. An alternative way to look at this metric is that 1% of the database traffic is experiencing values that are out of range. The lowest result represents the best latency. The test results are in milliseconds.

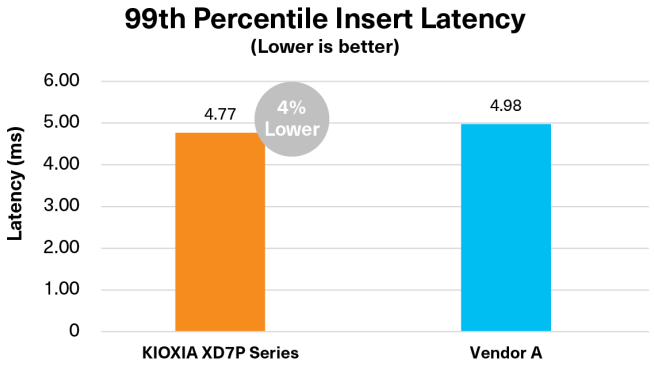

Test Analysis:

The KIOXIA XD7P Series SSDs demonstrated 4% lower latency to complete database insert operations at the 99th percentile versus Vendor A SSDs. Low latency positively affects database performance and application response times that can translate into better user experiences.

Workload 5: CPU Utilization

This workload test represents the percentage of the CPU cycles used for a given workload and measured to ensure that the server CPUs were efficiently processing the YCSB workloads. Low utilization means that the CPUs were not used efficiently, which could result in an underutilization of server capabilities and stranded compute resources. High CPU utilization indicates the CPUs are using its resources better. The test results are in percentages.

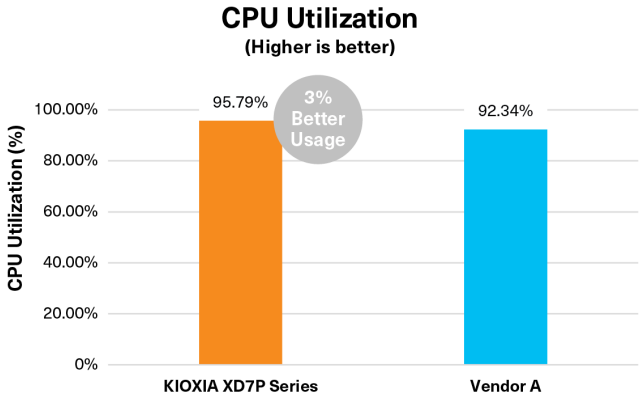

Test Analysis:

The KIOXIA XD7P Series SSDs demonstrated 3% better CPU usage versus Vendor A SSDs. With high CPU utilization, the benchmark workloads were efficiently processed and the CPUs used their resources better in the KIOXIA XD7P Series configuration.

Summary

Though a number of server and storage vendors have announced plans to support E1.S platforms, there are very few solutions available as of this publication date. The Supermicro SSG-121E-NES24R E1.S server and KIOXIA XD7P Series E1.S SSDs represent two available E1.S solutions that when combined, are a good match for database applications as showcased by the test results. This E1.S server/storage combination demonstrated fast database throughput, low database insert latencies and efficient use of the system CPUs better than the Vendor A E1.S SSD configuration as follows:

E1.S SSD Type	Number of Threads	Database Throughput *(in ops/sec) (higher is better)*	Avg. Insert Latency *(in ms) (lower is better)*	Avg. Insert Latency: 95th Percentile *(in ms) (lower is better)*	Avg. Insert Latency: 99th Percentile *(in ms) (lower is better)*	CPU Utilization *(in %) (higher is better)*
KIOXIA XD7P Series	128	122,093	1.10	2.07	4.77	95.79
Vendor A	128	99,087	1.29	2.36	4.98	92.34
XD7P Advantage		**23%**	**14%**	**12%**	**4%**	**3%**

Database throughput is important with respect to storage density, mostly for determining how many users that a single server can support. In other

words, higher throughput enables more operations per server with the same storage density. With the increased throughput of the XD7P Series SSDs (versus Vendor A SSDs) on an identical server, data center operators can reduce the number of servers used when handling the same database user workloads more efficiently than previous configurations, and thereby reducing costs.

Latency performance is another important database metric and often used in determining the end user experience. Since many NoSQL databases are the backend for web services, faster responses from the database make the applications appear snappier to end customers. The end user experience is a critical measure as a key revenue component for purchase-based websites. With fast database throughput, coupled with low database insert latency, the XD7P Series SSDs may enable more users per server, optimized resource utilization across the data center, and a better overall end user experience when compared to E1.S SSDs from Vendor A.

KIOXIA XD7P Series SSD Product Info

The KIOXIA XD Series product line aims to provide better total cost of ownership for customers. The latest generation KIOXIA XD7P Series PCIe 4.0 data center NVMe SSDs are 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 widths and supports up to 7,680 gigabyte[3] (GB) capacities at 1 Drive Write Per Day[5] (DWPD). The XD7P Series performance specifications[6] include:

KIOXIA XD7P Series SSD[7]: E1.S 15 mm

SPECIFICATION	Units	1,920 GB (up to)	3,840 GB (up to)	7,680 GB (up to)
Sequential Read *(128 KB; QD=32; 20W)*	MB/s	7,200	7,200	7,200
Sequential Write *(128 KB; QD=32; 20W)*	MB/s	3,100	4,800	4,800
Random Read *(4 KB; QD=256; 20W)*	IOPS	1,500,000	1,540,000	1,540,000
Random Write *(4 KB; QD=32; 20W)*	IOPS	95,000	180,000	200,000
Random Read Latency *(QD=1; 20W)*	µs	70	70	70
Random Write Latency *(QD=1; 20W)*	µs	10	10	10

MB/s = megabytes per second.
µs = microseconds.

Additional KIOXIA XD7P Series E1.S data center SSD specifications and information is available here.

Appendix A

Hardware/Software Test Configuration

Server Information	
Server	Supermicro SSG-121E-NES24R
No. of Servers	1
Number of CPU Sockets	2
CPU	Intel® Xeon® Gold 6444Y
No. of CPU Cores	16
CPU Frequency	3.6 GHz
Total Memory	128 GB DDR-5 DRAM
Memory Frequency	4.8 GHz
Operating System Information	
Operating System	Ubuntu®
Version	22.04.2
Kernel	5.15.0-76-generic

SSD Information		
Model	KIOXIA XD7P Series	Vendor A
Interface	PCIe 4.0 x4	PCIe 4.0 x4
Number of Devices	3	3
Form Factor	E1.S (15 mm)	E1.S (15 mm)
Capacity	3.84 TB	3.84 TB
DWPD	1 (5 years)	1 (5 years)
Active Power	up to 20 W	up to 16.6 W

Docker[8] Information	
Client Docker™ Engine Version	24.0.2
Server Docker Engine Version	24.0.2
containerd[9] Version	1.6.21
runC[10] Version	1.1.7
docker-init[11] Version	0.19.0

Apache Cassandra Information	
Version	4.1.2

Appendix B

Configuration Set-up / Test Procedures

Set-up

- A three-node Apache Cassandra cluster, with a replication factor of two, was set up by creating three containers through Docker.
- Each container deployed one KIOXIA XD7P Series SSD (E1.S, 15 mm) with 3.84 TB capacity to store its database on.
- A test database was created on the Cassandra application cluster using the YCSB synthetic benchmark tool. The database included three hundred million records.

Procedures

- The database tests were conducted using the three KIOXIA XD7P SSDs and three hundred million operations, and results were recorded for:
  - **Database Throughput** (in operations/second)
  - **Average Insert Latency** (in us)
  - **Average Insert Latency at the 95th Percentile** (in us)
  - **Average Insert Latency at the 99th Percentile** (in us)
  - **CPU Utilization** (in %)
- Set-up and test procedures were configured with the three Vendor A E1.S SSDs.
- The five database tests were run and the results were recorded.
- The KIOXIA XD7P Series SSD results were compared with the Vendor A SSD results.

NOTES:

[1] Developed by the Small Form Factor Technical Affiliate (SFF-TA) working group as part of the Storage Networking Industry Association (SNIA).

[2] A NoSQL database is a non-tabular database that stores data differently than traditional relational tables and provide flexible schemas that scale easily with large amounts of data and high user loads.

[3] Definition of capacity - KIOXIA Corporation defines 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.

[4] A 100% insert workload is common for such use cases as transaction logging, tracking systems, logistics, asset management, and others.

[5] Drive Write(s) per Day (DWPD): 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.

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

[7] The KIOXIA XD7P Series SSD product image may represent a design model.

[8] Docker is an open source platform that enables developers to build, deploy, run, update and manage containers.

[9] Containerd is a Docker-developed container runtime that manages the lifecycle of a container on a physical or virtual machine (i.e., a host). It creates, starts, stops and destroys containers.

[10] RunC is a lightweight, portable container runtime that includes all of the plumbing code used by Docker to interact with system features related to containers and designed with security in mind.

[11] Docker init is a command line interface (CLI) command that simplifies the initialization process of running new projects in a Docker container.

TRADEMARKS:

DISCLAIMERS: