



Application Brief

New KIOXIA XD7P Series E1.S SSDs Accelerate Database Storage Performance and Improve Density Challenges for Apache Cassandra® Databases

Tested on a Supermicro® SSG-1029P-NES32R E1.S Server with KIOXIA XD Series E1.S SSDs

With the recent availability of E1.S Enterprise and Datacenter Standard Form Factor¹ (EDSFF) enabled servers and SSDs, the data center objective of achieving higher performance and increased storage density is now more attainable. Today's data center applications generate tons of data that have increased storage requirements, and for applications such as databases, the data needs to be accessed in real-time to deliver high-quality user experiences. The massive amount of data generated in a data center has become more difficult to manage, more space-limited to store and more challenging to access quickly, creating a need for dense, fast and thermally-efficient storage capabilities which help to reduce exorbitant data center cooling costs. With the recent advent of the E1.S form factor, a server with E1.S SSDs is increasing the total capacity per server while providing a more dense and thermally-optimized solution.

E1.S-supported servers and SSDs are enabled by the PCIe® interface and NVMe® protocol – a combination proven to deliver fast data throughput and input/output operations per second (IOPS) performance, as well as low latency, when compared to legacy hard drive interfaces and protocols. Though many server and storage vendors are developing E1.S form factor platforms, there are few solutions available as of this publication date. The Supermicro SSG-1029P-NE32R server and KIOXIA XD Series SSDs represent two currently available E1.S solutions.

This application brief presents E1.S server/storage performance results relating to database throughput and latency as tested in a lab environment, and showcases CPU utilization benefits and how high performance is delivered without relying solely on system memory. Although many applications can benefit from this powerful E1.S solution, these results demonstrate potential system performance when running a NoSQL² Apache Cassandra database application driven by synthetic tests.

Test Environment

The focus of the test was to determine the database write performance in the SSDs. Apache Cassandra was the application used as it is a highly-regarded NoSQL database that delivers fast write performance. To test SSD write performance, write-based insert operations were used that adds new data to the database.

To demonstrate the write performance of KIOXIA XD Series E1.S SSDs in a Supermicro SSG-1029P-NES32R server, the system was tested with synthetic benchmarks from Yahoo! Cloud Serving Benchmark (YCSB) software. A test database was created on an Apache Cassandra application cluster consisting of three hundred million records and regarded as a 100% insert. It should be noted that a 100% inserted workload is common for such use cases as transaction logging, tracking systems, logistics, asset management, and others. The key test components included:

- Supermicro SSG-1029P-NES32R E1.S server
- Apache Cassandra NoSQL database software
- YCSB synthetic test software
- KIOXIA XD7P Series E1.S SSDs (2nd generation)
- KIOXIA XD6 Series E1.S SSDs (1st generation)

For testing purposes, one database test was created and a workload was generated that consisted of three hundred million insert operations to test the database performance of the Apache Cassandra NoSQL database. The test measured: (1) Operations throughput; (2) Average insert latency including at the 99th and 95th percentiles; and (3) CPU utilization. The E1.S SSDs tested included KIOXIA's first generation XD6 Series and second generation XD7P Series to demonstrate performance improvements between the generations. As the Supermicro SSG-1029P-NES32R E1.S server is PCIe 3.1-capable, XD6/XD7P Series SSDs were tested at PCIe Gen3 speed though both SSD families support PCIe Gen4 x4 performance.

Description of Benchmark Tests

The hardware and software equipment used for these benchmark tests included the following:

Test Criteria

- **Supermicro SSG-1029P-NES32R Server:** One (1) dual socket server with two (2) Intel® Xeon® Gold 622R processors, featuring 16 processing cores, 2.9 GHz frequency, and 192 gigabytes³ (GB) of DDR4
- **Operating System:** Ubuntu® v20.04.4 (kernel 5.4.0-124-generic)
- **Application:** Apache Cassandra v4.0.5
- **Benchmark Software:** Synthetic tests run through YCSB software (version 0.17.0)
- **Docker® Containers:**
Client Docker Engine (v20.10.17)
Server Docker Engine (v20.10.17)
'containerd' (v1.6.7)
'runc' (v1.1.3)
'docker-init' (v0.19.0)
- **Storage Devices (Table 1):** Three (3) KIOXIA XD7P Series E1.S 15mm SSDs with 3.84 terabyte³ (TB) capacities and three (3) KIOXIA XD6 Series E1.S 15mm SSDs with 3.84 TB capacities

Tested Specifications	XD7P Series	XD6 Series
Interface	PCIe 3.0 (x4)	PCIe 3.0 (x4)
Capacity	3.84 TB	3.84 TB
Form Factor	E1.S (15 mm)	E1.S (15 mm)
NAND Flash Type	BiCS FLASH™ TLC	BiCS FLASH™ TLC
Drive Writes per Day ⁴ (DWPD)	1 (5 years)	1 (5 years)
Operating Power	14W	14W

Table 1: Tested specifications for KIOXIA XD Series SSDs

Set-up & Test Procedures

The test system was configured using the hardware and software equipment outlined above. The Apache Cassandra application was configured to store both the log and database on each SSD tested. A three-node Apache Cassandra cluster, with a replication factor of 2, was setup by creating three containers through Docker software. Each container deployed one XD7P SSD to store its database. A test database was created on the Apache Cassandra application cluster using the YCSB synthetic test software that consisted of three hundred million records.

The database test was conducted and the workload of three hundred million operations was run against the Cassandra database cluster. The system recorded the following test results:

- Database Throughput in operations per second
- Average Insert Latency in microseconds (μ s)
- Average Insert Latency in μ s at the 95th percentile
- Average Insert Latency in μ s at the 99th percentile
- CPU Utilization in percentages


The tests above were repeated for the previous XD6 Series generation and compared to the XD7P Series results.

Test Results

The database throughput and latency benchmarks were conducted with the test results recorded. The CPU utilization benchmark was run in conjunction with the database throughput benchmark with results recorded to ensure that the server CPUs were efficiently processing the Apache Cassandra database workloads. For database throughput, the higher the value, the better the result. For the insert latency results, the lower the value, the better the result.

Database Throughput:

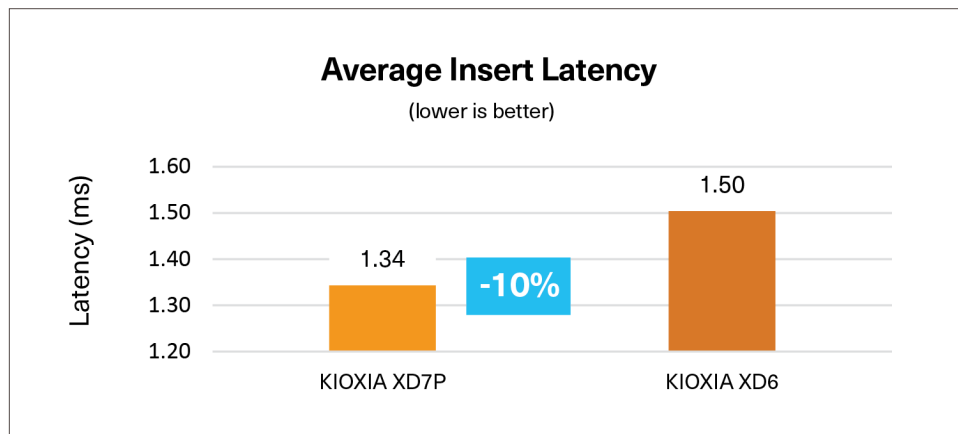
As the 300 million records were inserted into the database, these results represent the number of operations per second a system can complete on average. As it relates to database throughput, the results also measure how quickly a server can process incoming database operations and indicates if the number 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 impact application performance, as well as the user experience. The test results below are in operations per second.



As shown, the XD7P Series delivered 94,825 operations per second which is 11% faster than the previous XD6 Series generation.

Average Insert Latency:


These results represent the time it takes to perform an insert database operation. It includes the average time it takes for the YCSB workload generator to issue the insert operation and for the operation to successfully complete. Low latency positively impacts database performance and application response times that can translate into a better user database experience. The test results below are in milliseconds.



As shown, the XD7P Series delivered an average insert latency of 1.34 ms which is a little more than 10% faster in accessing written data than the previous XD6 Series generation.

Average Insert Latency at the 95th Percentile


These results represent 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 test results below at the 95th percentile are in milliseconds.



As shown, the XD7P Series delivered an average insert latency at the 95th percentile of 2.42 ms which is more than 36% faster in accessing written data than the previous XD6 Series generation.

Average Insert Latency at the 99th Percentile

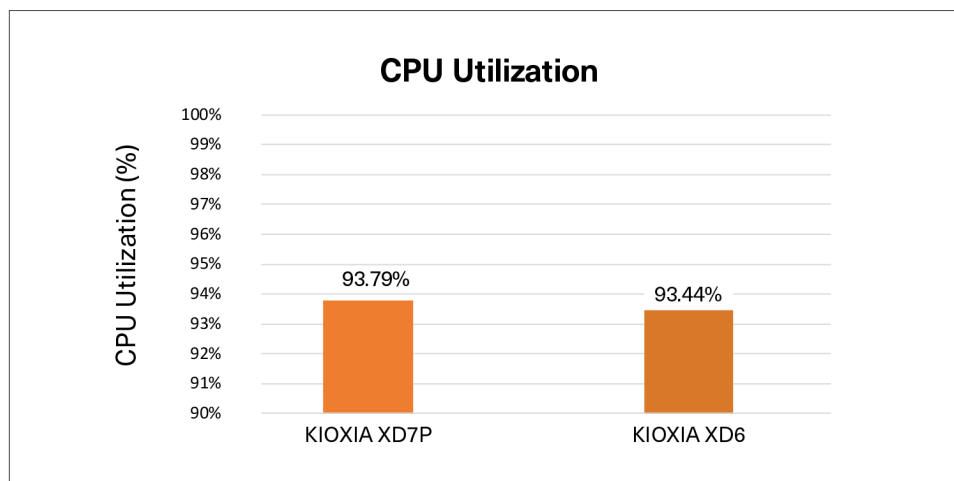
These results represent the time it takes to perform an insert database operation at the 99th percentile, or the point at which 99% of the database workload traffic is experiencing latency that is less than the 99% value. An alternative way to look at this metric is that 1% of the database traffic is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



As shown, the XD7P Series delivered an average insert latency at the 99th percentile of 4.92 ms which is more than 29% faster in accessing written data than the previous XD6 Series generation.

CPU Utilization

These test results represent a percentage of the CPU cycles that were being used for the Apache Cassandra workload and were measured to ensure that the server CPUs were efficiently processing the database workloads.



The CPU utilization for the XD7P Series was 93.79%, and the XD6 Series was 93.44%, indicating very little difference in CPU utilization while delivering higher throughput and lower latency.

NOTE: The test workload was a purely synthetic workload which dramatically increases the CPU utilization, whereas an actual production workload with a NoSQL database may have similar performance, but with much lower CPU utilization.

Test Analysis

The KIOXIA XD7P Series delivered higher write throughput and lower insert latencies than the previous XD6 Series, and both demonstrated that these E1.S SSDs are well-suited for database applications.

Data throughput is important for storage density, and 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 the previous XD6 Series generation) on an identical server, the test results show that data center operators can reduce the number of servers used when handling the same database user workload.

Latency performance is another important database metric and is often used in determining the end user experience. As many NoSQL databases are the back-end for web services, the faster the response time from the database, the snappier the application appears to the end customer. The end user experience is a critical measure that often influences how much time is spent by the users, but more importantly, is a key revenue component for purchase-based websites. With an increase in data throughput coupled with lower latency, the XD7P Series SSDs enable more users per server, optimized resource utilization across the data center, and a better overall end user experience when compared to the previous generation.

Summary

Though a number of server and storage vendors have announced plans to support E1.S platforms, there are very few solutions available today. The Supermicro SSG-1029P-NES32R E1.S server and KIOXIA XD Series E1.S SSDs represent an E1.S server/storage solution that is currently available and well-suited for database applications as showcased by these test results. This E1.S server/storage combination demonstrated very high database throughput and low insert latencies as follows:

XD Series (E1.S SSDs)	No. of Threads	DB Throughput (ops/sec) (higher is better)	Insert Latency (µs) (lower is better)	Insert Latency (µs) (lower is better)	Insert Latency (µs) (lower is better)	CPU Utilization (%) (higher is better)
XD7P	128	94,925	1.34	2.42	4.92	93.79
XD6	128	84,769	1.50	3.84	6.98	93.44
XD7P Advantage		+11%	-10%	-36%	-29%	~0%

This Supermicro and KIOXIA database solution can support up to 6 TB of system memory and up to 32 KIOXIA XD7P Series or XD6 Series SSDs. As the XD7P Series is available in 1,920 GB, 3,840 GB and 7,680 GB storage capacities, the solution can hold up to 96 TB in the 1U Supermicro SSG-1029P-NES32R E1.S server. Hot-swap capabilities are also featured making this E1.S solution a good match for hyperscale-class and enterprise-class applications.

Additional Supermicro SSG-1029P-NES32R E1.S server information is available [here](#).
 Additional XD7P/XD6 Series SSD information is available or soon to be available [here](#).

KIOXIA XD Series Specification Snapshot:



KIOXIA XD7P Series E1.S SSDs⁵

Sequential Read: up to 7,340 MB/s
Random Read: up to 1,500,000 IOPS
Sequential Write: up to 4,718 MB/s
Random Write: up to 200,000 IOPS

Capacities: 7,680/3,840/1,920 GB
 DWPD: 1 (5 years)

KIOXIA XD6 Series E1.S SSDs⁶

Sequential Read: up to 6,500 MB/s
Random Read: up to 880,000 IOPS
Sequential Write: up to 2,350 MB/s
Random Write: up to 90,000 IOPS

Capacities: 3,840 GB / 1,920 GB
 DWPD: 1 (5 years)

- 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 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 = 2³⁰ bits = 1,073,741,824 bits, 1GB = 2³⁰ bytes = 1,073,741,824 bytes and 1TB = 2⁴⁰ 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.
 - 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, for the specified lifetime. Actual results may vary due to system configuration, usage, and other factors.
 - ⁵ The KIOXIA XD7P Series E1.S SSD product image may represent a design model.
 - ⁶ The KIOXIA XD6 Series E1.S SSD product image may represent a design model.

TRADEMARKS:
 Apache and Apache Cassandra are registered trademarks of the Apache Software Foundation (ASF). Docker is a registered trademarks of Docker, Inc. in the United States and/or other countries. Intel and Xeon are registered trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. NVMe is a registered trademark of NVM Express, Inc. PCIe is a registered trademark of PCI-SIG. Supermicro is a registered trademark of Super Micro Computer, Inc. or its subsidiaries in the United States and other countries. Ubuntu is a registered trademark of Canonical Ltd. All other trademarks or registered trademarks are the property of their respective owners.

DISCLAIMERS:
 KIOXIA America, Inc. may make changes to specifications and product descriptions at any time. The information presented in this performance brief is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors. Any performance tests and ratings are measured using systems that reflect the approximate performance of KIOXIA America, Inc. products as measured by those tests. Any differences in software or hardware configuration may affect actual performance, and KIOXIA America, Inc. does not control the design or implementation of third party benchmarks or websites referenced in this document. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to any changes in product and/or roadmap, component and hardware revision changes, new model and/or product releases, software changes, firmware changes, or the like. KIOXIA America, Inc. assumes no obligation to update or otherwise correct or revise this information.

KIOXIA America, Inc. makes no representations or warranties with respect to the contents herein and assumes no responsibility for any inaccuracies, errors or omissions that may appear in this information.
 KIOXIA America, Inc. specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. In no event will KIOXIA America, Inc. be liable to any person for any direct, indirect, special or other consequential damages arising from the use of any information contained herein, even if KIOXIA America, Inc. is expressly advised of the possibility of such damages.

© 2022 KIOXIA America, Inc. All rights reserved.