KIOXIA

Solution Brief

Solving Storage Performance and Density Challenges with New E1.S Solutions: Supermicro® SSG-1029P-NES32R Server with KIOXIA XD6 Series SSDs

For MongoDB® and Applications Requiring High Performance at Very Low Latency and Dense Storage

Several applications generate large quantities of data that have increased storage requirements, and in turn require increased speed to access this data. IT administrators are tasked with delivering fast processing and storage performance to support these applications and workloads creating a need for dense and fast storage solutions that can keep up with the large amount of data generated. The higher volume of data created in a data center has become more difficult to manage, space-limited to store and more challenging to access quickly. With the recent availability of E1.S Enterprise and Datacenter Standard Form Factor[1] (EDSFF) servers and SSDs, the data center objective of higher performance and increased storage density has become more attainable.

E1.S form factor servers and SSDs are driven by the PCIe® interface and NVMe® protocol – a combination proven to deliver fast data throughput and input/output operations per second (IOPS) performance, and 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-NES32R is a new and available E1.S server solution that enables very dense storage installations with potentially greater density[2] versus media deployed with 2.5-inch or 3.5-inch[3] drives. It features KIOXIA XD6 Series PCIe 4.0 data center NVMe E1.S SSDs that are designed to the Open Compute Project (OCP) Datacenter NVMe SSD Specification and deliver fast read performance and a rich feature set.

This solution brief presents E1.S server/storage performance relating to latency and throughput, and showcases high performance 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 NoSQL[4] MongoDB database application driven by synthetic tests.

Performance and CPU Utilization with E1.S Form Factors

Database Creation, Operations and Performance Benchmarks

To demonstrate the performance of XD6 Series SSDs in a Supermicro SSG-1029P-NES32R server, the system was tested with synthetic benchmarks in a KIOXIA lab environment. A test database was created on the MongoDB application consisting of three hundred million records, as well as synthetic tests that were run through Yahoo! Cloud Serving Benchmark (YCSB) software.

There was one database test conducted where a workload of two billion operations (50% reads and 50% updates) was run against the NoSQL MongoDB database. A 50% read/50% update workload is fairly common for database applications. The one database test measured six metrics: (1) Run Time; (2) Average Read Latency; (3) Average Update Latency; (4) Operations Throughput; (5) Efficient CPU Utilization; and (6) Memory Utilization. As Supermicro SSG-1029P-NES32R E1.S server CPUs are PCIe 3.1-capable, the XD6 Series SSD was tested at PCIe Gen3 speed though they support PCIe Gen4 performance. The results of the synthetic tests yielded the following:

Run Time

This metric represents the total time required in hours to complete the workload consisting of two billion read and update operations against the MongoDB database. This measurement is heavily dependent on database size. As databases become larger, the amount of time it takes to successfully perform a read operation can increase as the system has to query the vast records stored in the database and find all of the data it needs.

For update operations, the system needs to traverse the entire database of records until it finds the record that needs to be updated, and then performs the necessary changes to the data. The low latency and high performance that the SSD provides enable the database to retrieve the records quickly from the SSD, which is critical in reducing the time it takes for each individual read and update operations to complete. With a database consisting of three hundred million records, the system took 18.33 hours to complete[5] two billion operations within the database workload.

Run Time
18.33 hours

Average Read Latency

This metric represents the time it takes to perform a read database operation. It includes the average time it takes for the YCSB workload generator to issue the read operation, and for the operation to successfully complete. This metric can positively impact database performance and application response times that can translate into a better user database experience. The average read latency delivered was 5.54 milliseconds (ms)[6] indicating very fast database performance.

Average
Read Latency
5.54 ms

Average Update Latency

This metric represents the time it takes to perform an update database operation. It includes the average time it takes for the YCSB workload generator to issue the update operation, and for the operation to successfully complete. This metric can also positively impact database performance and application response times that can translate into a better user database experience. The average update latency delivered was 11.34 ms[7] also indicating very fast database performance.

Average
Update Latency
11.34 ms

Operations Throughput

This metric represents the number of operations per second a system can complete on average. As it relates to database throughput, it also measures how quickly a server can process incoming database queries. It discerns if the number of incoming queries is much higher than the achievable database throughput. If this occurs, the server can overload creating longer wait times per query, and negatively impact application performance and the user experience. Operations throughput is especially important when a mix of operations, such as read and update operations from a large group of users, need to be simultaneously processed at sub-second response times. The operations throughput delivered was 30,316.44 operations per second[8] (ops/s) indicating fast database performance.

Operations
Throughput
30,316.44 ops/s

Efficient CPU Utilization

This metric represents a percentage of the CPU cycles being used for a given workload and was measured to ensure the server CPUs were efficiently processing the database workloads. Low utilization means that the CPUs were not being used efficiently, which could result in an underutilization of server capabilities and stranded compute resources. High CPU utilization indicates the CPU is better using its resources for higher application performance, however, most servers perform best when CPU utilization is under 80%. The results of this metric indicate that at 69.4% utilization, the CPUs were used heavily but not to the point of exhausting the available CPU resources.

Efficient
CPU Utilization
69.4%

Memory Utilization

This metric represents the percentage of memory being used for a given workload. The database size that was stored on the XD6 Series SSD was larger than the maximum memory allocated for MongoDB to use. This guaranteed that read and update operations could not be completed by simply retrieving a record from memory, and some data required for the read and update operations would be forced to access the XD6 Series SSD. IT administrators commonly assign 70% to 80% utilization of server memory which still provides available memory for urgent use. The results of this metric indicate that at 78% memory utilization, the database leveraged system memory appropriately and did not suffer major performance penalties or require any additional memory. This is largely due to the XD6 Series SSD's ability to deliver very low latency and high throughput performance.

Memory
Utilization
78%

Summary

Though E1.S platforms are being supported by a number of server and storage vendors, there are very few tested solutions available today. The Supermicro SSG-1029P-NES32R E1.S server and KIOXIA XD6 Series 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 low read and update latencies, fast workload throughput and high overall performance as follows:

MongoDB Test Results	Run Time	Avg. Read Latency	Avg. Update Latency	Operations Throughput	Efficient CPU Utilization	Memory Utilization
Supermicro E1.S Server: SSG-1029P-NES32R	**18.33 hrs.**	**5.54 ms**	**11.34 ms**	**30,316.44 ops/s**	**69.4%**	**78%**
KIOXIA E1.S SSD: XD6 Series						

This Supermicro and KIOXIA solution can support up to 6 terabytes[9] (TB) of system memory and up to 32 KIOXIA XD6 Series SSDs. As XD6 Series SSDs are available in 1 TB, 2 TB and 4 TB 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 well-suited for hyperscale-class and enterprise-class applications.

The system features heat sink options at varying sizes that are integrated directly into XD6 Series SSDs, enabling multiple SSD cooling options. Since XD6 Series SSDs fit directly within Supermicro SSG-1029P-NES32R E1.S servers without the use of carriers (as with M.2 drives), the system provides high storage density and efficient cooling (in compliance to the XD6 Series SSD operating specifications between 0°C and 70°C).

Additional Supermicro SSG-1029P-NES32R E1.S server information is available here.
Additional XD6 Series SSD information is available here.

ADDENDUM: TEST INFORMATION

Test Equipment

The hardware and software equipment used to perform the synthetic test and six metrics include:

- **Supermicro SSG-1029P-NES32R Server:**
  One (1) dual socket server with two (2) Intel® Xeon® Gold 6226R processors, featuring 16 processing cores, 2.90 GHz frequency, and 192 gigabytes[5] (GB) of DDR4 DRAM
- **Operating System:** Ubuntu® v20.04.3 (Kernel 5.4.0-89-generic)
- **Application:** MongoDB v5.0.3:
  Database size = 460.04 GB
  Maximum Allocated Memory = 150 GB
  Number of Connections / Threads = 256
  Number of Record = 300M
  Number of Operations = 2B
- **Test Software:** Synthetic tests run through YCSB software (version 0.17.0)
- **Storage Devices (Table 1):** One (1) KIOXIA XD6 Series PCIe 4.0 data center NVMe SSD with 3.84 TB capacity

Specifications	XD6 Series
Form Factor	E1.S
Package	9.5mm
Interface	PCIe 4.0
Capacity	3.84 TB
NAND Flash Type	BiCS FLASH™ 3D flash memory
Endurance Rating	1 Drive Write Per Day[10] (DWPD)
Endurance Time Frame	5 years
Power	<14W
DRAM Allocation	192 GB

Table 1: SSD specifications and set-up parameters

Test Set-up

The Supermicro SSG-1029P-NES32R E1.S server was configured with the Ubuntu v20.04.3 operating system and YCSB v0.17.0 test software. The YCSB software was used to create a database on the XD6 Series SSD to run a 50% read and 50% update YCSB workload against the MongoDB database. The 50%/50% mixed workload was used as it represents a common workload for many applications including database applications.

The Supermicro SSG-1029P-NES32R E1.S server Intel Xeon Gold 6226R processors are PCIe 3.0 capable, so the XD6 Series SSD was tested at Gen3 speeds. The XD6 Series SSD connects directly to the Supermicro server via an E1.S connector.

Test Procedures

The benchmark test performed on the XD6 Series SSD included the following metrics: (1) Run Time; (2) Average Read Latency; (3) Average Update Latency; (4) Operations Throughput; (5) Maximum CPU Utilization; and (6) Memory Utilization. The results of the six metrics were recorded.

NOTES:

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

[2] Source: Supermicro SSG-1029P-NES32R online product flyer - 6x E1.S drives = 2x U.2 drives in space utilization.

[3] 2.5-inch and 3.5-inch indicate the form factor of the SSD and not the drive's physical size.

[4] A NoSQL database stores data in a format other than relational tables.

[5] When tested internally by KIOXIA in a lab environment using a competitive E1.S SSD running the exact benchmark tests and configuration as presented in this Solution Brief demonstrated a Run Time of 41 hours. It should be noted that the competitive E1.S SSD was not supported on the Supermicro SSG-1029P-NES32R server at the time of this publication and its specifications are subject to change.

[6] When tested internally by KIOXIA in a lab environment using a competitive E1.S SSD running the exact benchmark tests and configuration as presented in this Solution Brief demonstrated an Average Read Latency of 1,154 ms. It should be noted that the competitive E1.S SSD was not supported on the Supermicro SSG-1029P-NES32R server at the time of this publication and its specifications are subject to change.

[7] When tested internally by KIOXIA in a lab environment using a competitive E1.S SSD running the exact benchmark tests and configuration as presented in this Solution Brief demonstrated an Average Update Latency of 1,211 ms. It should be noted that the competitive E1.S SSD was not supported on the Supermicro SSG-1029P-NES32R server at the time of this publication and its specifications are subject to change.

[8] When tested internally by KIOXIA in a lab environment using a competitive E1.S SSD running the exact benchmark tests and configuration as presented in this Solution Brief demonstrated Operations Throughput of 13,502 ops/s. It should be noted that the competitive E1.S SSD was not supported on the Supermicro SSG-1029P-NES32R server at the time of this publication and its specifications are subject to change.

[9] 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 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^{50} 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.

[10] 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.

TRADEMARKS:

DISCLAIMERS: