KIOXIA

Performance Brief

Improving Performance, Latency and System Utilization of Transactional Database Workloads with KIOXIA RM7 Series Value SAS SSDs versus SATA

Introduction

One of the most important applications in modern data centers are transactional databases. Many applications use them as they excel at reading and writing individual rows of data very quickly, while maintaining high data integrity. However, these databases are only as fast as their underlying storage, creating a need for high performance and low latency storage. SSD storage based on the SATA interface can become a performance and latency bottleneck for transactional database workloads.

The SATA interface is limited to a data transfer speed of 6 gigabits per second (Gb/s), with no planned performance improvements. As SATA and SAS SSDs connect to most server motherboards via a SAS interface card, utilizing the SAS protocol can enable a server to potentially double its data throughput from 6 Gb/s to 12 Gb/s.

With SATA SSD performance limitations, KIOXIA Corporation developed Value SAS SSDs - a unique class of SAS SSDs that deliver improvements in performance, capacity, reliability, manageability and data security over enterprise SATA SSDs, and priced to replace them. Utilizing a 12 Gb/s SAS interface (versus 6 Gb/s SATA), applications have access to higher throughput and lower latency storage. These benefits enable the server to utilize its hardware resources more efficiently when servicing data intensive workloads.

This performance brief presents a comparison of SSD performance, latency and system utilization using Microsoft® SQL Server™ transactional database workloads. It compares KIOXIA RM7-R Series Value SAS SSDs with SATA SSDs from a leading provider (Vendor A). A recently released rack server was used for the tests, as well as HammerDB[1] test software, which enabled the TPROC-C[2] online transaction processing (OLTP) workloads to run against the database.

The metrics recorded for both SSD configurations included Microsoft SQL Server database transactions, read/write throughput, read/write latency, power utilization, CPU utilization and total memory utilization. These metrics were recorded for each of three test runs for each SSD configuration to determine an average result.

The test results show that the combination of the 12 Gb/s SAS interface with KIOXIA RM7-R Series Value SAS SSDs delivered significant performance and latency improvements versus SATA SSDs (limited by the 6 Gb/s SATA interface).

The results include a brief description of each workload test, a graphical depiction of the test results and an analysis. Appendix A covers the hardware and software test configuration. Appendix B covers the configuration set-up and test procedures.

Test Results Snapshot

KIOXIA RM7-R Series Value SAS SSDs delivered the following results when compared with SATA SSDs from Vendor A:

Average SQL Server™ Database Transactions
(higher is better)
120% Higher

Average Read Throughput
(higher is better)
127% Higher

Average Write Throughput
(higher is better)
96% Higher

Average Read Latency
(lower is better)
82% Lower

Average Write Latency
(lower is better)
79% Lower

Average Power Utilization
+8% More

Average CPU Utilization
+18.8% More

Total Memory Utilization
Similar

Test Results[3]

The Average SQL Server™ Database Transactions Test

This test determined how the underlying storage impacted database performance and the user experience. It measured the number of transactions executed in the TPROC-C workload per minute. The HammerDB software, executing the TPROC-C transaction profile, randomly performed new order, payment, order status, delivery and stock level transactions. The test simulated an OLTP environment with a large number of users conducting simple and short transactions (that require sub-second response times).

There were three test runs conducted on each SSD configuration to obtain an average number of transactions that the server was able to execute. The results are in transactions per minute (TPM). The higher result is better.

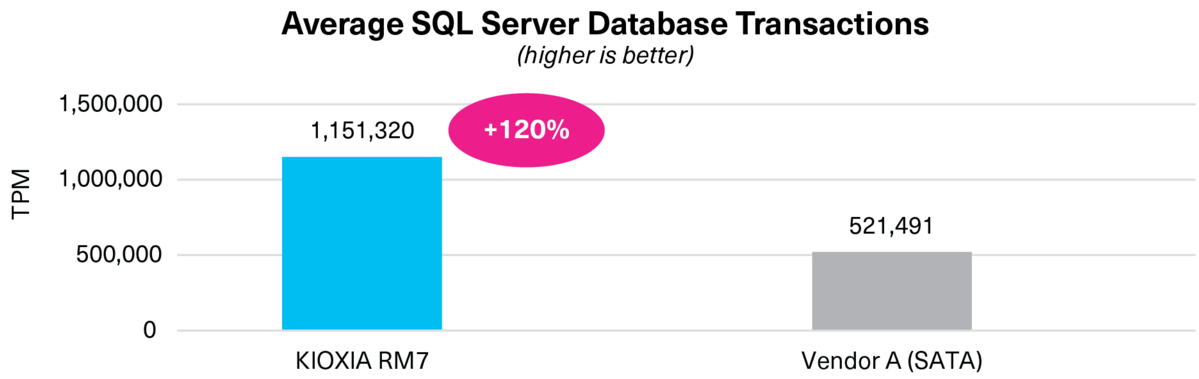

During this test run, additional performance metrics were collected to determine the impact that the TPROC-C transactional workload had on the server, including how much storage read and write throughput and latency that the system was capable of, as well as metrics associated with CPU, system memory and power utilization. These metrics included:

Test Metric 1: Average Read Throughput

This metric shows the average read throughput that each SSD configuration was capable of on the test server as it responded to incoming database queries. If the number of incoming queries is much higher than the achievable database throughput, the server can overload, creating longer wait times per query that can negatively affect application performance and user experience. The results show the average read throughput from three test runs for each SSD configuration. The results are in megabytes per second (MB/s). The higher result is better.

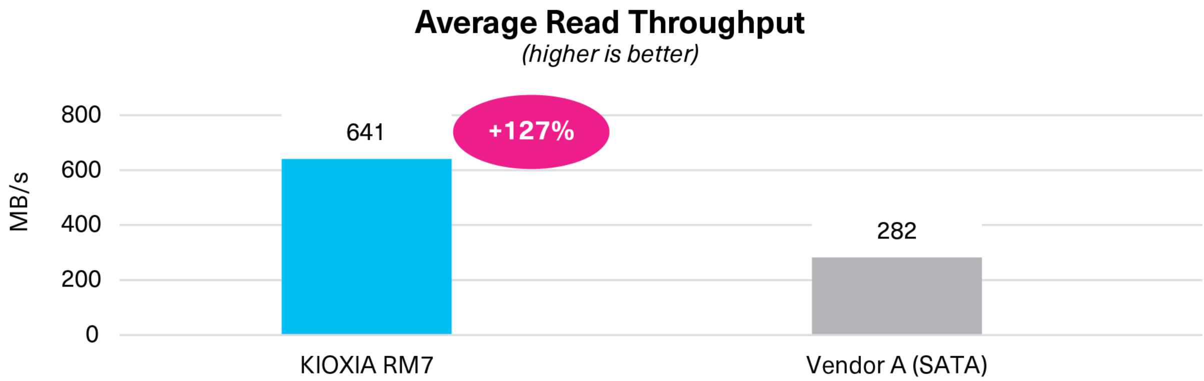

KIOXIA

Test Metric 2: Average Write Throughput

This metric shows the average write throughput of each SSD configuration and includes the average maximum data throughput delivered by the test server over an extended period of time. The write throughput issued to storage by the incoming database queries is generally the limiting factor for this test as the system must wait for the write operation to complete before issuing additional read or write queries. The process is known as IO-wait and can negatively affect application performance and the user experience. The results show the average write throughput from three test runs for each SSD configuration. The results are in megabytes per second (MB/s). The higher result is better.

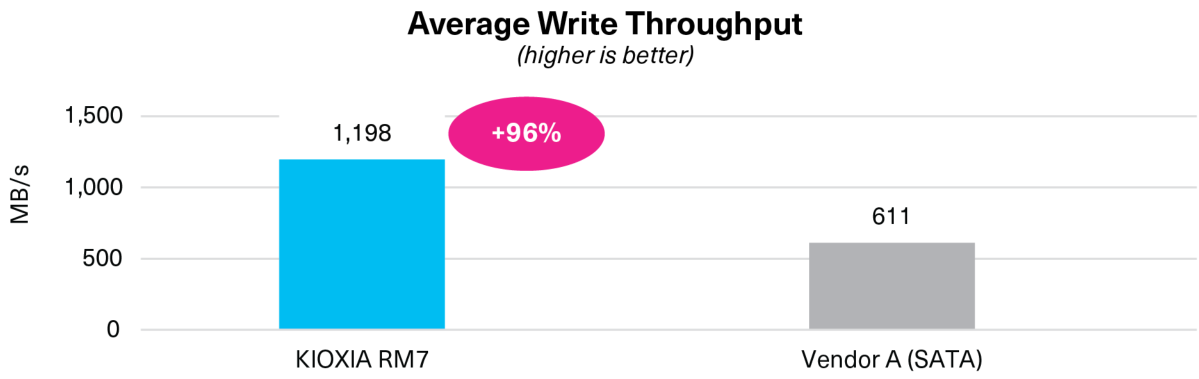

Test Metric 3: Average Read Latency

The average read latency is a metric of the time it took to perform a drive read operation. It includes the time it took for the workload generator to not only issue the read operation, but also the time it took to complete the operation and receive a 'successfully completed' acknowledgement. The following results show the average read latency from three test runs for each SSD configuration. The results are in microseconds (µs). The lower result is better.

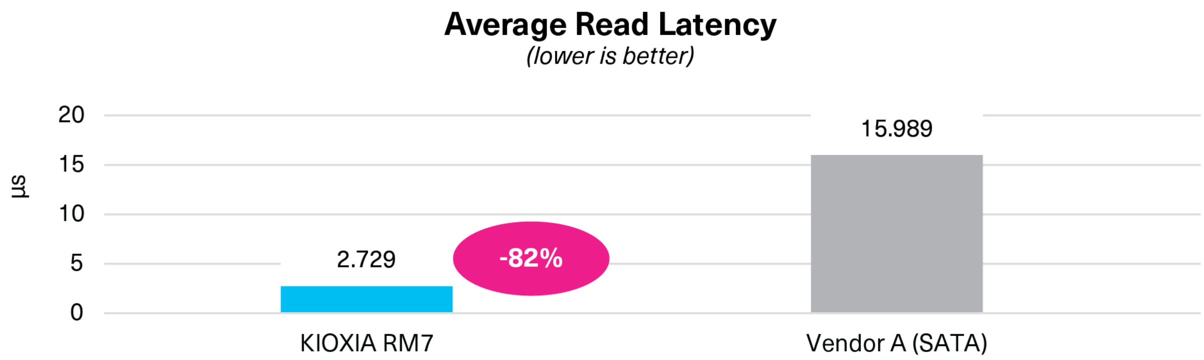

Test Metric 4: Average Write Latency

The average write latency is a metric of the time it took to perform a drive write operation. It includes the average time it took for the workload generator to not only issue the write operation, but also the time it took to complete the operation and receive a 'successfully completed' acknowledgement. The results show the average write latency from three test runs for each SSD configuration. The results are in microseconds (µs). The lower result is better.

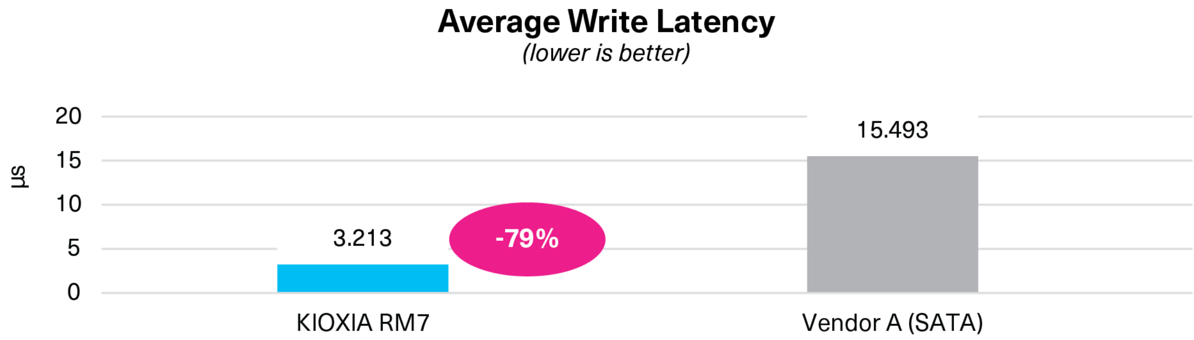

Test Metric 5: Average Power Utilization

The average power utilization is a metric of the average power drawn from the test system during the three test runs for each SSD configuration. The power draw is a component of all of the individual components that receive power from the power supply unit. For these tests, higher power could be required to deliver performance and latency improvements. The results are in Watts.

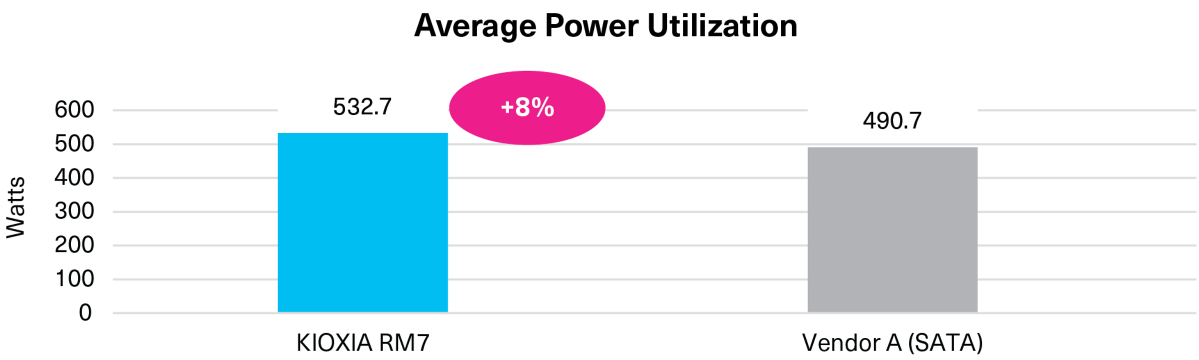

Test Metric 6: Average CPU Utilization

The average CPU utilization is a metric of the average CPU usage during the three test runs for each SSD configuration. For TPROC-C workloads, an increase in CPU usage correlates to faster storage since the CPU is not waiting for tasks. As a result, more transactions can be serviced making better use of the CPU resource. The results show the average CPU utilization from three test runs for each SSD configuration. The results are in percentage of use. For these tests, higher CPU utilization was the goal.

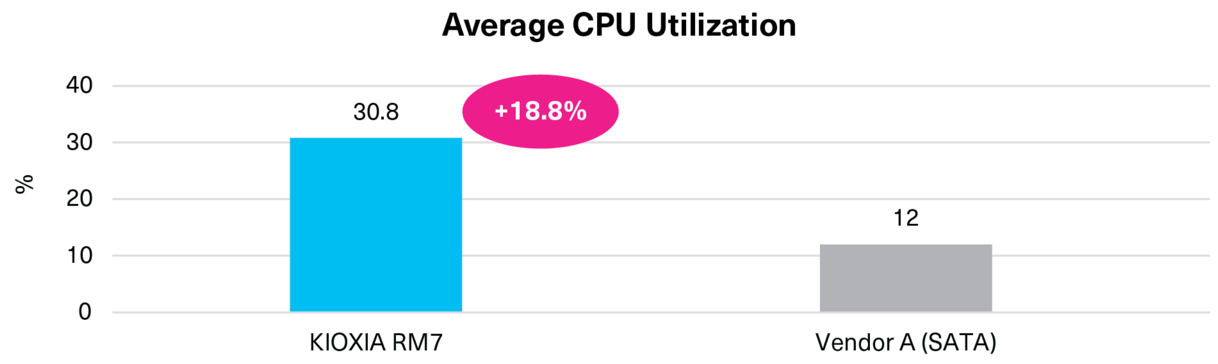

Test Metric 7: Total Memory Utilization

The total memory utilization is a measurement of the total memory (DRAM) utilization that the system used on average from three test runs for each SSD configuration. The Microsoft® SQL Server™ database was provided 118,096 megabytes[4] (MB) of DRAM. The results are in megabytes (MB). For these tests, similar system memory utilization was the objective.

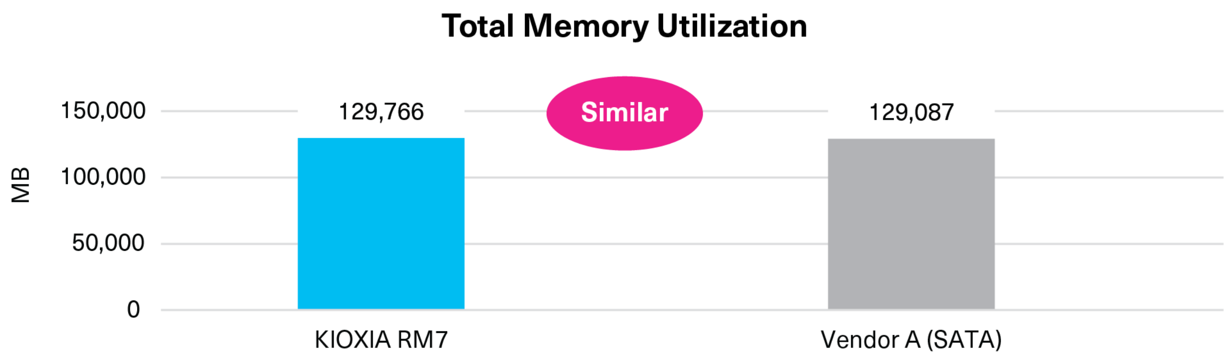

Analysis

KIOXIA RM7-R Series Value SAS SSDs deployed in the test server showed about twice the database performance (and more, in some cases) when compared with the SATA SSD configuration. Value SAS SSDs delivered on average 120% more transactions per minute, 127% higher read throughput and 96% higher write throughput, while reducing the server footprint needed to service these workloads.

From a latency perspective, it is especially important for the underlying storage to deliver low latency because slow storage can have a drastic impact on transaction queueing. The KIOXIA RM7-R Series Value SAS SSDs demonstrated significant latency improvements when performing a read or write operation when compared with the SATA SSD configuration, and delivered on average 82% lower read latency (or 2.87 µs better) and 79% lower write latency (or 3.25 µs better) .

The KIOXIA RM7-R Series Value SAS SSDs delivered about 8% higher overall system power draw on average than the SATA SSD configuration. The increase in power draw was required to deliver formidable performance and latency improvements.

High CPU utilization for these tests was the objective to demonstrate that the underlying storage can keep up with the Microsoft® SQL Server™ application. The KIOXIA RM7-R Series Value SAS SSDs demonstrated over 18% higher CPU usage when compared with the SATA SSD configuration. With the increase in performance delivered by KIOXIA RM7-R Series Value SAS SSDs, CPU usage increased and the server was able to efficiently use its processing cores to service more transactions. The lower CPU usage in the Vendor A configuration indicated that the drive performance required the CPU to wait for the SATA SSDs, leading to underutilization of the CPU.

Summary

This performance brief presented an SSD comparison using Microsoft SQL Server database workloads with KIOXIA RM7-R Series Value SAS SSDs and SATA SSDs from a leading provider (Vendor A). A recently released rack server was used for the tests, as well as HammerDB test software. Testing showed that KIOXIA RM7 Series Value SAS SSDs were able to deliver about twice the performance in measures of throughput, as well as significantly lower latency when compared with the SATA SSD configuration. KIOXIA RM7 Series Value SAS SSDs deliver a compelling value proposition and are an excellent drop-in replacement for existing SATA SSD deployments and for new or existing SAS servers.

KIOXIA RM7 Series Value SAS SSD Product Info

The latest generation KIOXIA RM7 Series Value SAS SSDs support the 12 Gb/s SAS-3 interface and are available in two 2.5-inch[5] configurations: RM7-R Series for read intensive applications (1 DWPD[6], up to 7.68 terabyte[4] (TB) capacities) and RM7-V Series for higher endurance mixed use applications (3 DWPD, up to 3.84 TB capacities). Security options[7] are available for both configurations.

Additional KIOXIA RM7 Series Value SAS SSD information is available at the KIOXIA 'Life After SATA' site.

KIOXIA RM7 Series SSD[8]

Appendix A

Hardware/Software Test Configuration

Server Information		
No. of Servers	1	
BIOS Version	BIOS-A42 v2.90	
CPU Information		
CPU Model	AMD EPYC™ 7543	
No. of Sockets	2	
No. of Cores	32	
Frequency (in gigahertz)	2.8 GHz	
Memory Information		
Memory Type	DDR4	
Memory Speed	DDR4-3200	
Total Memory Size	256 GB	
SSD Information		
SSD Model	KIOXIA RM7-R Series	Vendor A
Form Factor	2.5-inch	2.5-inch
Interface	SAS-3	SATA Revision 3.0
Interface Speed	12 Gb/s	6 Gb/s
No. of SSDs	4	4
SSD Capacity	3.84 TB	3.84 TB
Drive Writes per Day	1	0.8
Active Power	up to 9 watts	up to 3.6 watts
Operating System Information		
Operating System (OS)	Microsoft Windows® Server® 2022 Datacenter	
OS Version	10.0.20348	
Database Model	Microsoft® SQL Server™ 2022	
Database Version	16.0.1000.6	
Test Software Information		
Test Software Model	HammerDB	
Benchmark	TPROC-C	
Version	4.9	
No. of Virtual Users	256	

Appendix B

Configuration Set-up/Test Procedures

Configuration Set-up

The test server was installed with the Microsoft® Windows® Server® 2022 Datacenter operating system.

Four 3.84 TB[6] KIOXIA RM7-R Series Value SAS SSDs and four 3.84 TB Vendor A SATA SSDs were placed in a hardware RAID 10[9] set in the BIOS, using the default RAID controller settings.

Disk Management on the Microsoft Windows Server 2022 Datacenter operating system was initialized and a new volume was created by the RAID 10 hardware set.

Microsoft® SQL Server™ was installed on the newly created volume.

HammerDB test software was installed on the same volume that Microsoft SQL Server resided enabling the TPROC-C OLTP workloads to run against the database. To build a schema, 5,000 was entered for the number of warehouses. For the number of virtual users, 128 was entered. With 5,000 warehouses, the size of the database was 590.4 GB[6].

Once the schema was built, the Driver Script was edited by changing 'Minutes for Test Duration' to 15 minutes and selecting 'Use All Warehouses.' For virtual users, 'Options' must be opened and 'Virtual Users' changed to 256, 'User Delay' to 1, and 'Repeat Delay' to 1. Select three options: Show Output, Log Output to Temp and Log Timestamps. Once completed, select 'Run' to initiate the schema.

The HammerDB TPROC-C workload tests were started.

Test Procedures

The results were recorded when the TPROC-C workload was run against each SSD configuration covering:

- *Average SQL Server Database Transactions (in TPM)*
- *Average Read Throughput (in MB/s)*
- *Average Write Throughput (in MB/s)*
- *Average Read Latency (in µs)*
- *Average Write Latency (in µs)*
- *Average Power Utilization (in watts)*
- *Average CPU Utilization (in percentage of use)*
- *Total Memory Utilization (in MB)*

For each individual test, three total runs were performed and the average of the three runs were calculated and compared with each SSD configuration.

NOTES:

[1] HammerDB is benchmarking and load testing software that is used to test popular databases. It simulates the stored workloads of multiple virtual users against specific databases to identify transactional scenarios and derive meaningful information about the data environment, such as performance comparisons.

[2] TPROC-C is the OLTP workload implemented in HammerDB derived from the TPC-C™ specification with modification to make running HammerDB straightforward and cost-effective on any of the supported database environments. The HammerDB TPROC-C workload is an open source workload derived from the TPC-C Benchmark Standard and as such is not comparable to published TPC-C results, as the results comply with a subset rather than the full TPC-C Benchmark Standard. TPROC-C means Transaction Processing Benchmark derived from the TPC "C" specification.

[3] Read and write speed may vary depending on the host device, read and write conditions and file size.

[4] 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.

[5] 2.5-inch indicates the form factor of the SSD and not the drive's physical size.

[6] 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.

[7] Optional security feature compliant drives are not available in all countries due to export and local regulations.

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

[9] RAID 10 is a redundant array of independent disks configuration that combines disk mirroring and disk striping to protect data.

TRADEMARKS:

AMD, EPYC and combinations thereof are trademarks of Advanced Micro Devices, Inc. Microsoft, Microsoft SQL Server, Microsoft Windows and Windows Server and are registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. TPC-C is a trademark of the Transaction Processing Performance Council. All other company names, product names and service names may be trademarks or registered trademarks of their respective companies.

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 render inaccuracies 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. are advised of the possibility of such damages.