KIOXIA

Performance Brief

Upgrading Local Server Transactional Database Performance and Reliability with KIOXIA PM7 Series 24G SAS SSDs Versus SATA

Introduction

Transactional databases handle a broad range of applications in many industries (i.e., e-commerce sites, banks, stock exchanges, retail stores, etc.), and represent one of the most important workloads run in data centers today. As these workloads excel at reading and writing individual rows of data very quickly, they depend on high performance hardware. With a high dependency on the underlying storage, SATA SSDs can be a bottleneck for transactional database workloads.

The SATA Revision 3.0 interface supports up to 6 gigabits per second (Gb/s) of data transfer, which is lackluster when compared with today's interface standards and throughput requirements for modern computational applications. Additionally, SATA is without a roadmap for future performance improvements, which means it will not get faster any time soon, and likely never.

Most modern servers, even if they are purchased with SATA SSDs, ship with a SAS infrastructure and backplane, which allows both SAS and SATA drives (SSDs and HDDs) to be used in the same drive bay. Therefore, new and existing SATA servers have the ability to support SAS with pre-existing hardware, meaning, that as storage requirements change, new servers do not need to be purchased. Instead, slow SATA drives can be easily replaced with higher performing, more reliable and larger capacity SAS SSDs. With the availability of the 24G SAS interface, many servers have unrealized performance improvements that can be utilized very easily, and with great cost savings.

When compared with the SATA interface standard, 24G SAS can deliver multiple times the bandwidth (22.5 Gb/s versus 6 Gb/s). In a data center with random bursts or constant data throughput, 24G SAS SSDs can significantly improve both the number of transactions per minute that can be delivered, and the responses to incoming queries for read or write operations.

This performance brief presents an SSD performance and system utilization comparison using Microsoft® SQL Server™ as a database to simulate transactional workloads. It compares KIOXIA PM7-V Series 24G SAS SSDs with SATA SSDs from a leading provider (Vendor A). A Dell® PowerEdge™ R660 rack server was used for the tests and included a PERC (PowerEdge RAID Controller) 12 card that enabled both SAS and SATA drive configurations with local RAID (Redundant Array of Independent Disks) capabilities.

Database tests were performed on both SSD configurations and several metrics were recorded to show performance and system utilization. The metrics recorded were average database transactions, average read and write throughput, average read and write latency, average power utilization, average CPU utilization and total memory utilization. There were three total test runs, and the average of the three runs were calculated and compared with each SSD configuration. HammerDB[1] test software was used for both SSD configurations, which enabled the TPROC-C[2] online transaction processing (OLTP) workloads to be run against the database.

The test results show that the KIOXIA PM7-V Series 24G SAS SSDs delivered a 3.4x increase in transactions per minute (TPM) performance when compared with SATA SSDs which are limited by a 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 of the results. Appendix A covers the hardware and software test configurations. Appendix B covers the configuration set-up and test procedures.

Test Results Snapshot

KIOXIA PM7-V Series 24G SAS SSDs delivered improved performance and system utilization results when compared with Vendor A SATA SSDs:

Average SQL Server™ Database Transactions
243% more TPM

Average Read Throughput
281% faster MB/s

Average Write Throughput
223% faster MB/s

Average Read Latency
8.9x improvement

Average Write Latency
109.2x improvement

Average Power Utilization
14% more utilization

Average CPU Utilization
53% more utilization

Total Memory Utilization
Similar utilization

TPM = transactions per minute
MB/s = megabytes per second

Test Results[3]

The Average SQL Server™ Database Transactions Test

This test demonstrates how the underlying storage affects database performance, showing how many transactions that each SSD configuration is able to handle. It includes the number of transactions that were 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.

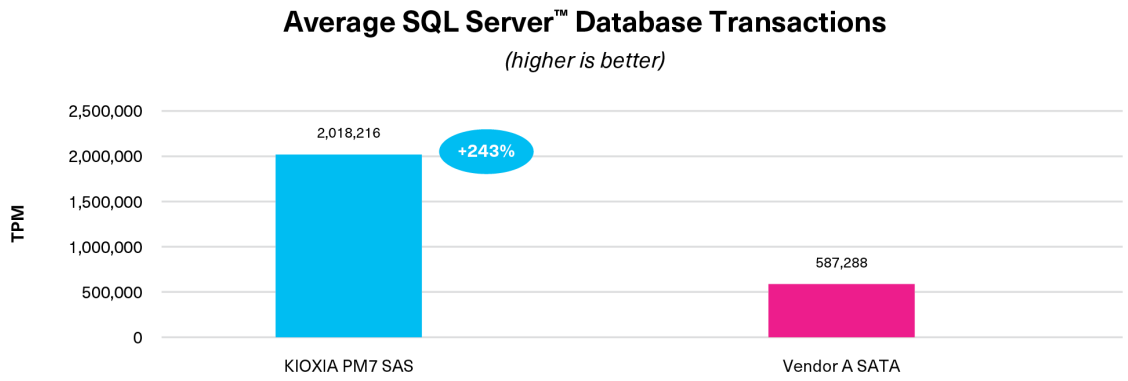

During this test run, additional performance metrics were recorded to determine the impact that the TPROC-C transactional workload had on the server as noted below:

Test Metric 1: Average Read Throughput

This metric measures the average read throughput that each SSD configuration was capable of achieving on the Dell® PowerEdge™ R660 server as it responds 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.

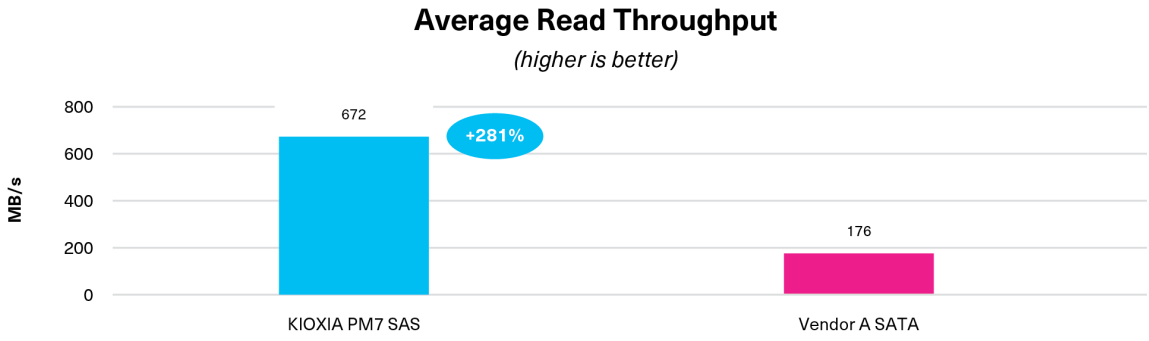

Test Metric 2: Average Write Throughput

This metric measures the average write throughput of each SSD configuration and includes the average maximum data throughput delivered by the Dell® PowerEdge™ R660 server over an extended period. 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, and more commonly known as IO-wait. 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.

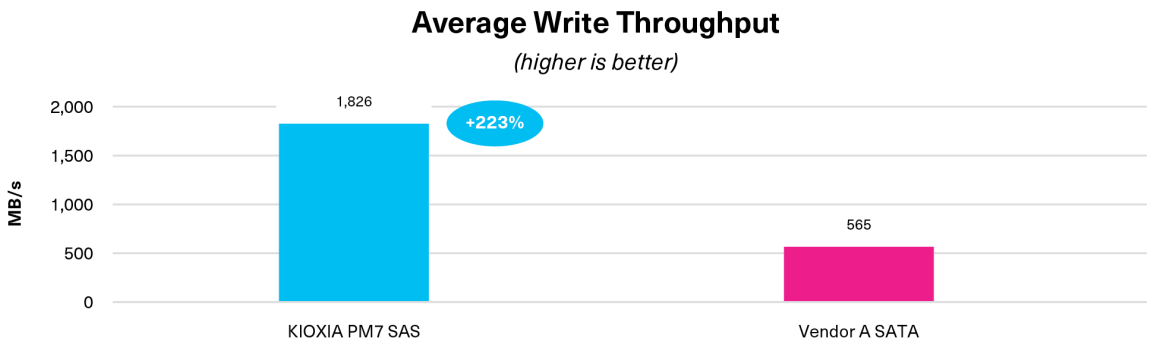

Test Metric 3: Average Read Latency

The average read latency is a metric that measures the time it takes to perform a drive read operation. It includes the time it takes for the workload generator to not only issue the read operation, but also the time it takes 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 milliseconds (ms). The lower result is better.

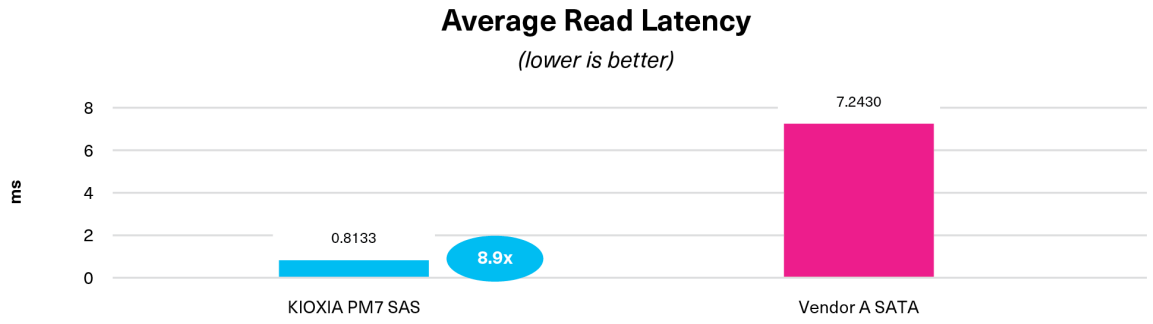

Test Metric 4: Average Write Latency

The average write latency is a metric that measures the time it takes to perform a drive write operation. It includes the average time it takes for the workload generator to not only issue the write operation, but also the time it takes 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 milliseconds (ms). The lower result is better.

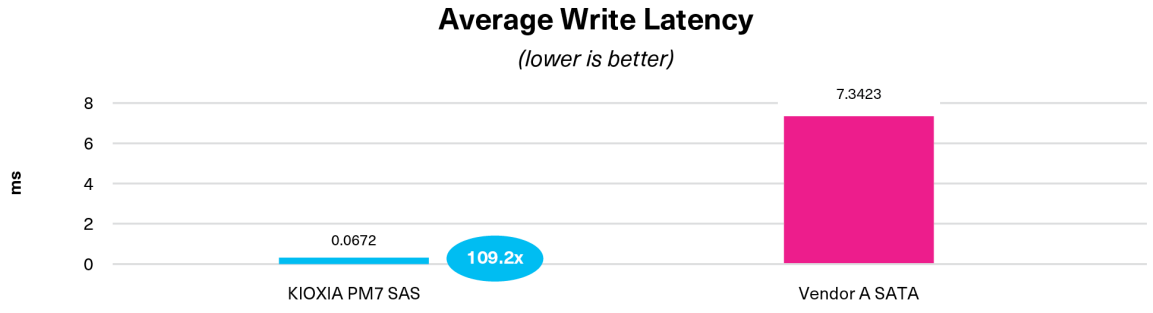

Test Metric 5: Average Power Utilization

The average power utilization is a metric that measures the average power draw from the test system during the three test runs for each SSD configuration. The power draw is a compilation of all of the individual components that receive power from the power supply unit. As the test system is identical for both configurations with the exception of the SSD, any difference in power consumption is caused by the storage. For these tests, higher power draw may be required to deliver performance and latency improvements. The results are in watts and includes four SSDs per test system.

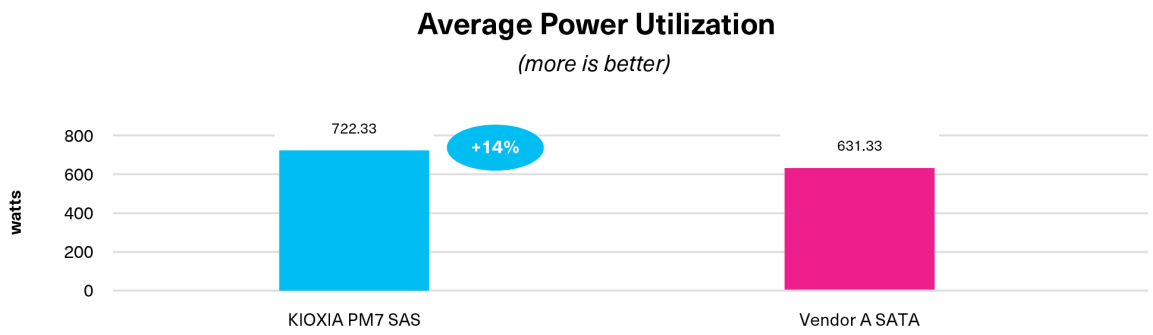

Test Metric 6: Average CPU Utilization

The average CPU utilization is a metric that measures the average CPU usage during the three test runs for each SSD configuration. For TPROC-C workloads, faster storage can improve CPU utilization. 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 correlates to more transactions that can be serviced.

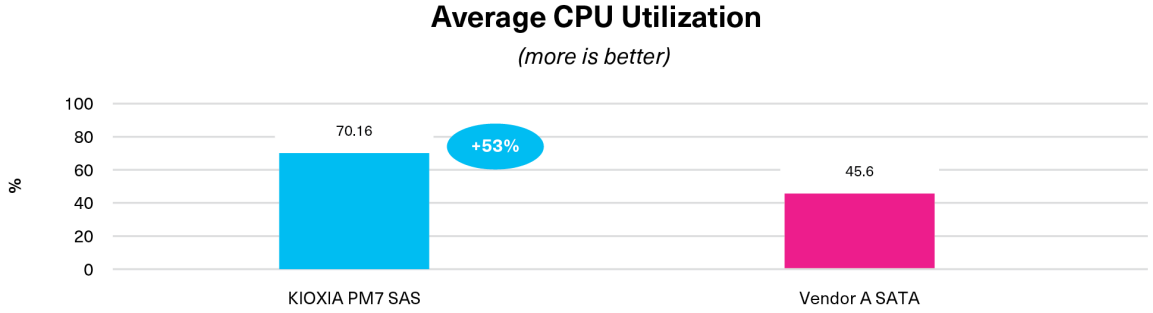

Note: The KIOXIA PM7-V Series SSD configuration delivered, on average, more than 1.5x improved CPU usage[4] when compared with the SATA SSD configuration; however, as seen earlier, the TPM increased even more, by 3.4x[5], which in turn, highlights the effects that faster storage has on CPU utilization.

Test Metric 7: Total Memory Utilization

The total memory utilization is a metric that measures the total DRAM utilization that the system uses on average from three test runs for each SSD configuration. For this test metric, the Microsoft® SQL Server™ database was provided 40,960 megabytes[6] (MB) of system DRAM. The results are in MB. For these tests, similar system memory utilization is the objective.

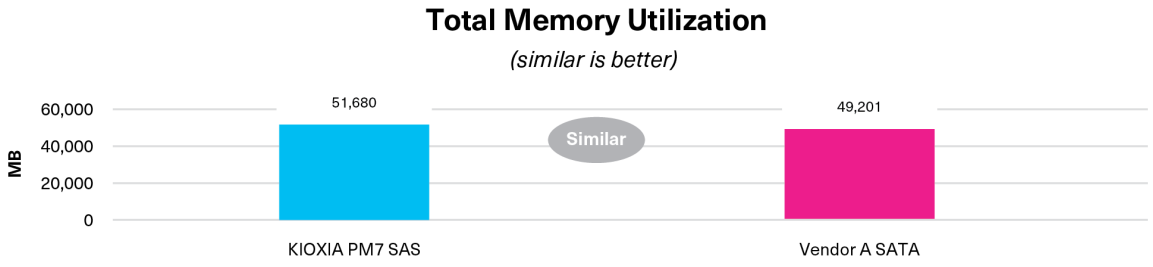

Analysis

KIOXIA PM7-V Series 24G SAS SSDs deployed in the Dell® PowerEdge™ R660 rack server was able to handle 3.4x more database transactions when compared with 6 Gb/s SATA SSDs. Additionally, the KIOXIA PM7 Series 24G SAS SSDs delivered 281% higher read throughput and 223% higher write throughput on average, while reducing the server footprint needed to service these transactional 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 PM7-V Series 24G SAS SSDs demonstrated considerable latency improvements when performing a read or write operation when compared with the SATA SSD configuration, and on average, delivered an 8.9x improvement in read latency and a 109.2x improvement in write latency. These latency improvements can breathe new life into a server that is experiencing slow or bottlenecked performance from deployed SATA storage. Applications, such as databases, cloud services, data analytics, artificial intelligence, machine learning and financial services, can garner a significant latency boost from 24G SAS deployments.

The KIOXIA PM7-V Series 24G SAS SSDs delivered about 14% higher overall system power draw on average when compared with the SATA SSD configuration. The increase in the power draw was required to deliver the increased performance and latency improvements.

High CPU utilization for these tests was the objective to demonstrate that the CPU was able to service more transactions utilizing the faster underlying storage, which in turn led to a higher TPM from the Microsoft® SQL Server™ application. The KIOXIA PM7-V Series 24G SAS SSDs demonstrated over 53% higher CPU usage when compared with the SATA SSD configuration. With the increase in performance delivered by these 24G SAS SSDs, CPU usage increased and the Dell® PowerEdge™ R660 server was able to more efficiently use its processing cores to service more transactions.

Summary

This performance brief presented an SSD performance and system utilization comparison using transactional database workloads with KIOXIA PM7-V Series 24G SAS SSDs and SATA SSDs from a leading provider (Vendor A). A Dell PowerEdge R660 server with a Microsoft SQL Server database benchmarked with HammerDB test software was used for the tests. Testing showed that KIOXIA PM7-V Series SSDs were able to deliver 3.4x the TPM performance when compared with the SATA SSD configuration. KIOXIA PM7-V Series 24G 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 PM7 Series 24G SAS SSD Product Info

The KIOXIA PM7 Series is the 2nd generation of 24G SAS SSDs that leverages 112-layer BiCS FLASH™ 3D flash memory technology in a 2.5-inch[7] form factor. The series is available in two configurations: (1) PM7-R Series for read-intensive applications and supports 1.92 terabyte[6] (TB) up to 30.72 TB capacities at 1 Drive Write Per Day[8] (DWPD); and (2) PM7-V Series for higher endurance mixed-use applications and supports 1.6 TB up to 12.8 TB capacities at 3 DWPD.

The PM7 Series represents KIOXIA's 7th SAS SSD generation that builds on the company's successes as a leading SAS SSD provider. More information available here.

KIOXIA PM7 Series SSD[9]

Appendix A

Hardware/Software Test Configurations

Server Information	
Server Model	Dell® PowerEdge™ R660
No. of Servers	1
BIOS Version	1.3.2

CPU Information	
CPU Model	Intel® Xeon® Gold 6430
No. of Sockets	2
No. of Cores	32
Frequency (in gigahertz)	2.1 GHz

Memory Information	
Memory Type	DDR5
Memory Speed	DDR5-4400
Memory Size	32 GB[6]
No. of DIMMs	4
Total Memory Size	128 GB

SSD Information		
SSD Model	KIOXIA PM7-V Series	Vendor A
Form Factor	2.5-inch	2.5-inch
Interface	24G SAS	SATA Revision 3.0
Interface Speed	22.5 Gb/s	6 Gb/s
No. of SSDs	4	4
SSD Capacity	12.8 TB	3.84 TB
Drive Writes per Day	3	1
Active Power	up to 5 watts	up to 3.6 watts

Operating System Information	
Operating System (OS)	Microsoft® Windows® Server® 2022 Datacenter
OS Version	8.0.1, 21813344
Database Model	Microsoft® SQL Server™
Database Version	16.0.1000.6

Test Software Information	
Test Software Model	HammerDB
Benchmark	TPROC-C
Version	4.9
No. of Virtual Users	176

Appendix B

Configuration Set-up/Test Procedures

Configuration Set-up

The Dell® PowerEdge™ R660 rack server was installed with the Microsoft® Windows® Server® 2022 Datacenter operating system.

Four 12.8 TB[6] KIOXIA PM7-V Series 24G SAS SSDs and four 3.84 TB Vendor A SATA SSDs were placed in a hardware RAID 10[10] set to hold the Microsoft® SQL Server™ database in the server.

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 online transaction processing (OLTP) workloads to run against the database. To build a schema, 1,000 was entered for the number of warehouses. For the number of virtual users, 176 was entered. With 1,000 warehouses, the size of the database was 126 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 176, 'User Delay' to 1, and 'Repeat Delay' to 1. 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 ms)*
- *Average Write Latency (in ms)*
- *Average Power Utilization (in watts)*
- *Average CPU Utilization (in percentage)*
- *Total Memory Utilization (in MB[6])*

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] The improved 1.5x CPU usage is determined by dividing the KIOXIA PM7-V Series average CPU utilization results of 70.16% by the Vendor A SATA SSD results of 45.6%.

[5] The TPM increase of 3.4x was determined by the Average SQL Server™ Database Transactions Test results as follows: 2,018,216 TPM (KIOXIA PM7 Series) divided by 587,288 TPM (Vendor A SATA SSDs).

[6] 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.

[7] 2.5-inch indicates the form factor of the SSD and not the drive's physical size.

[8] DWPD: Drive Write(s) Per Day. One full drive write 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] The product image shown is a representation of the design model and not an accurate product depiction.

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

TRADEMARKS:

Dell and PowerEdge are registered trademarks or trademarks of Dell Inc. Intel and Xeon are registered trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. 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.