KIOXIA

White Paper

Delivering Advanced Performance and Improved Server Utilization with Value SAS SSDs

A Unique Class of SAS SSD to Replace Enterprise SATA in Servers

Flash-based solid-state drives (SSDs) continue to replace hard disk drives (HDDs) in servers at a steady pace. As IT departments are tasked with choosing SSD solutions for their servers to support their varying workloads and applications, one of the decisions they face will be selecting the right SSD interface: SATA, SAS or PCIe®. This white paper presents a unique class of enterprise SAS SSD that can enable up to an 80% reduction in latency and up to a 54% improvement in performance over SATA SSDs in an Online Transaction Processing (OLTP) application – resulting in an increased host CPU utilization of up to 9%[1].

Industry Landscape

With the deployment of SSDs in servers, SATA SSDs were a ubiquitous choice for IT managers and data center architects due to their low cost and compatibility with SAS infrastructures. As more feature-rich and higher performing interfaces emerged over time, SATA still remained the most common interface for server-attached SSDs based on its low cost, despite performance limitations. Though the SATA interface speed is sufficient for hard drive implementations where disk rotations and seek latencies dominate performance contentions, flash-based SATA SSDs are generally bottlenecked by its 6 gigabits per second (Gb/s) performance ceiling.

The SATA SSD attach rate in servers is forecasted to decline as application and workload requirements are surpassing the capabilities of the SATA protocol (Figure 1).

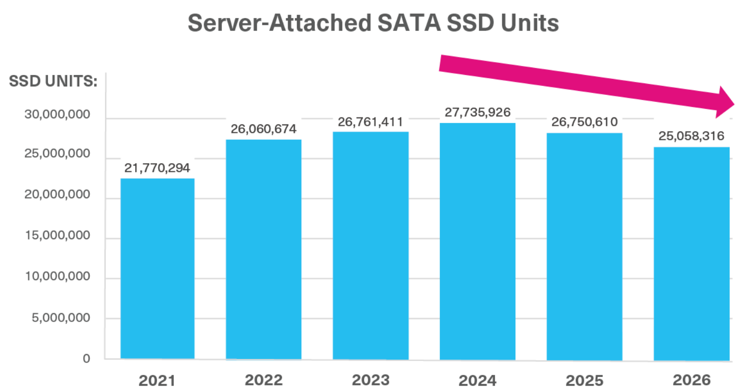

Figure 1: Depicts the decline of SATA SSDs in enterprise servers
Source: Forward Insights, SSD Insights Q2/22, Report #FI-NFL-SSD-Q222, Published May 2022

The SATA performance ceiling can become a bottleneck in I/O-intensive workloads that can result in underutilized server node compute capabilities and inefficient resource allocation, especially at data center scale. Since the SATA interface has no substantive improvements in the near-future roadmap, IT managers and solution vendors will seek alternative interface options for deploying SSDs, as well as realize the highest value from their flash investments.

Recognizing the limitations of SATA SSDs, KIOXIA Corporation (formerly Toshiba Memory Corporation) developed and introduced value SAS SSDs to the industry — a unique class of SAS SSD that delivers advancements in performance, capacity, reliability, manageability and data security over enterprise SATA SSDs, at a price designed to replace them. When value SAS SSDs are deployed in servers, applications have access to higher performing and lower latency storage, and no longer have to contend with a 6Gb/s performance ceiling. This improvement in storage performance enables a server node to efficiently utilize its CPU and DRAM resources while servicing I/O-intensive workloads. It also increases the server's load capacity enabling the node to support more users.

At data center scale, these system-level improvements can be translated into a reduction of nodes required to support a given workload. The end result is a decrease in cooling and power requirements that ultimately may reduce the infrastructure's total cost of ownership (TCO) necessary to deliver a certain performance level.

KIOXIA America, Inc. conducted a series of application benchmark tests using its RM Series of value SAS SSDs to demonstrate the performance benefits, as well as the improvements in server utilization, versus enterprise SATA SSDs. This white paper presents the criteria, results and analysis associated with these tests and showcases how value SAS SSDs deliver higher performance, lower latencies and improved server utilization when compared to enterprise SATA SSDs.

Challenges Associated with Enterprise SATA

Enterprise SATA SSDs utilize a 6Gb/s electrical interface which is half the bandwidth of the 12Gb/s SAS-3 interface. The SAS interface is also full-duplex which enables its full bandwidth to be utilized in both directions simultaneously (so up to 24Gb/s is actually supported). Conversely, SATA is half-duplex and limited to only 6Gb/s bandwidth in one direction at a time limiting its full performance to one quarter of the maximum bandwidth supported by 12Gb/s SAS.

The SATA interface is also limited by a low maximum queue depth of 32 versus up to 256 outstanding Inputs/Outputs (I/Os) in the queue supported by SAS. The relatively narrow bandwidth and shallow command queueing of SATA creates a bottleneck for SSDs, leaving much of the performance and parallelism of its flash memory underutilized.

Enterprise SATA SSDs are also driven by the Advanced Host Controller Interface (AHCI) instruction set designed to reduce development and deployment costs for consumer-grade hard drives. Unfortunately, this cost-minded optimization led to the omission of several enterprise-focused features that are supported in the SAS interface. Since the original deployment model for SATA drives included home PCs and workstations, signal strength and integrity were optimized for short cable lengths and there was not a need for in-depth health reporting, feature configuration, debug information, or data integrity verification. As such, the SATA interface does not support the same depth of features as SAS such as added value, reliability, availability and serviceability to the data center.

The performance, deployment and usability limitations of enterprise SATA SSDs may result in an inability to service today's performance-intensive applications or a growing user base. KIOXIA Corporation developed value SAS SSDs to improve on all facets of the SATA interface - from providing faster application performance, lower latencies and improved server utilization, to supporting more robust data integrity and management feature sets.

Description of Benchmark

Benchmark tests were conducted by KIOXIA America, Inc. in a lab environment that compared the system performance, transactional latency and system CPU utilization of a mainstream server platform configured with KIOXIA RM6 Series value SAS SSDs versus currently shipping enterprise SATA SSDs from a leading vendor (referred to as Vendor A). The tests utilized a TPROC-C workload[2] generated by HammerDB[3] against a SQL Server™ database to emulate an OLTP application.

Test Configuration:

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

- **Server:** One (1) HPE® ProLiant™ DL385 Gen10 Plus v2 dual socket servers with dual AMD EPYC™ 7543 CPUs featuring 32 processing cores, 2.8 GHz frequency, and 256 gigabytes[4] of DRAM
- **Operating System:** Microsoft® Windows Server® 2022 Standard
- **Application:** Microsoft SQL Server Standard (64-bit) version 15.0.2000.5
- **Storage Devices:** Four (4) KIOXIA RM6 Series value SAS SSDs with 3.84 terabytes[4] (TB) of capacity -
  Four (4) leading enterprise SATA SSDs (Vendor A) with 3.84 TB[4] of capacity

Specifications	RM6 Series	Vendor A
Interface	Value SAS	Enterprise SATA
Capacity	3.84 TB[4]	3.84 TB[4]
Form Factor	2.5-inch (15mm)	2.5-inch (7mm)
Drive Writes per Day[5] (DWPD)	1 (5 years)	0.8 (3 years)

Table 1: SSD specifications

Key Setup Parameters	RM6 Series	Vendor A
Data Warehouses[6]	1,000	1,000
Virtual User Count	512	512
InnoDB Buffer Pool Size	200 GB	200 GB

Table 2: Set-up parameters

- **Benchmark Software:** TPROC-C benchmark generated through the HammerDB test tool:
  - *TPROC-C is the OLTP workload implemented in HammerDB and includes a mix of five concurrent transactions of different types and nine types of tables with a wide range of record and population sizes – results are measured in transactions per minute (tpm).*
  - *HammerDB is benchmarking and load testing software used to test popular databases. The software simulates the stored workloads of multiple virtual users against the SQL Server database to identify transactional scenarios and derive meaningful information about the data environment, and in most cases, includes performance comparisons.*

Setup and Test Procedures:

The test system was set-up using the hardware and software equipment outlined above. One HPE ProLiant DL385 Gen10 Plus v2 server was used to host the SQL Server database and the HammerDB load testing software.

Prior to testing, the HammerDB software was configured with a test schema based on the TPROC-C benchmark (to emulate an OLTP environment). The SQL Server application was then loaded with 1,000 data warehouses[6] comprising of about 100 GB[4] of the server's storage capacity. For these benchmark tests, the memory allocation to the SQL Server buffer was set to 200 GB[4] to emulate a common data center configuration where approximately 80% of the available server DRAM was dedicated to database transactions while the remaining 20% was reserved for operating system functionality.

HammerDB was configured for 512 virtual users to simultaneously send query threads to get a response. The query response time was set to 1 millisecond (ms), showing the capability to achieve a near-immediate response time. HammerDB was also configured to timed testing runs using 512 virtual users where each test run included a 2-minute ramp up and one 60-minute test duration. Each test was run three times to record a score comprised of the average of the test runs.

The tests included: (1) Average transactions per minute (tpm); (2) Read latency; (3) Write latency; and (4) CPU utilization. The tests showcase how value SAS SSDs provide higher performance, lower latencies and improved server utilization when compared to enterprise SATA SSDs.

Test Results

The test results are divided into four test categories and were conducted with the average values presented based on three test runs. The following is a description of each of the four tests and associated results:

Test 1: Transactions per Minute

In an OLTP environment, tpm is the measure of how many new orders a system can support while it is executing additional transaction types such as payment, order status, delivery and stock levels. OLTP applications normally have a large number of users who conduct simple, yet short transactions that require sub-second response times and return relatively few records. The results are presented in Table 3 and Figure 2 respectively.

SSD Tested	Average Results (3 runs)	TPM Improvement
RM6 Series (value SAS)	2,257,996 tpm	+54%
Vendor A (enterprise SATA)	1,463,301 tpm	

Table 3: TPM comparison of value SAS SSDs versus enterprise SATA SSDs (higher is better)

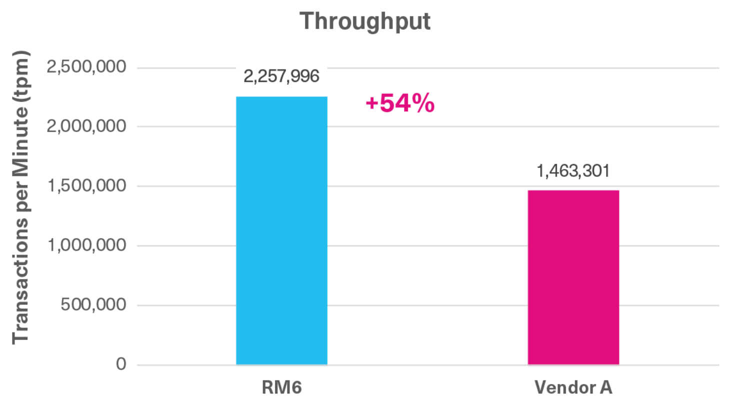

Figure 2: Graphical TPM comparison of value SAS SSDs versus enterprise SATA SSDs (higher is better)

Test 2: Read Latency

Latency is the measure of the time required for a sub-system or one of its components to process a single storage transaction or data request. The time it takes for the data to begin moving from one system to another may greatly affect application performance and the overall user experience. Read latency is the time delay before data is returned from a storage device following an instruction from the host for that transfer. The results are presented in Table 4 and Figure 3 respectively.

SSD Tested	Average Results (3 runs)	Read Latency Improvement
RM6 Series (value SAS)	1.69 milliseconds (ms)	-80%
Vendor A (enterprise SATA)	8.45 ms	

Table 4: Read latency comparison of value SAS SSDs versus enterprise SATA SSDs (lower is better)

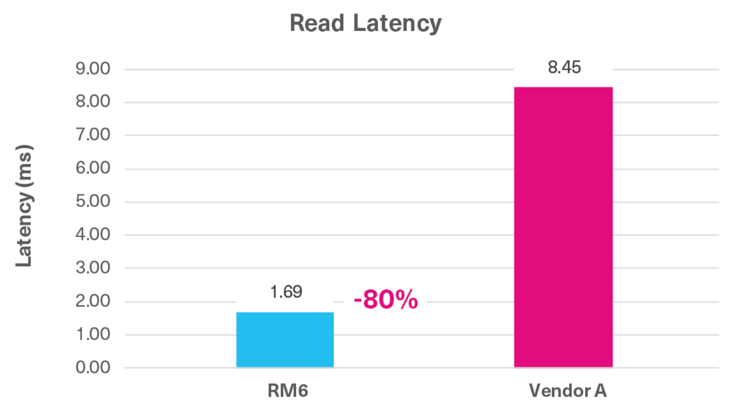

Figure 3: Graphical read latency comparison of value SAS SSDs versus enterprise SATA SSDs (lower is better)

Test 3: Write Latency

Write latency is similar to read latency but for writing data, and is the delay in time before a storage device completes writing the data following an instruction from the host for that transfer. The results are presented in Table 5 and Figure 4 respectively.

SSD Tested	Average Results (3 runs)	Write Latency Improvement
RM6 Series (value SAS)	13.81 ms	-62%
Vendor A (enterprise SATA)	36.82 ms	

Table 5: Write latency comparison of value SAS SSDs versus enterprise SATA SSDs (lower is better)

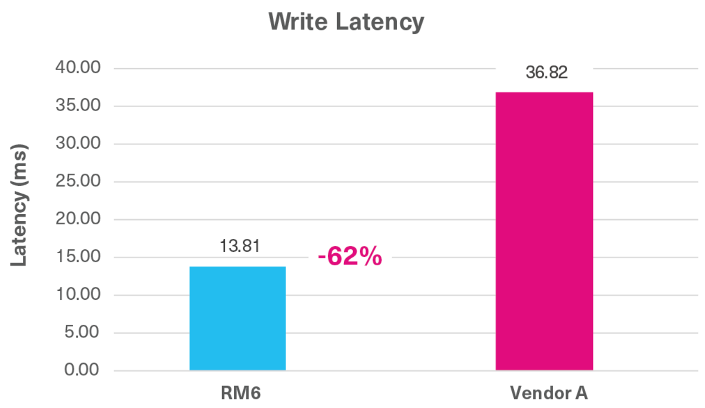

Figure 4: Graphical write latency comparison of value SAS SSDs versus enterprise SATA SSDs (lower is better)

Test 4: CPU Utilization

CPU utilization determines the amount of computing tasks that are being performed by the server's central processing unit. Low CPU utilization means that the processor is not being used efficiently. The results are presented in Table 6 and Figure 5 respectively.

SSD Tested	Average Results (3 runs)	CPU Utilization Improvement
RM6 Series (value SAS)	62.34%	+9%
Vendor A (enterprise SATA)	57.10%	

Table 6: CPU utilization comparison of value SAS SSDs versus enterprise SATA SSDs (higher is better)

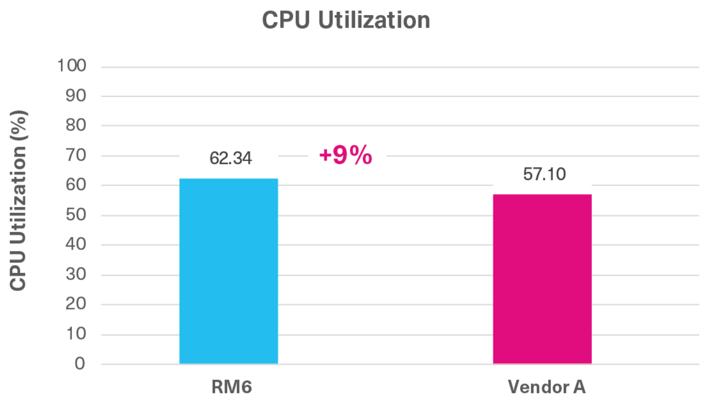

Figure 5: Graphical CPU utilization comparison of value SAS SSDs versus enterprise SATA SSDs (higher is better)

Test Results

From the benchmark results, value SAS SSDs enabled the server to deliver on average, 54% higher transactions per minute than when using comparable SATA drives under a TPROC-C workload. The full-duplex 12Gb/s bandwidth capability of value SAS enables a system to support more transactions and utilize more of the system's CPU, all while servicing transactions at up to 80% lower read latencies and 62% lower write latencies, on average.

When value SAS SSDs are installed in the mainstream server platform, they deliver faster, lower latency storage performance. This enables the server node to support up to a 54% higher application load versus when it was configured with comparable SATA drives. In summary, value SAS delivered the following advantages over enterprise SATA SSDs:

TEST RESULTS

54% higher average database performance per node

80% lower average latency per drive read

62% lower average latency per drive write

9% higher average CPU utilization

RM6 Series Overview

The RM6 Series is KIOXIA's 2nd generation of the value SAS product line that features improved performance from the initial offering, up to 7.68 TB[4] capacity, 1 DWPD[5] for read-intensive applications (RM6-R Series) and 3 DWPD[5] for mixed use applications (RM6-V Series), and a host of security options – all of which are geared to support a wide variety of workload requirements.

RM6 Series SSDs
12Gb/s Value SAS

Performance Specs[7]

SeqRead = up to 810 MB/s
RanRead = up to 155K IOPS[8]
SeqWrite = up to 635 MB/s
RanWrite = up to 52K IOPS[8]

Configurability

1 and 3 DWPD[5] options
960 GB to 7,680 GB capacities

Summary

The benchmark tests conducted by KIOXIA America demonstrate the performance improvements that the RM6 Series value SAS SSDs has over typical enterprise SATA SSDs. While enterprise SATA SSDs are cost-effective and easy to deploy, value SAS may be a better SSD option. It not only delivers improved application performance and lower latencies, but easily replaces SATA in SSD deployments built with a SAS infrastructure. RM6 Series value SAS SSDs are priced to replace enterprise SATA SSDs in server-attached configurations.

The results from the benchmark tests showcase that RM6 Series value SAS SSDs improved the performance of a server node, servicing an OLTP workload an average of 54% when tested against the same server node configured with SATA SSDs. The more storage bandwidth and IOPS[8] that a server can support, the more database transactions it can complete. This enables more users to be serviced at one time by the server and sets the foundation for delivering better Quality of Service (QoS) and potentially improving the TCO of a given infrastructure. The end result is better application performance, reduced IT costs and a better overall user experience.

For more information regarding the KIOXIA RM6 Series, visit the "life after SATA" website at https://business.kioxia.com/en-us/ssd/oem/hpe/rm6-value-sas.html.

NOTES:

[1] Based on system-level benchmarking conducted by KIOXIA America, Inc. included RM6 Series SSDs (3,840 GB capacities) and enterprise SATA SSDs (3,840 GB capacities), in an emulated OLTP environment. Results were based on three runs that tested average transactions per minute and overall disk latency. Read and write speed may vary, as well as latency, depending on the host device, read and write conditions and benchmark programs. Sequential read and write performance and latency results mentioned herein may vary from published RM6 Series SSD specifications.

[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] 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.

[4] 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 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 = 2^{30} bits = 1,073,741,824 bits, 1GB = 2^{30} bytes = 1,073,741,824 bytes and 1TB = 2^{40} 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.

[5] 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.

[6] The 1,000 data warehouses, in combination with the 200G buffer pool size and 512 virtual users represents a database configuration that creates a normal database size.

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

[8] IOPS, or Input/Output per Second, represents the number of I/O operations processed per second.

Trademarks:

AMD EPYC is a trademark of Advanced Micro Devices, Inc. HPE and ProLiant are either registered trademarks or trademarks of Hewlett-Packard Enterprise Company and/or its affiliates. Microsoft, SQL Server and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. PCIe is a registered trademark of PCI-SIG. 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: