

Delivering Advanced Performance and Improved Server Utilization with Value SAS SSDs

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

Ilya Cherkasov, Tyler Nelson, Scott Harlin



Introduction

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 paper will introduce a new class of SAS SSD that can enable up to a 50% reduction in latency and up to a 25% improvement in performance over SATA SSDs in an Online Transaction Processing (OLTP) application – all while increasing CPU utilization by up to 20%¹.

With the deployment of SSDs in servers, SATA SSDs have become 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 has remained the most common interface for server-attached SSDs, despite performance limitations, as a result of its low cost. Though the speed of the SATA interface is sufficient for hard drive implementations where disk rotations and seek latencies dominate performance contentions, flash-based SSDs are generally bottlenecked by SATA's 6Gb/s performance ceiling.

According to Forward Insights², the SATA SSD attach rate in servers is declining as application and workload requirements are surpassing the capabilities of the SATA protocol. This performance ceiling can become a bottleneck in I/O-intensive workloads and can result in underutilized server node compute capabilities, inefficient resource allocation, and stranded compute resources, especially at the data center scale. SATA has no substantive improvements in the near-future roadmap, so IT managers and solution vendors are looking to alternative interface options to deploy SSDs and realize the highest value from their flash investments.

Recognizing this limitation of SATA SSDs, KIOXIA America, Inc. introduced value SAS (vSAS) — a new class of SAS SSD that delivers advancements in performance, capacity, reliability, manageability and data security over SATA products at a price designed to replace enterprise SATA SSDs. 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 SATA's 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 the 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, which in turn, ultimately may reduce the total cost of ownership (TCO) of the infrastructure necessary to deliver a certain performance level.

KIOXIA conducted a series of application benchmark tests using its RM5 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 compared to SATA SSDs.

The SATA Problem

Enterprise SATA SSDs utilize a 6Gb/s electrical interface - only half of the maximum bandwidth of 12Gb/s SAS. While the SAS interface is full-duplex, meaning that the full bandwidth capability can be utilized in both directions simultaneously (up to 24Gb/s), SATA is only half-duplex and limited to only 6Gb/s of bandwidth in one direction at a time. Effectively, this limits 6Gb/s SATA to only a quarter of the maximum bandwidth supported by 12Gb/s SAS. Advancements in the SAS interface, such as MultiLink SAS™ wide port, enable two physical ports to be teamed together into a single logical port, doubling the interface bandwidth and further increasing the benefit over what the SATA interface can deliver.

Another limitation of the SATA interface is its low maximum queue depth of 32, versus up to 256 outstanding I/Os in the queue enabled in a typical SAS SSD. This relatively narrow bandwidth and shallow command queueing creates a bottleneck for the SSD, leaving much of the performance and parallelism of its flash memory underutilized.

SSDs utilizing the SATA interface support the Advanced Host Controller Interface (AHCI) instruction set. AHCI was 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 found 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 were inferior to the signal strength and integrity of the SAS interface. PC and workstation use cases had no need for in-depth health reporting, feature configuration, debug information, or data integrity verification, resulting in an interface that does not support the same depth of features available in the SAS specification that add value, reliability, availability and serviceability to the data center.

These performance, deployment and usability limitations mean that SSDs based on SATA may fall short of servicing today's performance-intensive applications or a growing user base. KIOXIA has designed value SAS 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 Benchmarks

Benchmark tests were conducted by KIOXIA in a lab environment that compared the system performance, transactional latency, and system CPU utilization of a mainstream server platform configured with RM5 Series value SAS SSDs and enterprise SATA SSDs. The tests utilized a TPC-C™-like workload generated by HammerDB against a SQL Server™ database to emulate an OLTP application.

Test Criteria

To measure the results as outlined in Section 3, the hardware and software equipment used for the benchmark testing included the following:

- **Server:** Dell™ PowerEdge™ R730xd dual socket servers with dual CPUs featuring four (4) processing cores, 3.4GHz frequency, and 32GB³ of DRAM
- **Operating System:** Microsoft® Windows Server® 2016
- **Application:** Microsoft SQL Server 2017
- **Storage Devices:** Four (4) KIOXIA RM5 Series SSDs with 960GB capacity (value SAS) Four (4) Intel® 4500 Series SSDs with 960GB capacity (enterprise SATA)
- **Benchmark Software:**
 - TPC-C benchmark generated through the HammerDB test tool
 - TPC Benchmark C is an OLTP benchmark that 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.

4a Set-up & Test Procedures

The test system was set-up using the hardware and software equipment outlined above. Separate Dell R730xd servers were used to host the SQL Server database and the HammerDB application to avoid contention from either application. Without the two-server testbed to conduct separate testing, the applications could have interfered with memory and/or CPU cycles, and could have compromised the test results.

Prior to testing, the HammerDB load testing software was configured with a test schema based on the TPC-C benchmark (to emulate an OLTP environment). SQL Server was then loaded with one thousand data warehouses that comprised about 100GB of the server's storage capacity. For these benchmark tests, the memory allocation to the SQL Server buffer was set to 10GB, so no more than 10% of the 100GB database could be cached at one time.

HammerDB was configured for 48 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 48 users where each test run included a 2-minute ramp up and two 5-minute test durations. Each of these tests were run three times to record a score comprised of the average of the test runs.

By recording and analyzing: (1) Average Transactions Per Minute (tpm); (2) Disk Read Latency; (3) Disk Write Latency; (4) Overall Disk Latency; and (5) 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 the five benchmark tests that were conducted with the average values presented in Chart 1 below (based on three test runs). A description of each test result is also included.

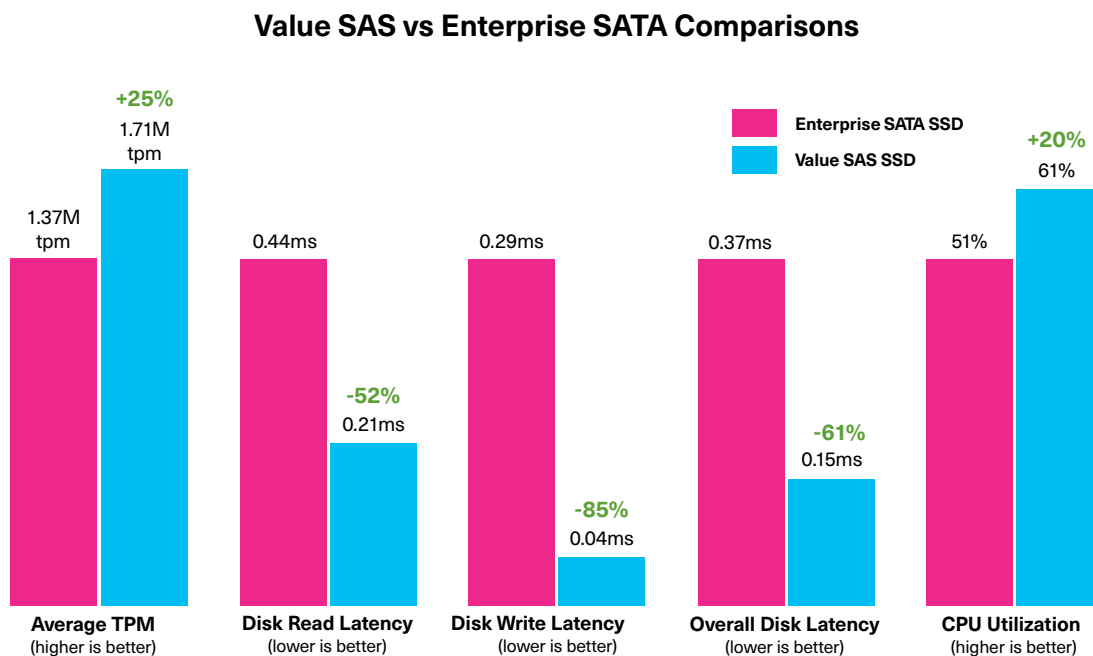


Chart 1: enterprise SATA SSD vs value SAS SSD benchmark comparisons in an OLTP environment

5a Transactions per Minute

In an OLTP environment, transactions per minute (tpm) is typically 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 that conduct simple, yet short transactions that require sub-second response times and return relatively few records. This type of transactional processing usually results in a 70% read and 30% write I/O split at the database level (Table 1).

SSD Tested	Average Results (3 runs)	Value SAS Advantage
Enterprise SATA	1.37M tpm	
Value SAS	1.71M tpm	+25%

Table 1: TPM comparison of enterprise SATA SSDs vs value SAS SSDs (higher is better)

5b Read Disk 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 the application's performance and user experience. Read disk latency is the time delay before the data is returned from a storage device following an instruction from the host for that transfer (Table 2).

SSD Tested	Average Results (3 runs)	Value SAS Advantage
Enterprise SATA	0.44ms	
Value SAS	0.21ms	-52%

Table 2: Read disk latency comparison of enterprise SATA SSDs vs value SAS SSDs (lower is better)

5c Write Disk Latency

The write disk latency has the same description as read disk latency but for writing data, and is typically the delay in time before a storage device completes writing the data following an instruction from the host for that transfer (Table 3).

SSD Tested	Average Results (3 runs)	Value SAS Advantage
Enterprise SATA	0.29ms	
Value SAS	0.04ms	-85%

Table 3: Write disk latency comparison of enterprise SATA SSDs vs value SAS SSDs (lower is better)

5d Overall Disk Latency

Overall disk latency is a representation of the average latency recorded throughout the workloads, including both read and write commands from the host (Table 4).

SSD Tested	Average Results (3 runs)	Value SAS Advantage
Enterprise SATA	0.37ms	
Value SAS	0.15ms	-61%

Table 4: Overall disk latency comparison of enterprise SATA SSDs vs value SAS SSDs (lower is better)

5e CPU Utilization

CPU utilization is used to determine 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 (Table 5).

SSD Tested	Average Results (3 runs)	Value SAS Advantage
Enterprise SATA	51%	
Value SAS	61%	+20%

Table 5: CPU utilization comparison of enterprise SATA SSDs vs value SAS SSDs (higher is better)

Test Analysis

From the results of the benchmarks, value SAS SSDs enabled the server to deliver on average, 25% higher transactions per minute than when using comparable SATA drives under a TPC-C workload. The full-duplex 12Gb/s bandwidth capability of value SAS enables a system to support more transactions and utilize much more of the system's CPU, all while servicing transactions at up to 85% lower 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 25% higher application load versus when it was configured with comparable SATA drives. A significant improvement like this

could mean up to a 20% reduction in TCO for an infrastructure supporting a TPC-C-like workload, where you could eliminate up to one fifth of the compute nodes by configuring your deployment with value SAS SSDs. In summary of the benchmark test results, value SAS delivered the following advantages over enterprise SATA SSDs:

- 25% higher averaged database performance per node
- 52% lower averaged latency per disk read
- 85% lower averaged latency per disk write
- 61% lower averaged latency per disk transaction
- 20% higher averaged CPU utilization

Summary

KIOXIA's benchmark tests demonstrate the performance advantages that the RM5 Series of value SAS SSDs have over typical enterprise SATA SSDs. While enterprise SATA SSDs are cost-effective and easy to deploy, value SAS is a better SSD option. It not only delivers improved application performance, but easily replaces SATA in SSD deployments. KIOXIA's RM5 Series is a new product family of value SAS SSDs that is priced to replace enterprise SATA SSDs and provides significant performance, reliability and latency benefits.

The results of KIOXIA's benchmark tests showcase that RM5 value SAS SSDs improved the performance of a server node servicing an OLTP workload an average of 25% over the same server node configured with SATA SSDs. The more storage bandwidth and IOPS 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 improving the TCO of a given infrastructure. The end result is better application performance, reduced IT costs and a better overall user experience.

In 2019, KIOXIA America, Inc. announced the availability of its RM5 Series of 12Gb/s value SAS SSDs utilizing KIOXIA's 64-layer BiCS FLASH™ TLC (3-bit-per-cell) 3D flash memory. RM5 Series SSDs have been adopted by some of the top server vendors in the world.

For details, visit the "life after SATA" website at <https://business.kioxia.com/en-us/ssd/life-after-sata.html>.

Disclaimer:

KIOXIA America, Inc. may make changes to specifications and product descriptions at any time. The information presented in this white paper 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 KAI 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.

¹ System-level OLTP benchmarking conducted by KIOXIA included RM5 Series SSDs (960GB capacities) and enterprise SATA SSDs (960GB capacities), in an emulated OLTP environment. Results were based on three runs that tested average transactions per minute (tpm) and overall disk latency (microseconds). Read and write speed may vary, as well as latency, depending on the host device, read and write conditions, and benchmark programs. KIOXIA defines a megabyte (MB) as 1,000,000 bytes and a kibibyte (KiB) as 2¹⁰ bytes, or 1,024 bytes. Sequential read and write performance, and latency, mentioned herein are reference data, and may vary from the RM5 product data in the datasheet.

² The decline of the SATA SSD attach rate in servers is available from Forward Insights, "SSD Forecast-SSD Insights" published Q1 2019.

³ Definition of capacity: KIOXIA 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 1 GB = 2³⁰ bytes = 1,073,741,824 bytes, 1 TB = 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.

Dell is a trademark of Dell, Inc. Intel is a registered trademark of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Microsoft, SQL Server and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. MultiLink SAS is a trademark of the SCSI Trade Association. PCIe is a registered trademark of PCI-SIG. TPC-C is a trademark of the Transaction Processing Performance Council. All other trademarks or registered trademarks are the property of their respective owners.