Life After SATA:
A New Strategy to Address SATA Limitations on CPU and Server Performance

WHITE PAPER

Choosing the Right SSD Increases Server Performance

The Serial Advanced Technology Attachment (SATA) is the most widely adopted storage interface used over the past several decades, serving storage requirements for a wide range of workloads due to its low-cost and simplicity. However, modern workloads now demand more and are driving storage administrators and infrastructure managers to develop new strategies that avoid stranding CPU cycles and DRAM utilization behind storage device bottlenecks.

Since its market introduction in 2000, the SATA interface has been the workhorse of server data storage. Beginning with hard drives, and more recently solid-state drives (SSDs), SATA-based servers have delivered 'good' performance—and at attractive price points—keeping many production servers humming even as data center computing workloads have expanded and diversified.

For data-intensive workloads such as online transaction processing (OLTP), hyperscale private clouds, artificial intelligence (AI), machine learning, analytics, data warehousing and more, even the most current SATA-based SSDs deployed within powerful servers will struggle to avoid being the performance bottleneck which slow CPU and server responses to unacceptable levels. Though SATA is still a good option for server workloads that don't require high IOPS or low latency, the last few generations of SATA have reached a practical interface plateau, limiting additional performance.

Given these SATA limitations, servers need to be re-architected and re-engineered to address the wide variety of data-intensive workloads that are now prevalent. With SATA on a decline, there are 'other' SSD types that can help organizations seamlessly get more out of their applications and utilize its servers to the fullest. Enter value SAS.

Value Serial-Attached SCSI (value SAS or vSAS)—an emerging class of SAS storage used for demanding workloads that need more performance and support to address higher application read/write rates and lower latencies than SATA, at price points that make it a direct SATA replacement.

DELL EMC KIOXIA TechTarget Custom Media

Delivering the Goods on I/O Operations per Second (IOPS)

Today's application workloads require higher disk reads and demand more Input/Output Operations per Second (IOPS) than ever. Without heightened SSD performance, bottlenecks can occur which throttle a user's ability to access critical data when they need it, and are a choke point for organizations.

Although SATA SSDs have provided an important performance boost over even the fastest hard drives, the steady influx of transaction-based workloads has raised the performance bar for infrastructure managers. The new breed of value SAS SSDs have created more performance-oriented platforms that deliver the high IOPS necessary to process and deliver critical data for read-intensive workloads.

This is particularly important for servers designed to address complex and high-performance workloads, such as Dell EMC™ PowerEdge™ servers. As these servers become the bedrock of modern data centers, enterprise workgroups and remote office/branch office (ROBO) environments, IT personnel have accelerated their search for IOPS-intensive storage solutions. High IOPS are essential for cloud and web-based workloads where many users need to be simultaneously serviced, resulting in an improved user experience.

As SATA SSDs offer significantly better performance when compared to hard drives, they do not meet the IOPS requirements for today's data-intensive applications. According to extensive laboratory testing[1] conducted by KIOXIA America Inc. **(KAI)** on separate Dell™ two-socket, rack mount PowerEdge servers, and emulating read-intensive workloads (such as OLTP), enterprise SATA SSDs can deliver an average of **1.37** million transactions per minute **(tpm)**. Under the same test configuration and environment, value SAS SSDs can deliver about **1.71** million transactions, or 25% more tpm than SATA, as depicted in Chart 1.

Value SAS vs Enterprise SATA Comparisons

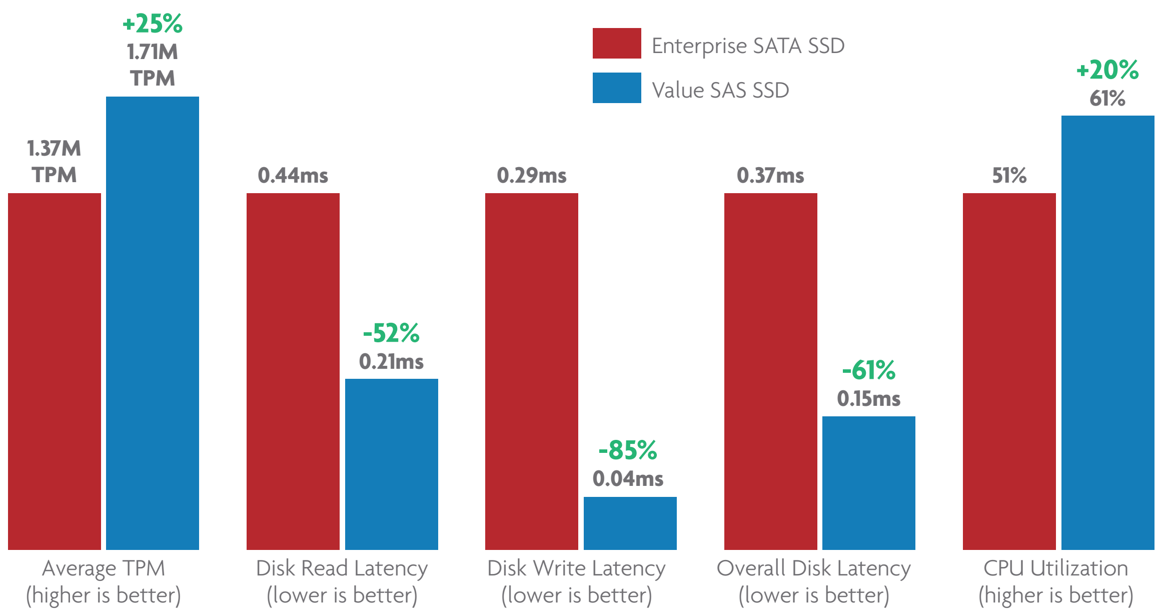

Chart 1: SSD comparisons between enterprise SATA and value SAS

Driving Down Latency

Low latency is another vital part of a modern storage solutions strategy that requires legacy SATA SSD replacement. For many data center workloads, latency can be a severely limiting factor since it can slow down data transfers, especially small block-sized transfers common with OLTP applications. The faster data can be moved, particularly in modern server configurations, the more transactions can be processed in a given session, and the more users that can be accommodated simultaneously.

From the KAI lab tests[1], enterprise SATA SSDs and value SAS SSDs were evaluated for latency where a lower value is better, and included:

- Read disk latency, or the time delay before data is read following an instruction for that transfer.
- Write disk latency, or the time delay before a data transfer begins writing the data following an instruction for that transfer.
- Overall disk latency, or the delay in time before a data transfer begins either reading or writing data following an instruction for that transfer and based on a mixed workload split covering 70% read operations and 30% write operations.

As it relates to disk read latency, enterprise SATA SSDs delivered an average of approximately 0.44 milliseconds (ms), while value SAS SSDs delivered 0.21ms of latency equating to a 52% improvement (Table 1). For disk write latency, enterprise SATA SSDs delivered an average of approximately 0.29ms, while value SAS SSDs delivered 0.04ms of latency equating to an 85% improvement (Table 2). For overall disk latency, enterprise SATA SSDs delivered an average of approximately 0.37ms, while value SAS SSDs delivered 0.15ms of latency equating to a 61% improvement (Table 3).

Read Disk Latency:

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

Table 1: Read disk latency comparison of enterprise SATA SSDs vs value SAS SSDs

Write Disk Latency:

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

Table 2: Write disk latency comparison of enterprise SATA SSDs vs value SAS SSDs

Overall Disk Latency:

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

Table 3: Overall disk latency comparison of enterprise SATA SSDs vs value SAS SSDs

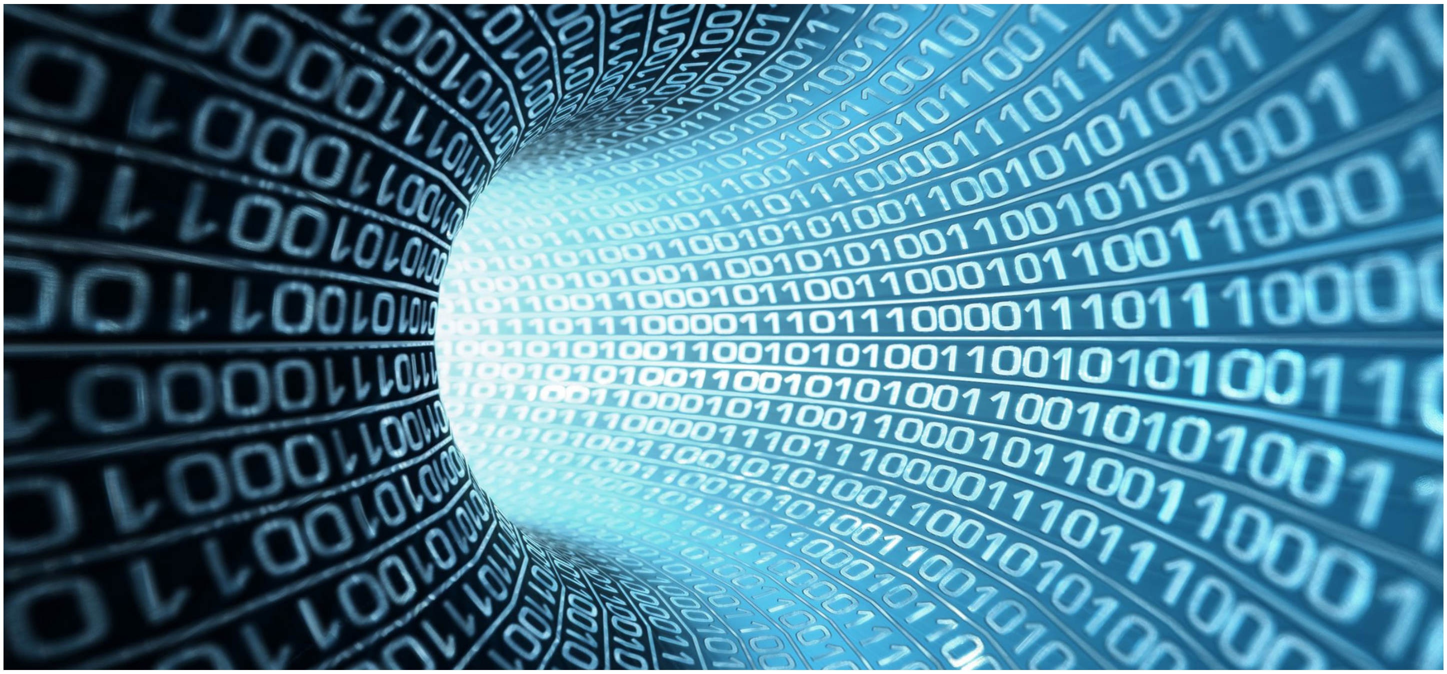

Obtaining More CPU Efficiency

The higher the CPU capacity utilization rate, the more transactions that are processed per CPU, resulting in a larger number of transactions and users that can be accommodated. CPU utilization directly correlates to performance and server efficiency so that additional servers don't need to be purchased and deployed in order to make up for performance limitations. This way, customers get the maximum performance and efficiency from their servers instead of leaving CPU cycles underutilized. Demanding workloads require high utilization of all the compute resources within a server and it is vital that SSDs be selected that enable workloads to efficiently utilize the available CPU capacity and obtain the most return on investment from the servers.

From the KAI lab tests[1], enterprise SATA SSDs and value SAS SSDs were evaluated to determine the impact of a slower storage interface on CPU usage in a read-intensive, I/O heavy workload. The enterprise SATA SSDs tested achieved CPU utilization of approximately 51% when tested on the Dell EMC PowerEdge server platform. When value SAS SSDs were deployed within the same server, the CPU utilization was tested at 61% (or about 20% higher CPU utilization (Chart 1). By breaking potential performance bottlenecks created by unacceptably low enterprise SATA CPU utilization, value SAS SSDs enable better use of all their server resources.

Summary of Results

From the test results, value SAS SSDs enabled the mainstream server platform under test to deliver 25% higher transactions per minute than comparable SATA drives under a TPC-C-like workload. The full duplex 12Gb/s bandwidth capability of value SAS enables a system to support higher transactions and utilize much more of the system's CPU, all while servicing transactions at up to 85% lower latencies. The result is that value SAS provides higher performance, lower latencies and improved CPU utilization that enable more application and server benefits, as well as a reduction in TCO.

As depicted in Chart 1, value SAS delivered the following advantages over enterprise SATA SSDs:

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

Supporting a SATA SSD Replacement Strategy

KIOXIA's performance testing validates the need for a storage-intensive workload strategy that replaces legacy enterprise-class 10K and 15K hard drives and SATA SSDs with value SAS SSDs. Its RM5 Series of value SAS SSDs are now available in Dell two-socket, rack mount PowerEdge servers to help organizations navigate to the best storage solutions for their present and future workloads.

For years, KIOXIA has been a long-standing market leader in SSD storage with the industry's broadest SSD product portfolio, a deep expertise and understanding of SSD technologies, and in developing strategies to meet customer storage needs. As such, they have helped many organizations develop "life after SATA" strategies to address SATA SSD limitations in I/O-intensive environments.

Conclusion

IT professionals, infrastructure decision-makers and data center managers need a long-term strategy to replace their SATA drives over time as I/O-intensive workloads demand higher performance and lower latency. While SATA drives still have their place for a number of workloads, their performance limitations and CPU underutilization begs for a better solution that not only addresses storage-intensive workloads, but also maximizes server capital expenditures. By planning, implementing and managing a long-term SATA SSD replacement strategy, IT decision-makers can help their organizations keep pace with rapidly changing workload requirements that are the basis of digital transformation.

To learn more about how KIOXIA can help you prepare for "life after SATA," visit https://business.kioxia.com/en-us/ssd/life-after-sata.html.

Footnotes

1 Test Criteria: The test system included separate Dell EMC two-socket, rack mount PowerEdge servers to host the Microsoft® SQL Server® 2017 database as well as the HammerDB application to avoid contention from either application. Without the testbed to conduct separate testing, the database workload would have interfered with memory and/or CPU cycles that could compromise database performance, as well as the test results. All benchmarks were tested in a Microsoft Windows Server® 2016 environment using four (4) Toshiba Memory RM5 value SAS SSDs with 960GB capacities and four (4) Intel® 4500 SSDs with 960GB capacities.

HammerDB load testing software was configured with a test schema based on a TPC-C benchmark (to emulate an OLTP environment). OLTP applications normally have many users that conduct simple, yet short transactions that require sub-second response times and return relatively few records, so the workload profile was designed for 70% read operations and 30% write operations. SQL Server was then loaded with one thousand data warehouses that comprised about 100GB of the server's storage capacity. Memory allocation was set to 10GB so no more than 10% of the 100GB database could be cached at one time. The HammerDB query response time was also reduced from a 500ms delay to 1ms, setting the demand generation delay to an immediate response time that provided the tested SSDs with substantially higher I/O demand. HammerDB was also changed from untimed testing to timed testing runs using 48 users where each test run included a 2-minute ramp up and two 5-minute test durations. Each test was run three times to record an average.