

Comparing SSD Interfaces in Servers

Which SSD Interfaces are Good, Better or Best?

Introduction

Given the ongoing quest to achieve faster data storage performance and heightened data reliability in the data center, flash-based solid-state drives (SSDs) are being used extensively in servers and external enterprise storage systems, replacing hard disk drives (HDDs) at a rapid pace. According to market research, total SSD revenue surpassed HDD revenue for the first time in 2017¹.


As the demand for higher data read and write performance and reliability increases, so does the need for more modern technologies that can keep up with today's data-centric storage and computing requirements. The most common SSD types are based on SATA (Serial Advanced Technology Attachment) and SAS (Serial Attached SCSI) interface protocols, but NVMe™ (Non-Volatile Memory Express), a newer protocol designed from the ground up for flash memory has risen to the occasion - supercharging SSDs to address today's data-intensive workloads.

SSDs continue to replace HDDs in servers requiring IT end users to decide what types will meet their varying workload and application demands. At present, there are three common SSD data transports used that connect servers to storage media and include SATA, SAS and PCIe® (Peripheral Component Interconnect Express). The NVMe specification uses PCIe as a transport for SSDs. So when tasked with choosing the right SSD for your server, one of the primary choices for selection will be the SSD interface itself.

SATA Overview

With the emergence of flash-based SSDs, SATA became a standard protocol and was soon positioned for direct HDD replacement. As other interface options emerged over time, SATA became the most common and cost-effective interface for SSDs despite some application performance delays. Though these delays were somewhat acceptable in hard drives where disk rotations and seek latencies are the norm, they may not be acceptable for SSDs that utilize flash-based memory cells.

SATA vs NVMe Stacks



When SATA SSDs are deployed within a server, I/O commands traverse through multiple layers (Figure 1) resulting in an inability to take complete advantage of the flash-based media and prevents users from leveraging the full performance capabilities of flash memory. As the SATA interface is half-duplex, it only uses one lane / one direction at a time for transferring data, so as CPUs get faster and DRAM bandwidth increases, the one lane can become a bottleneck. Servers with powerful, multicore processors and an abundance of RAM will be waiting for data transactions (reads or writes) to complete, and can result in an underutilization of compute resources that leave some of the capabilities of SSD flash on the table.

SATA SSDs have traditionally been popular being substantially less expensive than SAS or NVMe SSDs. The latest interface, SATA III, was launched in 2009 and can transfer data at 6Gb/s (or ~600MB/s). Though these SSDs are well-established, its adoption rate in servers is declining as they are being replaced by superior technologies.

Figure 1: Depicts the differences between SATA SSD and NVMe SSD Stacks

SAS Overview

SAS is a point-to-point serial protocol that moves data to and from storage devices using the SCSI command set. It is generally compatible with the SATA transport enabling SATA-based SSDs to connect to most SAS backplanes or controllers in servers. Since the majority of today's servers are equipped with a SAS infrastructure, SAS and SATA SSDs can be used in the same drive bay. SATA SSDs can be easily swapped out with SAS SSDs requiring no changes to the SAS-enabled server or infrastructure (unlike NVMe).

Generally, the SAS interface provides higher throughput than SATA and is geared toward applications that require data protection and high availability. It has a defined technology roadmap that currently supports larger capacities and higher performance capabilities with the new 24G SAS standard.

CPU Utilization

SAS SSDs enjoy several advantages over SATA SSDs as the SAS-3 full-duplex standard can transfer data bi-directionally at speeds up to 12Gb/s (or ~1,200MB/s) – performing reads and writes much faster than eSATA SSDs. It also supports hardware RAID (redundant array of independent disks) capabilities and multiple data paths to help enable high fault tolerance and data protection. SAS is also hot-swappable, has sophisticated and proven reliability, availability and serviceability (RAS), includes more over-provisioning options to prolong write life than SATA, and is designed to run in environments that require constant drive use. It also provides better software infrastructure management capabilities and better software interoperability with legacy operating systems in comparison to SATA. While a SATA drive could technically be used similarly to a SAS drive, performance cannot adequately support high data-intensive workload demands.

Introducing Value SAS

Value SAS (vSAS) is a new category of SAS SSD media that is expected to replace SATA SSDs in server applications. Initial SSD offerings deliver a theoretical throughput of 12Gb/s in a single port configuration with advancements in capacity, reliability, manageability and data security over SATA SSDs, at price points that compete favorably with them. Current offerings will utilize cutting-edge 96-layer BiCS FLASH™ 3D flash memory developed by KIOXIA, with capacities ranging from 960GB to 7,680GB². The new SAS category is also capable of delivering much faster bandwidth and input/output operations per second (IOPS) performance that include sequential reads at 840MB/s, sequential writes at 710MB/s, random reads at 160,000 IOPS and random writes at 50,000 IOPS. Value SAS also provides a very high level of data security with support for the Self-Encrypting Device (SED) standard and the Sanitize Instant Erase (SIE) standard. By replacing SATA drives with value SAS drives in servers, customers can recognize better system level performance per dollar.

PCIe/NVMe Overview

The PCIe interface can deliver higher data transfer rates than SATA or SAS interfaces due to more available channels (lanes) for data to flow. PCIe/NVMe SSDs are replacing SATA SSDs as the high-performance interface both in home computing and enterprise-level applications, and its big boost in IOPS helps to keep today's fast CPUs continually fed with data. These SSDs connect to servers and storage arrays through the PCIe interface bus which essentially serves as the backplane, effectively bypassing the host bus adapter (HBA) to make high-performance and low-latency its big advantages. Each PCIe Gen4 lane can move data at speeds up to 2GB/s, enabling 2.5-inch x4 lane PCIe SSDs that are capable of delivering data-intensive, low-latency content at exceptionally fast speeds.

The latest available PCIe SSDs feature support for the NVMe standard which goes well beyond the limited data transfer speeds that have plagued SATA SSD connections in the past. The NVMe specification provides efficient host processing of each command by supporting a queue depth of 64K commands in 64K queues (versus the SATA interface that uses a queue depth of 32 in a single command queue). By performing multicore processing of large I/O operations over multiple PCIe bus lanes, NVMe SSDs reduce bottlenecks and keep data flowing into and out of server CPUs, enabling users to tackle more demanding storage workloads quickly and efficiently.

“While SATA III theoretically provides 6Gb/s performance, and SAS provides 12Gb/s performance per lane, NVMe SSDs deliver data transfer performance of 2GB/s per lane, or upwards of 8GB/s in a typical x4-lane configuration”

While SATA III theoretically provides 6Gb/s performance, and SAS provides 12Gb/s performance per lane, NVMe SSDs deliver data transfer performance of 2GB/s per lane, or upwards of 8GB/s in a typical x4-lane configuration. SSDs based on NVMe should be more responsive to heavy workloads and more resilient against performance degradation due to too many I/O requests.

Introducing Data Center NVMe

Data center NVMe is a new category of PCIe/NVMe SSD media designed for read-intensive applications that delivers higher performance and capacities versus eSATA SSDs, but at a comparable price tag. These SSDs are expected to feature capacities up to 15.36TB with complete end-to-end data protection and emphasis on Quality of Service (QoS). Not all systems that support SATA will be able to convert to data center NVMe, but for systems that support NVMe, this option will likely drive the best system level performance per dollar.

Interface Performance

To better understand the performance benefits of each SSD interface, and provide a consistent comparison, the specified and practical performance capabilities of enterprise SATA, SAS and NVMe SSDs are now presented. Newer SSD categories, (e.g., value SAS and data center NVMe), are also included. The measurements assume server-side performance (versus communicating over a network) to provide the data locality benefits of direct-attached storage (high-performance and low-latency). Server-based storage houses storage media inside of servers and keeps the data local to the CPU for more immediate host processing.

Chart 1 depicts the interface limits and practical specifications associated with the three SSD protocols used to connect servers to enterprise SATA, SAS and NVMe SSDs. The **interface limit**³ represents the highest read or write performance that 'could' theoretically be obtained from an SSD and takes into account the physical limitations of the electrical connections. The **practical specification**³ accounts for typical performance at the product level and is based on benchmark testing in a lab environment.

“The measurements assume server-side performance to provide the data locality benefits of direct-attached storage (high-performance and low-latency)”

From the results in Chart 1, SSDs based on the SATA interface deliver the slowest read/write performance of the three protocols and may leave CPU resources underutilized and unable to realize their full performance potential from flash memory. To achieve the performance objectives that improve compute efficiencies, moving to enterprise SAS or NVMe SSDs provide heightened performance over SATA SSDs, but at a costlier price tag.

Chart 1: Enterprise SATA SSD Performance Comparison versus Enterprise SAS SSDs and Enterprise PCIe/NVMe SSDs

Read / Write Operation	6Gb/s SATA		12Gb/s Enterprise SAS		PCIe/NVMe Gen4x4 Enterprise	
	Interface Limit	Practical Specification	Interface Limit	Practical Specification	Interface Limit	Practical Specification
Sequential Read	600MB/s	550MB/s	1,200MB/s	1,100MB/s	8,000MB/s	6,900MB/s
Sequential Write	600MB/s	550MB/s	1,200MB/s	1,100MB/s	8,000MB/s	4,200MB/s
Random Read	146,000 IOPS	85,000 IOPS	293,000 IOPS	210,000 IOPS	1,923,000 IOPS	1,400,000 IOPS
Random Write	146,000 IOPS	50,000 IOPS	293,000 IOPS	90,000 IOPS	1,923,000 IOPS	350,000 IOPS

Chart 1: Interface limits and practical specifications comparing the read/write performance of enterprise SATA, SAS and PCIe/NVMe

A server with enterprise SATA SSDs deliver a practical read/write bandwidth of 550MB/s and 85,000 read IOPS. Enterprise SAS SSDs deliver a practical read/write bandwidth of 1,100MB/s and 210,000 read IOPS (or about 2x more bandwidth than SATA and 2.47x more IOPS). Enterprise PCIe/NVMe SSDs deliver a practical read bandwidth of 6,900MB/s, a practical write bandwidth of 4,200MB/s and 1,400,000 read IOPS (or about 12.54x more read bandwidth than SATA, about 7.63x more write bandwidth, and 16.47x more IOPS).

Charts 2 and 3 include two new categories of data center SSD media whose performance and price points are expected to replace SATA SSDs in server applications. The first is value SAS that deliver read/write performance advancements of 12Gb/s and can replace SATA with no changes to the server infrastructure. The second is data center NVMe where each PCIe Gen4 lane can move data at speeds up to a theoretical 8GB/s per drive in a typical x4 U.2 model. Value SAS is available today from KIOXIA America, Inc. and for those IT users ready to transition to NVMe, data center NVMe SSDs are being developed by KIOXIA as well.

Chart 2: Enterprise SATA SSD Performance Comparison versus Value SAS SSDs and Data Center PCIe/NVMe SSDs

Read / Write Operation	6Gb/s SATA		12Gb/s Value SAS		PCIe/NVMe Gen4x4 Data Ctr.	
	Interface Limit	Practical Specification	Interface Limit	Practical Specification	Interface Limit	Practical Specification
Sequential Read	600MB/s	550MB/s	1,200MB/s	840MB/s	8,000MB/s	6,700MB/s
Sequential Write	600MB/s	550MB/s	1,200MB/s	710MB/s	8,000MB/s	4,500MB/s
Random Read	146,000 IOPS	85,000 IOPS	293,000 IOPS	160,000 IOPS	1,923,000 IOPS	1,100,000 IOPS
Random Write	146,000 IOPS	50,000 IOPS	293,000 IOPS	50,000 IOPS	1,923,000 IOPS	180,000 IOPS

Chart 2: Interface limits and practical specifications comparing read/write performance of enterprise SATA, value SAS and data center PCIe/NVMe

Chart 3: Value SAS SSDs and Data Center PCIe/NVMe SSD Performance Gains versus Enterprise SATA SSD

Read / Write Operation	6Gb/s SATA	12Gb/s Value SAS		PCIe/NVMe Gen4x4 Data Ctr.	
	Practical Specification	Practical Specification	Performance Gain Over 6Gb/s SATA	Practical Specification	Performance Gain Over 6Gb/s SATA
Sequential Read	550MB/s	840MB/s	1.52x	6,700MB/s	12.18x
Sequential Write	550MB/s	710MB/s	1.29x	4,500MB/s	8.18x
Random Read	85,000 IOPS	160,000 IOPS	1.88x	1,100,000 IOPS	12.94x
Random Write	50,000 IOPS	50,000 IOPS	1.00x	180,000 IOPS	3.60x

Chart 3: Read/write performance gains of value SAS and data center PCIe/NVMe versus enterprise SATA

From the Chart 2 and Chart 3 results, a server with enterprise SATA SSDs deliver a practical read/write bandwidth of 550MB/s and 85,000 read IOPS. Value SAS SSDs deliver a practical read bandwidth of 840MB/s, a practical write bandwidth of 710MB/s, and 160,000 IOPS (or about 1.52x more read bandwidth than SATA, about 1.29x more write bandwidth and 1.88x more read IOPS). Data center NVMe SSDs deliver a practical read bandwidth of 6,700MB/s, a practical write bandwidth of 4,500MB/s and 1,100,000 read IOPS (or about 12.18x more read bandwidth than SATA, about 8.18x more write bandwidth, and 12.94x more read IOPS).

The results speak for themselves as enterprise SATA SSDs deliver **GOOD** read/write performance at cost-effective price points. However, value SAS delivers **BETTER** read/write performance and data center NVMe delivers the **BEST** read/write performance, both at comparable price points to enterprise SATA, representing viable replacement options.

Conclusion

These SSD interface bandwidth and performance results showcase how many operations or transactions that the server's CPU can process. The more bandwidth and IOPS that a server processes, the more throughput and operations/transactions that can be completed, as well as the number of users that can be serviced at one time, with better Quality of Service (QoS). The end result is more application performance, a reduction in IT costs and a better overall user experience.

SSD performance also affects CPU utilization, so when higher application performance is obtained, the server gets more out of its CPU. With SATA reaching its performance limits, SSDs based on SATA may become a bottleneck to the server, starving the CPU from performing more operations or transactions, thereby underutilizing its compute capabilities. Value SAS and data center NVMe break this paradigm and enable better use of all server resources leading to a better Total Cost of Ownership (TCO).

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

¹ Source: IDC Research, "Worldwide Enterprise SSD and HDD Combined Market Overview," 2018-2022, June 2018, document #US44076018, Slide 4.

² Definition of capacity: KIOXIA defines 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 1GB = 2³⁰ bytes = 1,073,741,824 bytes, 1TB = 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, such as Microsoft Operating System and/or pre-installed software applications, or media content. Actual formatted capacity may vary.

³ The specified and practical performance capabilities of enterprise SATA, SAS and PCIe SSDs were presented in charts on page 3, and included newer SSD categories such as value SAS and data center NVMe. In Chart #1 and Chart #2, the interface limit represents the highest read or write performance that 'could' theoretically be obtained from an SSD and takes into account the physical limitations of the electrical connections. The practical specification accounts for typical performance at the product level and is based on benchmark testing in a lab environment. An Online Transaction Processing (OLTP) application was used for measurement of server-side performance (versus communicating over a network) to provide the data locality benefits of direct-attached storage being high-performance and low-latency. The results showcase SSD interface bandwidth and performance and how many operations/transactions that a server's CPU can process.

The performance measurements were derived from existing KIOXIA America, Inc. SSD products and include the following:

- The 6Gb/s enterprise SATA SSD performance numbers are based on the KIOXIA HK6 Series of SATA SSDs
- The 12Gb/s enterprise SAS SSD performance numbers are based on the KIOXIA PM6 Series of SAS SSDs
- The 12Gb/s value SAS SSD performance numbers are based on the KIOXIA RM6 Series of SAS SSDs
- The Gen4 x4 enterprise PCIe/NVMe SSD performance numbers are based on the KIOXIA CM6 Series of NVMe SSDs
- The Gen4 x4 data center PCIe/NVMe SSD performance numbers are based on the KIOXIA CD7 Series of NVMe SSDs (in development)

NVMe is a trademark of NVM Express, Inc. PCIe is a registered trademark of PCI-SIG. All other trademarks or registered trademarks are the property of their respective owners.

© 2021 KIOXIA America, Inc. All rights reserved. Information in this best practices document, including product specifications, tested content, and assessments are current and believed to be accurate as of the date that the document was published, but is subject to change without prior notice. Technical and application information contained here is subject to the most recent applicable KIOXIA product specifications.