KIOXIA

Best Practices

Good, Better, Best SSDs for Servers

Comparing the Latest SSD Interfaces and Associated Benefits[1]

Introduction

Flash-based solid-state drives (SSDs) achieve faster data storage performance and provide more data reliability in the data center when compared with hard disk drives (HDDs). As the demand for higher data read/write performance and reliability increases, so does the need for advanced SSD technologies that can keep up with today's data-centric storage and computing requirements. Faster data transport interface technologies include PCIe® (Peripheral Component Interconnect Express) for both enterprise and data center applications, as well as 12 gigabits per second (Gb/s) value SAS (Serial Attached SCSI).

Before these faster SSD interface transports became available, the SATA (Serial Advanced Technology Attachment) interface was the common interface for SSDs and was based on a hard drive command set developed decades ago. It has an interface data transfer speed of up to 6 Gb/s, or approximately 550 to 600 megabytes per second (MB/s), and is much slower than today's PCIe and SAS interfaces. Additionally, when storing or retrieving data from SATA, redundant disk rotation and seek latency commands are performed which slow drive performance and limit the benefits of flash storage. With no planned SATA performance enhancements, IT end users should research the newer, faster and more reliable SSD technologies available for their server applications.

Featured within this Best Practices document is an overview of the NVM Express™ (NVMe™) protocol that includes an SSD command set designed specifically for flash memory storage and uses the PCIe interface for data transport. SSDs based on the PCIe interface and NVMe protocol represent the fastest SSD throughput and input/output operations per second (IOPS) to date. The latest PCIe 5.0 interface and NVMe 2.0 protocol enables SSDs to deliver bandwidths up to 14,000 MB/s. When compared with the bandwidth of a fast enterprise SATA SSD at 550 MB/s, enterprise NVMe SSD performance can be up to 25 times faster.

This Best Practices document also covers a new category of SAS SSDs called value SAS that delivers an interface data transfer speed up to 12 Gb/s, or up to twice the interface throughput when compared with 6 Gb/s SATA SSDs. IT end users can recognize better system level performance per dollar by replacing SATA drives with value SAS drives in servers.

Since the interface connects servers to storage media, choosing the right SSDs for server applications also requires choosing the right SSD data transport. The purpose of this Best Practices document compares the latest SSD interface data transports and their associated benefits to determine which SSD types are good, better or best for server configurations.

GOOD

Value SAS SSDs

Key Specifications

Interface	12 Gb/s SAS
Capacities	up to 7.68 TB*
Sequential Read	up to 1,100 MB/s
Sequential Write	up to 1,050 MB/s
Random Read	up to 190K IOPS
Random Write	up to 55K IOPS

KIOXIA RM7 Series SSD specifications[2]

BETTER

Data Center NVMe SSDs

Key Specifications

Interface	PCIe 4.0 / NVMe 1.4
Capacities	up to 15.36 TB*
Sequential Read	up to 7,200 MB/s
Sequential Write	up to 6,000 MB/s
Random Read	up to 1,250K IOPS
Random Write	up to 200K IOPS

KIOXIA CD8-R Series SSD specifications[3]

BEST

Enterprise NVMe SSDs

Key Specifications

Interface	PCIe 5.0 / NVMe 2.0
Capacities	up to 30.72 TB*
Sequential Read	up to 14,000 MB/s
Sequential Write	up to 7,000 MB/s
Random Read	up to 2,700K IOPS
Random Write	up to 310K IOPS

KIOXIA CM7-R Series SSD specifications[4]

Compared with Enterprise SATA SSDs

Key Specifications[5]

Interface	6 Gb/s SATA
Capacities	up to 3.84 TB*
Sequential Read	up to 550 MB/s
Sequential Write	up to 520 MB/s
Random Read	up to 98K IOPS
Random Write	up to 30K IOPS

**TB = terabytes*[6]

Enterprise SATA SSD Overview

When deployed within a server, SATA input/output (I/O) commands traverse through multiple layers in the SATA SSD stack to complete read or write operations. This results in an inability to take advantage of flash media and its capabilities. The SATA interface is also half-duplex and only uses one lane/one direction at a time to transfer data. As modern CPUs get faster and DRAM bandwidth increases, the one SATA lane can become a performance bottleneck. Servers with powerful, multicore processors and an abundance of RAM could be waiting for data transactions (reads and writes) to complete, resulting in an underutilization of CPU resources and SSD flash capabilities.

SATA SSDs are less expensive in acquisition cost than the faster SSD technologies. The latest SATA-3 interface was launched in 2009 with a data transfer rate of up to 6 Gb/s. Though SATA SSDs are a cost-conscious data storage option, their adoption rate in servers has declined and replaced by the newer and faster SSD interface technologies (Figure 1).

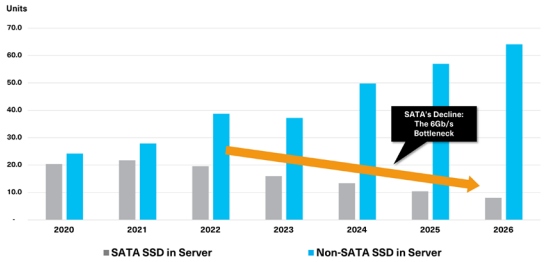

Figure 1: Decline of enterprise SATA SSDs in servers
(Source: Forward Insights, Q123, Enterprise Server Units)

Value SAS: SAS Reliability and Performance, Priced to Replace SATA

SAS is a point-to-point serial protocol that moves data to and from storage devices using the SCSI command set. It supports full-duplex bidirectional data transfers and performs read and write operations faster than enterprise SATA SSDs. In general, SAS is backward compatible with the SATA data transport enabling SATA SSDs to connect to most SAS backplanes or server controllers. Since the majority of today's servers are equipped with a SAS infrastructure, SAS and SATA SSDs can reside in the same drive bay and are easily swappable requiring no physical changes to the server or infrastructure. The ability to replace SATA drives with SAS drives in servers enables better system level performance per dollar.

SAS SSDs also support hardware redundant array of independent disks[7] (RAID) capabilities and multiple data paths that help enable high fault tolerance and data protection. These SSDs have sophisticated and proven reliability, availability and serviceability, include many over-provisioning options to prolong write life, and can run in environments that require constant drive use.

Value SAS SSDs represent a new category of SAS that delivers an interface data transfer speed of up to 12 Gb/s in a single port configuration and includes advancements in reliability, manageability and data security over enterprise SATA SSDs. Value SAS SSD price points enable them to compete favorably with enterprise SATA SSDs and deliver faster bandwidth and IOPS performance. Using KIOXIA RM7 Series value SAS SSDs[2] as an example, they deliver sequential read performance up to 1,100 MB/s, sequential write performance up to 1,050 MB/s, random read performance up to 190,000 IOPS and random write performance up to 55,000 IOPS. These value SAS SSDs also provide high-level data security with optional support for the Self-Encrypting Device[8] (SED) standard and the Sanitize Instant Erase[9] (SIE) standard. RM7 Series value SAS SSDs utilize BiCS FLASH™ 3D flash memory developed by KIOXIA Corporation and support capacities up to 7.68 TB.

KIOXIA RM7 Series SSD[10]

Speedburner SSDs: PCIe Snapshot

The PCIe interface delivers faster data transfer rates than SATA or SAS interfaces and has more channels (lanes) for data to flow. SSDs based on the PCIe interface are replacing SATA SSDs as the high performance interface in home computing and enterprise applications. Its big boost in IOPS performance helps keep today's fast CPUs continually fed with data. NVMe SSDs connect to servers through the PCIe interface bus that essentially serves as the backplane, effectively bypassing the host bus adapter to make high performance and low latency its big advantages. The PCIe 4.0 interface can move data through four lanes at speeds up to 8 gigabytes per second (GB/s), which enables 2.5-inch[11] NVMe SSDs to deliver data intensive, low latency content at exceptionally fast speeds. The PCIe 5.0 interface is approximately twice as fast and can move data through four lanes at speeds up to 16 GB/s.

The PCIe interface is the data transport for SSDs and the NVMe standard is the instruction set that provides capabilities 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. This enables users to tackle their more demanding storage workloads quickly and efficiently and be more resilient

against performance degradation due to too many I/O requests. NVMe SSDs include data center and enterprise classes with each having significant performance benefits when compared with enterprise SATA SSDs.

New Category: Data Center NVMe SSDs

Data center NVMe is a new category of SSD media designed for read intensive applications and delivers higher performance and capacities when compared with enterprise SATA SSDs, but at a comparable price tag. These SSDs feature capacities up to 15.36 TB with end-to-end data protection and emphasis on Quality of Service (QoS). At present, they typically use the PCIe 4.0 interface and NVMe 1.4 protocol.

KIOXIA Corporation recently launched its latest generation of CD8 Series data center NVMe SSDs positioned for general server, scale-out and cloud applications. The series is available in a 2.5-inch form factor with supported capacities up to 15.36 TB. The two configurations include KIOXIA CD8-R Series for read intensive workloads with 1 drive write per day[12] (DWPD) support and KIOXIA CD8-V Series for higher endurance mixed use workloads with 3 DWPD support.

KIOXIA CD8 Series SSD[10]

KIOXIA CD8 Series data center NVMe SSDs[3] deliver sequential read performance up to 7,200 MB/s, sequential write performance up to 6,000 MB/s, random read performance up to 1,250,000 IOPS and random write performance up to 200,000 IOPS. These SSDs optionally support the SED[8] and SIE[9] standards to provide high-level data security and utilize BiCS FLASH 3D flash memory developed by KIOXIA Corporation. They offer a balance of high performance, high reliability and low latency, and deliver outstanding system level performance per dollar.

Top Shelf: Enterprise NVMe SSDs

With the advent of the PCIe 5.0 specification revision, data can move through the PCIe interface approximately twice as fast when compared with the previous PCIe 4.0 generation. This enables supported SSDs to deliver I/O even faster than before. The speed upgrade is beneficial to data-intensive and computational enterprise applications such as cloud computing, databases, data analytics, artificial intelligence, machine learning, container orchestration and media streaming, to name a few. SSDs in this category push the speed limits of the PCIe 5.0 interface and deliver the highest performance, lowest latencies and highest capacities of the SSD classes. Key server and SSD vendors are developing solutions to comply with the PCIe 5.0 standard (and associated NVMe 2.0 protocol).

KIOXIA CM7 Series [10]

The PCIe 5.0 interface standard increases data transfer speeds from up to 16 gigatransfers per second (GT/s) to up to 32 GT/s per lane. In other words, the PCIe 5.0 interface can move data at approximately 4 GB/s per lane versus the almost 2 GB/s per lane supported by the PCIe 4.0 interface. As a result, PCIe 5.0 NVMe SSDs are able to deliver bandwidths of 14,000 MB/s. When compared to top-end enterprise SATA SSD bandwidths at 550 to 600 MB/s, PCIe 5.0 performance is significant.

The doubling of PCIe 5.0 performance enables SSDs to communicate faster with system CPUs and transfer large chunks of data with low latency. It also means that supported SSDs may be able to achieve the same throughput with fewer lanes, with more lanes available for other workloads. By utilizing half of the PCIe lanes, the PCIe 5.0 interface enables many more devices in a system versus the PCIe 4.0 interface. The increased performance gains of the PCIe 5.0 interface also have a positive effect on energy efficiency within a system.

KIOXIA Corporation recently launched its latest generation CM7 Series enterprise NVMe SSDs with support for 2.5-inch and new Enterprise and Datacenter Standard Form Factor (EDSFF) E3.S form factors and for the PCIe 5.0 specification and the NVMe 2.0 protocol. These SSDs are available in two configurations – KIOXIA CM7-R Series for read intensive workloads with 1 DWPD support and capacities up to 30.72 TB, and KIOXIA CM7-V Series for higher endurance mixed use workloads with 3 DWPD support and capacities up to 12.80 TB. Additional features include a dual-port design for high availability (HA) applications, flash die failure protection to maintain reliability in case of a die failure, and high-level data security with optional support for the SED and SIE standards. KIOXIA CM7 Series enterprise NVMe SSDs utilize BiCS FLASH 3D flash memory developed by KIOXIA Corporation.

Using KIOXIA CM7 Series enterprise NVMe SSDs[4] to demonstrate performance, they deliver sequential read performance up to 14,000 MB/s, sequential write performance up to 7,000 MB/s, random read performance up to 2,700,000 IOPS and random write performance up to 600,000 IOPS. These SSDs deliver the highest performance per dollar and highest performance per watt versus other KIOXIA SSD classes.

Summary

Today's data storage requires faster read/write performance and more data reliability than SATA can provide in order to keep pace with data-centric storage and large computing requirements. PCIe and SAS interface technologies are delivering significantly faster data transfers enabling higher SSD throughput and IOPS performance as presented in Table 1. SATA, being a legacy hard drive technology without performance improvements in the future, cannot compete with these faster SSD interface technologies. Relating to both enterprise and data center NVMe SSDs the performance gains versus enterprise SATA SSD performance are significant.

Read/Write Operation[13]	6 Gb/s SATA	12 Gb/s Value SAS	Value SAS Gains vs. SATA	Data Center NVMe	Data Center NVMe Gains vs. SATA	Enterprise NVMe	Enterprise NVMe Gains vs. SATA
SeqRead	550 MB/s	1,100 MB/s	2x	7,200 MB/s	13x	14,000 MB/s	25.4x
SeqWrite	520 MB/s	1,050 MB/s	2x	6,000 MB/s	11.5x	7,000 MB/s	13.4x
RanRead	98K IOPS	190K IOPS	1.9x	1,250K IOPS	12.7x	2,700K IOPS	27.5x
Ran Write	30K IOPS	55K IOPS	1.8x	200K IOPS	6.6x	310K IOPS	10.3x

Table 1: Performance comparison between faster SSD interface technologies and enterprise SATA SSDs

From Table 1, a server deployed with enterprise SATA SSDs[5] delivers about 550 MB/s read throughput, 520 MB/s write throughput, 98,000 read IOPS and 30,000 write IOPS. In comparison, a server deployed with value SAS SSDs delivers 2x higher read bandwidth, 2x higher write bandwidth, 1.9x higher read IOPS and 1.8x higher write IOPS. Value SAS SSDs deliver **GOOD** read/write performance at cost-effective price points.

When data center NVMe SSD performance is compared with enterprise SATA SSDs, they deliver exceptional read/write performance that includes 13x higher read bandwidth, 11.5x higher write bandwidth, 12.7x higher read IOPS and 6.6x higher write IOPS. These SSDs feature capacities up to 15.36 TB and balance high performance, high reliability and low latency at optimized price points making them **BETTER** than enterprise SATA or value SAS SSDs.

When enterprise NVMe SSD performance is compared with enterprise SATA SSDs, they deliver incredible read/write performance that includes a whopping 25.4x higher read bandwidth, 13.4x higher write bandwidth, 27.5x higher read IOPS and 10.3x higher write IOPS. These SSDs feature capacities up to 30.72 TB and possess the highest performance, highest reliability and lowest latency making them the **BEST** SSD option when compared with enterprise SATA SSDs, value SAS SSDs and data center NVMe SSDs.

General information for KIOXIA SSD products referenced in this paper:

KIOXIA RM7 Series Value SAS SSDs
KIOXIA CD8-R Series Data Center NVMe SSDs
KIOXIA CM7-R Series Enterprise NVMe SSDs

FOOTNOTES:

[1] This Best Practices document is an update from an original version published in January 2019.

[2] Value SAS specifications provided by KIOXIA Corporation using KIOXIA RM7 Series SSDs. Read and write speed may vary depending on the host device, read and write conditions, and file size.

[3] Data center NVMe SSD specifications provided by KIOXIA Corporation using KIOXIA CD8-R Series SSDs. Read and write speed may vary depending on the host device, read and write conditions, and file size.

[4] Enterprise NVMe SSD specifications provided by KIOXIA Corporation using KIOXIA CM7-R Series SSDs. Read and write speed may vary depending on the host device, read and write conditions, and file size.

[5] Enterprise SATA SSD specifications obtained from online and available published data supplied by a leading enterprise SATA SSD vendor. Read and write speed may vary depending on the host device, read and write conditions, and file size.

[6] Definition of capacity - KIOXIA Corporation 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 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.

[7] Redundant array of independent disks (RAID) is a technology that stores data in different places on multiple SSDs to protect the data in the case of a drive failure.

[8] Optional security feature drives are not available in all countries due to export and local regulations.

[9] Sanitize Instant Erase (SIE) is compatible with the Sanitize device feature set and the standard prescribed feature set by T10 (SAS) and T13 (SATA) committees of the American National Standards Association (ANSI) which makes it possible to invalidate data recorded on magnetic disks at a blink of an eye.

[10] The product image shown is a representation of the design model and not an accurate product depiction.

[11] 2.5-inch indicates the form factor of the SSD and not its physical size.

[12] 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, under the specified workload for the specified lifetime. Actual results may vary due to system configuration, usage, and other factors.

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

TRADEMARKS:

NVMe and NVM Express are registered or unregistered trademarks of NVM Express, Inc. in the United States and other countries. PCIe is a registered trademark of PCI-SIG. All other company names, product names and service names may be trademarks or registered trademarks of their respective companies.

DISCLAIMERS:

 Information in this Best Practices document, including product specifications, tested content, and assessments are current and believed to be accurate as of the publication date of the document, but is subject to change without prior notice. Technical and application information contained here is subject to the most recent applicable KIOXIA product specifications. Images within are for illustration purposes only.