KIOXIA

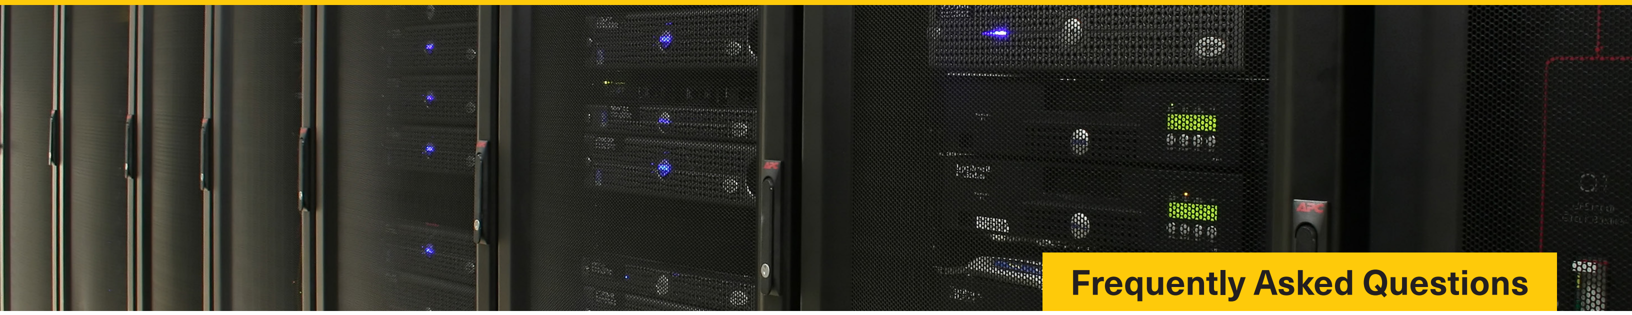

Frequently Asked Questions

Data Center NVMe™ SSDs: A New Category of SSD and Why You Need It

What is data center NVMe?

Data center NVMe™ is a newer category and subset of PCIe® NVMe solid-state drives (SSDs). Both data center and enterprise NVMe SSDs are supported by an industry consortium (NVM Express™) and an evolving NVMe specification that goes well beyond the limited data transfer speeds and capabilities that have plagued SATA SSD connections in the past. Most data center NVMe SSDs are designed for read-intensive applications and deliver much higher performance and larger capacities versus enterprise SATA SSDs, but within a comparable price tag.

Similar to enterprise NVMe SSDs, data center NVMe SSDs connect to a server through the PCIe interface bus where each PCIe lane is full duplex and can move data bi-directionally at speeds up to 1 gigabyte per second (GB/s). This enables these SSDs to perform multicore processing of large I/O operations over multiple PCIe lanes, thus reducing bottlenecks significantly while keeping data flowing into and out of server CPUs. Data center NVMe SSDs handle demanding storage workloads faster and a multitude of I/O requests more efficiently when compared to enterprise SATA.

In compliance with the NVMe specification, data center NVMe SSDs increase storage performance in servers by eliminating the SCSI command stack translation and direct-attached storage (DAS) bottlenecks typically associated with legacy SSD interfaces. Today's data center NVMe SSDs feature capacities up to 7.68 terabytes (TB)[1], provide complete end-to-end data protection with an emphasis on Quality of Service (QoS), and are positioned as a replacement to enterprise SATA SSDs in servers.

The entire NVMe SSD market segment, including enterprise and data center SSDs, is forecasted to represent about 45% of the total SSD consumption in the enterprise (in units) by the end of 2019[2]. The consumption is expected to grow to 78% by the end of 2021, and 91% by the end of 2023[2]. Conversely, enterprise SATA SSDs will represent about 42% of total enterprise SSD consumption by the end of 2019, and is projected to rapidly decline to 13% by the end of 2021, and 4% by the end of 2023[2].

Why is a new type of SSD needed?

Enterprise SATA SSDs are widely deployed in servers based on their cost-effectiveness, performance versus hard drives, and used mostly for application workloads that don't require exceptionally high performance or overly large capacities. The SATA interface itself has reached a performance plateau with no speed improvements planned in its future. For many applications, SATA SSDs may become a bottleneck to server CPUs, preventing them from reaching their operational or transactional potential. The end result may be an underutilization of the server's compute capabilities which can also limit the server's ability to service many clients simultaneously. An underutilization of the CPU reduces total cost of ownership (TCO) and application performance. SATA SSD consumption is expected to decline rapidly in servers as today's applications and workloads have far surpassed SATA's performance capabilities (Figure 1).

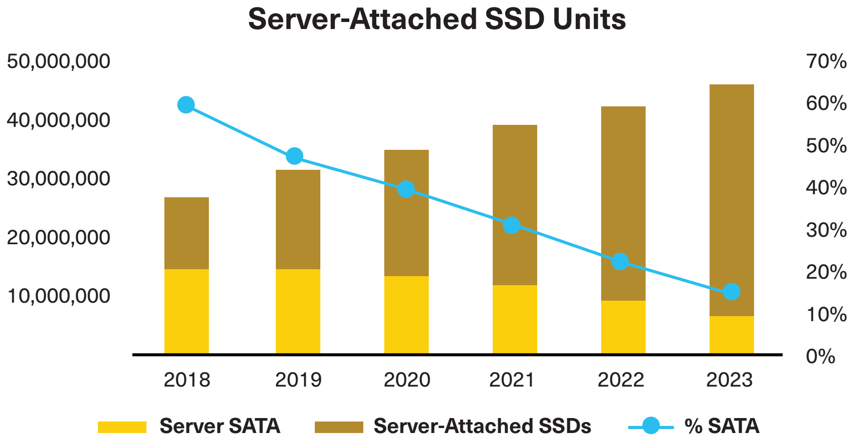

Figure 1: The decline of SATA SSDs in servers[3]
(source: Forward Insights, SSD Forecast-SSD Insights, Q3/19-FI-NFL-SSD-Q319)

In what areas are data center NVMe SSDs better than enterprise SATA SSDs?

Performance is the overwhelming benefit of data center NVMe SSDs when compared to enterprise SATA SSDs. For example, with PCIe Gen3, data center NVMe SSDs can transfer data at 1GB/s per lane, or upwards of 4GB/s in an x4-lane configuration. This has an immediate leg up on 6 gigabits per second (Gb/s) enterprise SATA SSDs relating to bandwidth and throughput. Performance comparisons between enterprise SATA and data center NVMe SSDs include both the ***interface limit*** (which represents the highest performance that 'could' theoretically be obtained from the SSD taking into account the physical limitations of the electrical connections), and the ***practical specification*** (that takes into account typical performance at the product level based on benchmark testing in a lab environment).

When tested in a KIOXIA lab environment[4], data center NVMe SSDs delivered 5.45x more read bandwidth than SATA, as well as 3.27x more write bandwidth, 5.88x more random read IOPS, and 1.43x more random write IOPS (Chart 1).

	Enterprise SATA SSDs 6Gb/s (1 DWPD*)		Data Center NVMe SSDs Gen3x4 (1 DWPD*)		Data Center NVMe SSD Advantage	
Read / Write Operation	**Interface Limit**	**Practical Specification**	**Interface Limit**	**Practical Specification**	**Interface Limit**	**Practical Specification**
Sequential Read	600MB/s	550MB/s	4,000MB/s	3,000MB/s	6.67x	5.45x
Sequential Write	600MB/s	550MB/s	4,000MB/s	1,800MB/s	6.67x	3.27x
Random Read	146,000 IOPS	85,000 IOPS	976,000 IOPS	500,000 IOPS	6.68x	5.88x
Random Write	146,000 IOPS	35,000 IOPS	976,000 IOPS	50,000 IOPS	6.68x	1.43x

Chart 1: Performance comparisons between enterprise SATA and data center NVMe SSDs
**Drive Write per Day*

The NVMe interface provides efficient host processing of each storage instruction by supporting a queue depth of 64K commands in 64K queues versus the SATA interface that supports a queue depth of 32 in a single command queue. Since the PCIe interface serves as the backplane for connecting data center NVMe SSDs to the server, the SAS or SATA host bus adapter (HBA) is eliminated, making high-performance and low-latency its key advantages (Figure 2).

In addition to the performance advantages, data center NVMe SSDs also include capacities that range from 960 gigabytes (GB[1]) to 7,680GB, where most SATA-based SSDs max-out at 3,840GB capacity. Multiple levels of data security are also available within NVMe SSDs that can include SIE (Sanitize Instant Erase), self-encrypting drive (SED) with TCG-Opal encryption and SED FIPS 140-2 (Level 2) support. The NVMe interface was designed from the ground up to support flash-based SSD implementations, while SATA was designed for consumer-grade HDDs.

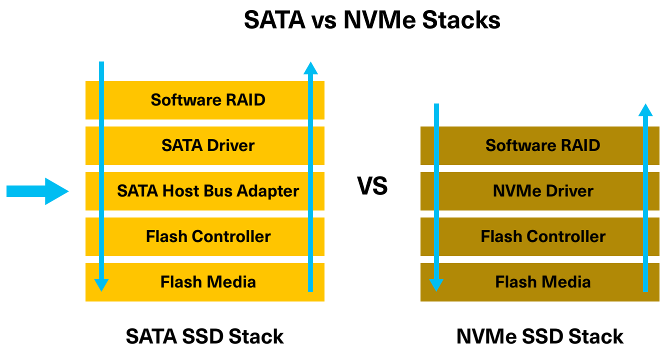

Figure 2: The host bus adapter is bypassed in the NVMe SSD stack

Are data center NVMe SSDs more effective than enterprise SATA SSDs?

Enterprise SATA SSDs utilize an instruction set originally developed for hard drives (AHCI). When SATA SSDs are deployed within a server, the I/O commands must traverse through a software stack (Figure 2) that was not designed for flash-based SSDs. Servers with powerful, multicore processors and an abundance of DRAM, may be left waiting for data transactions to complete, resulting in an underutilization of compute and storage resources.

Many servers today are designed with a SAS infrastructure that also supports SATA devices. This capability enables enterprise SATA SSDs to connect to SAS backplanes, HBAs or RAID controllers, allowing both SAS and SATA SSDs to be used in the same drive bay. When SATA SSDs use the server's SAS infrastructure and hardware RAID capabilities, protocol translation from SATA to SAS adds even more latency. Data center NVMe SSDs use native NVMe from end-to-end, eliminating the need for protocol translation, resulting in improved performance.

How do data center NVMe SSDs differ from enterprise NVMe SSDs?

Data center SSDs are high volume drives with little to no custom firmware or specialized features. If the drive fails, the user disposes it and replaces it with another. With fewer bells and whistles, and less performance and endurance, users expect to pay less for data center SSDs when compared to their more fully-featured enterprise counterparts.

Enterprise SSDs warrant a higher cost as they deliver the highest performance an SSD can currently achieve, and include features like data integrity checking, media wear reporting and error reporting. With these added capabilities, customers are willing to pay a premium for higher performance and improved data protection.

KIOXIA offers both data center NVMe SSDs in their CD5 Series and enterprise CM5 Series NVMe SSDs. Both NVMe SSD platforms are developed using 64-layer BiCS FLASH™ 3D technology and available in 960GB to 7,680GB capacities - a 15,360GB capacity CM5 Series SSD is also available. Performance differences[5] are shown in Chart 2.

	Data Center NVMe SSD CD5 Series (1 DWPD*)		Enterprise NVMe SSD CM5 Series (1 DWPD*)		CM5 Series Advantages	
Read / Write Operation	**Interface Limit**	**Practical Specification**	**Interface Limit**	**Practical Specification**	**Interface Limit**	**Practical Specification**
Sequential Read	4,000MB/s	3,000MB/s	4,000MB/s	3,350MB/s	1x	1.17x
Sequential Write	4,000MB/s	1,800MB/s	4,000MB/s	3,040MB/s	1x	1.69x
Random Read	976,000 IOPS	500,000 IOPS	976,000 IOPS	770,000 IOPS	1x	1.54x
Random Write	976,000 IOPS	50,000 IOPS	976,000 IOPS	80,000 IOPS	1x	1.60x

Chart 2: Performance comparisons between data center NVMe SSDs and enterprise NVMe SSDs
**Drive Write per Day*

In addition to its performance advantages, the CM5 Series has dual port support for high availability where the CD5 Series is single-ported and designed for DAS environments.

Are enterprise SATA SSDs easy to replace with data center NVMe SSDs?

Enterprise SATA SSDs can connect to SAS backplanes, HBAs or RAID cards enabling both SAS and SATA SSDs to be used in the same drive bay without any required changes to the server or infrastructure. A data center NVMe SSD cannot connect directly to a SAS RAID controller or a SAS/SATA backplane. Though data center NVMe SSDs offer even better drive performance than enterprise SAS or SATA, the server infrastructure and hardware RAID options available today are somewhat limited. To migrate over to data center NVMe SSDs will require purchasing new NVMe-enabled servers.

What is the future of data center NVMe SSDs?

The NVMe SSD segment, including enterprise and data center, is expected to dominate SSD consumption in the enterprise with 91% deployments by the end of 2023[2], supporting both on-premises and cloud data infrastructures. These architectures are general-purpose and utilize a pre-configured allocation of compute and storage resources to meet a variety of workload demands. As these data-intensive applications proliferate on a larger scale, the one-size-fits-all infrastructure comes up short relating to performance, capacity and scalability. Future cloud data centers require a disaggregation of compute, storage and networking resources so organizations can best utilize all of these resources. To enable these gains, these resources must be able to scale independently of each other.

Moving from a DAS architecture to a disaggregated shared storage model presents several benefits in resource utilization while delivering almost the same high-performance, low-latency advantages as if the NVMe SSDs were locally attached. Specialized NVMe over Fabrics (NVMe-oF™) software, such as KumoScale™ shared accelerated storage software from KIOXIA, is now available to better enable disaggregated and shared NVMe SSD storage to be used in an automated orchestration cloud environment. NVMe-oF has quickly become the network protocol of choice as it provides the ability to pool resources and provision the right amount of storage or compute for each application workload, and on each server within the data center. With NVMe-oF pooled flash storage, resources are accessed at much lower latencies for host sharing, and hosts will borrow compute resources from lower priority applications during peak workload demands if required.

When will data center NVMe SSDs be available?

Data center NVMe SSDs are available from select industry SSD and server suppliers. KIOXIA's CD5 Series of data center NVMe SSDs are available in 960GB to 7,680GB capacities, and deliver significantly more bandwidth and IOPS per drive versus enterprise SATA SSDs. These SSDs are now available from leading server companies including Dell EMC™ PowerEdge™ and HPE® ProLiant® (Chart 3).

CD5 Series SSD Capacity	Dell™ PowerEdge™ SKUs	HPE® ProLiant® SKUs
960GB	KCD5XLUG960G	P10208-B21
1,920GB	KCD5XLUG1T92	P10210-B21
3,840GB	KCD5XLUG3T84	P10212-B21
7,680GB	KCD5XLUG7T68	Contact HPE

Chart 3: CD5 Series availability through Dell™ EMC and HPE

Learn more about KIOXIA's CD5 Series and Life After SATA: https://business.kioxia.com/en-us/ssd/life-after-sata.html

Notes:

[1] Definition of capacity: KIOXIA Corporation 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 = 2^{30} bytes = 1,073,741,824 bytes, 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, such as Microsoft Operating System and/or pre-installed software applications, or media content. Actual formatted capacity may vary.

[2] Source: IDC. "Worldwide Solid State Drive Forecast, 2018-2023, Market Forecast-Table 12, Jeff Janukowicz, May 2019, IDC #43828819.

[3] Figure 1 forecasted numbers:

SSD Interface Type	2018	2019	2020	2021	2022	2023
SATA	15,617,032	14,686,974	13,638,225	12,107,294	9,418,066	6,813,239
SAS	447,941	500,146	456,533	351,385	251,014	179,450
PCIe/NVMe	10,448,244	16,369,353	20,635,845	26,319,461	32,373,792	38,655,059

[4] An Online Transaction Processing (OLTP) application was used for measurement of server-side performance to provide the data locality benefits of direct-attached storage (high-performance / 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 SSD products, tested at 1 DWPD (Drive Write per Day), and configured with 960GB capacities that included the HK4 Series of SATA SSDs and the CD5 Series of NVMe SSDs.

[5] An OLTP application was used for measurement of server-side performance to provide the data locality benefits of direct-attached storage (high-performance / 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 SSD products, tested at 1 DWPD, and configured with 960GB capacities that included the CD5 Series of data center NVMe SSDs and CM5 Series of enterprise NVMe SSDs.