Western Digital.

WHITE PAPER APRIL 2021

SAS to PCIe Transition

by Rohit Gupta | Scott Glenn | Jimmy Gomez | Praveen Midha

Non-Volatile Memory Express (NVMe™) SSDs

Legacy SSDs leveraged existing SAS/SATA physical interfaces, protocols, and form factors to minimize changes in the existing hard drive (HDD) slots and enterprise system architecture. However, SAS and SATA interfaces and protocols were not designed for high-speed, non-volatile storage media. Because of the interface speed, performance characteristics of the new storage media, proximity to the CPU and available lanes, PCI Express (PCIe) emerged as the next logical storage interface. PCIe slots directly connect to the CPU, providing memory-like access allowing software stacks to run very efficiently. *See figure 1.*

NVMe™ (Non-Volatile Memory Express) is a high-performance storage protocol for non-volatile storage media directly connected to the CPU via PCIe interface. The protocol is NUMA (Non-Uniform Memory Access) which is optimized and highly scalable, and brings many enterprise features compared to legacy protocols.

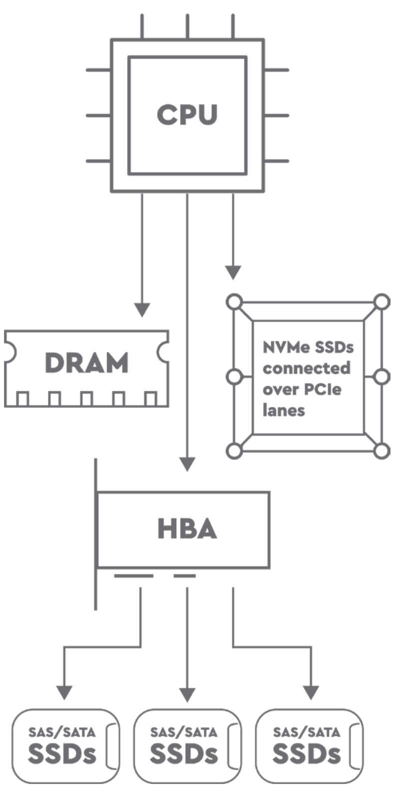

Figure 1: CPU connected with SSDs via PCIe interface vs. I/O Controller and HBA

Enterprise Data Centers

Data is driving and transforming businesses and economies at an unprecedented pace in the diverse yet globally connected world. Enterprises are increasingly storing and transforming more data to extract intelligence for better decisions, increase overall efficiencies, develop new products and maximize business profitability.

To serve new data-driven businesses, data centers are undergoing massive transformations such as:

- Accommodating more heterogeneous compute clusters
- Implementing distributed cloud architectures
- Storing data where it makes the most economic and strategic sense
- Enabling seamless dataflow across public cloud ↔ private cloud ↔ edge data centers
- Building cloud-native applications

The driving forces, explained above, will lead to the most complex and distributed architectures and require further improvement in efficiency, utilization, agility, elasticity and scalability of data centers. Clearly, next-generation data infrastructure is going to rely on shared and multi-tenant architectures, and some of the traditional storage protocols and physical interfaces (primarily SATA and SAS) will start exhibiting limitations and begin transitioning to PCIe (NVMe), an interface and protocol for high-performance, non-volatile data storage media.

SCSI to SAS

Small Computer System Interface (SCSI), a block storage and specification for parallel physical interface, is a decades-old set of parallel interface standards developed by the American National Standards Institute (ANSI) for physically connecting and transferring data between computer and disk drives and/or other peripherals, such as printers and scanners. SCSI, a bus, with bytes communicated in parallel on bulky cables, requires expensive connectors and limits data rates. Once the dataflow exceeds the threshold of megabytes per second on parallel wires, it was easier to send bits in serial at gigabits per second than to synchronize parallel wires carrying hundreds of megabits per second, and so Serial Attached SCSI (SAS) was created to support systems with large numbers of drives, such as data center systems.

Serial Attached SCSI (SAS) SSDs

Debuting in 2004, SAS protocol has evolved over a decade, delivering distinguishing and proven enterprise features, such as high reliability, fault tolerance (dual-port), speed and efficiency. SAS SSDs, primary storage or caching drives, are designed to fit in the same slot used for hard disk drives (magnetic drives) and use the same interface/protocol to connect to the host computer. 3.5" and 2.5" are the most common form factors across various SAS interface generations from 3 Gb/s, 6 Gb/s, 12 Gb/s, to 24 Gb/s, for both magnetic and semiconductor storage media.

Fast forward to today, across all data center systems, enterprise all-flash and hybrid arrays are the dominant users of SAS SSDs. Enterprise servers also deploy SAS SSDs but not to the scale of all-flash or hybrid flash arrays. SAS SSDs are primarily used with application workloads requiring high availability (HA), high input/output (I/O) and relatively lower latencies, as well as many other enterprise features, to deliver the highest level of infrastructure uptime, the most critical requirements for enterprise customers. Use cases for SAS SSDs include virtualization, online transaction processing, high-performance computing, enterprise-grade databases, data analytics, hyper-converged infrastructure (HCI) and software-defined storage (SDS), and many others.

SAS to PCIe Transition

So far, we have summarized that next-generation data centers will reduce (or eliminate) hardware boundaries; i.e., seamless data movement across the data infrastructures and increase overall system utilization and efficiencies. NVMe protocol capitalizes on parallel, low-latency data paths to the underlying storage media, similar to high-performance processor architectures. A comparison of SAS vs NVMe helps illustrate why the transition from SAS to NVMe SSDs makes both economic and strategic sense. In the following section, we will look at how NVMe is expected to play a pivotal role in achieving these objectives.

Comparison of SAS vs NVMe

	SAS	NVMe	Comments
Performance 4K block (IOPs)[1]	Random mixed 189K	Random mixed 401K	NVMe SSDs double random mixed IOPS and deliver the same performance with half the number of drives as that of SAS SSDs
Performance 4K latency (µs)[1]	Random read 150 µs	Random read 84 µs	NVMe SSDs offer 44% lower read latencies compared to SAS SSDs
Direct CPU connections	No direct CPU connection through HBAs or SAS expanders	Direct connection to CPUs	Reduce hopping, better QoS, faster data access
Queues	Single queue with 256 commands	NVMe specs allow 64K queues and 64K commands/queue	IOs on NVMe can have separate (dedicated) queues for multi-tenancy and command prioritization; NVMe can support long-outstanding IOs and process the chosen IOs for multi-threaded apps; Hosts can push more commands to the devices as compared to SAS protocol
Support for disaggregated or composable architecture	N/A	NVMe has roadmap, specs and products available today	NVMe supported infrastructure offers independent scaling of compute and storage without any significant overheads
New storage media enabler	Conventional device architectures to use low-endurance storage media	ZNS SSDs enable sequential operations and lower write amplification and therefore open the door for new low-endurance media	NVMe welcomes innovations and opens the door for new low-endurance/ high-capacity media
Management interface	SCSI Enclosure Services (SES) Supported	Further expanded on SES specs	Out-of-band/in-band support clearly makes NVMe SSDs a better choice
Form factor support	2.5" and 3.5" – although 2.5" has dominant presence	M.2, U.2, E1, E3 and various sub form factors i.e., dimensions	SAS SSDs have limited form factor footprint; NVMe SSDs offer many more form factors which allows for more features and customization options

[1] Performance values for example SAS and NVMe SSDs based on Ultrastar® DC SS540 and Ultrastar® DC SN840 SSDs with same capacity and endurance.

As shown above, NVMe SSDs offer significant advantages over SAS SSDs, including higher performance, lower latencies, newer form factors, rich management interface, telemetry support, and other enterprise features. Furthermore, innovations such as ZNS clearly position NVMe way ahead of current SAS SSDs.

All-Flash Arrays

All-flash arrays, home of SAS SSDs for caching or primary storage tiers, began with scale-up architectures and then spread into scale-out all-flash systems. Switching from SAS HDDs to SSDs began with operational simplicity (e.g., using an SSD in a slot previously used for an HDD) and continuation of similar system architectures (e.g., connection with HBAs/fan-out controllers/expanders). With this approach, all-flash arrays easily proved the total cost of ownership (TCO) benefits over all-HDD systems because of higher performance density and fewer storage controllers – fewer drives, offering same (or higher) storage capacity and higher performance, lower CPU requirements, lower data wait time for cores, less memory and networking, etc. The next phase of interface transition, from SAS to PCIe, was built on key advantages, such as direct CPU connections, performance density, scalability, enterprise features, existing and new form factors, and key technologies to enable low-endurance storage media, etc. Clearly, PCIe has a significant edge over SAS SSDs for both scale-up and scale-out all-flash storage systems. Also, moving forward, scale-out systems with disaggregated or composable architectures for independent scaling of compute and storage are also supported via NVMe-oF™ capabilities. Thus, the roadmap for SAS to PCIe transition is ready and enterprise server/storage system players have already started moving away from SAS infrastructure and prioritizing investments for all-NVMe systems.

Conclusion

Thus far, we have established that conventional storage protocols and physical interfaces were not designed for high-performance and non-volatile storage media. They have limited features and form factors, consume too many CPU cycles and delay delivery of data to applications. Also, IT infrastructure budgets are under tremendous pressure to maximize returns on infrastructure, both in storage and compute. Data centers are looking for protocols and interfaces that can handle rigorous application workloads with a smaller and more shareable infrastructure footprint and therefore allow organizations to reduce TCO and accelerate business growth.

NVMe SSDs, highly optimized protocol for non-volatile media, turned out to be a perfect fit. The new enterprise protocol, running on PCIe, eliminates many inefficiencies, improves performance and enables many more enterprise storage features and drive form factors. Also, there are many innovations quickly prioritizing NVMe over SAS protocol. Therefore, with superior benefits compared to SAS SSDs, NVMe is clearly a perfect storage protocol of choice to replace SAS SSDs in enterprise server and storage systems.

Western Digital Offerings

Western Digital offers its second generation of enterprise NVMe SSDs, the Ultrastar® DC SN840 SSD, for enterprise server and storage systems, allowing data infrastructures to embrace multi-cloud environments with SAS to NVMe transition.

Western Digital Zoned Storage Initiative

Western Digital launched Ultrastar® DC ZN540 ZNS SSDs[2] for enterprise data centers to deliver higher bandwidth and better QoS in a multi-tenant environment and further lower system level TCO.

[2] https://blog.westerndigital.com/zns-ssd-ultrastar-dc-zn540-sampling

Western Digital.

5601 Great Oaks Parkway
San Jose, CA 95119, USA
www.westerndigital.com