TOSHIBA

Life after SATA: how innovative NVMe SSDs will improve the business of MSPs

Introduction

The writing has been on the wall for a while: SATA has reached its performance limit. This leaves Managed Service Providers (MSP) in the position of reviewing their data center infrastructure and investigating the potential technologies to invest in for the future. Of course, the discussion here will not revolve exclusively around performance and capacity. There are also longer term, total cost of ownership (TCO) challenges to consider too. Here we review the background to this new set of challenges, the new XD5 NVMe low power SSD solutions from Toshiba Memory that are on offer, and the award-winning KumoScale™ software approach that delivers a higher return on investment for MSPs.

SATA has run out of steam

That SATA has run its course comes as no surprise to data center managers looking to capture customers demanding high-performance data storage Infrastructure-as-a-Service (IaaS) solutions. SATA drives, with their 6Gb/s interface, have become a bottleneck in servers, holding back CPUs from attaining their true potential, especially in applications that require high performance and low latencies. As more customers look to implement applications that benefit from high performance storage, it is predicted that the use of SATA SSDs in server applications will continue the decline first seen in 2018 (Figure 1).

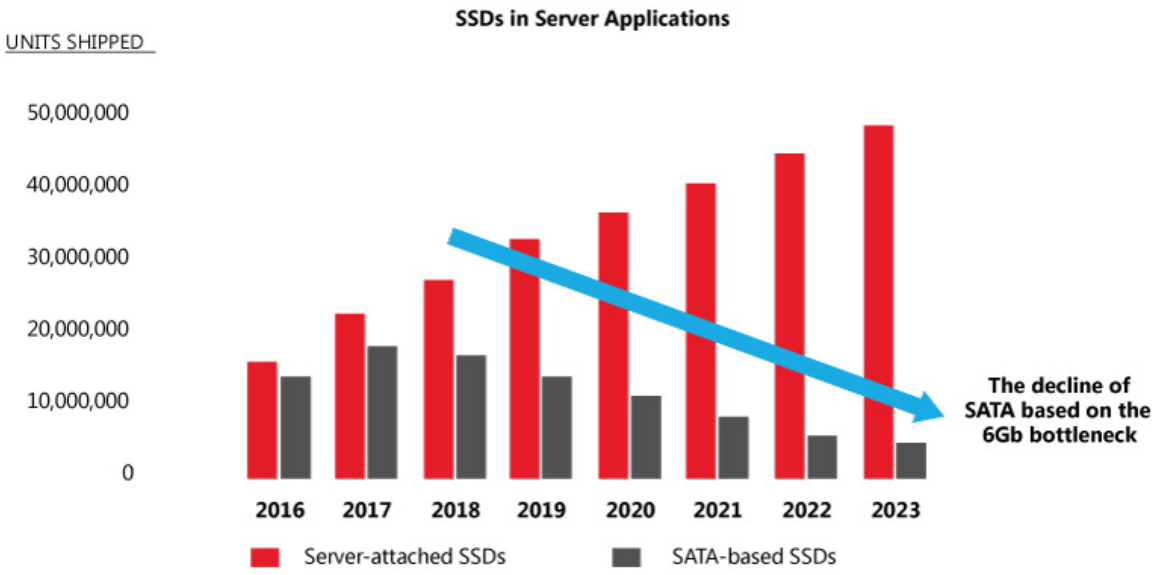

Figure 1: It is forecasted that the proportion of SATA SSDs will continue to decline in the coming years. (Forward Insights, SSD Forecast-SSD Insights, Q1 19)

The key replacement interfacing technology, Non-Volatile Memory Express (NVMe), provides multiple benefits thanks to both hardware improvements and optimizations that result in software. Firstly, NVMe has been developed from the ground up, allowing it to utilize the advantages that addressable non-volatile flash provides has over its spinning-platter predecessors. Using Peripheral Component Interface Express (PCIe) as a transport mechanism, NVMe provides support for multiple command queues and exploits the access parallelism inherent in an addressable solid-state memory. As the inventor of flash memory, Toshiba is able to deliver the full performance of its memory solutions over this interface.

Whereas SATA only supported a half-duplex interface, allowing data to be transferred in a single lane in a single direction, NVMe supports up to four full-duplex lanes. This provides a significant improvement both upon transfer rate and latency due to making multicore processing of I/O operations possible (Figure 2).

Figure 2: Data Centre NVMe SSDs display significant performance and latency improvement over SATA SSDs

Secondly, changes in the software stack to support NVMe mean that CPU accesses are simpler. Whereas access had to go through a driver layer and a host bus adapter layer for SATA SSDs to access flash storage controllers, the I/O protocol for NVMe SSDs simplifies this into a single layer (Figure 3). This change also contributes to the performance improvement of the system as a whole, especially latency.

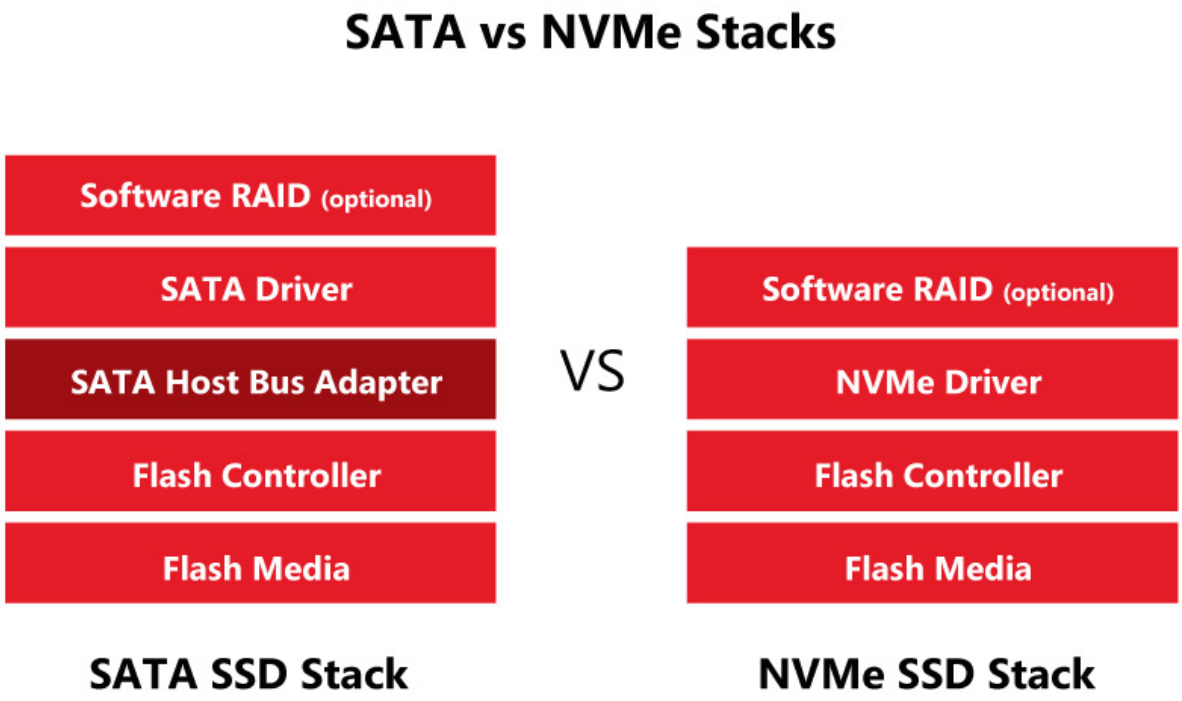

Figure 3: Differences between the software stacks to access NVMe SSDs compared to SATA SSD storage.

Looking beyond raw performance figures

As attractive as the performance and latency figures are, MSPs have a series of other concerns on their minds too. The capital expenditure (CAPEX) has to be evaluated, closely followed by a review of the operational expenditure (OPEX). Toshiba has been tackling precisely these issues in parallel to the performance challenges during the development of their latest data center NVMe SSD, the XD5 series.

Developed from the ground up, the XD5 series is designed with the challenges of data centers in mind. From the performance perspective, it was optimized to provide consistent performance and low latency; as well as keeping power consumption low. This has all been achieved at a comparable price point to equivalent SATA SSD storage options while surpassing them in terms of performance, sometimes several times over for some workloads. Thanks to this focus, the XD5 series is ideal for MSPs looking to offer NVMe performance in TCO driven applications.

Available in both an M.2 22110 and U.2 (2.5”), 7mm height form factors, capacities range from 960GB up to 3.84TB, covering a wide spectrum of data center applications (Figure 4). At a typical power consumption of 7.0W, they consume up to three times less power when compared to enterprise class NVMe drives, helping to reduce data center power demands, while the reduced heat dissipation results in lower cooling requirements. With a simultaneous upgrade of legacy servers, this allows also a redistribution of power budget from the storage to the compute. Volumetrically both form factors also enable the building of compact, data center class NVMe storage solutions, especially edge appliances.

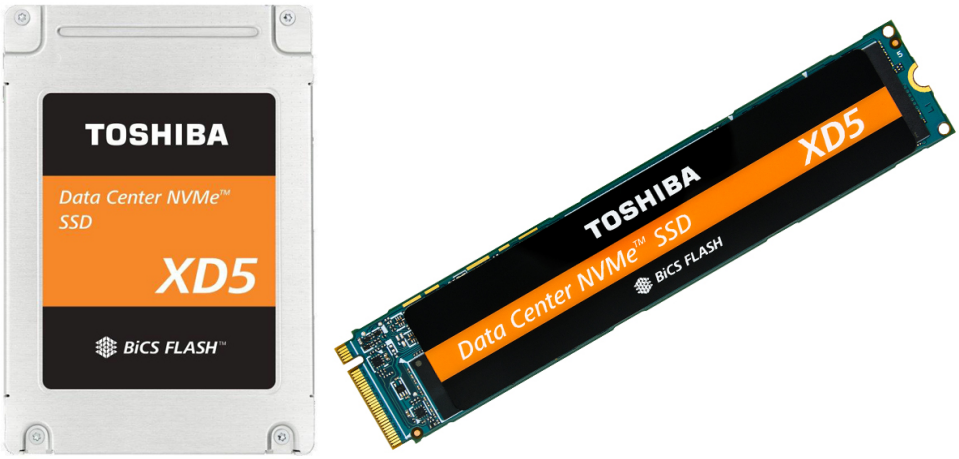

Figure 4: The XD5 Series is available in both 2.5" U.2 and M.2 22110 form factors.

When it comes to data accesses, the XD5 provides a typical random read latency of 150µs, and 45µs for random writes (4KiB QD=1). Sustained sequential reads can attain 2700MB/s and writes up to 895MB/s (128KiB) while for random reads this lies at up to 250KIOPS, or up to 21KIOPS for writes (4KiB).

https://business.toshiba-memory.com/en-emea/product/storage-products/data-center-ssd/xd5.html

Optimized for intensive read operations, the drives offer <1.0 and 1 DWPD respectively over the 5-year warranty of the drive. Use over shorter time periods, such as three years, allows a proportionally higher DPWD to be used. Naturally the XD5 also provides power loss protection together with a mean-time-to-failure (MTTF) of 2 million operating hours (2MOPH). The combination of low drive cost, power efficiency and high MTTF results in an improved TCO.

A little more effort for a much better TCO

One of the long-standing issues in data centers is achieving an optimal balance between computing resources and storage utilization. High performance processing push CPUs to their limits but leave storage underutilized, while the opposite is often true for storage intensive applications. The Direct-Attached Storage (DAS) topology is the issue here, leading to underutilized hardware, and lost operational agility and revenue.

Toshiba's award winning KumoScale™ software provides a standardized, fast, flexible, future-proof and optimized solution to this issue. Rather than rely on the NVMe SSD storage connected to individual CPUs, the storage and processors are disaggregated from one another, being linked back together using an NVMe over Fabrics™ (NVMe-oF) approach (Figure 5). This makes use of the high-speed networking on offer within todays data centers.

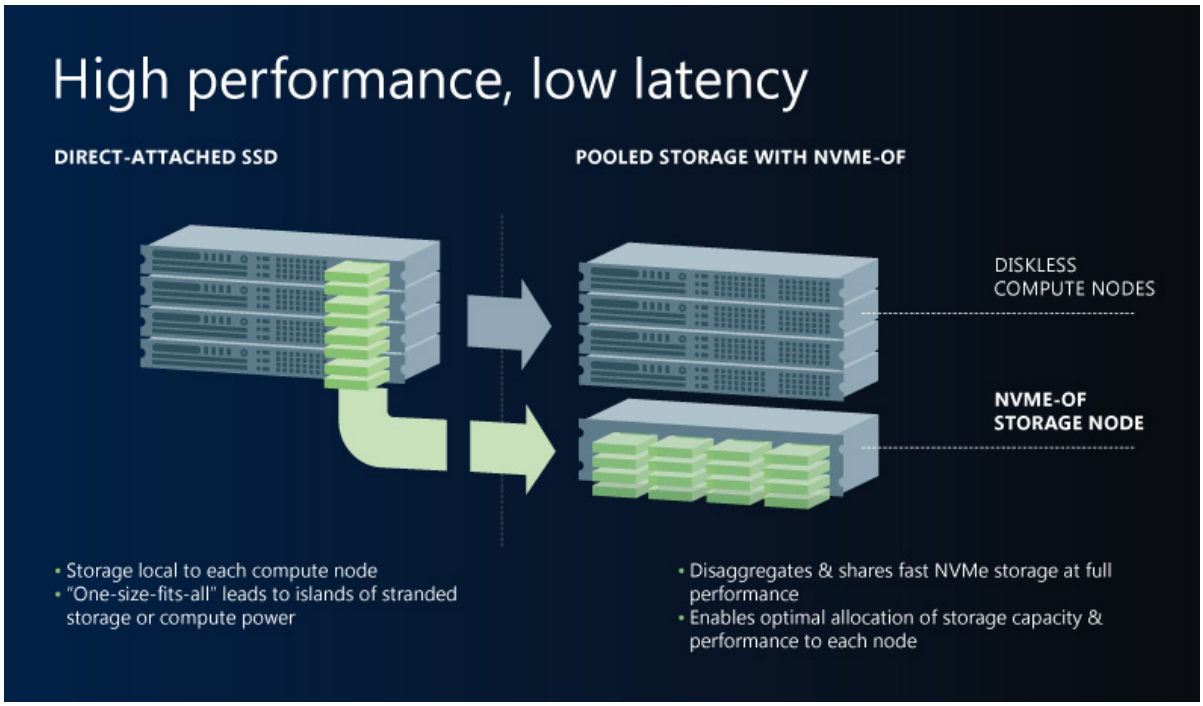

Figure 5: The move away from DAS to NVMe-oF when using NVMe SSD storage leads to further improvements in performance at low latencies for cloud applications

The result is an array of NVMe storage servers that provide ultra-high-performance block storage pools. These offer varying storage capacities and IOPS for sharing across thousands of job instances, outperforming DAS SSDs on both performance and latency. KumoScale nodes are provisioned, scheduled and managed just like compute nodes via a RESTful-API (Figure 6). Integrating with popular frameworks such as Kubernetes®, Intel® RSD and OpenStack® it also operates on the latest Linux® kernel release.

Overall, the NVMe-oF standard enables a better utilization of NVMe SSD storage capacity, leading to a lower per-node cost, while reducing stranded compute

resource through enabling server SKU consolidation. The system can also instantly adapt to changing workloads and peaks in demand, delivering greater IOPS and lower latency for cloud applications.

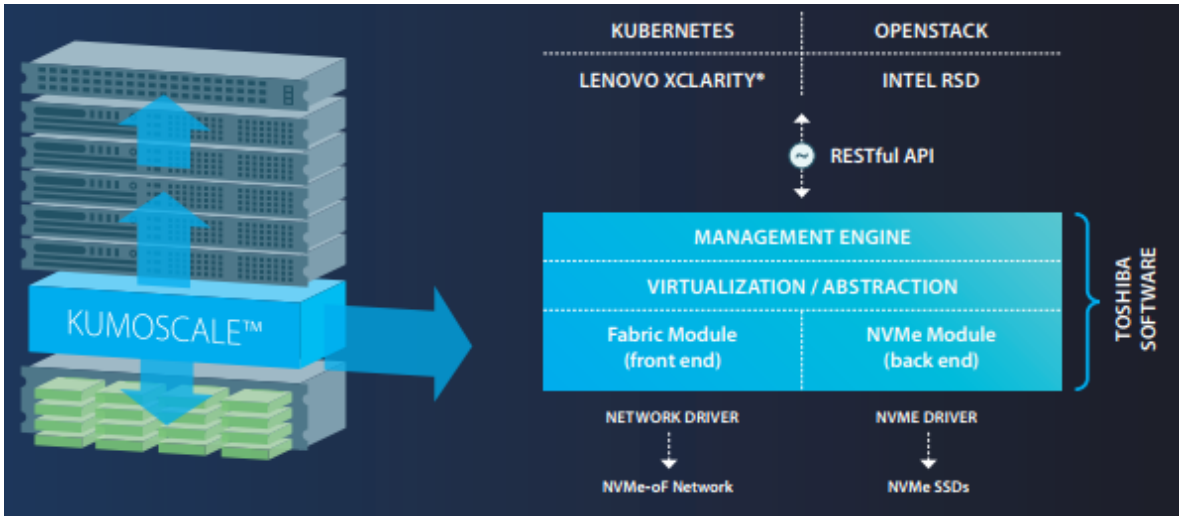

Figure 6: KumoScale nodes are provisioned and schedules just like compute nodes using standard frameworks over a RESTful API.

Summary

MSPs are facing significant changes in data center infrastructure as the performance of SATA SSDs hit their limits and fall below the expectations of their customers. NVMe is obviously the future-proof SSD interface that unleashes the performance that NAND flash can offer and fulfills the expectations that customers have of their applications. The XD5 NVMe SSDs Series offers performance superiority over SATA while maintaining a modest power consumption. Combined with investment costs comparable to SATA SSD, this drive helps Data Center Operators with a smooth transition from SATA to NVMe.

Use of Toshiba's KumoScale NVMe-oF solution also supports the move away from DAS approaches, while maximizing the advantages of NVMe SSD. This allows MSPs to offer customers better performance than today, while also achieving significantly higher levels of utilization of both storage and compute performance for improved TCO.

- E N D S -

References:

For specifications and features of the XD5 please follow this link:
https://business.toshiba-memory.com/en-emea/product/storage-products/data-center-ssd/xd5.html

Disclaimer:

Product specifications as well as product images are all subject to change without notice and may vary from those shown. Company names, product names and service names may be trademark of their respective companies.

All rights, changes and errors reserved.