Nanoseconds, Stubborn SAS, and Other Takeaways from the Flash Memory Summit 2019

Every year at the Flash Memory Summit held in Santa Clara, CA, attendees get a firsthand look at the technologies that will impact the next generation of storage. This year many of the innovations centered on forthcoming interconnects that will better deliver on the performance that flash offers today. Here are DCIG's main takeaways from this year's event.

Takeaway #1 – Nanosecond Response Times Demonstrated

PCI Express (PCIe) fabrics can deliver nanosecond response times using resources (CPU, memory, storage) situated on different physical enclosures. In meeting with PCIe provider, [Dolphin Interconnect Solutions](#), it demonstrated how an application could access resources (CPU, flash storage & memory) on different devices across a PCIe fabric in nanoseconds. Separately, [GigaIO](#) [announced](#) 500 nanosecond end-to-end latency using its [PCIe FabreX](#) switches. While everyone else at the show was boasting about microsecond response times, Dolphin and GigaIO introduced nanoseconds into the conversation. Both these companies ship their solutions now.

Takeaway #2 – Impact of NVMe/TCP Standard Confirmed

Ever since we heard the industry planned to port NVMe-oF to TCP, DCIG [thought](#) this would accelerate the overall adoption of NVMe-oF. Toshiba confirmed our suspicions. In discussing

its [Kumoscale](#) product with DCIG, it shared that it has seen a 10x jump in sales since the industry ratified the NVMe/TCP standard. This stems from all the reasons DCIG stated in a previous [blog entry](#) such as TCP being well understood, Ethernet being widely deployed, its low cost, and its use of existing infrastructure in organizations.

Takeaway #3 — Fibre Channel Market Healthy, Driven by Enterprise All-flash Array

According to FCIA leaders, the Fibre Channel (FC) market is healthy. FC vendors are selling 8 million ports per year. The enterprise all-flash array market is driving FC infrastructure sales, and 32 Gb FC is shipping in volume. Indeed, DCIG's research revealed 37 all-flash arrays that support 32 Gb FC connectivity.

Front-end connectivity is often the bottleneck in all-flash array performance, so doubling the speed of those connections can double the performance of the array. Beyond 32 Gb FC, the FCIA has already ratified the 64 Gb standard and is working on the 128 Gb FC. Consequently, FC has a long future in enterprise data centers.

FC-NVMe brings the benefits of NVMe-oF to Fibre Channel networks. FC-NVMe reduces protocol overhead, enabling GEN 5 (16 Gb FC) infrastructure to accomplish the same amount of work while consuming about half the CPU of standard FC.

Takeaway #4 — PCIe Will Not be Denied

All resources (CPU, memory and flash storage) can connect with one another and communicate over PCIe. Further, using PCIe eliminates the need for introducing the overhead associated with storage protocols (FC, InfiniBand, iSCSI, SCSI). All these resources talk the PCIe protocol. With the PCIe 5.0 standard formally ratified in May 2019 and discussions about PCIe 6.0 occurring, the future seems bright for the growing

adoption of this protocol. Further, AMD and Intel having both thrown their support behind it.

Takeaway #5 — SAS Will Stubbornly Hang On

DCIG's research finds that over 75% of AFAs support 12Gb/second SAS now. This predominance makes the introduction of 24G a logical next step for these arrays. A proven, mature, and economical interconnect, few applications can yet drive the performance limits of 12Gb, much less the forthcoming 24G standard. Adding to the likelihood that 24G moves forward, the SCSI Trade Association (STA) reported that the recent 24G plug fest went well.

Editor's Note: This blog entry was updated on August 9, 2019, to correct grammatical mistakes and add some links.