SOLID STATE STORAGE TECHNOLOGY CORPORATION

LITE-ON Storage is now Solid State Storage Technology Corporation

NVMe-oF™

The Future of Non-Volatile Media

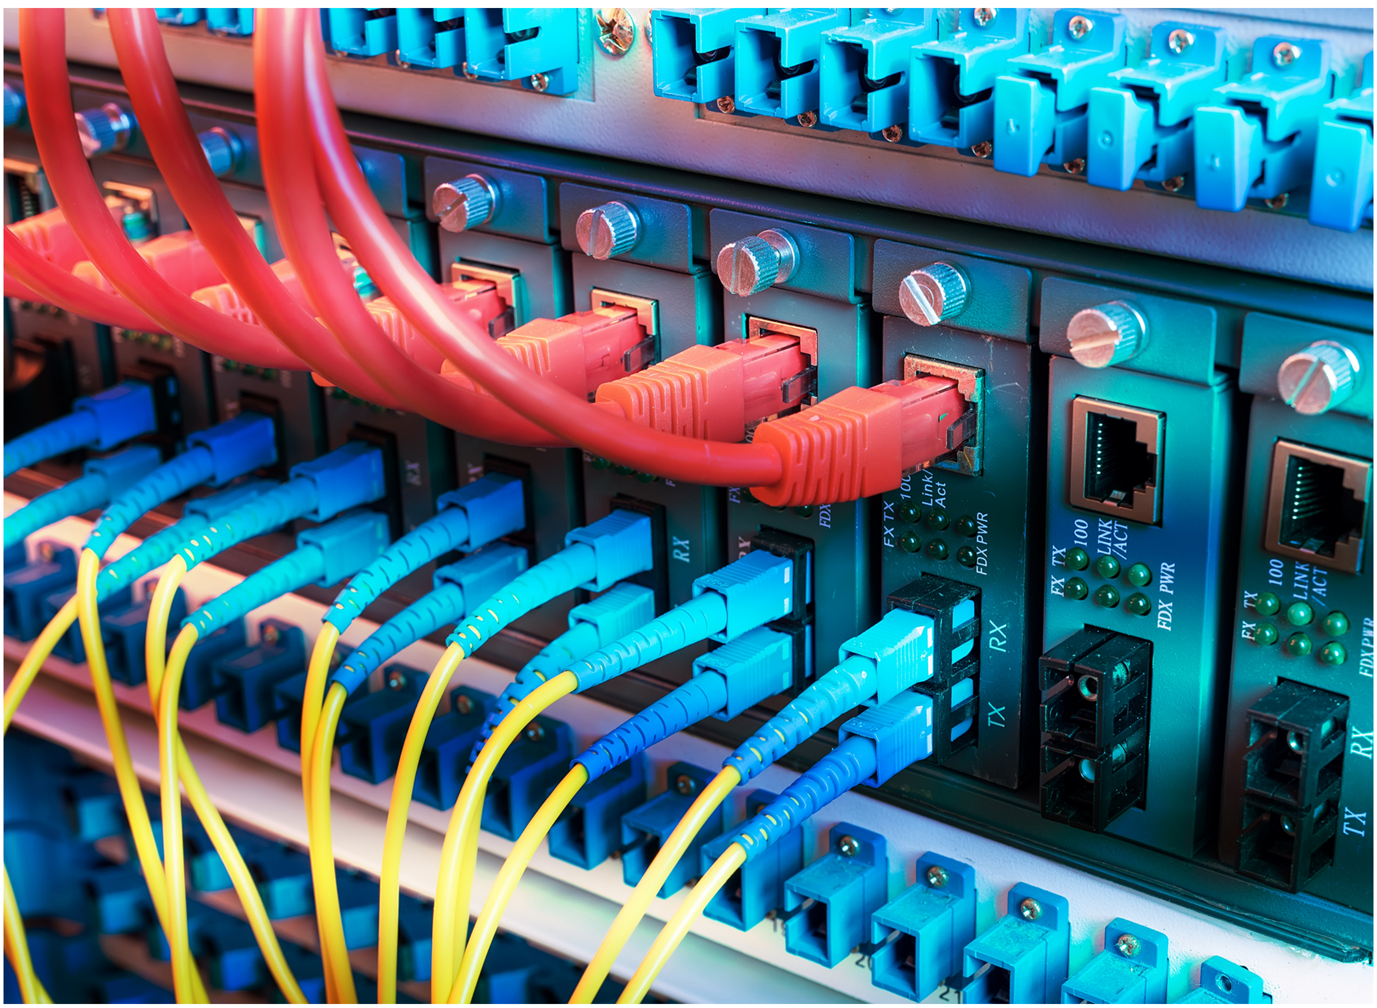

Introduction

Non-Volatile Media Express over Fabrics (NVMe-oF™) defines a way to leverage the NVMe™ protocol over a storage networking "fabric." NVMe-oF was introduced as a specification in 2016 and then more recently integrated as a technical proposal into NVMe 2.0 – please see our white paper on the NVMe 2.0 Specifications for more information. A networking fabric is defined as an any-to-any connection, extending NVMe protocols over, for example, Ethernet and Fibre Channel (FC). The goal is to improve storage scalability, especially over distance, utilizing or replacing existing networking hardware. Utilizing NVMe in this manner can provide performance near that of direct-attached storage (DAS) while allowing for an upgrade path.

Architecture

NVMe-oF™ as ratified can exist side-by-side with older protocols on a fabric like FC but the long-term goal is to unify and replace older protocols. This includes, for example, Internet Small Computer Systems Interface (iSCSI) and Serial-Attached SCSI (SATA). While NVMe was designed for a block protocol over the Peripheral Component Interconnect Express (PCIe®) bus standard, particularly with DAS, NVMe-oF seeks to create a wider Storage Area Network (SAN). Such a storage solution may also be known as a Network File System (NFS). The flexible connectivity offered by NVMe-oF also comes with only a minimal amount of latency overhead – no more than 10μs, as defined by specification – to maintain a high level of performance.

Connectivity

Aside from FC, NVMe-oF also works over Remote Direct Memory Access (RDMA). As with other forms of DMA, this allows storage communication to bypass the Central Processing Unit (CPU) which can improve performance and efficiency. RDMA comes in various flavors such as RDMA Over Converted Ethernet (RoCE), Internet Wide Area RDMA (iWARP), and InfiniBand. Ethernet is of course more convenient and as such, FC-NVMe also can work as Fibre Channel Over Ethernet (FCoE). In either case, NVMe for the Transmission Control Protocol (TCP) standard is the way forward.

Standards

Meeting the needs of current and future storage means is no easy task, and NVMe-oF™ must fill some requirements in order to do so. Cloud services, especially, require real-time provisioning, that is composable storage resources, to provide on-demand storage. Storage pools must be open, scalable, disaggregated, and extensible – flexibility is the name of the game. NVMe-oF can leverage either Just a Bunch of Flash (JBOF) or Ethernet-Attached Bunch of Flash (EBOF) with the mandatory Scatter Gather List (SGL) mechanism for NVMe command transfers. This extensibility means that NVMe-oF is multi-host, multi-port, and multi-path, supported by NVMe Input/Output (I/O) queue pairs to maintain performance with efficient messaging.

In addition to utilizing capsules for message-based transport, NVMe-oF has NVMe Qualified Name (NQN) addressing as an analogue to iSCSI Qualified Name (IQN). NVMe-oF also meets FC criteria via the INCITS/T11 tasks groups, which also cover the Intelligent Peripheral (IPI) and High-Performance Parallel Interfaces (HIPPI). While the FC Protocol (FCP) transported SCSI over fiber optics and copper cabling just fine for mechanical hard drives (HDD), solid state drives (SSD) and flash memory are bottlenecked by the protocol's limitations. The need for high-speed data transfer – that is, bandwidth – coupled with predictable, low latency, birthed NVMe-oF starting in 2014. Recent transitions to All-Flash Arrays (AFA) and the need for scalability over distance, particularly via Ethernet, led to the base NVMe specification being expanded for this type of command transfer.

Management

As part of the NVMe™ Specifications 2.0, NVMe-oF™ also works in tandem with the NVMe Management Interface (NVMe-MI). NVMe-MI provides an industry standard for managing NVMe devices across multiple deployments. NVMe-oF also works with Software-Defined Storage (SDS) and, moving forward, computational storage. The NVMe 2.0 specifications allow for easier development and implementation of such storage with the goal of improving efficiency by utilizing local computation resources, for example by offloading to storage controllers. Telemetry is also increasingly becoming important for SANs, given complex interconnections. NVMe-oF, therefore, benefits from the Storage Networking Industry Association's (SNIA) Swordfish and Distributed Management Task Force's (DMRF) Redfish specifications.

Benefits

NVMe-oF™ is a perfect fit for NVMe™ SSDs as it extends the base specification for use over distance within a networking framework. Organizing storage into pools with consistent performance allows for composable, disaggregated storage, which is both more flexible and scalable for use in enterprise and especially data centers. Leveraging bandwidth must come with minimal overhead and latency – which NVMe-oF provides, whether over FC or Ethernet – in order to ensure a DAS-like response for data requests. Servers often have limited resources – CPU time and memory (DRAM) – so efficiency is an important aspect of mass storage, making RDMA particularly attractive. This also paves the way forward for computational storage which allows for "smart" and Ethernet-integrated SSDs while allowing for Zoned Namespace (ZNS) and key-value implementations.

Storage must be consolidated while remaining agile – for example, with NVM Sets and Endurance Groups. NVMe and Ethernet are both widely-supported standards, as is TCP for the Internet, so it makes sense to encourage this aspect of NVMe-oF. Many organizations already rely on FC, however, for example with existing iSCSI implementations. The ability to run NVMe simultaneously over this hardware is to an advantage with long-term upgradeability in mind. The NVMe 2.0 Specification in fact even encourages the integration of HDDs, making NVMe-oF the perfect companion for storage solutions that cannot be entirely localized regardless of hardware type.

Native NVMe-oF™/Ethernet SSDs

Devices that can leverage this technology are native NVMe-oF™ SSDs, also known as Ethernet or Smart SSDs. These drives reduce the original storage stack – CPU, DRAM, NIC, and storage – to a drive with all of this embedded, including with a computational Application-Specific Integrated Circuit (ASIC). This simplifies both deployment and management. Costs are reduced also for the system and network with lower total cost of ownership (TCO) while maintaining high-bandwidth, low-latency performance. This unification under NVMe allows composable and virtualized storage with end-to-end benefits.

Summary

NVMe-oF™ takes the best of NVMe™ and lets it run over networking fabrics, even in legacy implementations. The new specifications help leverage the benefits of SSDs while providing for future innovations that will increase efficiency gains even further. Ensured is high-bandwidth, low-latency performance, all within a compatible package. NVMe-oF allows organizations the flexibility to organize remote storage into one universal implementation while maintaining the agility NVMe is known for, keeping storage responsive. The transition to Ethernet in particular is exciting.

NVMe-oF meets all existing standards and can be managed under Redfish, Swordfish, and NVMe-MI. The specification is future-proof as storage moves towards integrated solutions, including computational storage. This and RDMA will allow for better utilization of all server resources. The consolidation and organization or storage regardless of location, all with an ensured level of performance and shared specification, allows for a simplified storage stack. NVMe-oF, as developed alongside the NVMe Specifications 2.0, offers all of this while maintaining an upgrade path, providing a way to utilize NVMe everywhere.

Our SSD Solutions

PCIe™ - Our ED1 Series is a powerful, high performance SSD made for edge storage applications. It comes in M.2 and U.2 form factors.

SATA - Our ER2 SSD Series delivers affordability and performance with superior random read/write speeds of up to 90,000/45,000 IOPS. It comes in M.2 and 2.5" form factors.

Please contact our Solid State Storage Technology Corp. expert for more information.

*Specifications and features are subject to change without prior notice. Images are samples only, not actual products.

Request Full Specs Sheets

ABOUT US

A subsidiary of KIOXIA Corporation, **Solid State Storage Technology Corporation** is a global leader in the design, development, and manufacturing of digital storage solutions. We offer a comprehensive lineup of high-performance customizable SSDs for the Enterprise, Industrial, and Business Client markets. With various form factors and interfaces, our SSD solutions help businesses simplify their storage infrastructures accelerating variable workloads, improving efficiency, and reducing total cost of ownership.

Learn more at www.ssstc.com