LITE-ON Storage is now Solid State Storage Technology Corporation

The NVMe™ 2.0 Specifications

Refactoring, Unification, Flexibility, Ease

Introduction

The NVM Express™ (NVMe™) 2.0 specifications were announced on June 3rd, 2021, as an update to the older 1.4 base specification. Plurality here is an important distinction as there was a desire to create an entire family of specifications around the base, a move that NVM Express, Inc., calls restructuring or refactoring. This new library of specifications allows for a more flexible, modular approach, encouraging unification without the loss of backwards compatibility. Although NVMe™ was designed around non-volatile memory communication over PCI Express® (PCIe®), the NVMe™ 2.0 specifications also include NVMe™ over Fabrics (NVMe-oF™) along with the NVMe™ management interface (NVMe-MI™). Extending from this are other technical proposals (TPs) with a general delineation between Command Sets and Transport protocols.

The Command Sets include the Non-Volatile Media (NVM), Zoned Namespaces (ZNS), and Key Value (KV) specifications. While the first utilizes the traditional block-based addressing system, the other two rely on namespaces and key-value pairs, respectively. The Transport protocols include the PCIe®, Remote Direct Memory Access (RDMA), and Transmission Control Protocol (TCP) specifications. The first is the traditional transport for NVM while the other two allow for direct memory access or network transport, respectively. Together with NVMe-MI™ they make up a library that will work with existing hardware while allowing for simpler future development.

Goals

The transition from 1.4 to 2.0, rather than from 1.4 to 1.5, indicates more wide-sweeping changes. The reason for this is that there was a need to refactor the NVMe™ specifications with the goal of achieving the core values of “fast, simple, scalable.” This was a way of ensuring existing systems would not be majorly impacted while keeping open the pathway to nascent technologies, for example, computational storage. Having the base along with the other two major specifications, along with the Command Sets and Transport protocols, allows for a modular approach while leaving the door open for iterative TPs to arrive as minor rather than major updates. The NVM Express organization is especially focused on future SSD development with Hyperscale environments which must move beyond a block-based, PCIe®-only architecture.

One of the primary goals is unification so that storage management as a whole is more straightforward. This even includes hard drives (HDDs), with the desire to take the best of SCSI, SAS, iSCSI, and SATA, and allow devices to interact natively. SMR HDDs in many ways provided a blueprint for zoned namespaces through the SCSI Zoned Block Command (ZBC) and Zoned ATA Commands (ATA ZAC), for example. However, the non-zoned block basis is becoming obsolete especially for solid state as there is a transition to network- and DMA-based transport along with more efficient ways of addressing. Therefore, the NVMe™ 2.0 specifications seek to be flexible and future-oriented through a restructuring of storage organization.

Features

NVMe™ 2.0 has many new features, as covered above, but excepting smaller changes and additions there are four major ones: the ZNS specification, the KV Command Set, Namespace Types, and NVMe™ Endurance Group Management. The ZNS specification allows the host and SSD to communicate about data placement through the use of contiguous zones. The KV Command Set allows the SSD's controller and host applications to use key-value pairs instead of block addressing. Namespace Types provide "a mechanism to allow an NVMe™ SSD controller to support the different command sets … as part of NVMe™ 2.0," which in this case means variable namespace and command set pairing. NVMe Endurance Group Management allows media to be organized into different Endurance Groups, which allows greater control through NVM Sets and Namespaces.

As mentioned above, NVMe™ 2.0 allows for rotational media support – that is, hard drives – within a more general theme of backwards compatibility. Also added is Command Group Control which acts as a security measure to protect "the system from unintentional or malicious changes" with, for example, the Lockdown NVMe™ Admin Command. Within the framework are ways to adopt TCP, RDMA, etc., to prepare for hybrid cloud and Hyperscale architectures. NVMe™ 2.0 also seeks to supplant Open-Channel SSDs, for example, while including protocol support for things like SCSI's Protection Information (PI). As a whole, pre-existing features are more well-defined and integrated with clear paths forward, cleaning up the previous sprawl.

Advantages

ZNS works through sequentially-written zones with the assistance of header metadata from the host or application. As writes are sequential and more predictable, ZNS has lower write amplification, requires less overprovisioning, and also reduces the DRAM load for translation. Key Value also reduces translation overhead as it utilizes stored key-pairs, rather than addressing blocks, to retrieve data from SSD to host. Endurance Groups are primarily about flexibility – useful with NVM sets for IO determinism and efficient host-SSD division of labor – but especially capacity management. All of these can improve overall system efficiency while making the most of the advantages solid state has over rotational media.

NVMe™ 2.0, as mentioned above, was designed to refactor the NVM specifications for unification, easier management, scalability, and more flexible command sets and transport protocols. This allows organizations to tailor their storage around specific needs with clearer implementation. Scalability and flexibility are particularly important for the cloud, which has grown immensely in recent years. Improved management and unification – allowing fine-grained control from host and application, and over the network with various transport protocols – makes on-premises storage easier to consolidate and upgrade. NVMe™ 2.0's specifications work well to prepare for computational storage, which involves a balancing of compute workload between storage and host.

Summary

For solid state drives, the future is NVMe™ 2.0. The NVM Express organization carefully refactored its specifications and technical proposals around the intentions of NVMe™ 1.4. Fast – ZNS and KV reduce overhead and costs, improving efficiency which leads to faster storage regardless of transport protocol. Simple – the 2.0 base specification and the side specifications of NVMe-oF™ and NVMe-MI™ simply work, being backwards compatible, flexible for current systems, all while making it easier to develop for the future. Scalable – the separation of specifications and TPs plus the introduction Endurance Group Management ensure that NVMe™ is ready for Hyperscale arrays. Other additions, like Command Group

Control and NVMe™ Namespace Types, allow for better security and flexibility, as well.

As a whole, NVMe™ simplifies the storage interface while reducing total cost of ownership (TCO) and other costs, even for hard drives. The ability to remotely and transparently manage storage, including with help from the host and application, brings non-volatile memory into the future. The ability to use idle controller time allows computational work to be done on SSD, improving overall system efficiency. NVMe™ 2.0 embraces this concept of host-storage collaboration and brings some of the Open-Channel SSD ideals to fruition. The ultimate unification of storage, along with easier technical development, will allow future networks to continue growing unabated.

*All product and company names may be trademarks or registered trademarks of their respective holders.

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