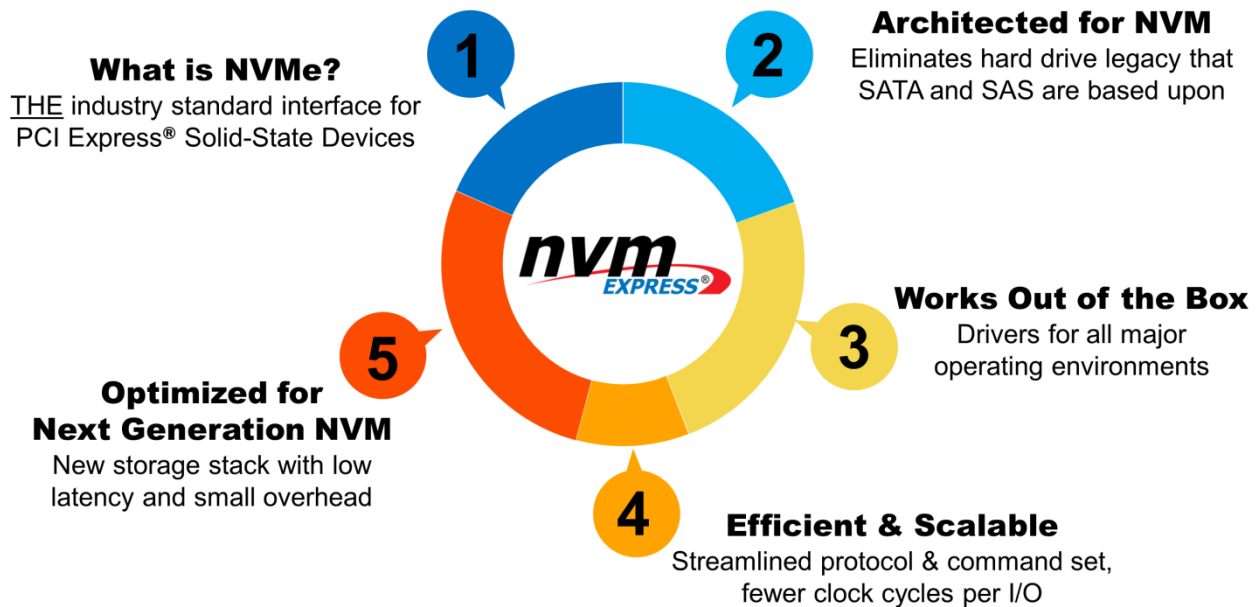


Introduction

NVM Express® (NVMe™) is an optimized, high-performance scalable host controller interface designed to address the needs of Enterprise and Client systems that utilize PCI Express®-based solid-state storage. Designed to move beyond the dark ages of hard disk drive technology, NVMe is built from the ground up for non-volatile memory (NVM) technologies. NVMe is designed to provide efficient access to storage devices built with non-volatile memory, from today's NAND flash technology to future, higher-performing, persistent memory technologies.

There are several performance vectors that NVMe addresses, including bandwidth, IOPs, and latency. For example, the maximum IOPs possible for Serial ATA was 200,000, whereas NVMe devices have already been demonstrated to exceed 1,000,000 IOPs. By supporting PCI Express and Fabrics such as RDMA and Fibre Channel, NVM Express can support much higher bandwidths than SATA or SAS (e.g., a PCI Express Gen3 x4 delivers 4 GB/s). Finally, next generation memory technologies may have read access latency under a microsecond, requiring a streamlined protocol that enables an end-to-end latency of under 10 microseconds, including the software stack.

NVMe is a completely new architecture for storage, from the software stack to the hardware devices and systems.





History

The original NVM Express Work Group was incorporated as NVM Express, Inc. in 2014 and is the consortium responsible for the development of the NVM Express specification. The organization currently has more than 100 member companies. Those interested in joining should visit www.nvmexpress.org/join-nvme.

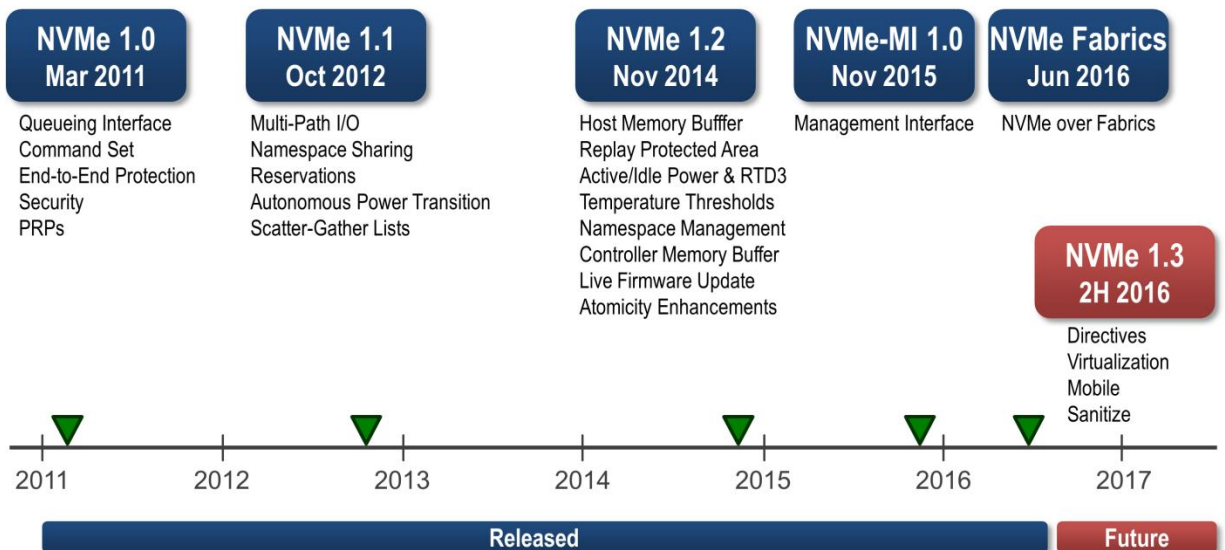
Version 1.0 of the NVM Express specification, developed cooperatively by the member companies from across the industry, was released on March 1, 2011. Version 1.1 of the specification was released on October 11, 2012, and version 1.2 was released on November 3, 2014.

In November 2015, the NVM Express Management Interface specification (NVMe-MI) was released, providing out-of-band management for NVMe components and systems. NVMe-MI provides a common baseline management feature set across all NVMe devices and systems, and a consistent method for implementing optional features. Commands include inquiring and setting the configuration, getting the health of the subsystem, firmware management, namespace management, security management, etc.

The NVM Express organization is currently working on version 1.3 of the NVMe specification, and this specification is expected to be completed during calendar year 2016. Version 1.3 addresses the needs of mobile devices, with their need for low power consumption and other technical features. When version 1.3 is completed, NVMe will be the only storage interface available for all platforms from mobile devices through data center storage systems.

Seeing the need to share storage devices, scale to large numbers of devices and extend over longer distances, work on the NVMe Express over Fabrics specification was begun in 2014. Completed in 2016, the NVMe over Fabrics specification extends NVMe onto fabrics such as Ethernet, Fibre Channel and InfiniBand®, providing not only access to individual NVMe devices but also to NVMe storage systems.

NVMe History & Roadmap




NVMe Protocol Advantages



NVMe as a storage protocol is designed from the ground up to work with non-volatile memory (NVM), including the current NAND flash technology and next-generation NVM technologies. As such, it does not have the same limitations as storage protocols designed for hard disk drives.

NVMe protocol supports multiple deep queues, which is an advancement over traditional SAS and SATA protocols. Typical SAS devices support up to 256 commands and SATA devices support up to 32 commands in a single queue. These were adequate for hard disk drive technologies, but cannot take full advantage of current and next-generation NVM technologies.

In contrast, NVMe supports 64K commands per queue and up to 64K queues. These queues are designed such that I/O commands and responses to those commands operate on the same processor core and can take advantage of the parallel processing capabilities of multi-core processors. Each application or thread can have its own independent queue, so no I/O locking is required. NVMe also supports MSI-X and interrupt steering, which prevents bottlenecking at the CPU level and enables massive scalability as systems expand.



In addition, NVMe has a streamlined and simple command set that uses less than half the number of CPU instructions to process an I/O request that SAS or SATA does, providing higher IOPS per CPU instruction cycle and lower I/O latency in the host software stack. NVMe also supports enterprise features such as reservations and client features such as power management, extending the improved efficiency beyond just I/O.



Form Factors and Power Consumption

NVMe is focused on storage protocol and can work on a variety of form factors. These include:

- M.2 form factor
- U.2 2.5-inch drive form factor (formerly known as SFF-8639)
- Add-in PCIe card (AIC)

The NVMe power range spans full-power enterprise devices down to low-power mobile devices. NVMe also takes advantage of processor power states.

Information on NVMe products from NVM Express, Inc. member companies is available at <http://www.nvmexpress.org/products/>.

Operating System Support

NVMe is supported in the following environments.

	Driver Available
Chrome OS	Chrome OS
Linux	RHEL 6.5, 6.6, 6.7, 7.0, 7.1 SUSE 11 SP3, 12 Ubuntu 13, 14
Windows	Windows 7, 8, 8.1, 10 Windows Server 2008 R2, 2012, 2012 R2
UEFI	UEFI
UNIX	FreeBSD Solaris 11.2 and higher
VMware	ESXi 5.5, 6.0

Work is ongoing to expand the compatibility list; updated information on NVMe drivers is available at <http://www.nvmexpress.org/drivers/>.

Storage Security

NVMe addresses today's data privacy concerns by supporting a tunneling protocol that provides security features produced by the Trusted Computing Group (TCG) and other related communities.

Features planned for NVMe devices and systems include simple access control, data at rest protection, crypto-erase, purge-level erase and others.

Additional information on NVMe storage security is available in the *Trusted Computing Group and NVM Express Joint White Paper: TCG Storage, Opal and NVMe* available at <http://www.nvmexpress.org/white-papers/>.

Conclusion

NVMe is gaining rapidly in mindshare among consumers and vendors. Some industry analysts are forecasting that NVMe will become the dominant storage interface over the next few years. With its high-performance and low-latency characteristics, and availability for virtually all platforms, NVMe is a game changer.


For the first time, storage devices and storage subsystems have a fundamentally different way to operate with host computers, unlike any previous storage protocol. The streamlined instructions, lower latency software stack, parallelism of queues and large queue depths, plus the design for non-volatile memory, provides previously unheard of I/O and throughput rates and the lowest latencies for storage ever seen.



Architected for Performance

NVM Express and the NVM Express logo are registered trademarks, and NVMe is a trademark of NVM Express, Inc.

All other names mentioned may be trademarks or registered trademarks of their respective holders.

A decorative footer graphic consisting of horizontal, multi-colored streaks in shades of green, yellow, and blue, mirroring the design of the NVM Express logo.