Technical white paper

NVMe technology for SSDs

HP Workstations

Table of contents

Purpose

This whitepaper is an introduction to NVMe (Non-Volatile Memory express) technology and the capabilities of NVMe to deliver the throughput for high bandwidth storage applications. Designed from the ground up for use with non-volatile memory, NVMe is a departure from the incumbent storage protocols used today, namely the ATA/AHCI command sets used in SATA and the SCSI block commands as used in SAS. This whitepaper will focus on new features of NVMe as compared to the features of the ATA command set.

What is NVMe?

According to the NVMe working group, NVMe is the standardized high performance host controller interface for PCIe SSDs. In fact, this innovative technology was built for non-volatile memory storage devices. Instead of the multitude of commands that were designed into AHCI for the sole purpose of being backwards compatible with the ATA command set, an NVMe device can complete storage operations with minimal commands and minimal overhead.

Below are the main differences between NVMe and ATA/AHCI:

Feature	NVMe	ATA/AHCI
Un-cacheable register reads per I/O	0	4
Number of I/O queues	Up to 64K	1
Number of commands per I/O queue	Up to 64K	32
MSI-X and interrupt steering	Yes	No
Number of mandatory I/O commands	13 (10 admin commands and 3 I/O commands)	18
Total number of available I/O commands	25	81
Windows 7 driver support	3rd party or Microsoft QFE[1]	Inbox
Multiple namespace support	Yes	No
Asynchronous event notification	Yes	No
Parallelism and multiple threads	No locking – Submission and completion doorbell registers for each I/O queue	Requires sync lock to issue command

Why NVMe?

As non-volatile memory (e.g. SSDs) storage technology evolves, the price and performance of the technology make it an attractive alternative to traditional magnetic media (e.g. HDDs). The non-volatile memory storage is based on NAND memory devices. These NAND devices have made great strides in boosting performance in the storage subsystem. However, as NAND technology developed, the 6 Gb/s SATA 3 bus began to hinder further advances in performance. Storage devices then moved to the PCIe bus, which is capable of up to 500 MB/s per lane for PCIe Generation 2 and up to 985 MB/s per lane for PCIe Generation 3. HP Workstations will use four PCIe lanes for individual storage devices to achieve bandwidth up to roughly 4 Gb/s.

The first PCIe attached storage devices for the client market were four lane devices compliant to the PCIe Generation 2 standard. These devices used the AHCI command set and were fully compatible with existing drivers for SATA HDDs and SSDs. AHCI was created for HDDs, addressing the need for multiple commands to read the data. SSDs do not have this need. Because the first PCIe devices used the AHCI command set, they were burdened with the overhead that comes with AHCI. In particular, an un-cacheable register read consumes 2000 CPU cycles and there are 4 un-cacheable register reads per command. This translates into 8000 CPU cycles, or roughly 2.5 µs of latency per command. Thus, to complete a single command with an AHCI controller, there are multiple small delays which add up to a large total latency. In order to reach the full potential of NAND, as well as the next generation NVM (Non-Volatile Memory), a new storage command set was required.

Performance

The benefits of the low latency associated with the NVMe protocol is clearly illustrated in the following chart.

AHCI vs. NVMe sequential performance

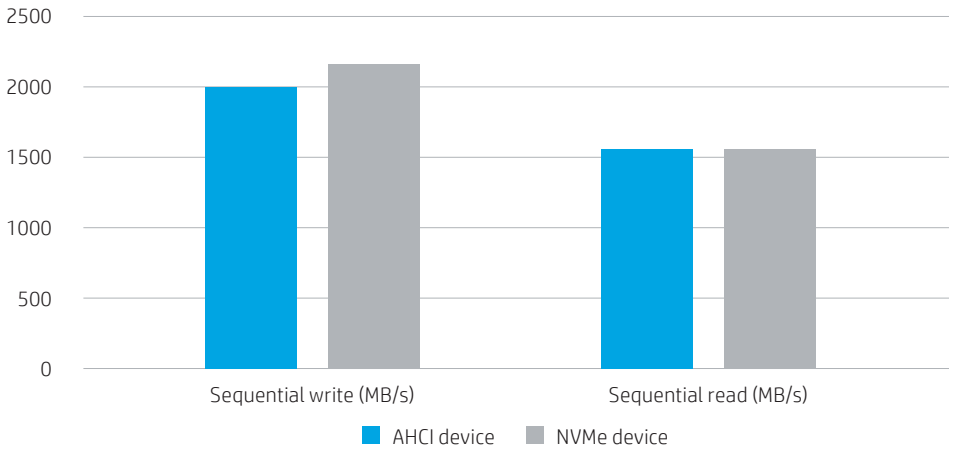

AHCI vs. NVMe random performance

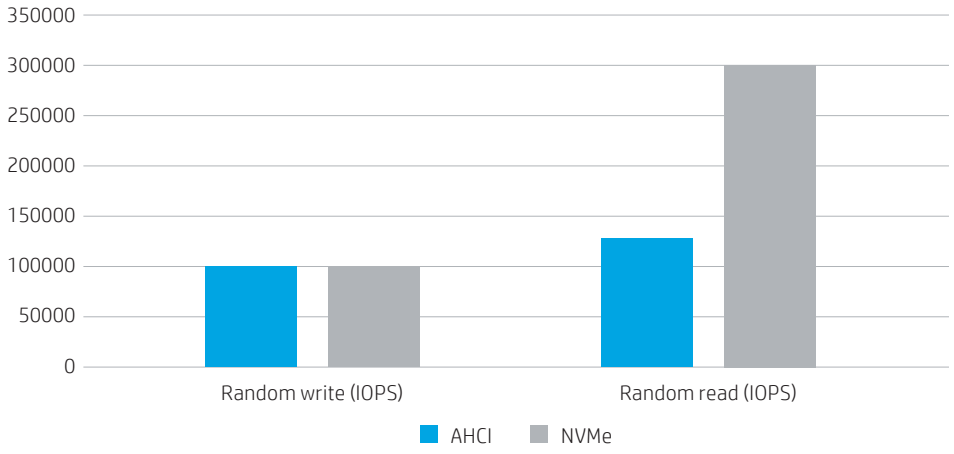

As noted previously, one of the design goals for NVMe was very low latency compared with AHCI. The benefit of this lower latency is illustrated in the random read performance shown in the following chart. Random write performance (currently limited by NAND) will show similar gains compared to AHCI as non-volatile memory technologies continue to evolve.

Compatibility with operating systems

There are inbox drivers for NVMe in Windows 8 and Windows 10, and Linux® support for NVMe has been available since the release of the 3.3 kernel. However, because NVMe is a newer storage standard, there are not standard (class) inbox drivers for NVMe devices in Windows 7. HP has been working with hardware vendors to enable NVMe on Windows 7 through vendor-developed NVMe drivers. In addition, Microsoft has provided a Quick Fix Engineering (QFE) update that enables NVMe support on Windows 7.[1] HP has thoroughly tested NVMe devices with both the Microsoft QFE driver and vendor-developed drivers. Please check the HP driver website for the recommended driver for your HP computer. Similarly, many Linux® distributions have backported NVMe support into older kernel revisions. Please check with your Linux® provider for NVMe support in your favorite distribution.

Integrating NVMe devices into your environment

The NVMe devices that currently ship on HP Workstations are M.2 modules in PCIe carriers, branded as HP Z Turbo Drive G2. NVMe devices can also be found in the 2.5" U.2 form factor and standard add-in PCIe cards, but they are typically less cost-effective than M.2 devices.

NVMe devices can co-exist with devices using other storage standards, such as SATA HDDs and SSDs, as well as SAS HDDs. HP has performed integration testing with NVMe drives as the boot devices with both SATA and SAS drives as data devices, and vice-versa. Additionally, NVMe devices from HP do not carry a cost premium over SATA SSDs and are offered in 256 GB and 512 GB capacities.

HP believes that the IT industry is moving to enable this new NVMe technology for PCIe SSD products. It has the performance benefits without a cost premium.

Resources, contacts, or additional links

RAID FAQs

HP Z Turbo G2 Quick Reference Guide

White papers with more depth on the capabilities and benefits of HP Workstations
hp.com/go/whitepapers

Information about HP Workstations
hp.com/go/workstations

[1] Microsoft has made a QFE available that enables NVMe support for Windows 7. Some HP systems may rely on this QFE provided driver to enable NVMe on Windows 7.

Sign up for updates
hp.com/go/getupdated

Share with colleagues

Rate this document

4AA6-0751ENW, August 2015