State of NVMe: Perceptions and Misconceptions

SURVEY REPORT

IN PARTNERSHIP WITH

ActualTech Media

Western Digital®

Key Findings and an NVMe Primer

In our first-ever study around knowledge and perceptions of Non-Volatile Memory express (NVMe), ActualTech Media and Western Digital sought to assess people's thoughts, opinions, and understanding of this evolving protocol. This survey and report are not intended to gauge market share and interest, but rather to better understand how much people know about NVMe. To this end, we received completed responses from 399 IT professionals and decision makers and analyzed the results in several ways. For more information about our overall methodology, please see "Appendix: Methodology" at the end of this report.

KEY FINDINGS

There are three key findings we want to highlight:

Finding 1: NVMe Hype Appears Well Deserved

It's clear from this survey that NVMe benefits are probably not being overhyped. With many new technologies, there is a lot written about how they will transform entire industries. Unfortunately, many of those technologies fail to live up to sky-high expectations. In some cases, people that are just considering new technology have great expectations compared with those that have deployed it, only to find out that it doesn't live up to the hype.

With NVMe, our results show just the opposite. Those who have yet to deploy NVMe are cautious, while those with direct experience are reporting even more positive outcomes.

Finding 2: Respondents Understand NVMe's Prime Workloads

In general, respondents seem to understand that not all workloads need or will benefit from NVMe, at least from a performance perspective. Respondents identified real-time workloads as the biggest beneficiaries of an NVMe performance boost. Given how NVMe impacts workload performance, this assessment is spot on. In other words, people are grasping the outcomes achieved from NVMe deployment and correctly applying those outcomes to potential workload benefits.

Finding 3: For Large Businesses, Current Flash Storage Solutions Are Poised to Become Too Slow

Flash is fast, but it's only as fast as all the components that sit between it and applications. NVMe aims to eliminate interconnect bottlenecks and unleash the full power of flash. For larger organizations, this transformation cannot come soon enough. Just 26% of enterprise respondents say that flash is exceeding their expectations. Sixteen percent of the respondents from large enterprise companies say that flash meets their needs today, but are concerned that this may not be the case in the future. Nearly 3% of respondents say that flash has failed to meet their ongoing performance needs.

There are some caveats to keep in mind. Most importantly, flash usage may be as part of a hybrid-flash system, not an all-flash system, a factor that could impact results. Regardless, it's clear that larger companies are ahead of smaller ones in terms of potentially exhausting the performance limits of their flash-based storage systems.

AN NVME PRIMER

NVMe is quickly emerging as a replacement for legacy interconnect protocols between compute and storage. It promises to shatter performance barriers that have become limiting since the mainstream adoption of flash. As flash became a more viable option for organizations, storage vendors moved quickly to add this fast interface to their products. In some cases, entire new companies were born, focusing on flash-first approaches to storage.

However, a number of entrants into the flash market provided products that were originally architected in a disk-based world, which operates under far different physics than flash, and performs orders of magnitude slower. The architectural design decisions that go into a system built around spinning disks are vastly different than the ones that go into a flash system.

Disk-based protocols such as SAS (Serial Attached SCSI) and SATA (Serial ATA) work well for spinning disks, but develop serious constraints when deployed in flash systems. As an example, these disk-based protocols were developed to handle seek time induced latency issues, which do not exist in flash systems. The NVMe protocol is highly optimized for simultaneous reads and writes.

For all storage controllers, there are two physical constraints to understand:

- **Number of queues.** A queue is exactly what it sounds like. It's the lineup of commands that can be sent to storage. SAS and SATA storage controllers each have just a single queue. Think of it like a road. You have a single lane of traffic and vehicles are not allowed to pass each other. When you consider how spinning storage works, this limitation makes sense. When a read or write request is sent to a disk drive, the drive head needs to move to the correct disk platter location and then spin the disk to access the data. Due to this mechanical limitation, it's impossible to read from or write to multiple locations simultaneously. Therefore, a single queue is adequate.

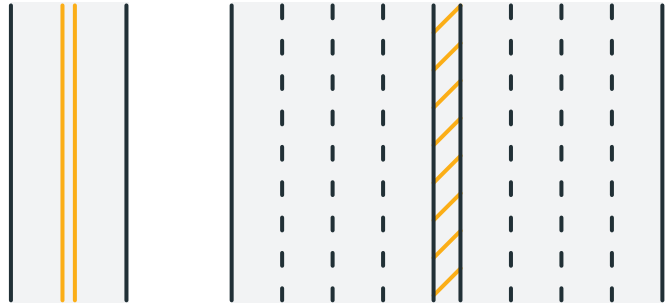

Number of queues can be compared with the number of lanes on a road. SAS and SATA have just a single queue, like a road where vehicles aren't allowed to pass each other. NVMe dramatically increases the number of queues available, like a multi-lane super highway.

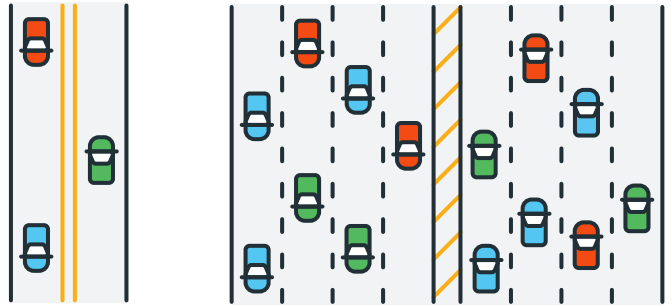

Queue depth equates to the total number of cars on the road at the same time. SAS and SATA have queue depths of 32 and 256, respectively, while NVMe is capable of supporting up to 64K commands.