

NVMe: Setting Realistic Expectations for 2018

Non-volatile Memory Express (NVMe) has captured the fancy of the enterprise storage world. Implementing NVMe on all-flash arrays or hyper-converged infrastructure appliances carries with it the promise that companies can leverage these solutions to achieve sub-millisecond response times, drive millions of IOPS, and deliver real-time application analytics and transaction processing. But differences persist between what NVMe promises for these solutions and what they can deliver. Here is a practical look at NVMe delivers on these solutions in 2018.

First and foremost, NVMe is an exciting and needed breakthrough to deliver on the performance characteristics as of early 2018. Unlike the SCSI protocol that it replaces which was designed and implemented with mechanical hard disk drives (HDDs) in mind, NVMe comes to market intended for use with today's flash-based systems. In fact, as evidence of the biggest difference between SCSI and NVMe, NVMe cannot even interface with HDDs. NVMe is intended to speak flash.



speaking flash, NVMe no longer concerns itself with the limitations of mechanical HDDs. By way of example, HDDs can only handle one command at a time. Whether it is a read or a write, the entire HDD is committed to completing that one command before it can start processing the next one and it

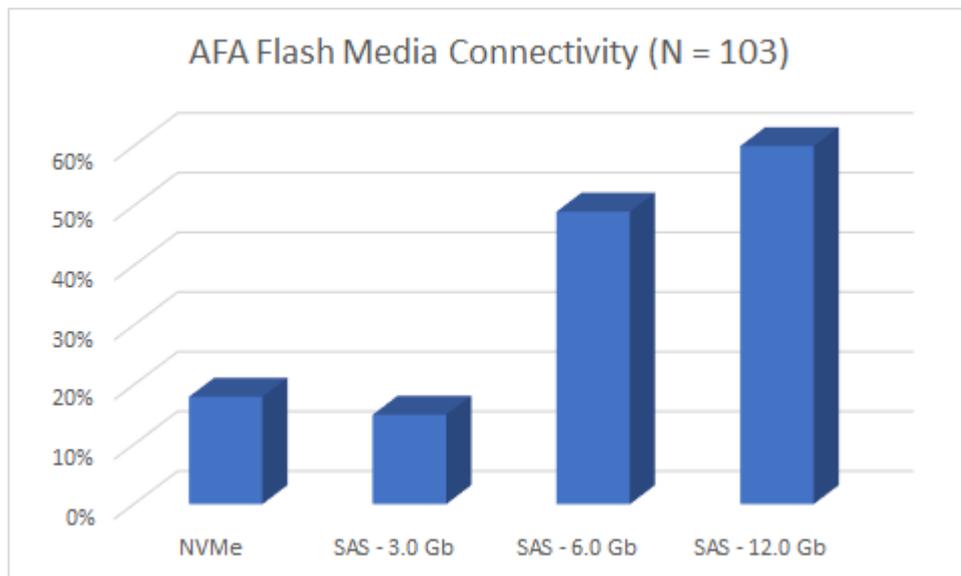
only has one channel delivering that command to it.

The limitations of flash, and by extension, NVMe, are exponentially higher. In the case of NVMe, it can support **65,535** queues into the flash media and stack up to 64,000 commands per queue. In other words, **over 4 billion commands** can theoretically be issued to a single flash media at any time.

Of course, just because NVMe can support over 4 billion commands does not mean that any product or application currently even comes close to doing that. Should they ever do so, and they probably will at some point, it is plausible that published IOP numbers might be in the range of tens or hundreds of millions of IOPs. But as of early 2018, everyone must still develop and mature their infrastructure and applications to support that type of throughput. Further, NVMe as a protocol still must continue to mature its interface to support those kinds of workloads.

So as of early 2018, here is what enterprises can realistically expect from NVMe:

1. If you want NVMe on your all-flash array, you have a short list from which to choose. NVMe capable all-flash arrays that have NVMe interfaces to all SSDs are primarily available from [Dell EMC](#), [Huawei](#), [Pivot3](#), [Pure Storage](#), and [Tegile](#). The number of all-flash arrays that currently support NVMe remains in the minority with only 18% of the 100+ all-flash arrays that DCIG evaluated supporting NVMe connectivity to all back end SSDs.



Source: DCIG

The majority of AFAs currently shipping support a 3, 6, or 12 Gb SAS interface to their backend flash media for good reason: few applications can take full advantage of NVMe's capabilities. As both applications and NVMe mature, expect the number of AFAs that support NVMe to increase.

2. Your connectivity between your server and shared storage array will likely remain the same in 2018. Enterprises using NAS protocols such as CIFS or NFS or SAN protocols such as FC or iSCSI should expect to do so for 2018 and probably for the next few years. While new standards such as NVMe-oF are emerging and provide millions of IOPs when implemented, such as evidenced by early solutions from providers such as [E8 Storage](#), NVMe is not yet well suited to act as a shared storage protocol between servers and AFA arrays. For now, NVMe remains best suited for communication between storage array controllers and their backend flash media or on servers that have internal flash drives. To use NVMe for any other use cases in enterprise environments is, at this point, premature.

3. NVMe is a better fit for hyper-converged infrastructure solutions than AFAs for now. Enterprises expecting a performance boost from their use of NVMe will likely see it whether they deploy it in hyper-converged infrastructure or

AFA solutions. However, enterprises must connect to AFAs using existing storage protocols such as listed above. Conversely, applications running on hyper-converged infrastructure solutions that support NVMe may see better performance than those running on AFAs. Using AFAs, protocol translation over a NAS or SAN must still occur over the storage network to get to the NVMe enabled AFA. Hyper-converged infrastructure solutions negate the need for this additional protocol conversion.

NVMe will improve performance but verify your applications are ready. Stories about the performance improvements that NVMe offers are real and validated in the real world. However, these same users also find that some of their applications using these NVMe-based all-flash arrays are not getting the full benefit that they expected from them because, in part, their applications cannot handle the performance. Some users report that they have uncovered their applications have wait times built into them because the applications were designed to work with slower HDDs. Until the applications themselves are updated to account for AFAs by having those preconfigured wait times removed or minimized, the applications may become the new choke point that prevent enterprises from reaping the full performance benefits that NVMe has to offer.

NVMe is almost without doubt the future for communicating with flash media. But in early 2018, enterprises need to set realistic expectations as to how much of a performance boost NVMe will provide when deployed. Sub-millisecond response times are certainly a realistic expectation and maybe almost a necessity at this point to justify the added expense of using an NVMe array since many SAS-based arrays may achieve this same metric. Further, once an enterprise commits to using NVMe, one also makes the commitment to only using flash media since NVMe provides no option to interface with HDDs.