451 Research | Business Impact Brief

# NVMe: The Welcome Aftershock of the Flash Revolution

## The 451 Take

It is no exaggeration to say that flash has revolutionized datacenter storage. Over the past several years, the rapid adoption of flash in datacenters has transformed the performance of enterprise storage systems. The process began with the deployment of hybrid storage systems that combine disk with small amounts of flash, and rapidly progressed to all-flash storage systems. Use of all-flash systems is now mainstream and continuing to grow (see figure below). This has heavily boosted overall application performance, and rescued storage from its long-standing status as the slow link in IT infrastructure.

### Adoption of All-Flash Arrays

Source: 451 Research's Voice of the Enterprise: Storage, Workloads and Key Projects 2019

Q. What is your organization's current adoption status for each the following infrastructure technologies? - All-flash array

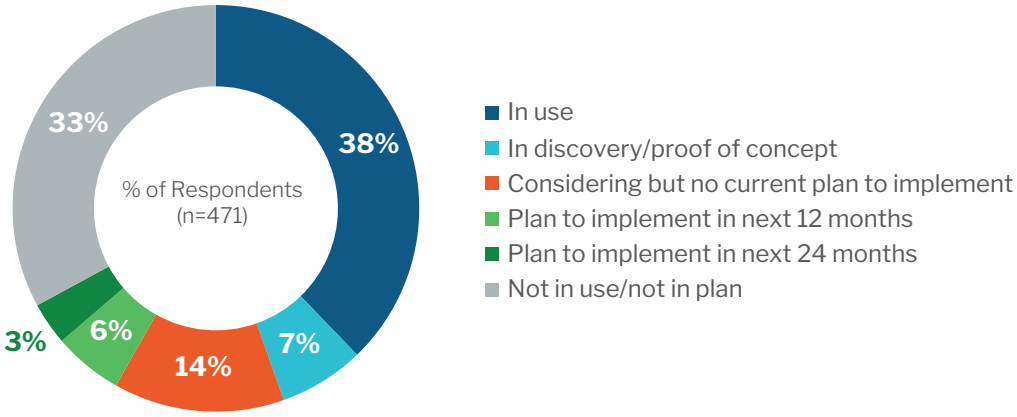


But during this time, the rest of IT has not stopped developing. Processing power has continued to grow in line with Moore's Law, and has been boosted by GPUs and other specialist processors. Network speeds have also increased dramatically, and the loads applied to IT infrastructure have not stopped growing. New business-critical and data-intensive analytics and machine-learning (ML) applications are placing especially heavy demands on storage systems to provide yet faster access to data.

Fortunately, storage is enjoying an aftershock of the flash revolution, in the form of an emerging system interface called NVMe, and its networked variant NVMe over Fabrics (NVMe-oF). Together they are providing a second major shot in the arm for storage performance. Designed specifically to maximize performance of flash and other solid-state storage technologies, the new interfaces are replacing the disk-era SAS and SATA protocols that until now have been used to access data in flash, and have limited its performance. NVMe and NVMe-oF were both developed by a widely backed cross-vendor body, and while NVMe is already well on its way to universal adoption, NVMe-oF is now following in its footsteps.

However, CIOs should be aware that most of the NVMe-based flash storage systems currently on the market are modified versions of devices that were designed during the disk era, and as a result they do not allow the full performance benefits of NVMe to be realized. To maximize flash performance using NVMe, storage systems need to be based on new internal architectures. IT organizations should also know that the NVMe-oF interface delivers a second boost to storage performance, and can be implemented over a variety of existing networks – including a very new option that allows it to be implemented at low cost, with zero disruption.

451 Research is a leading information technology research and advisory company focusing on technology innovation and market disruption. More than 100 analysts and consultants provide essential insight to more than 1,000 client organizations globally through a combination of syndicated research and data, advisory and go-to-market services, and live events. Founded in 2000 and headquartered in New York, 451 Research is a division of The 451 Group.

WWW.451RESEARCH.COM

451 Research® | Business Impact Brief

## Business Impact

**FASTER ACCESS TO DATA.** The IT industry has never ceased to demand greater infrastructure performance. Among new applications, analytics and ML are currently the highest profile among a number of emerging and existing workloads that need yet faster access to data. They will drive the adoption of NVMe and NVMe-oF, while NVMe and NVMe-oF will in turn drive adoption of analytics and ML by allowing IT to meet business needs for real-time (rather than batch) analytic runs that can take hours to complete, and by shortening development cycles for ML models.

**NVME MORE THAN DOUBLES FLASH PERFORMANCE.** The majority of all-flash storage systems deployed over the last five years are powered by flash drives that use SAS and SATA interfaces. Those two protocols were designed for use with mechanical spinning disk, which is not only far slower than flash but also has very different characteristics. The NVMe interface was specifically created to maximize the performance of solid-state storage media such as flash, and exploit their ability to allow data to be accessed in parallel. As a result, NVMe flash drives offer more than double the performance of SAS and SATA equivalents.

**NOT ALL NVME SYSTEMS ARE EQUAL.** There is a major challenge in translating the doubled performance of NVMe flash drives into a doubling of overall storage system performance. Modifying legacy systems to use NVMe drives creates bottlenecks elsewhere in those systems, heavily limiting the overall performance gain. A handful of pioneering storage suppliers have developed innovative architectures that address this issue, but the majority of these approaches have compromised the key infrastructure concept of storage disaggregation. That concept is universally practiced in data-centers, as it drives efficiency and reduces costs.

**NVME-OF PROVIDES A SECOND BOOST TO STORAGE PERFORMANCE.** The networks that link storage systems to application servers apply an overhead to application performance, and are currently also still based on disk-era protocols. NVMe-oF solves this problem by extending NVMe into those networks. NVMe-oF can run over multiple types of underlying networks, such as Fibre Channel and high-performance RDMA-enabled versions of Ethernet. This choice of underlying network has recently been extended with an option to use TCP/IP. This is important because TCP/IP is a part of plain vanilla Ethernet, which allows NVMe-oF to be implemented at low cost, and without modifications to existing application servers or the networks themselves.

## Looking Ahead

By boosting the performance of stand-alone or shared, disaggregated storage, NVMe and NVMe-oF are already reducing opex and capex penalties imposed by local or DAS storage – which until now has been required to meet the performance needs of some data-intensive applications. This process will continue.

The suppliers that are pioneering new storage architectures to maximize NVMe flash performance will continue to gain traction. Just as with the first wave of all-flash storage 5-6 years ago, the incumbent suppliers will not lead this technology trend. Storage systems that have been purposed-designed to fully exploit NVMe performance while not compromising the concept of disaggregation will offer major benefits to IT organizations.

NVMe-oF adoption will be slower than NVMe adoption, but the network protocol will eventually become mainstream. NVMe-oF will be implemented over a range of underlying networks. In the short term, Fibre Channel and the RDMA-enabled RoCE variant of Ethernet are set to be the most popular choices, while over the longer term, NVMe-oF will fill out the picture as a simpler and lower-cost choice. Few IT organizations will ignore the option to boost network performance, and thus overall application performance, when presented with the penalty-free option to use NVMe-oF over TCP/IP.

PAVILION

A the leader in NVMe-oF storage, Pavilion has adopted an agnostic policy for Ethernet, InfiniBand, TCP, RoCE, NFS, even iSCSI. In fact, with Pavilion's Hyperparallel Flash Array, it is possible to have multiple fabrics and protocols operating concurrently. This allows organizations to easily migrate existing workloads into the array and deliver blazing-fast NVMe-oF performance while deploying new capabilities like persistent containers. To learn more, visit www.pavilion.io.