



Are NVMe™ Flash Arrays Viable for Traditional Workloads?

By George Crump





Are NVMe Flash Arrays Viable for Traditional Workloads?



NVMe Flash Arrays promise an unprecedented level of performance thanks to the higher command count, queue depth and PCIe connectivity. Most NVMe Arrays are boasting IOPS statistics of close to one million IOPS and latency in the low hundreds of milliseconds. Workloads like Artificial Intelligence (AI), Machine Learning (ML), Deep Learning (DL) and high-velocity analytics are ideal NVMe Flash Array use cases, but most organizations are only at the beginning phases of implementing these projects. If they are to invest in NVMe Flash Arrays, these organizations need to see value in using them first with their traditional workloads like VMware®, Oracle®, and MS-SQL while they wait for the AI, ML, DL, and high-velocity analytics workloads to move into production.

The challenge is that the early NVMe All-flash Arrays are being delivered mostly by start-ups. These start-ups focus on delivering extremely high-performance systems that most traditional data centers don't need, and which also require investing in NVMe over Fabrics networking for which many data centers still aren't ready. The focus on extreme performance comes at a higher than expected cost, an "NVMe Tax." It also comes with feature deprivation. Many NVMe Flash Systems lack enterprise-class data services.

The result is many organizations, not ready for the AI/ML transition resort to buying traditional SAS-based all-flash arrays. While SAS-based all-flash arrays still have their place, many enterprise data centers can benefit from the additional performance that NVMe provides but they need NVMe arrays that are affordable, provide feature parity and ideally don't require their storage administrators to learn an entirely new storage interface.

Designing NVMe for Traditional Workloads

An NVMe system designed for traditional workloads should use standard networking protocols, (Fibre Channel (FC), iSCSI or IP, for filesystems like NFS and SMB. Standard networking protocols enable the organization to use the network infrastructure it has in place instead of acquiring new networking interface cards. NVMe flash drives are the critical component of the upgrade, but vendors who simply put NVMe drives in their current SAS AFA systems are not likely to see much of a performance boost. Because of NVMe Flash's low latency, storage system vendors need to equip the storage system hardware with more powerful processors and more RAM so that the storage software can keep up with the NVMe Flash

Drives. While NVMe over Fabrics network interface cards may not be required, the system should leverage higher speed, traditional network interface cards like 16GB FC and 40 GE Ethernet.

The NVMe Flash Array vendor should also optimize their software to perform in a much lower latency environment. First, since these systems leverage high-count multi-core processors, with some systems having as many as eighty cores, the software needs to be further optimized to take advantage of all of these cores. Second, the software needs further optimization to reduce as much latency as possible when

communicating with the NVMe flash drives. Features like inline deduplication and compression, thin provisioning, snapshots, clones and remote replication all need further examination by developers to ensure optimal usage in such a low latency eco-system.

Finally, the NVMe Flash Array should also include broad protocol support. Many of the early incarnations were block only or required a proprietary file system. The NVMe Flash Array should provide FC, iSCSI and File System (NFS and SMB) support. FC and iSCSI are essential to support performance hungry VMware, Oracle, and MS-SQL workloads. NFS in particular is vital to support future AI and ML workloads since many of them count on a file sharing protocol like NFS.

“

...need to equip the storage system hardware with more powerful processors...

”

The ROI of NVMe on Traditional Workloads

The ROI of NVMe on traditional workloads is scale. Server computing power is rarely fully utilized in virtualization because storage IO is the limiting factor. An NVMe Flash Array, when supporting a VMware environment, removes the IO limits and enables the CPU to be fully utilized. As a result, the environment can support dozens of more virtual machines per physical host. Greater virtual machine density means a reduction in physical hosts, potentially by a factor of ten. With the reduction in hosts comes the important reduction in networking interface cards, network switch ports and simplification of the VMware cluster.

Oracle and MS-SQL Server benefit in similar ways. Again, these environments seldom tap into the full potential of the host CPU. They are continually waiting for IO response. That waiting leads organizations to add more hosts to the environment, which increases hardware costs, software licensing costs, and management complexity. An NVMe All-Flash Array enables the organization to scale the number of users per physical host much further thus reducing and simplifying the environment. Some organizations may find that they cannot only scale the number of users further but that they can also consolidate database instances without impacting performance, once again reducing software licensing costs.

Transitioning NVMe Flash Arrays to Modern Workloads

The final ROI is, even though NVMe Flash Array vendors design their systems for traditional workloads the NVMe Flash Array is also more than adequate for the vast majority of modern AI/ML workloads which enables the organization to use the array to solve its most immediate problems first (traditional workloads) while laying the foundation for future workloads. Again, NFS is important as AI/ML workloads slowly move into production; NFS is often the first choice in file systems, and an NVMe Flash Array gives them the performance they need. As the workload moves into full production, the customer may decide a more AI/ML specific file-system is required. Again, the block (fibre and iSCSI) protocol support of the NVMe Flash Array enables these file systems to be easily supported when the organization is ready.

StorageSwiss Take

Almost every large enterprise can benefit from a transition to NVMe Flash Arrays. That transition, however, needs to come without an “NVMe-Tax” and deprivation in features. If enterprises can make a move to NVMe for a small price increase while maintaining the features they need then the ROI on an NVM Flash Array is very compelling.



The Firm

Storage Switzerland is the leading storage analyst firm focused on the emerging storage categories of memory-based storage (Flash), Big Data, virtualization, and cloud computing. The firm is widely recognized for its blogs, white papers and videos on current approaches such as all-flash arrays, deduplication, SSD's, software-defined storage, backup appliances and storage networking. The name "Storage Switzerland" indicates a pledge to provide neutral analysis of the storage marketplace, rather than focusing on a single vendor approach.



About Our Partner

Based in Silicon Valley, Tintri by DDN is a wholly owned subsidiary of DataDirect Networks (DDN), the data-at-scale powerhouse and world's largest privately held storage company. Tintri solutions serve the needs of enterprise data centers with intelligent infrastructure. Our product portfolio includes advanced technologies from Tintri (acquired September 2018), Nexenta (acquired May 2019), and IntelliFlash (acquired from Western Digital November 2019). This portfolio offers a fresh choice and a new experience to elevate your user experiences through automation, analytic insight, and tight integration with application workflows for your enterprise applications, including mission-critical databases and business applications, virtualization, and flexible software defined solutions.

Experience different! For more information on how Tintri's portfolio of solutions can turbo-charge your business success through a simple, intelligent infrastructure, visit <https://tintri.com>.

The Analyst



George Crump is the founder of Storage Switzerland, the leading storage analyst firm focused on the subjects of big data, solid state storage, virtualization, cloud computing and data protection. He is widely recognized for his articles, white papers, and videos on such current approaches as all-flash arrays, deduplication, SSDs, software-defined storage, backup appliances, and storage networking. He has over 25 years of experience designing storage solutions for data centers across the U.S.