IDC TECHNOLOGY SPOTLIGHT

Sponsored by: NetApp

While NVMe-based local storage has been available for some time, 2018 marks the year that large, established enterprise storage providers started fielding true enterprise-class arrays built around NVMe.

NVMe and NVMe over Fabric: The Future of Primary Enterprise Storage

July 2018

Written by: Eric Burgener, Research Vice President, Storage

Introduction

As new enterprises continue to pursue digital transformation, the need to accommodate next-generation applications (NGAs) is driving a new set of storage requirements. These NGAs generally service social media, mobile computing, big data analytics, and cloud-based workloads, and they put high demands on storage infrastructure to consistently deliver low latency and high throughput in the face of varying I/O loads. NVMe is a new storage technology designed specifically for flash media, delivering lower latencies and much higher bandwidth and throughput than SCSI, which makes it a much better fit for these higher-performance workloads. NVMe will ultimately become the core storage technology around which enterprise-class storage arrays are built. IDC believes that by 2021, more than 50% of external, primary storage revenue will be generated by NVMe-based systems. Almost all these systems will include NVMe over Fabric (NVMf) host connections.

AT A GLANCE

KEY STATS

End-to-end NVMe-based enterprise-class arrays will deliver 50% lower latencies and up to an order of magnitude higher throughput than their SCSI-based counterparts.

By 2021, more than 50% of external, primary storage revenue will be generated by NVMe-based systems.

As the installed base of enterprise-class arrays upgrades to NVMe technology, Fibre Channel will be the dominant transport for NVMe over Fabric.

Local storage based around NVMe (PCIe) has been used in enterprises for several years, but some customers felt limited by capacity scalability constraints and an inability to efficiently share this high-performance storage across multiple servers. The first NVMe-based storage appliances began to appear in 2016, and while the systems provided excellent performance, they lacked the enterprise-class data services (inline data reduction, thin provisioning, RAID, snapshots, encryption, quality of service, replication, etc.) needed to support mixed enterprise workload consolidation. Established enterprise storage providers have begun to integrate NVMe into their flagship storage platforms, and they are offering customers the performance of NVMe in scalable storage platforms that include a comprehensive set of mature enterprise-class data services, "five nines plus" availability, and the proven multitenant management features needed for mixed enterprise workload consolidation.

The availability of NVMe in enterprise-class systems can help companies meet several goals. The higher performance of NVMe addresses the latency and throughput requirements of NGAs, allowing both legacy and next-generation workloads to be cost-effectively consolidated onto a single platform. The higher infrastructure density enables more workload consolidation, driving a lower total cost of ownership (TCO) on a per-application basis. In addition, because NVMe was specifically designed for flash storage, it makes much more efficient use of flash performance and capacity while positioning customers to easily migrate to newer storage technology such as storage-class memory that is incompatible with SCSI.

Definitions

Enterprise storage vendors have been shipping all-flash arrays (AFAs) based on SCSI for the past seven years, and these systems clearly outperformed hard disk drive (HDD)-based arrays. However, they did not efficiently deliver the full performance value of persistent flash storage media. NVMe is a remote direct memory access (RDMA) protocol developed specifically for flash media, and it is optimized for reading and writing flash rather than spinning disk media. SCSI is a serial protocol, a factor that severely limits performance scalability as the underlying media gets faster. Unlike SCSI, NVMe supports massive parallelism. NVMe can handle up to 64,000 simultaneously outstanding requests against backing media; this feature, combined with the flash-optimized I/O protocol and the much higher bandwidth enabled by the NVMe specification, allows NVMe to fully unlock the high performance of flash-based storage devices.

NVMe-based arrays include storage controllers, backplanes, and storage devices that use NVMe technology rather than SCSI technology. These types of systems support significant performance within an array: SCSI-based AFAs generally could accommodate up to several million IOPS, but NVMe-based arrays will be able to deliver an order of magnitude higher throughput with noticeably lower latencies. NVMf, which effectively extends the RDMA capability of NVMe over a switched fabric, delivers this performance over a storage area network (SAN) back to server-resident applications. NVMf can potentially run over Fibre Channel (FC), Ethernet, or InfiniBand networks, giving customers options to leverage networking expertise already resident in their IT organizations. As enterprise storage transitions to NVMe, storage systems will exhibit NVMe technology end to end: NVMe storage controllers, backplanes, and storage devices in a single system that can be shared across hundreds of servers connected over a high-speed, NVMf-based switched fabric. NVMf technology will also be required to accommodate newer persistent storage technologies (such as storage-class memory) in the future, providing a nondisruptive bridge to the higher performance they will offer.

Benefits

NVMe technology as implemented in enterprise-class arrays delivers many different benefits to enterprises undergoing digital transformations. Benefits include:

» **Performance.** NVMe delivers lower latencies and higher bandwidth and throughput than SCSI-based arrays, making it feasible to configure much more powerful systems. The lower latencies better enable real-time NGAs and can foster higher revenue, faster response, and better customer service for latency-sensitive workloads that directly drive bottom-line business benefits. The higher bandwidth and throughput allow for much better data mobility, enabling rapid workload rebalancing, easier data migration, faster rebuilds for very large capacity storage devices, and the ability to better support big data analytics.

- **Efficient sharing of high-performance storage.** Enterprise-class arrays built around NVMe that also support NVMf host connections enable the sharing of this high-performance storage across a switched fabric, which can support literally hundreds of servers. Relative to local storage, an enterprise-class array delivers much higher capacity, much more efficient sharing of that capacity, and potentially a full set of data services that allow NVMe performance to be used for dense workload consolidation (which can now include NGAs that have very stringent latency requirements).
- **Easy transition to NVMf.** Because NVMf runs over the existing FC, Ethernet, or InfiniBand transports, it will be very easy for customers to transition to this higher-performance host connection. Given that FC is more often used for the type of latency-sensitive primary workloads where NVMe will be first deployed, IDC expects NVMe over FC to dominate as the installed base of enterprise-class arrays upgrades to NVMe technology. NVMf can run on 8Gb FC networks, but depending on workload requirements, customers may wish to upgrade to the newer Gen 6 (16Gb/32Gb) FC technology.
- **Lower TCO.** The increased infrastructure density of NVMe-based storage platforms, boosted by the protocol's ability to support massive parallelism, allows more workloads to be consolidated onto smaller platforms that use less energy, take up less floorspace, and user server-side CPU cycles much more efficiently (reducing the number of servers needed to meet any performance requirement). The significantly better reliability of flash-based storage (relative to HDDs) means administrators spend far less time replacing failed devices, particularly in the multi-petabyte configurations that are increasingly used for big data workloads. All these benefits translate to a lower TCO, particularly in larger, higher-capacity configurations.
- **Future proofing.** Given that flash is the future of enterprise storage, future solutions in this area will assume NVMe, not SCSI. A storage protocol such as NVMe that is built specifically for flash will use flash-based performance and capacity resources most efficiently, and NVMf will be required to most effectively deliver that utilization to server-side applications. NVMf support "future proofs" storage systems for the impending release of these newer technologies, enabling them to be nondisruptively added to systems as they become available.

Considering NetApp

NetApp is an enterprise storage vendor that provides a wide range of storage platforms and makes solutions available through all five enterprise storage consumption models: storage appliance, software only, converged infrastructure, hyperconverged infrastructure, and cloud. In 2015, the company entered the AFA market, one of the fastest-growing segments in the overall enterprise storage market. Since then, NetApp has achieved the number 2 market share position (by revenue) in the AFA space. NetApp's initial AFA offerings were built around SCSI technology, but the company has already leveraged NVMe technology as a caching layer for some of its all-flash systems.

In May 2018, NetApp announced its first enterprise-class array built entirely around NVMe — the NetApp All-Flash FAS A800. This new high-end storage platform supports FC-based NVMf host connections, and it offers performance that makes it as appropriate for real-time big data analytics workloads such as artificial intelligence (AI) and machine learning (ML) as it is for dense mixed enterprise workload consolidation (or both at the same time). While some other storage vendors have announced the availability of NVMe over InfiniBand (including NetApp in an October 2017 announcement for its E/EF-Series platforms), NetApp's NVMe over FC for the All-Flash FAS is a true enterprise-grade NVMf solution. Customers can transition to this NVMf implementation without having to replace their existing FC storage networking infrastructure.

NetApp also introduced a new storage operating system (OS) version — ONTAP 9.4. This release is required for NVMe and NVMf support. An upgrade to ONTAP 9.4 also enables existing A300, A700, and A700s systems to use NVMf as a host connection protocol to achieve improved performance (note that this requires an NVMf-capable storage network fabric and appropriate host software).

ONTAP is an enterprise-class clustered OS that supports unified storage (block/file), RAID, multipathing, synchronous replication for disaster recovery (MetroCluster), adaptive quality of service, and a complete suite of OnCommand systems management tools. Inline storage efficiency technologies include compression, deduplication, thin provisioning, and compaction —a new feature that NetApp expects will drive the average data reduction ratio across its entire installed base to roughly 5:1 (the current ratio is 4.75:1). Extensive snapshot technology support includes FlexClone, SnapMirror, SnapRestore, SnapManager, and SnapCenter — all features that improve snapshot-driven workflows such as application-specific backup and restore as well as test and development. WORM (SnapLock), data-at-rest encryption (NetApp Volume Encryption), and automatic data tiering to the cloud (NetApp FabricPool) are also supported.

Challenges

While NetApp is not the first enterprise storage vendor to ship a truly enterprise-class array built around NVMe technology, it is among the first to provide the NVMe over FC connection that many enterprises will require to deliver NVMe performance all the way to existing server-side applications. NetApp has been working with NVMe in its enterprise arrays for several years, however, so this is not a new technology for the company. As NVMe-based enterprise-class arrays either are or soon will be available from many of NetApp's established competitors, the company will not be able to rely solely on support for NVMe as a differentiator. It will need to clearly distinguish its offerings based on other key advantages. The NetApp Data Fabric vision and supporting technologies shipping today are a strong competitive differentiator for NetApp and should continue to be as the company builds out its support for the hybrid cloud environments that Data Fabric enables (and that over 72% of enterprises are already using).

Conclusion

NVMe is the future of enterprise storage, but in the near term, customers should let their deployment decisions around NVMe be driven by workload requirements. Clearly there are workloads, particularly among NGAs, that directly benefit from the lower latency and higher bandwidth and throughput that NVMe enables in ways that directly drive bottom-line benefits. Enterprise storage vendors that do not offer NVMe technologies risk losing out to other vendors as their customers deploy AI/ML and other real-time workloads, but there will still be many workloads that can get all the performance they need from SCSI-based AFAs at a potentially lower cost. A good strategy for vendors for the foreseeable future is to offer both NVMe and SCSI-based enterprise storage platforms and let customers choose the platforms that best meet their workload requirements.

In the near term, customers should let their NVMe deployment decisions be driven by workload requirements.

This is the strategy that NetApp is following. It offers the high-end, NVMe-based A800 and the low-end, SAS-based A220 as well as a range of platforms in between that support both NVMf and SCSI host connections. Customers that require NVMe performance can select an end-to-end NVMe-based array that supports NVMe over FC host connections today (the A800). Customers that need to support a nondisruptive upgrade path to NVMe and NVMf over the next 12–18 months can select from the A300, A700, and A700s, and customers for which SCSI meets the performance requirement over the life cycle of the array can select the A220.

IDC believes that revenue for enterprise-class NVMe-based arrays will start to have an appreciable impact on overall AFA revenue by the end of 2018 — growth that will only accelerate as more enterprises add real-time big data analytics, AI/ML workloads, and other very latency-sensitive NGAs to their application mix. With its new NVMe-based array lineup, NetApp offers compelling value for customers undergoing digital transformation: These systems deliver high performance, offer a full suite of the enterprise-class data services necessary to enable dense mixed workload consolidation (including new latency-sensitive NGAs), and offer nondisruptive upgrade paths to future storage technologies such as storage-class memory and NVMf (if they do not already support them right out of the box). IDC believes that the market for NVMe technology will continue to grow and that NetApp is well-positioned to leverage the benefits of NVMe technology in its customers' datacenters. To the extent that NetApp can address the challenges described in this paper, the company has a significant opportunity for success.

About the analyst:

***Eric Burgener**, Research Vice President for Storage, Enterprise Platforms, Storage & Network Infrastructure Group*

Eric Burgener is Research Vice President for Storage in the Enterprise Platforms, Storage & Network Infrastructure Group. Mr. Burgener's core research coverage includes storage systems, software and solutions, quarterly trackers, and end-user research as well as advisory services and consulting programs. Based on his background coverage of enterprise storage, Mr. Burgener's research includes an emphasis on flash-optimized arrays as well as software-defined storage.

IDC Custom Solutions

IDC Corporate USA
5 Speen Street
Framingham, MA
01701, USA
T 508.872.8200
F 508.935.4015
Twitter @IDC
idc-insights-community.com
www.idc.com