White Paper

NVMe: The Key to Unlocking Next-Generation Tier 0 Storage

Sponsored by: Dell EMC

Eric Burgener
April 2018

IDC OPINION

Over the past several years, persistent flash storage has changed the game for primary workloads and dominates enterprise storage revenue in this area. Denser consolidation and new types of workloads are starting to demand more than today's SAS-based all-flash arrays (AFAs) can deliver, and these increasing workload performance requirements will drive the penetration of NVMe technologies into mainstream enterprise storage platforms over the next several years. When end-to-end NVMe-based systems begin to dominate primary storage revenue in 2021, vendors that have already amassed several years of production use will deliver more mature systems that are better suited to the dense mixed workload consolidation that information technology (IT) organizations seeking to improve infrastructure efficiencies will be pursuing.

Over the next three years, the race will be on between start-ups that have recently introduced end-to-end NVMe systems and established vendors that will be offering NVMe options based on existing product lines. The start-ups will be looking to add enterprise storage functionality and prove the platform resiliency needed for more mainstream use, while the incumbents will look to orchestrate their customers' nondisruptive migration to NVMe while leveraging their strengths in enterprise functionality and proven resiliency. With PowerMax, Dell EMC has moved the comprehensive storage functionality and the proven resiliency of its flagship VMAX array to a new tier 0 enterprise storage platform that has been re-architected around an end-to-end NVMe design. With PowerMax, Dell EMC can offer customers all the performance of NVMe in a platform that provides the proven functionality and resiliency needed for dense mixed workload consolidation and for next-generation low-latency applications.

IN THIS WHITE PAPER

As next-generation applications that require new levels of performance increasingly become part of mainstream computing, enterprise workloads are also pushing the limits of storage performance. Because NVMe helps unlock new levels of performance to address this challenge, it has begun to appear in storage solutions and will spread quickly over time. This white paper discusses the evolving primary storage market, highlighting customer needs that will drive the rapid adoption of NVMe into mainstream storage solutions over the next 12–24 months. It also explores the design considerations customers should be examining in next-generation enterprise storage platforms. It then reviews Dell EMC's PowerMax array, discussing Dell EMC's strategy for integrating NVMe technologies into this platform and how that maps to customer requirements.

SITUATION OVERVIEW

IT transformation and digital transformation are under way in most organizations. Next-generation workloads that leverage mobile computing, social media, big data analytics, and/or cloud-based platforms are becoming increasingly important and driving new performance requirements for enterprise infrastructure, particularly in the storage arena. These workloads include new data management applications, such as financial and other transaction-oriented databases, unstructured data analytics used for data access, analysis, and delivery, and unstructured content analytics such as cognitive and artificial intelligence (AI) platforms. In fact, IDC forecasts these types of latency-sensitive and real-time analytics-oriented applications to be the fastest-growing storage workloads over the next five years. Today's storage platforms need the performance, functionality, and flexibility to be able to host these next-generation workloads along with traditional workloads that are demanding more power as well (such as relational databases, collaboration platforms, and enterprise applications).

Over the past several years, these performance requirements have resulted in a rapid evolution toward AFAs as the primary storage platforms of choice. In 2017, AFAs will generate over 80% of all primary external storage revenue, and that percentage will continue to increase over time. But as more demanding next-generation workloads in the data management area appear more often in enterprises, it is driving the need for even higher performance than today's AFAs can deliver. This increasing performance demand has generated considerable interest in NVMe technologies. In 2016, several niche storage vendors produced end-to-end NVMe-based systems, but IDC expects that as the need grows, more established vendors will enter this space. In 2017, we have seen several vendors announce AFAs intended for general-purpose mixed workload consolidation that are based around the latest NVMe technology.

Most AFAs in the market today leverage a SAS (SCSI-based) protocol, which worked well when enterprise storage systems were built primarily from hard disk drives (HDDs), but as persistent flash usage has become more mainstream, the protocol's limitations are becoming more apparent, particularly in higher-performance environments. With the proliferation of multicore CPUs and the denser storage consolidation enabled by virtualization, SAS' lack of parallelism is limiting the ability to get the most out of native flash media performance.

NVMe running over a PCIe interconnect addresses these shortcomings in several ways. The NVMe protocol stack is specifically developed for the much higher performance of flash and emerging next-generation media (storage-class memory [SCM]) and is therefore much more efficient for read/write operations. The leaner NVMe protocol can deliver significantly lower latencies and higher throughput than SAS can when used against the very same flash media. Because NVMe supports massive parallelism, it can maximize the use of multicore CPU architectures and modern media, removing the bottleneck from the storage array. Therefore with NVMe, storage systems can offer better performance and support even more consolidation of traditional and many next-generation workloads.

The use of NVMe opens up the opportunity to build more highly efficient enterprise storage systems that make much better use of IT resources such as compute and capacity. As storage vendors evolve to broader use of NVMe, however, other architectural changes are required to get the most out of the higher-performance media types (flash, SCM, etc.). To maximize the performance and efficiencies of NVMe, vendors will need to optimize an array's hardware and software. First, the hardware must be upgraded to NVMe-based controllers, backplanes, and devices. Then the storage operating system should be optimized for NVMe-based hardware, which today includes NVMe solid state disks (SSDs) but will soon include NVMe-based SCM devices such as Intel Optane. Finally, the system should also support NVMe over Fabric (NVMeoF) for either Ethernet and/or Fibre Channel (FC) host connection options.

Similar to the initial AFA offerings, the early end-to-end NVMe-based systems have focused purely on performance and do not include many basic data services such as snapshots and replication or multifaceted resiliency and data protection features. Many of these systems are also very limited in their capacity scalability, relegating these niche systems to a relatively small market. At the same time, we have also seen NVMe appear in more proven enterprise arrays — initially deployed to enable a high-performance path for a cache tier. However, we have reached an inflection point where we are seeing broader use of NVMe technologies within enterprise storage platforms as persistent storage and network fabric for host connections.

As established enterprise storage vendors introduce NVMe technology into their flagship platforms, data services, resiliency, availability, and multitenant management capabilities will differentiate those offerings from the early NVMe systems from start-up vendors. As enterprises adopt broader use of NVMe with more mainstream workloads, the scalability and maturity of these enterprise-class features will become strong differentiators between storage platform types. The early NVMe-based systems were generally used for custom, real-time big data analytics–oriented workloads and/or extremely high-performance databases, but they lacked the scalability, availability, and management features to serve as general-purpose consolidation platforms. As established vendors infuse NVMe technologies into their proven, mature platforms, they will clearly be going after a consolidation play that includes both legacy and next-generation workloads. This focus depends heavily on the availability and management capabilities necessary to handle mixed workloads that include at least some mission-critical applications. Key to the success of the established vendors will be a migration strategy that is nondisruptive to application workloads.

The higher infrastructure density that NVMe enables brings with it some caveats. With higher-performance devices and controllers, a failure can potentially have a larger impact. Moving to consolidate more workloads onto a single storage platform potentially offers cost and efficiency advantages but can also increase fault domain risks. With mainstream workloads that include at least some mission-critical applications, platform resiliency becomes even more important as customers increase their workload density per platform. "Five-nines plus" availability requires well-thought-out resiliency strategies that include features such as integrated data integrity checking, dual-parity RAID, snapshots, and advanced replication. Increased consolidation will drive larger average configurations in terms of capacity, potentially introducing "at scale" challenges for performance, management, and other areas. Platform maturity is an important consideration here, as are enhancements that specifically address the problems of dense multitenant management, such as the ability to select data services by application, quality-of-service (QoS) controls, and predictable performance at scale.

Over the next few years, the industry will see a slow migration in enterprise-class storage platforms from SAS toward NVMe. In 2021, NVMe-based systems will drive over 50% of all primary storage sales, signifying the dominance of NVMe-based systems for these workloads. CIOs are already evaluating technology refreshes in part around infrastructure density, and the newer systems that better exploit flash and SCM will all be based around NVMe technologies. This higher infrastructure density is already a strong focus among service provider customers, but it will become increasingly important in the enterprise as well.

The New Dell EMC PowerMax: Tier 0 Storage with End-to-End NVMe

Dell EMC has been the market share leader (by revenue) in the enterprise storage market for 15 years. Over the past several years, Dell EMC has built out a strong, all-flash storage portfolio with VMAX as the flagship offering serving the requirements of medium-sized to large enterprises. Differentiating features of VMAX include its ability to scale performance and capacity independently across a wide range; its comprehensive and mature set of enterprise-class data services, capped by proven reliable replication options; and its ability to deliver "six-nines" availability.

With the introduction of PowerMax, Dell EMC has created its first mainstream NVMe-based array designed for general-purpose use with mixed workloads. PowerMax builds on the core proven enterprise features of VMAX combined with brand-new end-to-end NVMe hardware and a myriad of software innovations delivering tier 0 storage.

PowerMax delivers excellent performance, with up to 10 million IOPS and 150GBps bandwidth, making it 50% faster than and providing roughly 3 times the performance density of VMAX All Flash. The building block of PowerMax is an NVMe-optimized PowerBrick, which includes an engine with redundant controllers and a Drive Array Enclosure (DAE). Each PowerBrick leverages NVMe from the engine to the DAEs, which each support up to 24 drives in 2U. The array scales anywhere from a single PowerBrick up to eight PowerBricks, providing 4 petabytes (PB) of effective capacity (based on 3:1 data reduction). In addition, PowerMax effectively doubles the storage density, enabling customers to get up to 2PB of effective storage capacity in a single floor tile.

PowerMax offers a variety of industry-standard NVMe SSD drive options. The system will initially ship with NVMe-based flash drives, and in early 2019, it will also support NVMe-based SCM drives (specifically Intel Optane). VMAX All Flash already delivers sub-500 microsecond latencies, and PowerMax will offer up to 25% better response times with NVMe-based flash drives and up to 50% better response times with SCM drives.

NVMe flash and SCM drives can be mixed in a single system, and Dell EMC has incorporated machine learning (ML) algorithms to optimize data placement. The built-in ML engine automatically recognizes I/O profiles and moves data to the appropriate media type in real time. This design approach ensures that PowerMax can offer the fastest performance without any storage management overhead, easily adapting to a varied and evolving workload mix over time and providing many options for customers to cost effectively support a wide variety of mixed workloads with different I/O requirements within a single PowerMax array.

PowerMax is NVMeoF ready and will offer a nondisruptive upgrade to a new NVMeoF I/O module, which will support 16/32Gb FC host connections when it ships in early 2019. The NVMeoF option will allow customers to create an end-to-end NVMe-based array that delivers the latencies of local NVMe storage with all the efficiencies, scalability, and proven storage management functionality of Dell EMC's flagship enterprise storage platform. The difference between NVMeoF and the older protocols mirrors the difference between the NVMe and SAS protocol stacks discussed previously. NVMeoF is more efficient, supporting much lower latencies and much higher throughput. In IDC's opinion, NVMeoF is the host connection protocol of the future for enterprise storage, and NVMeoF for FC will enable customers to take advantage of it quickly with their current infrastructure.

PowerMax includes significant new software innovations as well. Dell EMC has added inline deduplication that is hardware assisted, running on the same card that handles inline compression. These data reduction technologies are selectable together at the storage group level for maximum

flexibility and are fully compatible with the use of all data services including encryption (implemented at the storage controller level after the data reduction has been performed and as the data is written to persistent storage). PowerMax also adds a QoS feature that allows customers to segment performance by workload using latency limits. The use of these three data services technologies together (compression, deduplication, and encryption) does not impact the array's ability to consistently deliver sub–500 microsecond latencies at scale. Combined with PowerMax's ability to deliver "six-nines" availability, all these features provide a strong foundation for dense, cost-effective workload consolidation. Like the other all-flash platforms in Dell EMC's enterprise storage offerings, PowerMax is fully covered by Dell EMC's Future-Proof Storage Loyalty Program, which includes a three-year satisfaction guarantee, a 4:1 storage efficiency guarantee, inclusive software bundling (for data services), a data migration guarantee, hardware investment protection, and guaranteed fixed maintenance costs (at a component level) for the life of the array.

The new PowerMax family includes the PowerMax 2000 and the PowerMax 8000. PowerMax systems are sold in an appliance-based package with bundled software. The Essentials software package with the base system includes PowerMaxOS with SnapVX, inline compression and deduplication, nondisruptive migration, QoS, and iCDM Basic. The included migration tools enable customers to move nondisruptively from any VMAX or migrate data simply from most third-party arrays to a PowerMax. The Pro software package extends the advanced capabilities of the platform, including SRDF, encryption, and eNAS (for a unified storage solution) as well as host multipathing and iCDM Advanced (via AppSync). The PowerMax 8000 also supports mainframe and IBM i storage workloads. This new offering provides NVMe performance, scalability, and efficiency in a mainstream platform that includes all the proven availability, reliability, security, data protection, and manageability that Dell EMC is known for in enterprise storage.

FUTURE OUTLOOK

Current enterprise arrays based around 12Gb SAS are providing sufficient performance for most mainstream workloads, although there are certain applications, in particular very latency-sensitive databases, where the performance of NVMe is needed today. However, as more next-generation workloads based on real-time analytics — where microseconds translate into additional revenue or competitive advantage — are brought into production, NVMe will become increasingly important. IDC predicts that by 2020, 60–70% of Fortune 2000 companies will have at least one mission-critical workload that leverages real-time-oriented (versus batch-oriented) big data analytics, and a number of them will have several. NVMe will provide the ability to run these extremely demanding workloads on an IT shop's primary storage platform, eliminating the need for high-performance storage silos. Dense consolidation along these lines, however, underlines the importance of having a reliable platform. As customers consider moving in this direction to increase the efficiency of operations, they should focus on platform availability and recoverability.

NVMe is an enabler of the future of enterprise storage, and IDC expects that NVMe-based arrays will start to dominate primary storage platform revenue in 2021. Vendors should already be discussing strategies with their customers about how they support not only NVMe SSDs but also next-generation SCM technologies and NVMeoF options. Vendors that introduce NVMe support earlier will have more mature platforms once NVMe usage goes mainstream, which is a good reason for vendors to start offering these systems earlier while maintaining SCSI-based options. For the next year or two, SCSI-based systems may offer some cost advantages, but as IT organizations become comfortable with denser consolidation, the infrastructure density of NVMe-based systems will ultimately offer a better economic model.

CHALLENGES/OPPORTUNITIES

Customers will need to carefully consider which option is most appropriate for the next technology refresh of their primary storage platforms — SAS or NVMe. A particular customer's strategies for dense storage consolidation and adding more next-generation applications to the mainstream workload mix will be key factors in this decision. Customers can also consider how other benefits of NVMe besides low latency and efficiency can help them achieve their overall goals. For example, the higher bandwidth of NVMe can enable much greater data mobility, opening up options in data protection, data migration, and potentially, composability that may move customers in desirable directions given their strategies. It will be up to vendors that introduce NVMe-based systems to justify why customers might need to move to them now versus purchasing another SCSI-based system. For at least the foreseeable future, the best option for vendors that can afford to do so is to offer both types of systems and let customers choose.

When selecting an NVMe-based array, customers should ask about specific storage operating system enhancements that optimize the system for modern media, including flash and SCM, rather than spinning disk media. For the next year or two, how well the system is optimized for NVMe will likely be a meaningful differentiator in the areas of performance, efficiency, and cost — considerations that are significant for most enterprises. Vendors with longer experience with NVMe in production use among their customers will likely do better on this metric.

Managing the transition presents both a challenge and an opportunity. Customers should demand that whatever path they take to get to NVMe be nondisruptive to their application environments and end users. Systems designed to ultimately support NVMe end to end is where the industry is headed, but by following a strategy that lets customers incrementally move to NVMeoF or SCM technologies as needed, vendors offer maximum flexibility. Effectively navigating the NVMe migration over time is one way vendors can differentiate themselves from their competitors.

In the past, the breadth of Dell EMC's all-flash portfolio allowed the company to offer choice to its customers. The addition of a new enterprise array completely re-architected around NVMe and its ability to support emerging flash and SCM technologies is a strong addition to that portfolio. This is clearly an opportunity for Dell EMC to differentiate itself from vendors with more limited all-flash offerings.

CONCLUSION

Next-generation workloads will increasingly require the level of performance that NVMe unlocks. Customers should familiarize themselves with their enterprise storage vendors' strategies to incorporate NVMe technology into their offerings. Combining their own strategies around storage consolidation and next-generation workload deployment with their understanding of their vendors' NVMe direction will allow them to make intelligent choices about enterprise storage platforms.

Customers looking to leverage NVMe-based infrastructure density to further consolidate storage will do well to focus on key issues such as storage solution resiliency, multitenant management capabilities, platform maturity, and vendor migration strategies. Incumbent vendors with proven resiliency at scale and a broad set of comprehensive storage capabilities may have an edge here if they intelligently integrate NVMe technologies into their platform in a manner that preserves a nondisruptive migration strategy for their customers. Workload requirements will determine whether customers choose a SAS-based array or an NVMe-based array at their next technology refresh, and vendors that have both options will offer their customers better choice.

With the introduction of PowerMax, Dell EMC is adding a key forward-looking platform to its already strong all-flash portfolio. The NVMe technology is intelligently fused into a platform that delivers a full suite of enterprise-class data services and proven "six-nines" availability, significantly increasing the system's storage performance, infrastructure density, and efficiency. With its performance profile, the PowerMax not only competes with the niche NVMe systems introduced in 2016 to host extremely latency-sensitive next-generation workloads but also goes beyond what those systems provide with proven enterprise storage functionality, making it a preferred alternative for dense mixed workload consolidation.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com