White Paper

# How Broad All-Flash Vendor Portfolios Help IT Customers

Sponsored by: Dell EMC

Eric Burgener
May 2017

## IDC OPINION

In the current 3rd Platform computing era, enterprises of all sizes are being asked to continue to support legacy applications while providing support for next-generation applications (NGAs) in mobile computing, social media, big data/analytics, and cloud environments. At the same time, most information technology (IT) organizations are consolidating around virtual infrastructure for both legacy and new applications. This has significantly changed the I/O profiles that storage infrastructure has to deal with, and it has become clear that for most primary storage workloads, and even now for some secondary workloads, hard disk drives (HDDs) no longer adequately meet requirements. As a result, we are seeing overall enterprise storage revenue shift away from HDD-based designs toward newer flash-based designs with the following results:

- **Flash is rapidly coming to dominate primary workloads as a persistent storage medium.** For latency-sensitive workloads, flash offers performance, capacity, storage density, energy and floor space consumption, efficiency, and reliability benefits that are compellingly better than the benefits systems based around HDDs can offer for 3rd Platform computing workloads, and 76% of enterprises plan to deploy or evaluate all-flash array (AFA) platforms in the next 12 months.
- **Differing workload requirements are driving the rise of new storage architectures.** In the past, most enterprise storage designs were based around a scale-up, dual-controller model, but in the past several years, new architectures have emerged, which include scale-out, software-defined, converged and hyperconverged infrastructure, and cloud designs. These new architectures offer benefits for many NGA workloads that scale-up, dual-controller designs could not, and all-flash versions of each of these architectures are now available from enterprise storage providers to meet increasingly stringent performance requirements.
- **Decreasing flash costs are enabling its use in workloads that are not just about low latency.** Traditionally, flash had been targeted for use with mission-critical, latency-sensitive primary workloads where better performance translated directly to improved bottom-line business benefits. However, as flash costs have continued to decrease, the use of flash with less latency-sensitive workloads, justifiable based on the secondary economic benefits of flash deployment at scale (need for fewer devices, lower energy and floor space consumption, higher infrastructure density, increased efficiencies, and better reliability), has emerged as a market. This significantly increases the total available market that vendors can address with their all-flash products, expanding the use of flash to unstructured (file/object based) workloads as well as the traditional block-based workloads.

Customers with varied storage needs will have their requirements best met by a range of all-flash offerings, driving increased revenues for those vendors that offer broad all-flash portfolios.

## IN THIS WHITE PAPER

As more and more primary workloads are moving to all-flash storage configurations, those vendors that offer a broad portfolio of all-flash options, targeted for different customers and workloads, find themselves in the best position to meet today's enterprise storage requirements. This white paper discusses the evolving all-flash market, highlighting the different customer needs and the all-flash array market segments that have emerged to meet them. This document then turns to an overview of Dell EMC's broad all-flash portfolio, identifying unique system differentiators and the customer and workload targets for each platform.

## SITUATION OVERVIEW

Enterprise datacenters have experienced significant changes in workloads over the past 10 years. Virtualization has become a mainstream computing platform, more efficiently enabling the denser consolidation of workloads in a much more flexible management environment. Data streams (the I/O profile that today's enterprise storage infrastructure must accommodate) have become much more random, exhibiting a variety of read/write ratios, a wide distribution of block sizes, a high ratio of data reducibility for certain workloads, and I/O bands (i.e., hotspots) that drift over time. High customer expectations about low latency and predictably consistent performance are paired with an increasing demand for application services that are always available. These requirements have driven the emergence of new storage architectures and media. One of the most notable of these is the broad use of flash as a persistent primary storage medium.

The performance benefits of flash are well known — a single flash device like a solid state disk (SSD) offers on average 10 times lower latency and 10–100 times the throughput of a 15,000rpm HDD, and it offers very consistent and predictable performance (unlike an HDD) regardless of workload spikes. But there are a number of other benefits that IDC collectively refers to as the "secondary economic benefits of flash deployment at scale" that allow all-flash systems to deliver a compellingly lower total cost of ownership (TCO) than HDD-based systems — the need for far fewer devices to deliver required performance, much lower energy and floor space consumption, higher CPU utilization driving the need for fewer application servers (since each server gets much more work done) and possibly lower software licensing costs (because of the need for fewer CPU cores), much lower administration costs (time spent tuning to address storage performance issues tends to go to zero), and much better device reliability (flash is a solid state device rather than a mechanical device). It is this combination of benefits that enterprises moving from HDD-based systems to all-flash systems find so attractive.

It is important to note that the much lower latency of flash enables the use of inline storage efficiency technologies like compression, deduplication, thin provisioning, pattern recognition, write minimization, and space-efficient snapshots and replication with latency-sensitive primary storage workloads, something generally not possible with HDD-based systems for performance reasons. On average, enterprises using these technologies running mixed primary workloads achieve data reduction ratios of at least 3:1, and many achieve a much higher ratio. In 2017, the cost of enterprise flash stands at $0.43 per gigabyte, while the cost of performance-intensive 15,000rpm HDDs stands at $0.27 per gigabyte — this means that customers achieving only a 2:1 data reduction ratio with their all-flash systems are enjoying a lower effective dollar-per-gigabyte cost for flash, and that is before any of the other secondary economic benefits of flash deployment at scale are taken into account.

The industry's move to the use of all-flash systems for mixed primary workload consolidation is undeniable. In 2016, when the AFA market generated $4.94 billion in revenue (versus overall external storage market revenue of $24.53 billion), over 70% of all primary (versus secondary) storage spend was driven by AFA purchases. 27% of information technology organizations are already committed to an "all flash for primary storage" strategy, and another 50% let flash deployment be driven on a case-by-case basis (which means it is frequently deployed for latency-sensitive primary storage workloads). The AFA market is the most dynamic external storage segment by far and is forecast by IDC to grow at a 26.2% compound annual growth rate (CAGR) through 2020 when it will hit $9.39 billion. Every customer that IDC has spoken with about its all-flash deployment experiences wants to move more workloads to it over time because of its combination of performance, ease of use, and compelling economic benefits.

The primary flash market has driven almost all AFA revenue over the past five years. The primary flash segment targets latency-sensitive primary workloads. In 2016, IDC noted the emergence of an additional AFA market segment we call "big data flash." Built exclusively around scale-out designs, these systems optimize more for high storage density and capacity, high bandwidth to handle very large data sets, and low cost — at the expense of low latency. These systems still deliver very consistent performance in the face of varying workloads, like primary flash arrays do, but their lower cost of raw flash capacity makes them much better suited for petabyte (PB)-scale workloads where certain secondary economic benefits of flash deployment at scale — high density, low energy and floor space consumption, and high reliability — become increasingly important. Initial products were shipped in this space in 2016, and IDC expects this market to surpass $1.05 billion in revenue by 2020. For those vendors willing and able to field platforms in this space, this opens up yet another market for all-flash products to help drive future all-flash revenue growth while allowing customers to extend the benefits of flash to additional workload types.

## The Importance of the All-Flash Portfolio

The first enterprise-class AFAs shipped in 2011, and these early systems were deployed as dedicated storage platforms for generally a single application environment like a high-performance database or virtual desktop infrastructure (VDI) environment. After three years of use, enterprises started to want to move more workloads to flash, and end-user purchase criteria started to shift away from just performance and acquisition cost and more toward the ability of these systems to support mixed workloads and integrate into preexisting datacenter workflows. Vendors responded, and by the end of 2015, there were a number of AFA platforms that, in addition to performance, included extended capacity, inline data reduction, enterprise-class data services, and extensive API support for workflow integration. By the end of 2016, the top AFA vendors offered very similar product feature-based value propositions on a per-platform basis, and differentiators like go-to-market strength and an ability to sell enterprise flash effectively began to become strong determinants of market success.

Enterprises must accommodate a variety of different workloads with a variety of needs — there are block-, file-, and object-based data sets; low latency and high throughput are critical for some workloads, while for others, capacity scalability (and the cost of that capacity) is much more important, and some require "six-nines plus" availability, while others do not. As enterprises moved toward the replacement of HDD-based systems with AFAs and adopted more flash-oriented datacenter strategies, customers wanted platforms that made the benefits of flash available across a much broader spectrum of workloads than in the past. Vendors responded, and it is possible today to buy flash-optimized storage platforms for block-, file-, and/or object-based workloads — those optimized to deliver low latency and high overall throughput and bandwidth as well as those optimized for low cost and high infrastructure density in petabyte-scale environments. A single all-flash platform cannot

handle all these needs — multiple platforms are required — and it is becoming increasingly clear that a broad all-flash portfolio that meets a wide spectrum of customer requirements is a strong differentiator driving overall all-flash revenue growth on a per-vendor basis.

In the past, most enterprise storage systems were sold as storage appliances — turnkey solutions including both hardware and software built around scale-up, dual-controller architectures. But as NGAs proliferated in mobile computing, social media, big data/analytics, and cloud environments, other architectures arose that offered benefits for not only new workloads but also new customer types. Today, there are five different consumption models — storage appliances, software-only, converged infrastructure, hyperconverged infrastructure, and cloud — that target different customers and different workloads. Definite customer preferences have developed for certain of these models, and vendors looking to maximize their ability to meet customer requirements are making their intellectual property available through at least several and, in some cases, all five of these consumption models. This requires a broad all-flash portfolio, and vendors that do not offer their products through a certain consumption model will not be able to go after customers that prefer that model. This puts single-product AFA vendors at a distinct disadvantage to competitors that offer a broader portfolio.

## The Dell EMC All-Flash Portfolio

Dell EMC sells a broad all-flash portfolio that offers block, file, and unified storage platforms; includes both scale-up and scale-out architectures; has solutions that target everything from small to midmarket customers to large enterprises and cloud providers; and covers different consumption models.
Dell EMC was an early entrant into the all-flash space back in 2013 with the XtremIO product, the first AFA to break the $1 billion revenue mark. XtremIO is a block-based storage solution built around a scale-out architecture that is sold as an appliance, primarily to medium to large enterprises. However, as market interest in all-flash platforms grew, Dell EMC made other all-flash offerings available to address the requirements of different customers. Today, Dell EMC's all-flash portfolio includes seven types of different enterprise storage platforms and covers four of the five consumption models. All-flash options are not available as part of Dell EMC's cloud storage offerings (Virtustream, VMware Cloud Foundation), but hybrid storage configurations that can include flash are available as part of these cloud storage services.

Given the strong need of customers for all-flash performance with their enterprise workloads, Dell EMC has infused its flagship VMAX platform with a number of optimizations to make it run better using flash as the persistent storage medium. Already, this system generates roughly 80% of all VMAX revenue. With its mature and proven HYPERMAX operating system environment, the VMAX All Flash includes a complete range of enterprise-class data services, "six-nines plus" availability, and support for the consolidation of open systems (e.g., Oracle, SQL, SAP), mainframe, and IBM iSeries workloads on a single platform. It also includes Dell EMC's broadest ability to integrate with a variety of different datacenter tools and workflows through a complete set of APIs, and an ability to scale both performance and capacity independently. Inline compression supports improved performance, increases flash endurance, and lowers the effective cost of flash capacity, and it comes with a 4:1 data reduction ratio guarantee. An ability to migrate from prior VMAX generations to the VMAX All Flash nondisruptively helps support the system's "six-nines availability" through technology refresh. Targeted mainly at medium to large customers looking for a proven enterprise storage platform for mixed workload consolidation, a VMAX All Flash offers a comprehensive and mature set of enterprise-class data services while delivering flash-optimized performance that scales to well over 6 million IOPS and over 4PB of effective storage capacity in a single floor tile. EMC Unisphere, the management interface, provides EMC's most complete set of multitenant administrative tools and natively supports a number

of advanced configurations, such as stretched clusters, based around proven, scalable replication. The VMAX All Flash is packaged as a storage appliance and is often used by customers to consolidate both legacy and newer applications.

In May 2017, Dell EMC announced the next-generation XtremIO platform, the X2. Based around a scale-out architecture, XtremIO is the only AFA in the industry that uses a metadata-centric system architecture that significantly improves the efficiency of both real-time and data services operations by executing them in main memory. XtremIO's metadata operations are content aware, operate 100% in memory, and are fully distributed across all controllers in a system to ensure consistent performance even as a system scales. The system's snapshot (XtremIO Virtual Copy [XVC]) facility is in particular differentiated in its performance, scalability, and efficiency because of its metadata-centric implementation. Because of the low latency with which all operations are executed, XtremIO excels at delivering consistently low latencies even while numerous data services, such as inline data reduction, space-efficient read/write snapshots, and encryption, are all in simultaneous use with heavy real-time I/O workloads. XtremIO is a block-based, highly efficient storage platform with a full complement of enterprise-class data services that are designed to operate at the scale necessary for mixed workload consolidation. Packaged as a storage appliance, the platform is targeted for use in medium to large enterprise virtual environments for virtual desktop infrastructure, high-performance databases, and snapshot-intensive workloads (DevOps, etc.). Dell EMC stands behind this platform with a 4:1 data reduction ratio guarantee as well.

The Dell EMC Unity platform is a unified storage offering supporting block and file together on a single system. Targeted at midmarket customers, Dell EMC Unity is available as an all-flash appliance, a hybrid flash appliance, or a virtual (i.e., software-only) appliance. Leveraging off-the-shelf Intel-based x86 servers with internal storage, Dell EMC Unity offers inline data reduction (compression) backed by a 4:1 data reduction guarantee, point-in-time snapshots; synchronous and asynchronous replication; built-in encryption; tiering to the cloud; deep ecosystem integration with VMware, Microsoft, and OpenStack; and cloud-based predictive analytics through CloudIQ and is fully integrated with EMC data protection options. Dell EMC Unity supports multiple SSD options, including a 3D NAND triple-level cell (TLC)–based 15.3TB drive that supports very high storage densities with low power consumption and floor space utilization. Target customers include midmarket enterprises looking for a simple, modern architecture that will be used to consolidate mixed block and file workloads where consistent flash performance and low latency are required.

For midmarket customers looking for hybrid arrays that offer the best value, Dell EMC offers the SC Series platforms. All-flash configurations of this platform are available and start at under $25,000, but the SC Series is really optimized for hybrid deployments that include a mix of flash and spinning disk. The SC Series includes efficient auto-tiering technology that moves data between media types based on activity and other user-defined policies, inline compression and deduplication backed by a 4:1 data reduction guarantee, space-efficient snapshots, replication, and built-in multiarray federated management. Midmarket customers that need the most flash optimizations, higher performance, and scalability will be better served by the Dell EMC Unity offerings, but those customers simply looking for the lowest cost in a flash-enhanced system will find better value with the SC Series offerings. These storage appliances include a good mix of data services and are targeted for use as a mixed workload consolidation platform for small to medium enterprises, remote or branch offices, or other distributed environments.

For those customers interested in converged infrastructure offerings, Dell EMC offers several flavors of VxBlock and VBlock systems. Converged systems include servers, storage, and networking, all packaged into a single pre-configured, pre-wired rack that is purchased under a single SKU. These systems are easy to purchase and easy to deploy, can be managed through an integrated administrative GUI, and are supported under a comprehensive support plan with a single point of contact. Dell EMC offers three types of VxBlock and VBlock systems, giving customers a choice between purchasing a version with Dell EMC Unity storage, Dell EMC XtremIO storage, or Dell EMC VMAX All Flash storage. VxBlock systems include Cisco Unified Computing System (UCS) and support VMware NSX or Cisco Application Centric Infrastructure (ACI) network virtualization, whereas VBlock systems also include Cisco UCS but support only Cisco ACI virtual networking. These systems are generally targeted at medium to large enterprises looking specifically for ease of purchase and deployment and integrated management, although workload targets vary based on the type of integrated storage.

VxBlock and VBlock System 350, which include Dell EMC Unity storage, support a broad spectrum of mixed block and file workloads, offering the scale needed for virtualized and cloud-enabled deployments. VxBlock and VBlock System 540, which include Dell EMC XtremIO storage, are targeted at workloads that demand the highest throughput at the lowest latency — applications like online transaction processing (OLTP), business analytics, and end-user computing. VxBlock and VBlock System 740 deliver enterprise scale for mission-critical applications and consolidated mixed workloads that include SAP, Oracle, and Microsoft platforms, regardless of whether they run on open or mainframe systems.

For those customers interested in all-flash web-scale infrastructure, Dell EMC offers ScaleIO and VMware vSAN options, which support block-based storage; the XC Series, which supports block and file storage; and Isilon All-Flash, which supports file-based storage. ScaleIO, a software-defined storage product that supports mixed server brands, operating systems (both physical and virtual), and storage media types (HDDs, SSDs, and PCIe and NVMe flash), can be deployed as either a traditional storage-only architecture or a modern hyperconverged architecture. Storage-only architectures allow storage capacity to be scaled independently from compute resources and support common IT infrastructure and roles. ScaleIO brings web-scale economics and simplified storage life-cycle management to this disaggregated storage model. Hyperconverged architectures scale compute and storage together for an easier deployment model, providing all the benefits of a SAN but with web-scale economics, and offer cost efficiencies (since the same server is running both the application and the storage services). ScaleIO is more flexible and efficient than traditional SAN infrastructures, allows for resources to be increased or decreased "on the fly" with no downtime, and supports performance that scales linearly to the 10M+ IOPS range with predictable, sub-millisecond latency.

ScaleIO can be purchased as a software-only product to be deployed on customer-selected server and storage hardware. Dell EMC also sells what are called ScaleIO Ready Nodes — pre-validated, tested, and optimized Dell PowerEdge server nodes, which include the ScaleIO software preloaded — all supported by a single vendor (Dell EMC). ScaleIO can also be purchased as part of a converged infrastructure offering called VxRack FLEX. Offering the classic benefits of converged infrastructure — easy purchase and deployment, integrated networking, and a single point of customer support — VxRack FLEX is a rack-scale, hyperconverged system that uses Dell EMC ScaleIO as the foundation for a fully functional software-defined enterprise-class storage platform using Dell PowerEdge servers, is hypervisor agnostic (as well as supports physical servers), and provides the scalability to start small and grow to 1,000+ nodes. Target customers include those looking for web-scale, software-defined storage solutions to support a mix of workloads, and VxRack FLEX is often used as the basis for cloud-based infrastructure (public and/or private). Any solution based on Dell EMC ScaleIO supports all-flash as well as hybrid storage configurations, depending on customer preference.

Dell EMC VxRail is another hyperconverged offering that supports both all-flash and hybrid storage configurations. Powered by VMware vSAN and offered with Dell PowerEdge servers, VxRail is the industry's only fully integrated, pre-configured, and pretested VMware hyperconverged infrastructure appliance family on the market. Sold as storage appliances, VxRail includes Dell EMC's mission-critical data services such as replication, backup, and cloud tiering at no additional charge, and it is the only hyperconverged appliance featuring kernel-level integration between VMware vSAN and the vSphere hypervisor for improved performance and efficiency. Target customers include existing VMware shops running mixed virtual workloads looking to easily expand their infrastructure, regardless of whether that supports noncloud or cloud environments, as well as new customers looking to install their first VMware deployment.

The XC Series combines Dell PowerEdge server hardware with the Nutanix Virtual Computing Platform (VCP) to create a hyperconverged appliance based on a web-scale architecture. Supporting both block- and file-based access, this platform supports multiple hypervisors (VMware ESXi, Microsoft Hyper-V, and Nutanix's own Acropolis Hypervisor), easy "one click" operations, data locality management for higher performance, and all-flash as well as hybrid storage configurations. Nutanix VCP is a software-defined storage product that features fully distributed data and metadata operations for predictable, high performance; features self-healing, fully automated recovery from failures, and fully nondisruptive upgrades for firmware as well as for expansion and/or reconfiguration; heavily leverages automation and analytics to help administrators make better informed management decisions; and supports enterprise-class data services like compression, RAID, snapshots, replication, quality of service, and REST APIs for integration with third-party management tools. Administrators manage virtual machines (VMs) rather than LUNs, volumes, or RAID groups, allowing them to easily apply the XC Series' policy-based management at the application level. Target workloads include mixed business-critical workloads including VDI, private cloud, database, OLTP, and data warehouse as well as virtualized big data analytics deployments.

Isilon is a scale-out NAS platform, targeted at unstructured data and file workloads that are very demanding in terms of overall scalability (both performance and capacity). Based around the proven, mature OneFS operating environment, Isilon offers massive scalability; file-based deduplication backed by a 4:1 data reduction guarantee; automated tiering, which optimizes data placement; broad multiprotocol support; comprehensive data protection including space-efficient snapshots and data replication, SEC 17a-4–compliant WORM; and role-based access controls. Systems can be expanded quickly and easily, without downtime, manual data migration, or application logic reconfiguration — all handled automatically by Isilon's AutoBalance feature. The Isilon All-Flash platform, which started shipping in 2017, extends the benefits of the OneFS environment to unstructured workloads with the most demanding low latency, predictability of performance, and storage density requirements. Isilon All-Flash solutions scale up to 33PB in a single cluster with up to 9 million IOPS and 540GBps throughput performance. Workload targets include digital media (broadcast, real-time streaming, rendering, and postproduction), electronic design automation (simulation, verification, and analysis), life sciences (DNA/RNA sequencing and genomics), and enterprise data lakes to support big data analytics workflows across a variety of industries.

With this broad offering, Dell EMC is well positioned to meet customer needs for all-flash platforms across many key areas. Its all-flash products cover the primary flash market (VMAX All Flash, XtremIO, Dell EMC Unity, SC Series, and ScaleIO), the big data flash market (Isilon All-Flash configured with the 15.3TB SSDs), structured workloads (VMAX All Flash, XtremIO, Dell EMC Unity, and ScaleIO), unstructured data workloads (Dell EMC Unity and Isilon All-Flash), and unified storage (Dell EMC Unity and XC Series) as well as four of the five different consumption models (storage appliances, software-only, converged infrastructure, and hyperconverged infrastructure). Dell EMC does offer cloud-based storage offerings through its Virtustream division as well as through VMware, its wholly owned subsidiary, but flash is used as a caching tier rather than persistent storage in these services.

## FUTURE OUTLOOK

The use of flash as a persistent storage medium will come to dominate primary storage infrastructures in the next two to three years. With enterprise flash costs dropping at 26.0% per year through 2020, the lower cost of raw flash capacity will help that medium make strong inroads into secondary storage environments that have any degree of performance sensitivity as well. While many of these secondary workloads — backup, disaster recovery, content and media streaming, big data/analytics, and so forth — are not particularly latency sensitive, they often have to deal with the movement of extremely large data sets, and the benefits that flash brings to the table in terms of throughput and bandwidth can make it quite interesting for these applications. Given the multi-petabyte scale of many of these secondary storage environments, the savings associated with high flash storage density, low energy and floor space consumption, and high reliability also make it extremely interesting. Note that spinning disk technology is not likely to support a general-purpose device that will offer more than 12TB of raw capacity, while SSD vendors like Samsung and Seagate are already talking about delivering 64TB-class devices within the next year. The same benefits are very much of interest for many unstructured workloads as well, and flash that is sufficiently affordable for even small portion of this $32.8 billion market (by 2020) opens up a lot of new opportunities for all-flash offerings in these markets.

Because of competition from flash, revenue from performance-optimized storage devices like 10,000rpm and 15,000rpm HDDs is dropping precipitously. HDDs are not going away, however. 7,200rpm HDDs today only cost about $0.03 per gigabyte, and deep archive environments where data is written, must be maintained for many years, is rarely if ever accessed, and must be kept online are a perfect fit for these types of storage devices. Flash is not likely to be cheap enough to compete effectively with HDDs for these types of workloads that have almost no performance sensitivity for many years, if ever. But we will see flash become affordable and used for more and more secondary environments as enterprise flash costs continue to come down over the next five years. This makes it extremely important for any enterprise storage provider to cater to the all-flash requirements of a number of different customers and workloads.

## CHALLENGES/OPPORTUNITIES

Despite the overwhelming benefits of flash for mixed primary workload consolidation to those that really understand its multifaceted benefits, IDC surveys still indicate that the single biggest reason IT organizations give for not purchasing AFAs is cost. It is critical for vendors to comprehensively explain the benefits that flash brings to the table for all types of workloads and make sure in particular that finance officers understand the secondary economic benefits of flash deployment at scale. To that end, IDC has published a document for finance organizations entitled *Justifying Investment in All-Flash Arrays* (IDC #US41646416, August 2016), which reviews the economic case for all flash. In a worldwide survey taken by IDC in mid-2016, only 15% of enterprises had AFAs deployed in production environments, but 76% of survey respondents had plans to deploy and/or evaluate one in the next 12 months. IDC has, in fact, noted an acceleration in AFA revenue growth since that survey was taken. More and more IT organizations are coming to understand the comprehensive picture, but there is still a lot of work to do around getting the message out about not only flash performance but also flash economics.

As customers consider how the benefits of flash apply differently to primary and secondary workloads, it can be confusing to sort out which AFAs are best suited for their own needs. Vendors should take care in positioning their all-flash offerings to make this as clear and easy as possible for customers to make the right choice.

## CONCLUSION

Highly scalable, truly enterprise-class all-flash platforms are here and are being widely used and considered as a general-purpose platform for mixed primary workload consolidation. When legacy enterprise storage platforms come up for technology refresh, IT organizations should absolutely be evaluating AFAs as replacements. As flash costs continue to drop, more and more enterprises will be moving to "all flash for primary storage" strategies and will be considering their use more often for secondary storage environments that can benefit from the high density, low energy and floor space consumption, and high reliability of flash. Across the variety of different storage needs that enterprises have — block, file, and/or object; low latency; low cost of capacity; dense mixed workload consolidation; high density; high reliability; and the ability to move large data sets quickly and easily — IT organizations will need to deploy multiple storage platforms. Vendors with a broad all-flash portfolio are better positioned to serve the needs of these customers than vendors that just offer a single AFA platform, regardless of how impressive it is.

Dell EMC was an early entrant into the enterprise AFA market, and since that entry in 2013, it has expanded its all-flash portfolio to offer solutions for block-, file-, and object-based environments; mission-critical workloads that require low latency, high throughput, and "six-nines plus" availability; mainframe support; more capacity and cost-conscious workloads; and both midmarket and large enterprise customers. Flash optimizations made available on Dell EMC's initial platforms have, in many cases, made a rapid transition to other all-flash platforms as Dell EMC brought those to market, and today, the company offers a highly flash-optimized enterprise-class all-flash portfolio that can meet the needs of customers across a number of different primary and secondary storage workloads.

## 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

---