



White Paper

NetApp All Flash FAS: Flash-Optimized Storage for Mixed Workload Consolidation

Sponsored by: NetApp

Eric Burgener
July 2016

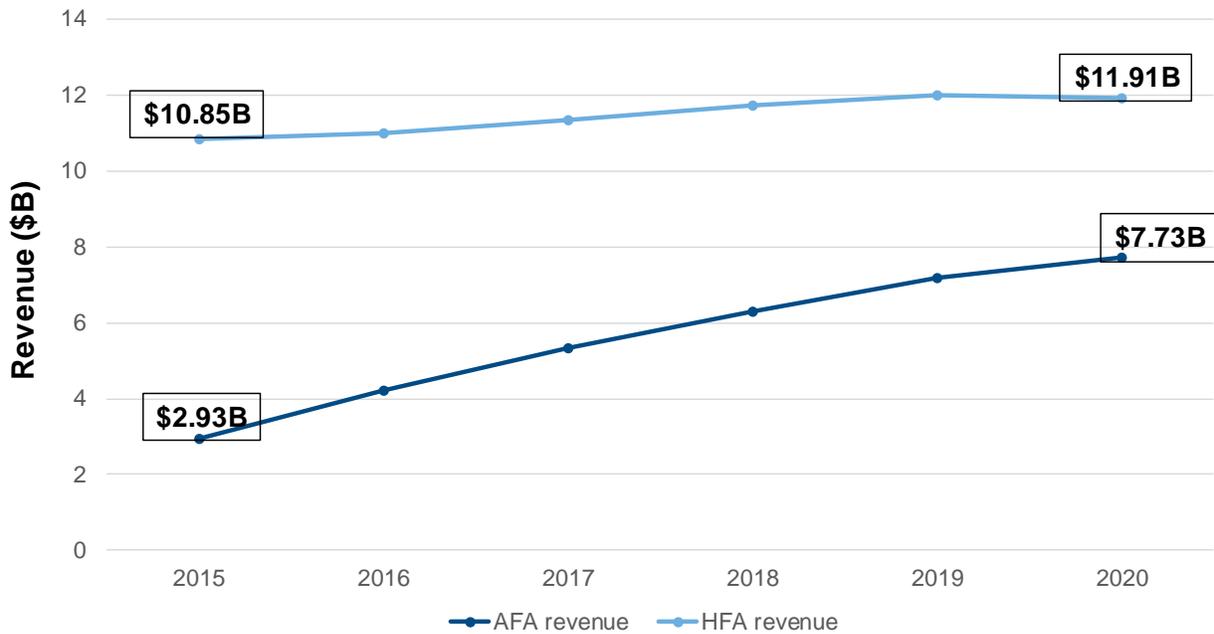
SITUATION OVERVIEW

Flash storage is taking the enterprise market by storm. Datacenters are now consolidating legacy application environments, such as relational database, messaging, and collaboration platforms, with next-generation applications built to meet the needs of social media, mobile computing, big data/analytics, and cloud infrastructure. New, flash-optimized storage architectures are required to cost-effectively deliver the performance necessary as these mixed workloads are virtualized. As legacy enterprise storage workhorse platforms come up for technology refresh, organizations are looking for new systems that deliver comprehensive and mature enterprise-class functionality in a flash-optimized design that maximizes the performance, endurance, reliability, storage density, and cost-effectiveness achievable with flash media.

Shared network storage provides significant benefits in terms of proven resiliency, enterprise-class data services, and resource sharing efficiency. Hybrid flash arrays (HFAs) and all-flash arrays (AFAs) both share this architecture, and the combined revenue from these two platforms already dominates external enterprise storage spend. In 2016, revenue from these two platforms should generate \$15.18 billion, growing at a 5.3% compound annual growth rate (CAGR) over the next four years to crest \$19.64 billion in 2020 (see Figure 1). HFAs include a mix of flash and spinning disk media, providing options to house both performance-sensitive primary applications and more cost-sensitive secondary applications on the same platform, while AFAs use only flash media and tend to be deployed more for a mix of performance-sensitive primary applications. Both platforms can exhibit significant levels of flash optimization, depending on the vendor, providing organizations with a choice between two options that can deliver full flash performance for at least selected applications.

FIGURE 1

External Enterprise Storage Market, 2015-2020



Note: Data includes historical (2015) restatement based on new AFA taxonomy introduced in June 2016.

Source: IDC, 2016

IDC expects that over the next five years, AFAs will grow at a faster rate than HFAs and, by 2020, should drive over 70% of all primary storage spend. As flash costs continue to plummet, we will see AFAs become cost justifiable for more and more workloads, although hard disk drives (HDDs) will clearly not be going away anytime soon. Growth rates for 15,000rpm and 10,000rpm HDD shipments have dropped significantly as these HDDs have been replaced with flash-based options for performance-sensitive applications, but the extremely aggressive cost per gigabyte (GB) of 7,200rpm HDDs guarantees their continued use for more cost- and capacity-sensitive environments such as backup, disaster recovery, and archiving. The combination of primary and secondary storage requirements will have AFA and HFA revenues slowly converging over time despite the fact that primary application environments typically make up only 20-30% of an organization's storage capacity requirements (with the remainder being dedicated to what are considered secondary applications). Over the forecast period, however, the AFA market will grow at almost six times the rate of the HFA market.

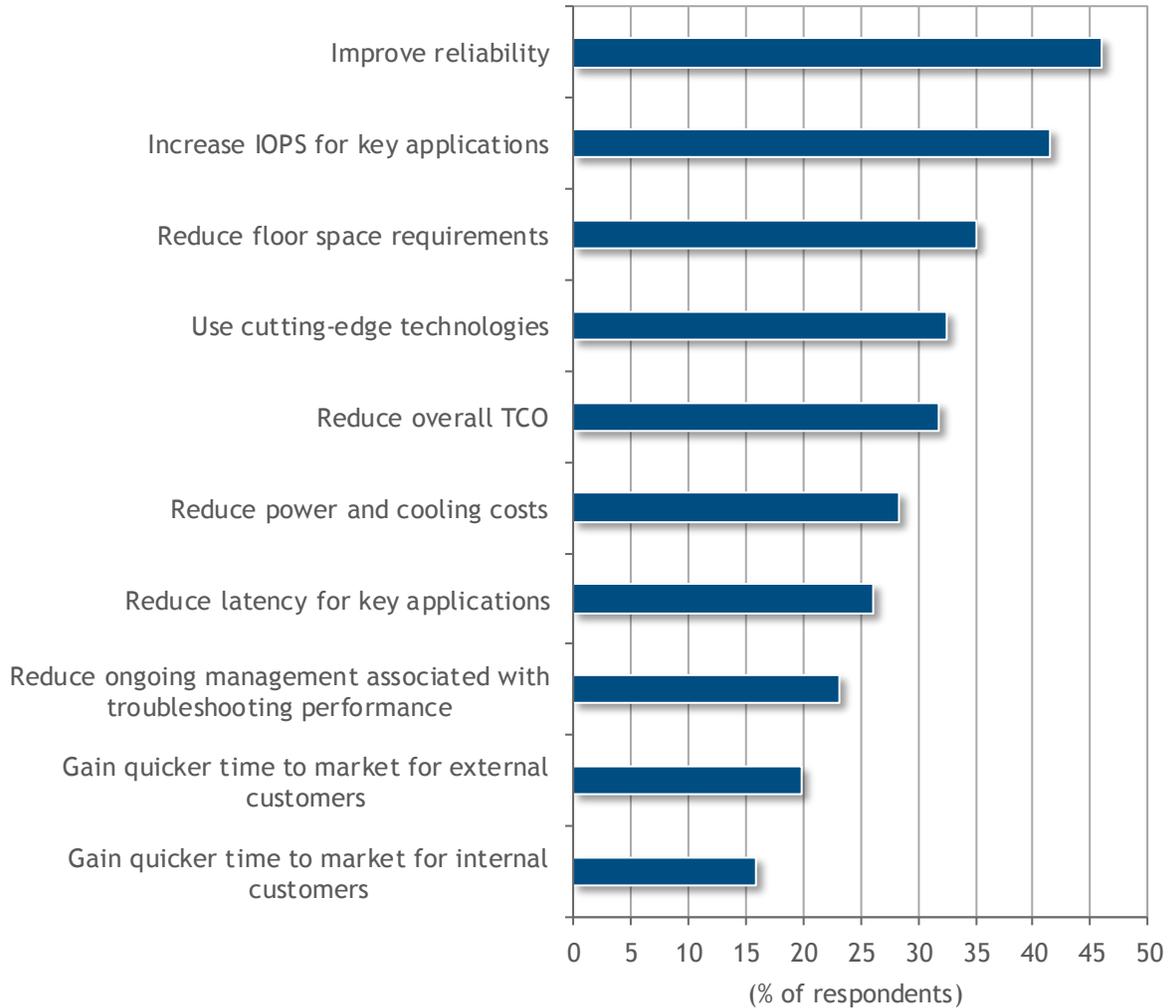
Several key factors are driving the higher AFA growth. IDC survey data from 2015 indicates that improved reliability, higher performance, and increased infrastructure density are the top 3 reasons why organizations move to AFAs when their legacy storage platforms come up for renewal. On a capacity basis, solid state disk (SSD) devices have 3x better reliability than HDDs, and as more and more datacenters find themselves managing hundreds of terabytes (TB) of raw storage capacity, this becomes increasingly important. Flash performance significantly outruns the performance available from HDDs – latencies can easily be an order of magnitude lower and will be much more predictably consistent despite varying workloads, and throughput can range from 10x to 100x better with flash.

The ability to shrink infrastructure when moving from HDDs to SSDs reduces floor space and energy requirements, allowing organizations to better utilize their datacenter space and energy budgets. Figure 2 shows the top benefits of all-flash arrays.

FIGURE 2

Top Benefits of All-Flash Arrays

Q. What does your organization see as the three most significant benefits of using AFAs?



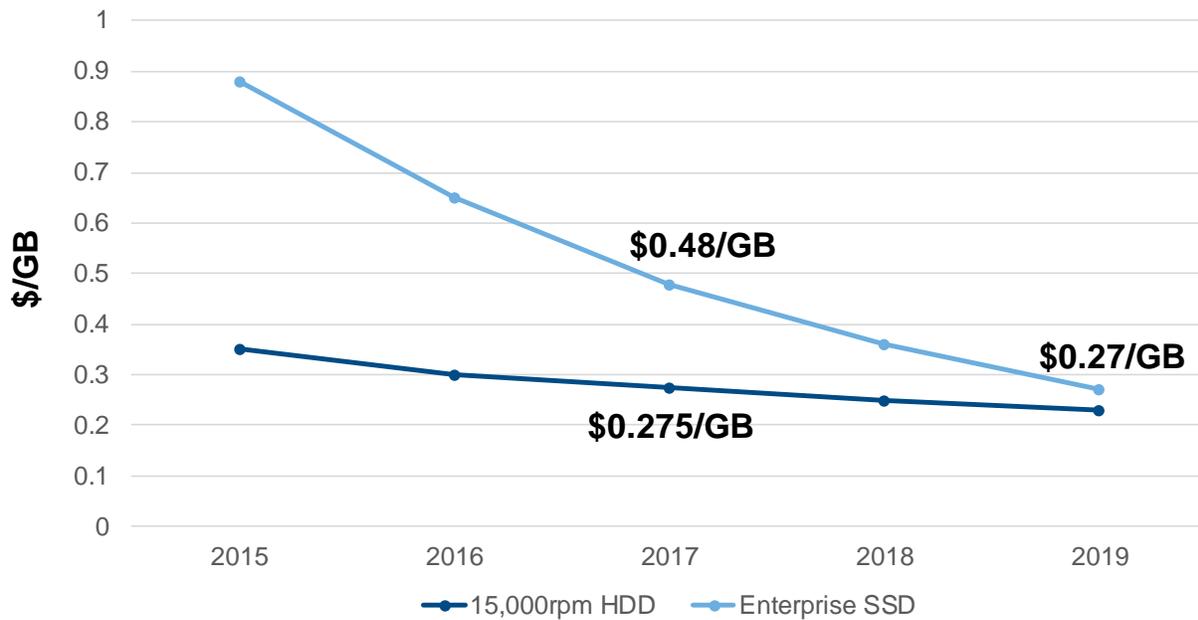
Source: IDC's survey of AFA users, March 2015

Although enterprise SSDs still cost more than 15,000rpm and 10,000rpm HDDs on a raw capacity basis, flash enables the use of inline data reduction and other storage efficiency technologies with performance-intensive applications that significantly lower their effective cost per gigabyte (see Figure 3). Data reduction ratios vary widely by workload type but can range from 2:1 for relational database environments to as much as 10:1 or more for virtual desktop infrastructure (VDI) and other clone-intensive environments. Because of the significant impact of inline data reduction technologies, raw cost

per gigabyte is no longer a relevant way to compare enterprise storage platform options – comparisons based on effective cost per gigabyte and total cost of ownership (TCO) are much more relevant. For most primary application environments, the effective cost per gigabyte of flash is already below that of HDDs because of the use of these technologies, and those TCO advantages bring with them all the reliability, performance, and infrastructure density advantages of flash. Over the forecast period, flash costs are expected to continue to decrease, only widening the existing financial advantage to flash deployment for performance-sensitive primary application environments.

FIGURE 3

Flash Costs Continue to Plummet: Enterprise 15,000rpm HDD Versus Enterprise SSD Dollar-per-Gigabyte Costs, 2015-2019



Source: IDC, 2016

In 2015, IDC noted a marked change in AFA deployment models. From 2011 through 2014, many AFAs were used as dedicated storage platforms for a single, very demanding application environment such as VDI, Oracle, or SQL Server, but as familiarity with these platforms increased, many enterprises wanted to move more workloads onto them. AFA vendors noted this change in deployment models and, over the past two years, have focused on better enabling their platforms to support mixed workload consolidation. How well an AFA platform supports mixed workload consolidation will clearly be the competitive battleground for these systems going forward. When customers want to support a mix of primary application environments on a single platform, AFAs are often deployed. When customers want to consolidate both primary and secondary applications on a single storage platform, HFAs are often chosen. In these latter configurations, performance-intensive applications can be pinned to flash storage to provide flash performance all the time, while less demanding applications can be hosted on spinning disk. Vendors that offer both AFAs and HFAs in their portfolio provide customers with the option to choose the platform that best meets their individual requirements.

NetApp: Flash-Optimized Enterprise Storage for Every Need

NetApp offers a broad portfolio of flash-optimized enterprise storage platforms. Hybrid offerings include the FAS and E-Series systems, while all-flash offerings include the EF-Series, the All Flash FAS (AFF), and SolidFire. This portfolio includes flash performance, scalability into the multipetabyte range, a wide set of data services that support easy administration and rapid recovery in a variety of scenarios, nondisruptive operations that give customers "five-nines plus" availability options, and storage offerings that allow customers to cohost block- and file-based workloads on the same storage platform. With tens of thousands of storage systems installed across Fortune 5000 customers, NetApp has proven its ability to reliably host customers' mission-critical workloads over several decades.

Since the introduction of the AFF products in August 2015, NetApp's all-flash revenue has skyrocketed. IDC observed that NetApp – with a year-over-year growth rate of 238.2% (1Q15 to 1Q16) for its all-flash business – took over the number 2 market share by revenue spot in 1Q16 and grew its revenue at almost three times the overall AFA market's growth rate over that same period (1Q15 through 1Q16). The strength of the AFF offering is its ability to reliably host mixed workloads cost-effectively, and NetApp's introduction of this platform coincided with the AFA market's move toward mixed workload consolidation.

Going forward, NetApp has a multifaceted strategy to meet customer requirements for flash-optimized performance on mission-critical and other workloads. The EF-Series delivers unadulterated flash performance for dedicated application environments that will primarily be leveraging the data services provided by the application (e.g., Oracle's DataGuard). The AFF is targeted for mixed workload consolidation, bringing with it a set of proven, mature data services and excellent datacenter ecosystem integration on a highly scalable platform that delivers "five-nines plus" availability. The SolidFire platform offers the benefits of scale-out architectures, including an ability to scale both performance and capacity in a balanced manner over time as well as the agility and cost economics inherent in these architectures, for customers looking to host workloads on Webscale infrastructure. The NetApp portfolio of storage platforms also offers data mobility across AFAs, HFAs, and cloud-based virtual storage arrays (ONTAP Cloud), which increases IT flexibility when creating disaster recovery and other tiered storage configurations.

NetApp All Flash FAS: Flash-Optimized Enterprise Workload Consolidation

The NetApp AFF is an enterprise-class storage platform built around an active-active dual controller design that leverages NetApp's proven and mature ONTAP operating environment. ONTAP is enterprise-grade data management software that supports a full suite of data services, including inline data reduction and other storage efficiency technologies; a number of RAID options supporting single, dual, and triple parity; multipath I/O; storage quality-of-service (QoS) controls; space-efficient snapshots and clones with integrated data protection capabilities; AES-256 encryption; and multiple replication options including support for zero-downtime stretched clusters – all in a platform that supports unified storage. Support for nondisruptive operations enables online firmware and software updates, hot pluggable components such as SSDs, power supplies and fans, and transparent data migration between NetApp platforms that helps simplify maintenance and technology refresh tasks.

The AFF is available in multiple flavors that range from entry level to massively scalable. The largest system currently available can support up to 1,920 SSDs, available in a variety of sizes. NetApp was the first enterprise storage vendor to introduce the 3D NAND, TLC-based 15TB SSD for use in enterprise-class storage systems, and using this drive, high-end AFF configurations can support tens to hundreds of petabytes of raw storage capacity (which can be increased significantly depending on the data reduction ratios achieved in actual production usage). File-based environments can scale

from 1 node to 24 nodes (12 high-availability [HA] pairs), while block-based environments can scale from 1 node to 12 nodes (6 HA pairs) in clustered configurations. All nodes in these configurations can be managed through a single pane of glass. Host connections include 16Gb Fibre Channel (FC) and 10GbE, while supported protocols include FC, FCoE, iSCSI, NFS, pNFS, and CIFS/SMB.

Performance and Scalability

Every AFF is highly flash optimized to deliver the most out of purchased flash capacity. A group of flash-specific optimizations, marketed under the name ONTAP Flash Essentials, ships with every ONTAP system and includes enhancements specifically designed to improve the performance, reliability, endurance, and cost-effectiveness of available flash capacity. With Flash Essentials, writes are acknowledged in NVRAM-based WAFL file system write caches and then reordered before being written to persistent flash media to improve performance, minimize garbage collection operations, and reduce the number of writes required to handle any given workload. This promotes not only low latency but also very consistent performance even under heavy workloads and maximizes the endurance (and therefore the reliability) of flash media.

Built for performance, the ONTAP WAFL file system distributes both data and metadata across all resources in a system to maintain performance consistency across varying workloads and maximize flash media endurance and reliability. Interestingly, WAFL's wide-striping approach to writing data spreads the write load evenly across flash media, further improving its endurance. Hybrid and all-flash systems can coexist in the same cluster, and the NetApp Data Fabric allows workloads to be nondisruptively moved between on- and off-premises platforms for workload balancing, to perform technology transitions without impacting application services, or to assist in scheduling maintenance operations and other tasks. A NetApp all-flash cluster can support roughly 1 million IOPS at sub-millisecond latencies and hundreds of petabytes of raw capacity.

Storage Efficiency Technologies

ONTAP includes a number of inline storage efficiency technologies that not only save space but also improve performance and increase flash endurance. In the AFF, those features include "always on" inline deduplication, inline compression and zero-block elimination, thin provisioning, FlexClones (space-efficient, high-performance clones), and replication technologies based around delta differentials (SnapMirror). Inline data reduction features like this are critical in reducing the effective cost per gigabyte for flash capacity and supporting the high storage densities that make for more efficient workload consolidation in smaller footprints (in terms of both floor space and energy). In mixed workload environments, many AFF customers achieve data reduction ratios in the 5:1 to 10:1 range (although these ratios are very dependent on workload mix).

Return on Investment

Return on investment (ROI) should be one of the major considerations in storage system selection for organizations of all sizes. For all-flash systems such as the AFF, the relevant calculation will reflect lower capital expenses because of better infrastructure density, lower operating expenses because of easier provisioning and administration with no time spent on performance tuning, reduced floor space and energy costs, the need for fewer servers to drive required storage performance (because flash latency significantly increases server CPU utilization), and lower software license costs (as a result of the need for fewer servers). Consistent high performance with predictable low latencies leads to increased revenue generation, an improved customer experience, and faster, better business decisions. Compared with legacy environments hosting a variety of mixed workloads, the AFF will

deliver a TCO that can easily be 50-70% lower than the TCO of HDD-based systems, and that advantage will only increase as flash costs continue to plummet in the coming years.

Mission-Critical, Enterprise-Class Functionality

What sets the AFF apart from other AFAs on the market is its mature ONTAP operating environment. Proven in hundreds of thousands of production installations over the past several decades, ONTAP offers the comprehensive features needed to host dense, multitenant environments that include a variety of different workloads, mission critical and otherwise. We've reviewed the enterprise-class data services, but ONTAP also supports a wide range of APIs that provide seamless integration into existing datacenter workflows: VMware's VAAI, VASA, VADP, VVOLs, and SRM; Microsoft's VSS; Oracle's RMAN; a REST API; and management APIs such as SNMP and SMS. Important features typically not available on other AFA offerings include:

- **Integrated data protection.** While many other AFAs still rely on ISV offerings to meet data protection requirements, ONTAP includes a full suite of proven tools that provide reliable recovery options: space-efficient snapshots, FlexClones, SnapVault for disk-to-disk backup, SnapMirror for replication and disaster recovery, SnapCenter for application-consistent data protection, and MetroCluster for continuous data availability (through stretched clusters).
- **Unified storage.** With customers avidly using AFAs as general-purpose primary storage platforms, any features that provide more consolidation flexibility provide extra value. All NetApp AFF models support both SAN and NAS on the same platform with a number of different block and file protocols and support the use of all data services across both block- and file-based storage. Although block-based storage was the initial focus of AFAs, more customers are looking to leverage flash performance across file-based storage as well, and the AFF is well positioned in this area.
- **Quality of service.** QoS is a critical feature in mixed workload environments because it can be used to avoid "noisy neighbor" problems that might otherwise arise. Many of today's workloads have widely varying I/O patterns, and it is important that I/O spikes that occur on one workload do not impact the performance of other workloads cohosted on that same array. This was a problem that had to be actively managed in legacy environments, but the AFF's combination of flash performance and QoS controls for IOPS ensures that storage objects (regardless of whether they are block or file based) do not consume more than their allotted resources in the face of varying workloads.
- **Multitiered storage.** NetApp's enterprise storage portfolio includes a variety of all-flash and hybrid flash platforms, offering options for customers to deploy the solution that best meets their particular needs for performance, scalability, and cost. The NetApp Data Fabric supports the easy, nondisruptive migration of workloads between platforms, including both on- and off-premises options such as cloud, as needed for test and development, production, and data protection operations. All of this storage can be managed from a single pane of glass, which makes it easy to determine and execute optimal data placement. The AFF makes up the performance-oriented tier in these environments.

CHALLENGES/OPPORTUNITIES

With AFAs now available from all the major enterprise storage vendors, they are quickly becoming a mainstream general-purpose deployment platform for mixed workloads. While some customers still want to purchase AFAs for dedicated application deployments, more customers are looking to use AFAs in a more mainstream manner and care about the multitenant capabilities of their AFAs. Most AFA vendors

offer a wide range of data services that provide the functionality necessary for mixed workload consolidation, but the maturity of these offerings vary. Vendors that have flash-optimized preexisting enterprise storage platforms that have been successfully deployed in tens of thousands of production environments clearly have a more proven, mature offering with a more comprehensive set of data services than vendors that have built AFAs from scratch in the past four to five years. For certain workload mixes, these platforms may be of more interest. There are clearly customers, though, that specifically want to deploy newer Webscale architectures, whether for certain next-generation application workloads or because the organization has a strategic plan to move more toward this kind of infrastructure over time. Vendors with a portfolio that includes all three types of AFAs, along with HFA options, provide more flexibility to meet customer needs than vendors that offer just a single platform type.

As AFAs evolve to become the enterprise storage workhorse of record over the next several years, platform maturity, enterprise-class data services, and the ability to easily integrate into preexisting datacenter workflows will become competitive differentiators. NetApp's acquisition of SolidFire, a Webscale platform that has very strong multitenant capabilities, combined with the company's existing EF-Series (for dedicated application deployment) and All Flash FAS (for the consolidation of mission-critical and other enterprise workloads), gives NetApp the flash-optimized platform portfolio to meet a wide variety of 3rd Platform computing infrastructure requirements. This provides NetApp with a market opportunity relative to those storage vendors that can offer only a single flash-optimized platform type.

CONCLUSION

Flash technology has reshaped the face of enterprise storage in just a few short years, and it will come to dominate primary storage platforms by 2020. Enterprises have a variety of needs, and it is difficult for a single platform type to provide the performance, scalability, functionality, and maturity to meet all workload requirements. Vendors with a broad portfolio of all-flash and hybrid offerings give customers the most flexibility in building the right storage infrastructure for their own environments. With its array of storage products, NetApp offers platforms that meet a variety of primary and secondary storage requirements:

- The AFF for the consolidation of mixed enterprise workloads that require comprehensive and proven data services, good datacenter integration, and "five-nines plus" reliability
- The SolidFire array for customers specifically looking for a highly scalable, flash-optimized Webscale architecture
- The EF-Series for dedicated application deployments that require the lowest latencies
- A wide range of proven HFAs at entry-level, midrange, and high-end price points that give customers flexibility in deploying multitiered storage environments under centralized management

As the crown jewel in this portfolio, the AFF is a highly flash-optimized, enterprise-class network storage platform that transcends the abilities of AFA point products with its extensive data services, integrated data protection, unified storage with multiprotocol support, and native ability to support the requirements of a multitiered storage environment that can grow nondisruptively to handle the needs of the largest enterprises.

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

Copyright Notice

External Publication of IDC Information and Data – Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2016 IDC. Reproduction without written permission is completely forbidden.

