



WHITE PAPER

NETAPP ALL FLASH FAS: PERFORMANCE OF FLASH WITH FEATURE-RICH ENTERPRISE STORAGE FUNCTIONALITY

Sponsored by: NetApp

Natalya Yezhkova
June 2015

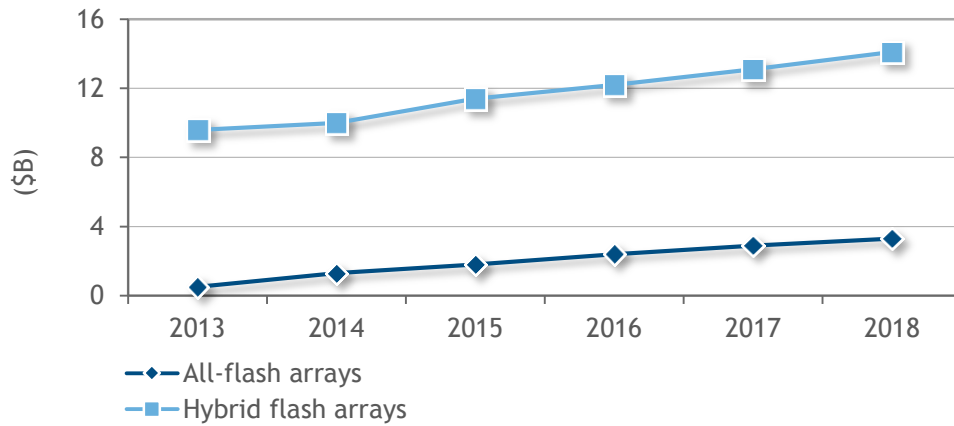
SITUATION OVERVIEW

Numerous IDC studies indicate that the overall adoption of flash storage is fast, yet there are varying degrees of adoption within individual organizations. The spectrum of adoption ranges from organizations that use flash selectively for data caching to organizations that have deployed flash as the first tier within a multitiered storage implementation. Less common are organizations that use flash for persistent storage for primary data, secondary copies, and even archives.

The forms in which flash storage is delivered to end users also vary. In its "simplest" implementation, flash is available via the PCIe flash cards typically seen on servers and on some storage arrays. Another form of flash is found in hybrid storage arrays (e.g., those designed for supporting both hard disk drives [HDDs] and solid state drives [SSDs]). These arrays typically represent the next generation of well-established storage platforms rich in functionality and optimized at different levels to include SSD in the storage media mix. Typically, hybrid arrays have a flexibility that will allow for configuration with mixed levels of HDD/SSD, ranging from just a few solid state drives to all-flash configurations. Going one step further, some vendors create a whole new line of all-flash products grown from existing well-established storage platforms. Finally, a new group of all-flash arrays (AFAs) designed from the ground up exclusively for flash. This market segment currently includes product offerings from several top market players and a small group of start-ups building their future on the benefits of flash technology. IDC projects that the AFA market will grow at a healthy CAGR of 46.1% from 2013 to 2018 and will reach \$3.3 billion by the end of 2018 (see Figure 1). The hybrid flash array market will grow at a CAGR of 8.1% during the same period. Comparatively, the overall enterprise external storage systems market is expected to grow at a low single-digit CAGR during the same period.

FIGURE 1

Worldwide All-Flash Array and Hybrid Flash Array End-User Spending, 2013-2018



Note: For details, see *Worldwide All-Flash Array and Hybrid Flash Array 2014-2018 Forecast and 1H14 Vendor Shares* (IDC #252304, November 2014).

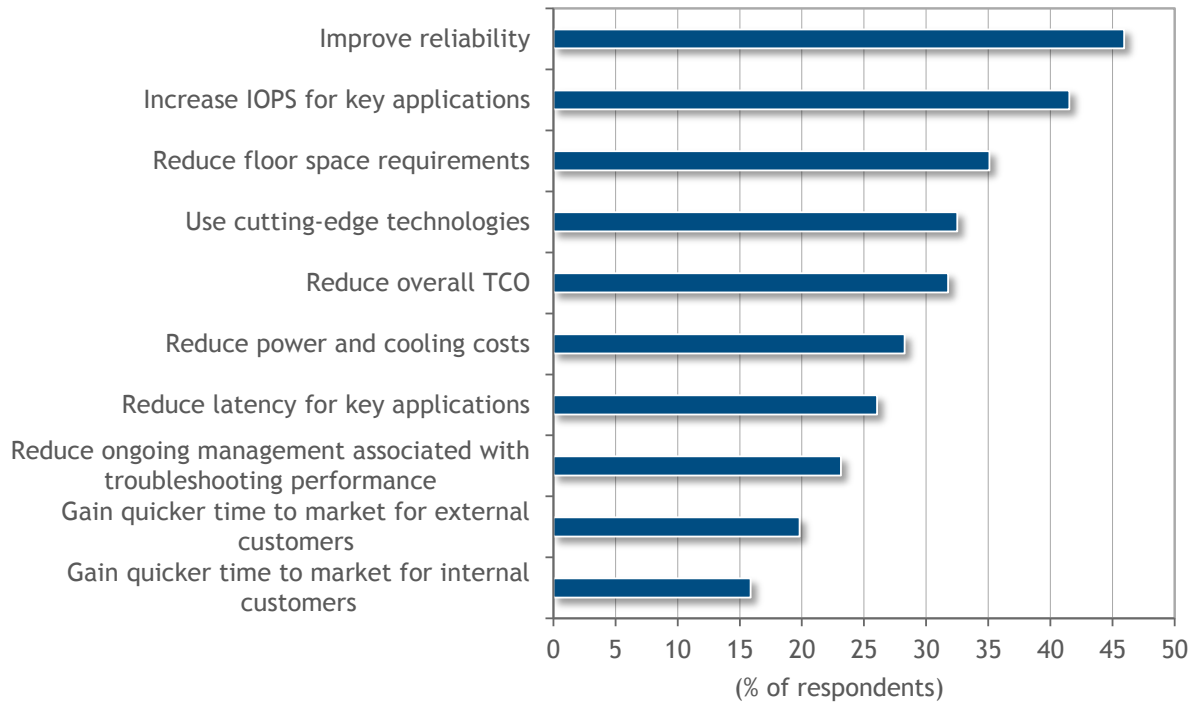
Source: IDC, November 2014

What are the major forces behind this high growth? What drives end-user interest in all-flash arrays? In a recent IDC survey of 300 end users in the United States, respondents cited performance, price, and scalability as the top 3 purchasing criteria for all-flash arrays. While performance is an obvious major reason for choosing all-flash arrays (and flash overall), price and scalability represent an important development in end users' criteria for flash. For years, the industry used a number of metrics to compare storage systems of different classes, of which the most common were IOPS for performance comparisons and raw dollar per gigabyte for price comparison. While AFAs win big time on the former, they also lose big time on the latter. Flash media is expensive, even compared with 15K performance-optimized HDDs, but, as IDC and many industry experts argue, applying the same dollar-per-gigabyte metric when comparing flash- and HDD-based systems is misleading. One of the basic benefits of flash articulated early on is its superior performance, which previously could be achieved only with multiple performance-optimized HDDs, and even then latencies were significantly higher. On top of that, most current AFA products provide built-in inline data compression, which optimizes the amount of data that needs to be processed, transferred, and stored. Taking this into account leads to a more appropriate comparison of the total cost of ownership (TCO) of flash-based systems versus HDD-based systems, an estimate that should include a variety of metrics such as dollar per IOPS, dollar per usable terabyte, and operating cost benefits. It is little wonder, therefore, that reduction in overall TCO was ranked high as a benefit of AFAs by survey respondents, after reliability and performance (see Figure 2).

FIGURE 2

Top Benefits of All-Flash Arrays

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



n = 300

Source: IDC's *Survey of AFA Users*, March 2015

Even with all the benefits, all-flash arrays are not a panacea for every storage need and every storage end user. The same survey and other IDC research suggest that the majority of end users still don't rush to adopt AFAs. There are multiple reasons that users are slow to adopt AFAs:

- Despite TCO benefits, many AFAs are still expensive. Most AFA products play in the upper midrange/high-end space with average selling prices of more than \$100,000. While the high price tag is justified for certain environments in which performance is critical to business operations, it's not easily digestible by all users who are looking for a storage system performance boost. This might change, however, because some of the vendors have started offering products priced under \$100,000 and even as low as \$25,000.
- The AFA market is not yet mature enough for a mainstream end user to risk spending storage budgets on systems (or vendors) that still don't have enough history.
- Most importantly, "built from the ground up", all-flash arrays utilize newly developed operating systems and software, thus often resulting in a solution lacking many advanced data management features that "optimized for flash" platforms have. Performance-demanding mission-critical applications, which are primary targets for AFAs, don't run in isolated environments. These applications require enterprise data life-cycle management features. Some end users can tolerate lack of such features in currently available all-flash array

offerings, but for many end users, the richness of proven enterprise-level features becomes a deciding factor in choosing how to bring flash storage into their storage environments. For them, all-flash offerings built on existing mature storage software platforms are an optimal solution that brings together the benefits of flash with the reliability of established system functionality.

There is a common misperception that established storage platforms can't compare with newly designed all-flash arrays because of the lack of "native support" for flash. Specifically, end users express concerns about the following:

- **Performance.** Can existing storage platforms deliver the full performance of flash?
- **Data optimization.** Can existing platforms deliver high effective capacity?

In the past twelve months, several storage systems suppliers have "flash optimizing" their storage systems to remain competitive among new all-flash array offerings. NetApp is one such vendor that has worked to enhance the all-flash array offering based on Data ONTAP. With its new line of products – All Flash FAS (AFF) – NetApp strives to provide end users with benefits that are above and beyond those available in native AFA products.

NetApp All Flash FAS: High-Performance Flash-Based Storage Solution with Rich Enterprise Features

The NetApp portfolio of storage system products consists of the E-series (hybrid) and EF-series (all-flash) families of block-oriented storage, hybrid FAS systems running Data ONTAP, and the extended family of All Flash FAS systems. The latter is the result of extensive flash optimization of Data ONTAP. Over the years, the FAS family has evolved to deliver very competitive solutions to the market. As a result, NetApp is the second largest supplier of external enterprise storage systems for the past two years running (according to IDC's Worldwide Quarterly Disk Storage Systems Tracker, 4Q).

The Data ONTAP operating system is the foundation of the FAS product family. It is a solid software platform that provides a broad set of enterprise features. NetApp continuously works on enhancing Data ONTAP functionality and adding innovative features to maintain its competitiveness in the market and appeal to end users. This includes flash optimization to ensure the FAS systems fully realize the benefits of solid state technology. So how do the resulting offerings from NetApp address end-user concerns related to hybrid flash arrays configured as all flash? The following sections discuss those concerns one by one.

Performance

NetApp has created an all-flash storage solution by optimizing its existing storage software (and systems) to provide flash performance while complementing mature enterprise capabilities. These enhancements available in Data ONTAP FlashEssentials and include the following:

- Coalesced writes to free blocks maximize performance and flash media longevity.
- Random read path built from the ground up and optimized for flash media.
- Parallelized processing delivers low latency.

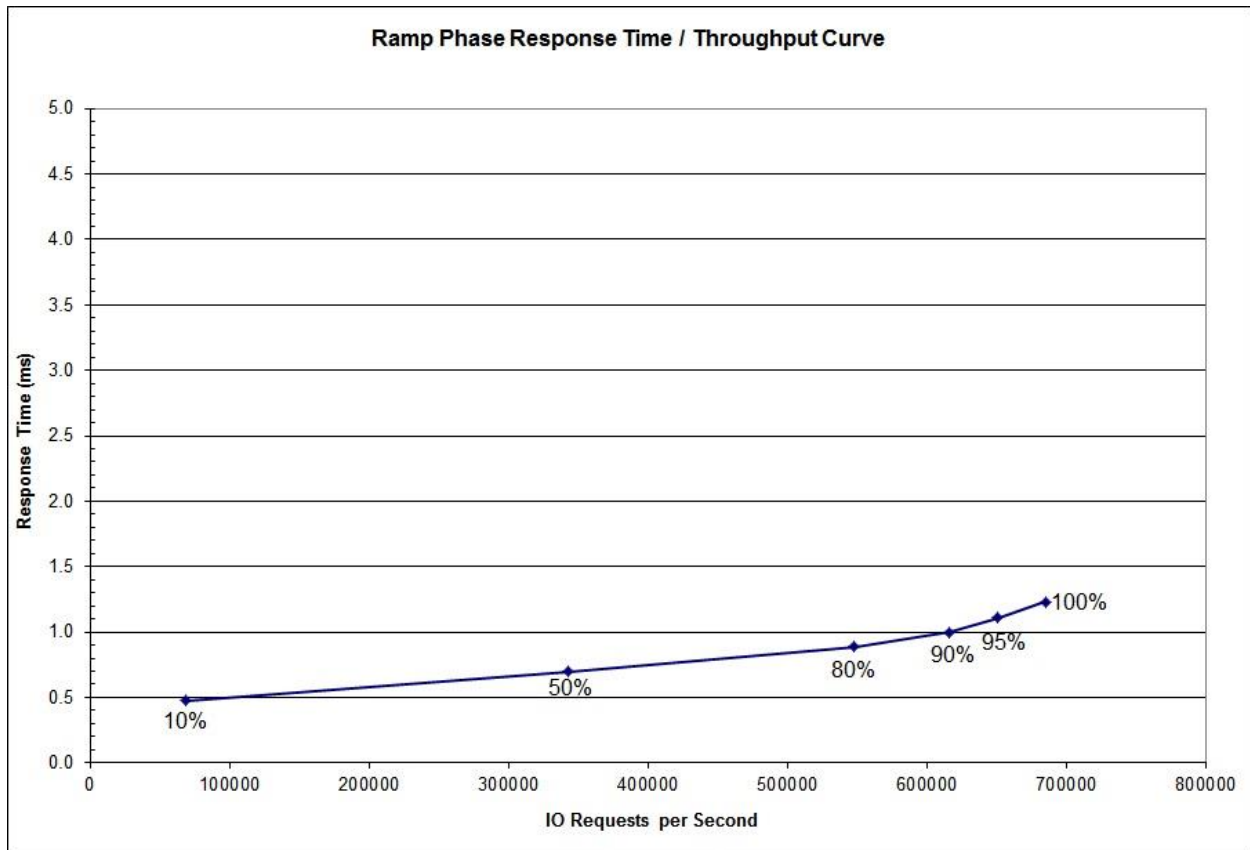
Some of these functional points are delivered through NetApp's Write Anywhere File Layout (WAFL). WAFL logs writes in NVRAM and acknowledges them immediately, resulting in very fast write response times. Random writes are then reordered and written out sequentially to a disk or an SSD. The contiguous placement gives a boost to read performance. Write performance of All Flash FAS

systems benefits significantly from this "fast write" feature provided by WAFL. The benefit of this approach becomes apparent for write-intensive workloads such as VDI and write-intensive benchmarks such as SPC-1.

Figure 3 illustrates results of the SPC-1 test of the average response time of NetApp's All Flash FAS8080EX with 4 high-availability (HA) controller pairs and 384 SSDs at various levels of I/O requests per second.

FIGURE 3

NetApp All Flash FAS8080EX SPC-1 Test Average Response Time and I/O Requests per Second



Source: NetApp, 2015

Another interesting measure of latency is least response time (LRT), the response time at 10% load. Many AFA users try to run their arrays in this "low end of the curve" to achieve minimum latency. In the SPC-1 Benchmark All Flash FAS showed 0.48ms LRT -under half a millisecond.

Data Optimization

Data compression is one of the biggest differentiators of all-flash systems. This not only leads to better utilization of available storage capacity and a fair analysis of system price (which goes beyond a straight dollar per raw terabyte) but also enables savings in operating expenses. The less storage capacity needed to store data, the less energy consumed and floor space required for a storage

system. In addition, less data is easier to manage. All these are big pluses when hundreds of terabytes or even petabyte-scale storage environments are involved.

NetApp is not a rookie in data reduction. Data ONTAP has a range of options for data reduction:

- Data ONTAP supports inline zero-block deduplication, and always on deduplication. NetApp states that it can deliver data reduction rates comparable to those provided by AFA systems.
- NetApp introduced support for inline data compression in its Data ONTAP in 2010. Since then, NetApp has been enhancing the functionality of its data compression algorithms and minimizing the burden on system resources. One of the challenges associated with data compression is assurance that system performance doesn't suffer when the compression is performed. In its upcoming Data ONTAP 8.3.1 release, NetApp will introduce the next version of the data compression functionality, which further reduces the performance impact associated with compression and expects to achieve near-zero performance impact.
- NetApp states that typical level of data reduction achieved by All Flash FAS users is 5x-10x when using inline compression, deduplication, and other space saving technologies.

Return on Investment (ROI)

Return on investment (ROI) should be one of the major considerations in the storage system selection process for organizations of all sizes. There is no single metric when measuring ROI; calculating ROI takes into account a variety of components, including monetary savings (both capex and opex), business benefits, and end-user experience. NetApp's All Flash FAS includes everything necessary to satisfy the needs of performance-demanding environments and workloads to delivery high ROI:

- Consistent high performance and low latency lead to increased business and end users productivity.
- All Flash FAS storage performance improves business operations and increases server CPU utilization, enabling consolidation of server and software license costs for substantial savings that offset investments in flash.
- Economic benefits in the form of improved total cost of ownership come from multiple areas, including reduction of data and floor space and energy savings – all adding up to the bottom line.

All Flash FAS ROI Example - SQL Server

SQL Server is one of the leading workloads for which end users deploy all-flash arrays. NetApp's internal testing of All Flash FAS8080 with Data ONTAP 8.3.1 compared to a legacy HDD-only configuration on the same SQL Server workload demonstrated remarkable performance and data efficiency:

- 4X increase in IOPS and 20X decrease in latency.
- 4X increase in server CPU utilization. Besides a direct impact on performance, this outcome also leads to the end-user ability to decrease number of servers and database licenses. In this testing, NetApp cut the number of servers and licenses needed in half.

In calculations performed over a 3-year period, NetApp estimated that deployment of the All Flash FAS resulted in **65% ROI** and **6 month payback**. The calculation includes the cost of the AFF8080 EX, eliminates the costs associated with the legacy storage system, and cuts the total number of database servers from 10 to five. This reduces SQL Server licenses and licensing costs by 50% - the biggest

cost savings factor. In addition, space savings reduced energy consumption for power and cooling. The total cost saving compared to the legacy storage system was over \$1 million.

Another important distinguishing factor for All Flash FAS systems is the ability to utilize the Data ONTAP storage operating system with all its enterprise-class functionality. This results in a short adoption cycle for current NetApp users and superior feature set for both current and new users.

The "Above and Beyond" Bonus: Enterprise-Class Functionality

NetApp offers a solid portfolio of enterprise-class features with its Data ONTAP software, which makes its All Flash FAS systems one of the most feature-rich all-flash storage systems on the market. A list of the most important features not typically available on "pure" all-flash array solutions is as follows:

- **Integrated data protection.** The job of storing production data cannot be completed without a proper data protection strategy. Many of the current AFA systems are still lacking in this area or provide limited data protection functionality that relies on the offerings of ISV partners. This works for some customers, especially those with deep technical knowledge and experience of integrating product offerings from multiple vendors.
For other enterprise customers, this approach is less preferable, particularly when reliance on a new feature for mission-critical applications is beyond their risk tolerance. For these customers, NetApp offers both off-the-shelf performance of all-flash arrays and proven data protection tools including Snapshot, SnapVault for disk-to-disk backup, SnapMirror for replication, SnapManager software for application-consistent data protection, and MetroCluster for continuous data availability. All Flash FAS also includes FlexClone technology which allows any file, LUN, or volume to be space-efficiently cloned. This has significant advantages in desktop and server virtualization, software development, and analytics.
- **Unified support for SAN and NAS.** IDC research indicates that most end users examining all-flash arrays consider deploying them for a single workload. Databases, virtual desktop infrastructure (VDI), and virtual servers are the top 3 workloads for which AFAs are deployed or considered for deployment. However, as end users have more hands-on experience with AFAs and the AFA market itself matures, deployment scenarios will quickly expand to other workloads as well, making workload consolidation a more common use case for all-flash arrays in the future. As this trend continues, end users will need broader sets of system protocols to support a variety of workloads and data use cases on all-flash arrays. A majority of current AFA solutions currently support block protocols – FC, iSCSI, and/or FCoE – but not NAS and object protocols. While demand for multiprotocol support can be a more forward-looking conversation for most environments given that most performance-driven workloads are largely block oriented, NetApp is already well prepared to address this demand.
- **Quality of service (QoS).** QoS is critical for environments in which storage systems are used for mixed virtual workload consolidation or in multitenant environments (e.g., in service provider environments). These environments become a more common target for AFA solutions. NetApp has provided storage QoS workload management since 2013, starting with clustered Data ONTAP 8.2. This feature allows the NetApp storage cluster to be divided into storage virtual machines that have their own level of permissions and rights. The service levels are defined through policies governing individual files, volumes, LUNs, or an entire storage virtual machine. QoS guarantees the delivery of an established performance level to a particular workload or tenant.
- **Scale-out architecture.** Scale out is a type of storage architecture in which performance and capacity can be expanded in a nondisruptive manner. Scale out is a fast-growing segment of

the storage market in general, and some of the all-flash array offerings, including the NetApp FAS with clustered Data ONTAP, have scale-out capabilities,. The market for scale-out architecture is poised for growth. Scale-out architecture is widely adopted in service provider environments, and adoption is also spreading in enterprise environments. All Flash FAS can scale up to 24 nodes (or 12 high-availability pairs) in a cluster, delivering millions of IOPS at sub-millisecond latency and a total of nearly 5PB of solid state storage capacity. In internally run tests, NetApp was able to achieve 4 million IOPS in 100% random reads with the All Flash FAS8080EX with 12 high-availability pairs.

- **Support for multitiered environments.** The storage needs of end users can vary dramatically depending on workload setup and business requirements. While there are companies that build storage environments on, or plan to move their storage environments to, all-flash arrays, a much larger group of end users will stick to having multiple tiers of storage that provide the best price/performance setup for various types of data. To design environments for the optimal placement of data, end users need some tools that help them identify the right storage tier for a particular data set at a particular time, depending on policies or data profile. To address these needs, NetApp provides support for multiple tiers within its own product portfolio, with All Flash FAS being the highest level of the performance-oriented tier, Dynamic data migration between flash and HDD tiers in clusters that mix all-flash and hybrid FAS systems is simple and nondisruptive. NetApp also supports integration with cloud-based storage to move data between on-premises storage and the cloud. The live workload migration between flash and HDD tiers as well as between on-premises and cloud tiers is delivered as part of Data ONTAP FlashEssentials mentioned previously.

CHALLENGES/OPPORTUNITIES

Flash-based storage is no longer a completely new concept on the market. However, it's still a young segment for many organizations that must continue to learn about the benefits of flash. These organizations are researching whether and how flash fits into their environments. For vendors, it is a perfect time to bring their all-flash or hybrid products on the market and be ready to capture the opportunity as adoption of flash storage and, in particular, all-flash arrays, spreads from the early-adopter stage to mainstream.

Deployment of AFAs delivers flash benefits, but this approach also involves the risks associated with the deployment of a new system – in many cases, from a start-up vendor. It's no surprise that adoption of AFAs comes in stages.

IDC believes that the market for all-flash arrays is about to leap to its new stage. Currently, a typical scenario for AFA deployment involves a single performance-demanding workload. But the situation is changing as end users gain experience with all-flash arrays and all-flash arrays prove their capabilities. The next stage for both new and experienced AFA users is adoption of AFA for mixed workload environments. Addressing such environments is a major opportunity for NetApp with its All Flash FAS arrays running clustered Data ONTAP.

This opportunity extends into multitiered environments. NetApp benefits from the breadth of its product portfolio, which supports a variety of storage media and cloud-based storage with integrated data migration tools, enabling dynamic movement of data between tiers. This and other enterprise features delivered by Data ONTAP are key to NetApp's competitive position in the all-flash market. For many users, the maturity of the product portfolio is critical in choosing a storage product.

As the AFA market continues to evolve, so will the richness of AFA product functionality. Current AFA vendors are working aggressively to add features to their storage platforms. This means that it is inevitable that NetApp will face stronger competition from AFA vendors and products.

The other area of opportunity for NetApp comes from its broad presence in various customer and regional segments. Some of these segments are not yet on the road map of most AFA start-ups, which mainly target large enterprises and service providers, thereby limiting competition in the near term to a few large incumbents. NetApp should use this opportunity to secure its All Flash FAS message in these market segments.

CONCLUSION

Flash-based storage is among a number of technologies that are shaping and transforming datacenters. The industry is still in the beginning of the flash adoption journey. Choosing the right partner that will provide strong support along the way is critical for long-term success. As end users start experiencing the benefits of all-flash storage solutions, they should look at the entire ecosystem of product features and data services a vendor can provide. Storage vendors such as NetApp have a solid portfolio of proven data services yet are innovative; vendors that can optimize the functionality of their systems to accommodate the latest technologies – like flash – should be on a short list for end users to consider.

NetApp's All Flash FAS arrays provide the performance end users expect from all-flash arrays while also delivering mature data services that new AFA products can't yet provide, such as integrated data protection, multiprotocol support, and native support for multitiered environments. Clustered Data ONTAP environments can encompass All Flash FAS and hybrid FAS storage deployed on-premises in corporate datacenters and the cloud. In addition, NetApp offers its customers not only a broad portfolio of its own solutions but also leading application and ecosystem integration with VMware, Microsoft, Oracle, Citrix, and other leading technology partners with solutions that can be found in any datacenter.

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-insights-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 2015 IDC. Reproduction without written permission is completely forbidden.

