



IDC MarketScape

IDC MarketScape: Worldwide All-Flash Array 2015-2016 Vendor Assessment

Eric Burgener
Eric Sheppard

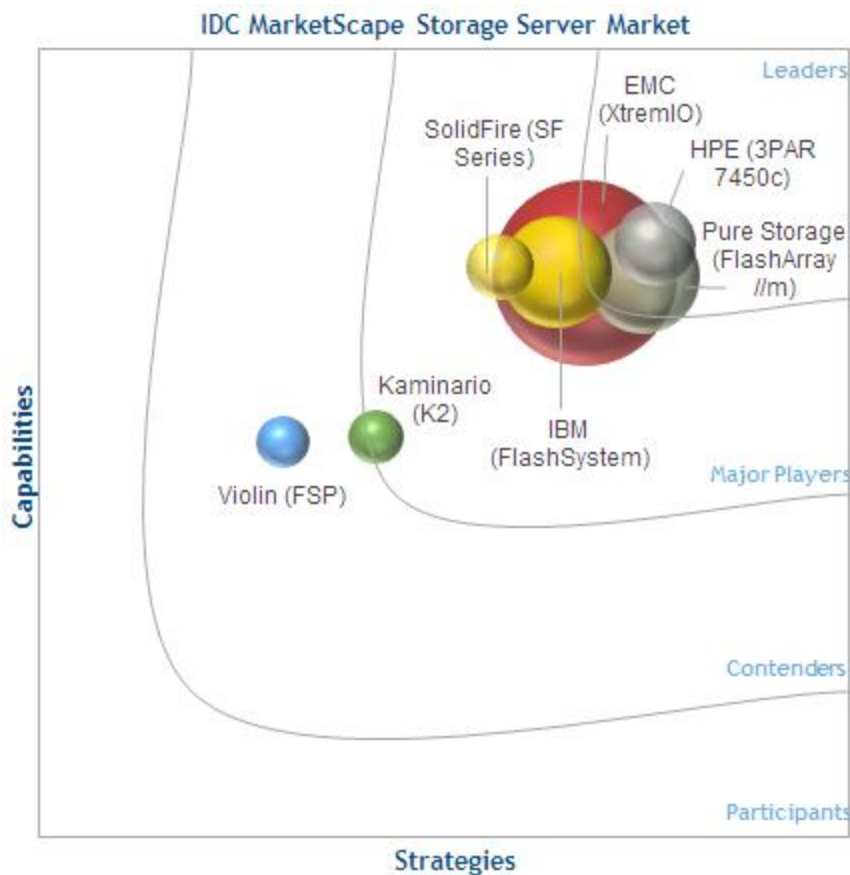
Ashish Nadkarni

THIS IDC MARKETSCAPE EXCERPT FEATURES: PURE STORAGE

IDC MARKETSCAPE FIGURE

FIGURE 1

IDC MarketScape Worldwide All-Flash Array Vendor Assessment



Source: IDC, 2015

Please see the Appendix for detailed methodology, market definition, and scoring criteria.

IN THIS EXCERPT

The content for this excerpt was taken directly from IDC MarketScape: Worldwide All-Flash Array 2015-2016 Vendor Assessment (Doc #US40721815). All or parts of the following sections are included in this excerpt: IDC Opinion, IDC MarketScape Vendor Inclusion Criteria, Essential Guidance, Vendor Summary Profile, Appendix and Learn More. Also included is Figure 1.

IDC OPINION

The all-flash array (AFA) market is beginning to transition from an emerging market to a more mature one. Among the relevant vendors evaluated in this study, a relatively comprehensive set of enterprise-class capabilities is broadly available, and it is clear that customers are thinking more and more about these platforms' ability to serve as general-purpose primary storage platforms hosting a variety of mixed workloads. All participating vendors offer solutions that easily support hundreds of thousands of IOPS, consistently deliver sub-millisecond latencies, and can support at least half a petabyte (PB) of effective storage capacity. That said, there are certain design decisions and/or characteristics that differentiate vendors, particularly in the areas of system architecture (scale-up, scale-out, or some combination of the two), media packaging (custom flash modules [CFMs] or solid state disks [SSDs]), multitenancy, support for replicated configurations, and the ability to integrate well into existing datacenter environments. There is less of a spread between vendors' existing platform capabilities today than there is between their road maps and their vision of how the companies will succeed over the next several years. Regardless of how they position their products externally, in IDC's view, each vendor has uniquely distinguished itself in a certain area. All of the AFAs from the vendors in this study can confidently be deployed today as general-purpose primary storage arrays to deliver flash performance and enterprise-class availability/reliability, providing opportunities for customers to transform both their IT infrastructure and their business opportunities. Prospective customers should look to the following second-order considerations to select the solution that best meets their requirements and make sure that they evaluate the consonance of vendors' road maps with their own vision of future enterprise storage requirements:

- The scalability and consistent high performance of EMC XtremIO's flash-optimized snapshot implementation for copy data management combined with its unprecedented rapid growth that have propelled it to the prominent market share position by revenue
- The data services comprehensiveness and maturity of HPE's 3PAR StoreServ 7450c combined with its best-in-class ability to integrate seamlessly with other platforms for virtualization, data movement, backup, disaster recovery (DR), and archive under a single consistent management environment
- The low latency, storage density, and power efficiency of the IBM FlashSystem based on its use of a CFM-based design (called IBM MicroLatency modules)
- The best-in-class combination of scale-up and scale-out storage architecture technology to deliver consistent performance and flexibility in the Kaminario K2
- The comprehensive ability to cost-effectively and nondisruptively perform in-place upgrades across technology generations of the Pure Storage FlashArray//m along with its unparalleled ability (from a quantitative point of view) to please its customers as demonstrated by its net promoter score (NPS)

- The quality-of-service (QoS) capabilities (implemented as they are in a scale-out storage architecture), the ability to support multitenancy, and the top-end scalability of the SolidFire SF Series
- The industry-leading single-node performance, storage density, power efficiency, and native replication portfolio implementation of Violin Memory's Flash Storage Platform (FSP)

IDC MARKETSCOPE VENDOR INCLUSION CRITERIA

To be included in this IDC MarketScape analysis, an AFA vendor must meet the following criteria:

- Have revenue generated from an AFA product line in 1H15 (prior to July 1, 2015)
- Use a multi-controller architecture or a scale-out design with nondisruptive failover (systems that require HA node pairs are excluded)
- Meet IDC's definition of an AFA: the system must have hardware that is unique to the all-flash configuration within the vendor's product line (other than just the choice of populating the array with all flash media) and the array's inability to support hard disk drives (HDDs) must be more than just a marketing limitation; note that all-flash configurations of HFAs (HFA/A) do not qualify
- Only use either CFMs or SSDs to meet both performance and capacity requirements within the array
- Provide at least two customer references to whom IDC can speak

Note that a number of vendors market their HFA/A products as AFAs, but because those products do not meet IDC's definition of an AFA, they have not been included. IDC's review of unpublished performance testing data using more real-world workloads (not "hero" tests) combined with simultaneous data services indicates that there is a noticeable performance difference, particularly with respect to consistent latencies, as a system is scaled up to its maximum throughput range between AFAs and HFAs. This area is evolving with several key vendors starting to narrow the performance gap between their HFA/A offerings and what IDC defines as true AFA products, and IDC will continue to keep a close eye on developments in this area.

ESSENTIAL BUYER GUIDANCE

The AFA market has undergone a significant evolution in the past year, moving away from dedicated application deployment models and more toward mixed workload consolidation. Over the past year, the major players in this space have all enhanced their offerings to provide the set of enterprise storage management features necessary for these arrays to be used as general-purpose primary storage arrays hosting many applications. There are still a few vendors that are missing a feature here or there – lack of deduplication or lack of a native replication capability – but these features are on their near-term road maps. Those vendors successfully selling AFAs to small and medium-sized enterprise have already seen their systems commonly used as general-purpose primary storage arrays, but more and more large enterprises are evaluating AFAs' multitenant capabilities prior to initial purchase.

Table stakes in the AFA market include the following functionalities, which are either available on the systems already or in the near-term road map:

- Stellar flash performance, easily providing hundreds of thousands or more IOPS at consistent sub-millisecond latencies combined with effective flash capacity of at least half a petabyte

- Wide striped data layout for improved performance consistency, better reliability, and faster recovery in the event of storage device failures
- Use of a variety of storage efficiency technologies that do not preclude systems from still delivering consistent sub-millisecond performance across their entire throughput range, including inline data reduction, thin provisioning, delta differentials for snapshot-based replication, and broad support for VAAI-based *unmap* functionality, among others
- Flash-optimized RAID 6 implementations for write minimization and capacity maximization
- All or most data services included with the array base price (except for, in some cases, encryption and replication)
- Effective price per gigabyte for most vendors, ranging from \$1 per gigabyte to \$2 per gigabyte (assuming data reduction ratios of 4:1), supported by the increasing use of cMLC flash media moving toward 3D NAND TLC in the near term
- Broad API support to aid in enterprise infrastructure integration – including VAAI, VADP, SRM, VSS, RMAN, SNMP, and RESTful APIs and/or scriptable CLIs for management – with VVOLs and Microsoft Hyper-V ODX generally being a near-term road map item (if not already available)
- Field-proven five-nines-plus reliability closing in on six-nines based on real-time monitoring of system uptime and usage statistics, which will increasingly be utilized not only by vendors but also customers directly

Despite the broad commonality in enterprise storage management capabilities, there are still a few areas where vendor offerings show significant differentiation. Storage architecture is one of those areas. There are classic dual-controller architecture systems (such as Pure Storage FlashArray//m and Violin Memory Flash Storage Platform), and there are systems built around classic scale-out designs (EMC XtremIO and SolidFire SF Series). As noted in *Evaluating Scale-Out and Scale-Up Architectural Differences for Primary Storage Environments* (IDC #256932, June 2015), more and more storage systems across all areas of enterprise storage are blending aspects of both storage architectures to improve scalability and flexibility. IBM (FlashSystem) can cluster up to four dual-controller nodes, allowing workloads to be spread across all nodes in the cluster (although workload redistribution when nodes are added requires manual intervention). Kaminario (K2) exhibits the elements of both scale-up and scale-out, allowing up to eight dual-controller nodes (each of which supports internal capacity expansion) to be clustered together under a true single-system image, complete with fully distributed data and metadata. HPE can scale a single system, like a 3PAR StoreServ 7450c, in terms of both controller nodes (up to four) and capacity within a single-system image and then can further cluster nodes together under a single management (but not system) image.

Other architectural design decisions in which systems differ include the use of CFMs instead of SSDs and whether a system uses active/active or active/passive controllers. Vendors using CFMs argue that increased visibility at the individual flash cell level enables more efficient and less impactful garbage collection operations, but with most systems in production use in the field still running at under 30% utilization, it is difficult to prove this based on real-world customer experience. IDC notes that those AFAs that use CFMs tend to lead the market in storage densities (TB/U) and power efficiency. Higher storage density can make for more compact systems (less floor or rack space required), but all evaluated systems easily support at least half a petabyte or so of effective capacity with near-term growth paths in the petabyte range and beyond. Active/active or active/passive controller implementations will have different performance impacts in post-failure scenarios.

Another area where significant differences exist today between vendors is in quality of service (QoS) controls. A premier vendor in this space, SolidFire, sets the bar by providing a complete set of controls on

the array itself (ability to set maximum, minimum, and burst limits for IOPS, QoS evenly balanced across all available system resources, and admission control) as well as extensive integration with QoS controls on third-party platforms, like VMware vSphere. Most AFAs in use today run at relatively low levels of utilization and with few applications, minimizing "noisy neighbor" concerns. However, as customers start to use these systems more and more as general-purpose primary storage arrays, contention for system resources will become a real concern. As a group, service providers using AFAs tend to run them at the highest levels of utilization, and SolidFire's strong position in QoS is due in large part to its initial focus on the service provider market (which has since been expanded to more broadly target enterprise users as well). Other vendors' road maps include QoS features as part of a general set of enhancements to improve these arrays' ability to support dense multitenancy while continuing to meet performance and security requirements. Note that the ability to apply data services selectively at the volume level (instead of just at the system level) is another important multitenancy requirement.

Over the past three to four years, most AFAs in the enterprise were initially purchased specifically for a single application – most often a database or a virtual desktop infrastructure (VDI) environment. Customers love the performance and ease of use that AFAs bring to the table (storage performance tuning time immediately drops to zero), and 100% of the customers interviewed by IDC about their AFA experiences expressed an interest in moving more workloads to flash over time. This has driven the interest in array-based data services, and replication is key. AFAs are in widespread use with mission-critical application environments, and customers need disaster recovery solutions for these configurations. The replication technology portfolio is a strong differentiator amongst vendors today with some only supporting replication through the purchase of separate products, some offering just the basic snapshot and/or asynchronous replication, and others offering a full replication portfolio – including snapshot, continuous asynchronous, and synchronous replication as well as stretched cluster support, data migration between disparate systems, and other features. A vendor's portfolio strategy is an important consideration, and those vendors that offer an ability to replicate between their AFAs (at a primary site) and other less-performant but potentially less-expensive HFAs (at secondary sites) have an important differentiator.

When making buying decisions, it is important to understand the systems' future intended use. All systems can deliver extremely high performance and low latency for a single application, but a system's ability to support key multitenant features should be considered for customers that plan to host multiple workloads, thereby maximizing the return on investment (ROI) that derives from the secondary economic benefits of flash deployment at scale. IDC believes that mixed workload consolidation is the future of AFAs and, by 2019, they will dominate primary storage spend in the enterprise. We also expect those systems that best support the features and flexibility that multitenancy requires to be the near-term winners.

In closing the "guidance" comments, it is important to take note of how the AFA vendors got to where they are now. Six of the seven (EMC, IBM, Kaminario, Pure Storage, SolidFire, and Violin) started with a blank sheet of paper to create systems that were specifically optimized for flash and, in the process, they were able to create some interesting differentiators: EMC's flash-optimized snapshot implementation for copy data management, IBM's excellent power efficiency, Kaminario's combination of scale-up and scale-out in a single system, Pure's comprehensive ability to cost effectively and nondisruptively perform in-place upgrades across technology generations, SolidFire's QoS, and Violin's excellent single-node performance. HPE came at this problem from a different point of view, taking the more mature 3PAR operating environment (which interestingly had originally been designed around the wide data striping that all the other vendors starting with a blank sheet of paper went with as well) and choosing to flash optimize that over time. HPE's strength is in the comprehensiveness and

maturity of its data services offerings, allowing customers to take advantage of flash performance while continuing to work within the familiar 3PAR operating environment for management. For some enterprises, particularly those that care about data services maturity and the ability to integrate with and move data between AFAs and other non-AFA storage platforms, these could be important features. For a further discussion along those lines, see *Strategic Considerations for Choosing Between All-Flash Arrays and Hybrid Flash Arrays* (IDC #258876, September 2015).

VENDOR SUMMARY PROFILES

This section briefly explains IDC's key observations resulting in a vendor's position in the IDC MarketScape. While every vendor is evaluated against each of the criteria outlined in the Appendix, the description here provides a summary of each vendor's strengths and challenges.

Table 1 shows the list of vendors, the product evaluated for each, and their classifications.

TABLE 1

List of Vendors and Their Classifications

IDC MarketScape Category	Participating Vendor	Product Evaluated
Leaders	EMC	XtremIO
	HPE	3PAR StoreServ 7450c
	Pure Storage	FlashArray//m
Major Players	IBM	FlashSystem 900/V9000
	SolidFire	SF Series
	Kaminario	K2
Contenders	Violin Memory	Flash Storage Platform

Source: IDC, 2015

Pure Storage

Pure Storage was one of the earlier entrants in the AFA space and the first vendor to realize and capitalize on the importance of inline data reduction. The company first shipped its FlashArray platform in 3Q11, right around the same time that Violin Memory shipped its own; however, Pure was positioning itself more as a general-purpose storage solution even then while Violin (with its excellent performance but limited data services) was clearly pursuing the "dedicated application" deployment mode at that time. Pure uses a scale-up, active/active (on the front end), dual-controller architecture in a system that can deliver over 1 million IOPS (assuming the standard 100% 4K random read hero test) and scale to support 136TB of raw storage capacity (or over half a petabyte of effective capacity assuming a 4:1 data reduction ratio). Interestingly, Pure does not quote performance at 4K random

reads but uses a more realistic 32K block size (given that block sizes in mixed virtual workload environments tend to vary significantly but average, according to IDC's research, in the 24K-32K range), claiming to offer up to 300,000 IOPS.

Pure primarily sells its products through partners, using a high-touch, channel-fulfilled model that has scaled well over time. Pure's solutions can be purchased either as a storage appliance or as part of a converged infrastructure stack called FlashStack that is jointly built on components from Cisco, Pure, and VMware. The FlashArray//m – Pure's latest product announced in June 2015 – offers a full complement of data services, including inline data reduction, snapshots, clones, AES 256-bit encryption, asynchronous replication, and QoS capabilities. The inline data reduction leverages a number of different technologies, including variable compression, 512-byte aligned variable deduplication, pattern removal, metadata copy reduction, thin provisioning, and a secondary "deep reduction" post-process compression. All data services are included as part of the base array purchase price.

Pure does offer a sub-\$100,000 price point, but primarily sells into the medium-sized and large enterprise. Pure customers, even early on, were hosting multiple applications on their FlashArrays, and Pure is sold across a wide range of verticals to support databases (both relational and NoSQL) and other applications, mixed virtual server workloads, and VDI environments.

Strengths

Pure has been the most successful of the AFA start-ups that did not get acquired; although, this past summer, it did complete an IPO. It is a well-funded company that – despite its size – is still growing the overall business very rapidly, and it is doing that while maintaining an unprecedentedly high level of customer satisfaction as evidenced by its extremely high Net Promoter Score. Pure does not share this score publicly, but it is the highest NPS this analyst has seen for any enterprise storage vendor ever by a wide margin, even putting the scores of companies such as Apple (that are the darlings of their customer base) to shame. Pure has earned this high rating by how customers are treated during the sales process, how easy its arrays are to deploy and use, the extremely high levels of availability its systems provide in production use, how responsive the company is to customer support issues, and its ability to deliver value across generational upgrades for its customers. Other vendors offer guarantees around performance and data reduction ratios that will be achieved in actual usage, but only Pure offers a Love Your Storage Guarantee that allows the unconditional return of a purchased system within 30 days, no questions asked.

Pure handles generational upgrades and investment preservation like no other company and, in fact, has made specific hardware enhancements in its FlashArray//m platform that enables nondisruptive, in-place upgrades of all components to newer generations – not just controllers and SSDs but also NVMe cache cards (which are uniquely hot pluggable) and internal array backplane bandwidth. This program is called Evergreen Storage, and it offers unique value to Pure's customers by delivering flat and predictable maintenance pricing throughout the multigenerational life of the system, free controller upgrades every third year, and ongoing maintenance and support for all components (including flash media) at guaranteed prices for as long as the customer chooses to continue to use the system. With Evergreen Storage, customers avoid the hassle and cost of forklift upgrades to move to new generations of technology.

Pure delivers AFA industry-leading management capabilities with Pure1, a global cloud-based storage management platform fully instrumented with telemetrics tracked at the individual system and the entire installed base level that are available for both customers and Pure employees in engineering, support, and product management. While this platform is not unique in the enterprise storage industry

(Nimble Storage, an HFA vendor, offers similar capabilities with its InfoSight Portal), it was an early provider in the AFA space and has prompted AFA competitors to emulate this kind of functionality.

Challenges

Today, the FlashArray//m only supports snapshot-based replication. Customers interested in deploying replicated configurations using continuous asynchronous, synchronous, or stretched clusters will need to wait until Pure introduces these capabilities in the near future. Pure's data services are also "always on," which can be an issue in certain mixed workload environments where customers may want to deploy, for example, encryption for use with some applications but not with others. Pure's adaptive data reduction tracks data reducibility in real time, and if no benefits are achieved, then it can automatically shut off the inline compression and deduplication at the volume level. This adaptive capability can conserve CPU cycles and reduce latencies; although, even with data services in the "always on" mode, the system is still able to produce consistent sub-millisecond latencies in mixed workload environments. Pure also ranks low on the storage density metric, supporting a maximum today of 12.2TB/U on the FlashArray//m. Storage density is expected to double, however, in the next several months.

Pure is the only true AFA vendor today that does not support clustering. While the system supports enough performance and capacity for near-term requirements, support for clustering would provide a long-term growth path for customers that does not exist today.

APPENDIX

Reading an IDC MarketScape Graph

For the purposes of this analysis, IDC divided potential key measures for success into two primary categories: capabilities and strategies.

Positioning on the y-axis reflects the vendor's current capabilities and menu of services and how well aligned the vendor is to customer needs. The capabilities category focuses on the capabilities of the company and product today, here and now. Under this category, IDC analysts will look at how well a vendor is building/delivering capabilities that enable it to execute its chosen strategy in the market.

Positioning on the x-axis, or strategies axis, indicates how well the vendor's future strategy aligns with what customers will require in three to five years. The strategies category focuses on high-level decisions and underlying assumptions about offerings, customer segments, and business and go-to-market plans for the next three to five years.

The size of the individual vendor markers in the IDC MarketScape represents the market share of each individual vendor within the specific market segment being assessed.

IDC MarketScape Methodology

IDC MarketScape criteria selection, weightings, and vendor scores represent well-researched IDC judgment about the market and specific vendors. IDC analysts tailors the range of standard characteristics by which vendors are measured through structured discussions, surveys, and interviews with market leaders, participants, and end users. Market weightings are based on user interviews, buyer surveys, and the input of a review board of IDC experts in each market. IDC analysts base individual vendor scores, and ultimately vendor positions on the IDC MarketScape, on detailed surveys and interviews with the vendors, publicly available information, and end-user experiences in

an effort to provide an accurate and consistent assessment of each vendor's characteristics, behavior, and capability.

Market Definition

In 2015, AFAs are expected to generate \$2.53 billion in revenue and grow at a compound annual growth rate (CAGR) of 21.6% to crest \$5.5 billion in 2019. At that point, AFA revenue is expected to account for 60-70% of all external primary storage spend. IDC defines AFAs as enterprise-class storage arrays that must:

- Only use either CFMs or SSDs to meet both performance and capacity requirements within the array
- Include hardware that is unique to the array (other than just the choice of populating the array with all-flash media) and is not shared with HFAs (if present in the vendor's portfolio)
- Include a specialized operating environment that is specifically optimized to support and get the most from a performance, endurance, and reliability point of view out of flash media

LEARN MORE

Related Research

- *External Storage Systems and the Impact of Flash: Comparing Growth of Different External Array Types* (IDC #259426, October 2015)
- *Oracle Enters the High-Growth All-Flash Array Market with an Oracle-Optimized Solution* (IDC #259344, September 2015)
- *Tegile Enters the All-Flash Array Market with a Super Density Offering* (IDC #259253, September 2015)
- *Strategic Considerations for Choosing Between All-Flash Arrays and Hybrid Flash Arrays* (IDC #258876, September 2015)
- *Executive Interviews on Flash Array Deployments in the Enterprise, 2015* (IDC #258124, August 2015)
- *Evolving Flash-Optimized Architectures* (IDC #256994, June 2015)
- *Evaluating Scale-Up and Scale-Out Architectural Differences for Primary Storage Environments* (IDC #256932, June 2015)
- *Managing the IT Infrastructure Balance Point for Improved Performance and Efficiency* (IDC #256754, June 2015)
- *IDC's Worldwide Flash Storage Solutions in the Datacenter Taxonomy, 2015* (IDC #255995, May 2015)
- *Kaminario Exhibits Major Market Momentum in the Hotly Contested AFA Market* (IDC #255495, April 2015)
- *IBM and Micron's One-Two Punch Around the New IBM FlashSystem Launch* (IDC #254862, March 2015)
- *SolidFire Poised for a Breakout Year in 2015* (IDC #254631, March 2015)
- *HPE 3PAR Architecture Demonstrates Performance Milestone for Flash Arrays* (IDC #254327, February 2015)

Synopsis

This IDC study represents a vendor assessment model called the IDC MarketScape. It's a quantitative and qualitative assessment of the characteristics that assess vendors' current and future success in the AFA market segment and provide a measure of the vendors' ascendancy to become a Leader or maintain a leadership position. IDC MarketScape assessments are particularly helpful in emerging markets that are often fragmented, have several players, and lack clear leaders.

The AFA market segment, which is part of the overall external enterprise storage market, is an example of an emerging market. In this IDC MarketScape, IDC attempts to assess the capabilities and strategies of key vendors of AFA storage platforms. IDC expects that market forces such as fierce competition and buyer demand will accelerate the evolution of this market into a mature market with only a few dominant vendors. In all likelihood, the only survivors in this market may ultimately be vendors with broad portfolio offerings that enable seamless integration of AFAs intended for use as primary storage platforms with other types of existing and emerging storage platforms, particularly secondary storage platforms. A key question in this market over the next several years concerns the viability of AFAs as cost-effective secondary storage platforms.

"The IT infrastructure and business transformation that AFAs enable will be strategic in the successful companies of the future," said Eric Burgener, research director, Storage Systems. "IDC expects AFAs to dominate primary storage market spend by 2019, opening up a key subsequent phase of evolution in enterprise storage markets: how secondary storage platforms will evolve to better meet the real-time needs of 3rd Platform computing."

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

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, telebriefings, and conferences. Visit www.idc.com to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit www.idc.com/offices. Please contact the IDC Hotline at 800.343.4952, ext. 7988 (or +1.508.988.7988) or sales@idc.com for information on applying the price of this document toward the purchase of an IDC service or for information on additional copies or Web rights. [trademark]

Copyright 2015 IDC. Reproduction is forbidden unless authorized. All rights reserved.

