



## IDC MarketScape

# IDC MarketScape: Worldwide All-Flash Array 2017 Vendor Assessment


Eric Burgener

**THIS IDC MARKETSCAPE EXCERPT FEATURES PURE STORAGE**

### IDC MARKETSCAPE FIGURE

**FIGURE 1**

## IDC MarketScape Worldwide All-Flash Array Vendor Assessment



Source: IDC, 2017

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 2017 Vendor Assessment (Doc #US43310017). 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 has undergone significant maturation over the past two years. A high percentage of customers have already committed to an "all flash for primary storage" strategy, and every customer interviewed for this study was among those. In 2017, AFAs will drive over 80% of all primary storage revenue. All of the established storage vendors have entered this space, and there are several start-ups with over \$100 million in revenue. With this level of market maturation, multiple segments have developed within the primary flash array space. There are systems targeted for dedicated application deployment, there are systems specifically for web-scale applications, and there are systems intended for dense mixed workload consolidation. These latter systems are driving most of the AFA revenue, and they aspire to become the primary storage platforms of record for enterprises of all sizes. This study evaluates the suitability of 10 vendors' AFA platforms for dense mixed enterprise workload consolidation that includes at least some mission-critical applications.

The successful vendors in this space have all basically determined what is required for dense mixed workload consolidation and, for the most part, have delivered it. Any of the AFAs evaluated in this study is appropriate for this use case, but some are better geared for specific requirements (dual-controller or scale-out designs, block or file storage or both at the same time, virtual only or mixed physical/virtual, etc.). They all offer adequate performance (in terms of throughput and latency), capacity, and functionality to consolidate workloads, although a few are still missing support for advanced replication configurations like stretch clusters. Platform maturity is another important consideration, and several of the vendors entered the AFA market in only the past couple of years (although their storage operating systems may have been shipping for several years before that on hybrid flash arrays [HFAs]).

The areas where the most differentiation between vendors was noted were in their strategies around NVMe and cloud-based predictive analytics, how they track and manage customer experience (CX), and how they handle technology refresh. For most customers looking for a reliable platform for mixed enterprise workload consolidation, these may all be secondary considerations today, but the more forward-thinking vendors in these areas are driving industry transformation in a way that will significantly benefit customers over the long run.

While Pure Storage was rated highly in this analysis, three other vendors (HPE, Dell EMC, and NetApp) also fared well. While all four of these vendors were very similar in terms of capabilities, Pure Storage's strategies in NVMe, cloud-based predictive analytics, CX, and technology refresh (with the Evergreen Storage program) made the difference for the company. Other vendors that were rated as Major Players each exhibit some strongly differentiating features and/or reputations: Hitachi Vantara for its reliability (backed by the industry's only 100% availability guarantee), Kaminario for its excellent blend of scale-up and scale-out capabilities and innovative composable storage strategy based around

NVMe over Fabric (the latter capability was not announced in time to be incorporated into its rating), IBM for its strong FlashCore technology (its use of custom flash modules [CFMs] across multiple AFA platforms), and Tegile for the value it provides based on its pricing models. Huawei, which is rated as a Major Player, has product functionality on par with all the others in the market but exhibited more conventional thinking in its telemetrics, CX, and technology refresh approaches. Tintri, a smaller vendor that shows up as a Contender, has the strongest virtual machine (VM)-aware offering across all vendors, a factor that differentiates it in terms of ease of use. Although Fujitsu did not meet the revenue requirement for inclusion, as its platform (the ETERNUS AF) had just recently started to ship, functionally the company is on par with many of the vendors in this study, but it did not fare as well in some of the newer ways that forward-thinking enterprise storage vendors are driving value for their customers (CX management, cloud-based predictive analytics, and technology refresh).

## IDC MARKETSCOPE VENDOR INCLUSION CRITERIA

---

This IDC study assesses the capabilities and business strategies of key suppliers in the AFA market, the fastest-growing segment of the external enterprise storage market. This evaluation is based on a comprehensive framework and a set of parameters that gauge the success of a vendor in delivering an enterprise storage platform intended for dense mixed workload consolidation that includes at least some mission-critical applications. This study includes analysis of the 10 most notable players in this space today, which (in alphabetical order) are Dell EMC, Hitachi Vantara, Hewlett Packard Enterprise (HPE), Huawei, IBM, Kaminario, NetApp, Pure Storage, Tegile, and Tintri.

To make this list, vendors need to have a platform that meets the following criteria:

- **Meets IDC's definition of an AFA.** IDC defines an AFA as any external storage array that supports only all-flash media as persistent storage and is available under a unique SKU. Systems that can support a mix of solid state disks (SSDs) and hard disk drives (HDDs) are considered by IDC to be hybrid flash arrays, even if they can ship in all-flash configurations. These arrays are not included in this IDC MarketScope. Candidate arrays must be targeted for primary storage workload consolidation even if in practice customers may also host some secondary storage workloads on them. For more information on the AFA taxonomy, see *IDC's Worldwide Flash in the Datacenter Taxonomy, 2017* (IDC #US42214317, January 2017).
- **Is developed and owned by that vendor.** The vendor needs to be the owner of the intellectual property (IP) associated with the candidate array and is responsible for developing the platform going forward.
- **Was generally available by June 30, 2017.** System functionality for this study was compared based on what was actually shipping by June 30, 2017. Candidate arrays must also have generated at least \$50 million in end-user revenue by June 30, 2017, for the 2017 calendar year to be considered for this study.

IDC recognizes at least three different kinds of AFAs: primary flash arrays, big data flash arrays, and rackscale flash arrays. This study focuses on primary flash arrays, but vendors of all primary flash array types (types 1, 2, and 3) were considered before the 10 most relevant vendors were selected. The purpose of this IDC MarketScope is to specifically evaluate candidate arrays for their ability to be used as storage platforms for the dense consolidation of mixed enterprise workloads that include at least some mission-critical applications. With this focus, the breadth of enterprise-class functionality is important, but so are overall availability and reliability, scalability, platform maturity, manageability, and ability to integrate into preexisting datacenter workflows (among others). Other areas of strategic importance include the vendor's approach to cloud-based predictive analytics, the availability of the IP

through multiple consumption models, programs around standardized customer experience management, and the vendor's plans with respect to NVMe.

Note that some vendors offer a single AFA, while others have a very broad portfolio of AFAs. For this study, vendors needed to select a single AFA platform that would be evaluated for its ability to host mixed enterprise workloads. There can be significant value, however, in a vendor offering a portfolio of AFAs, each of which may be targeted for different deployment models (e.g., mixed workloads, dedicated single workload, and web-scale workloads). For that reason, some weight is given to the availability of an AFA portfolio in the strategies criteria.

In the AFA market, there are at least 30 different AFA products, although not all are specifically targeted for dense mixed workload consolidation. There were a number of available AFA products that did not make the list because they did not meet one or more of the selection criteria.

## ADVICE FOR TECHNOLOGY BUYERS

---

This study evaluates candidate arrays against a specific use case: dense mixed workload consolidation that includes at least some mission-critical workloads. While all AFAs provide much higher performance than HDD-based arrays, the AFAs in this study span a wide range of performance and capacity scalability, exhibit different architectures, and take different approaches to the efficiency of resource utilization. They all offer a pretty complete set of enterprise-class data services, good datacenter integration capabilities, and multitenant management features. It is safe to say that any of these AFAs could be used to handle enterprise workload consolidation, but there are differences between them of which customers should take note. Some of the questions customers can ask themselves to help determine which of these AFAs might be a better fit for them are:

- Do I want to consolidate block- and file-based workloads on the same platform, or do I want to use separate platforms for those workloads? Do I want to consolidate both physical and virtual workloads on the platform, or am I just interested in virtual workloads?
- Do I want to get the next generation of an incumbent array (because of preexisting management expertise or other issues), or do I want to move to a new vendor that shows distinct differences in focus from the incumbent?
- Do I need synchronous replication and/or stretch cluster support now, or do I just need to know that the AFA will support it in the near future?
- How important is platform maturity to me?
- What is the proven reliability in actual usage based on customer references? All vendors pretty much claim to support at least "five-nines plus" availability, and many of them guarantee that, but anecdotal evidence indicates that reliability can still be a significant differentiator between platforms.
- Am I looking for vendors that are strong proponents of modernized ways of doing business (cloud-based predictive analytics, standardized CX management programs, and an updated technology refresh model), or have I been happy with existing approaches?
- How important are design attributes to really fine-tune the efficiency of resource utilization?
- In considering failure domains, what are my "performance" and "storage density" sweet spots? For example, to hit a requirement for 5 million IOPS, some customers want to purchase multiple smaller platforms, while others want a single platform for that.

One other area to consider is your comfort level with "self-driving storage." Some AFAs are very far down this path, and while they provide minimal opportunity to manually tune the array, they are much easier to manage, a consideration very important for customers that have delegated storage management tasks to information technology (IT) generalists (like more and more shops of all sizes are doing). Others offer all the manual tuning capabilities of the arrays of the past but layer templates and automated workflows into their GUIs to improve ease of use.

In the area of design efficiencies, these features tend to have less of an impact on smaller capacity configurations but more of an impact on larger ones. If you are installing a 150TB AFA, then you would likely care less about these issues than if you're installing a 1.5PB AFA. These features include the obvious considerations, like architecture, storage density (TB/U), and energy consumption, but they also include a number of other less obvious considerations:

- Does the AFA use fixed or variable block sizes? In widely varying workloads, the use of variable block sizes can lower the IOPS required to handle any given mix of workloads by as much as 10-15%.
- How is metadata handled? This can affect capacity consumption, addressable storage on the back end, latencies (particularly at scale), and other considerations.
- Do I have different RAID options so that I can select the data protection scheme that best balances my need for low latency, rapid recovery, and capacity overhead with my workload?
- Does the system use redirect-on-write or copy-on-write snapshots? If snapshots will often be written to (like is common in evolving copy data management use cases), redirect-on-write snapshots are more efficient and can be higher performing.
- Are data services selectable at a low level of granularity, or are they always on?
- For block-based arrays, has the vendor actually shipped Virtual Volumes (VVOLs) integration, which improves the efficiency of space and bandwidth utilization for operations involving snapshots and replication, or does the vendor just support the VVOLs API? VVOLs integration can also have significant ease-of-management benefits, depending on exactly how the vendor did the integration.
- For vendors that support cloud-based predictive analytics, exactly how are they using the telemetric data they collect to drive value for their customers?
- Does the vendor natively support replication to HFA platforms (which can be less expensive remote location targets for disaster recovery purposes), or does the vendor only support replication to another AFA of the same type?
- What are the total cost of ownership (TCO) implications of the vendor's technology refresh model?

This is not an exhaustive list, but it does point out the types of less obvious distinctions IDC noted between the AFAs evaluated in this study.

One of the key criteria in selecting AFA vendors for this study was whether or not they were pursuing a "flash first" sales strategy and when that started. From IDC's point of view, a vendor makes the shift to a "flash first" sales strategy when it markets either a type 1 or a type 2 AFA, leads with AFA offerings for primary storage opportunities, and has trained both its own direct sales force and channel providers on that sales strategy. For each of the established enterprise storage providers covered in this document that also sell HFAs (Dell EMC, Hitachi Vantara, HPE, IBM, and NetApp), IDC noted a distinct upward surge in AFA revenue growth for that vendor in the wake of its commitment to a "flash first" sales strategy. At this point, all the majors have made that leap, and several of the smaller AFA

vendors that have evolved over time out of HFA backgrounds (Tegile and Tintri) have done that as well. The "pure play" AFA vendors (Kaminario and Pure Storage) have obviously been following "flash first" sales strategies since their inception. For prospective buyers, it may be of interest to take note when a vendor initially committed to a "flash first" sales strategy as part of the vendor evaluation process.

There is no "best" AFA for all needs. Factors that positively contributed to an AFA vendor's position on the IDC MarketScape figure included the vendor's range of scalability, the extent of data services support, platform design efficiencies, ability to integrate into preexisting datacenter workflows, overall platform maturity, and the vendor's strategies around cloud-based predictive analytics, CX management, and technology refresh. Once a customer has created a short list based on higher-level considerations like architecture (dual controller or web scale), block and/or file, synchronous replication/stretch clusters or not, performance and storage density sweet spots, and incumbent vendor or new, any of the remaining AFAs will likely be a good fit. As an example, Huawei offers a very scalable, very high-performance AFA with a very comprehensive set of enterprise-class data services, but the platform is relatively new and uses more traditional approaches to telemetrics, CX management, and technology refresh. On a pure feature comparison, that system would rate very high, but platform maturity and other considerations could limit its appeal with some customers.

## A Note on Failure Domains

Although we are looking at overall system scalability, some customers are specifically interested in limiting the size of their failure domains and are willing to take on a little bit more management complexity to achieve their goals in this area. Note that while this analysis does comment on AFA scalability, many vendors offer multiple models to let customers determine what size of platform best meets their needs. For example, for NetApp, while much of the analysis is done around the All Flash FAS (AFF) A700 platform, NetApp fields smaller platforms that use that same architecture but offer lower entry price points and smaller failure domains (the A200, the A300, and the A700s). This flexibility is taken into account as part of a vendor's portfolio strategy.

Storage density (TB/U) is another consideration here. The highest storage density, while it may result in the most efficient energy and floor space consumption, also has potentially the biggest impact with the failure of an entire disk shelf. Vendors are clearly consciously aware of this, and some even claim to have limited their storage densities because of working with customers to define a "sweet spot" storage density. Many vendors offer multiple SSD or CFM device sizes, letting the customer mix and/or match to meet its own density requirements. This study does take storage density into account, but the most flexible systems in it will offer customers multiple device size options to hit their own storage density requirements rather than just offering the highest storage density period. For those systems on the lower end of storage density, however, this detracts from the value the system provides. Clearly, one of the benefits that flash media-based devices offer relative to HDDs is an ability to achieve much higher storage densities with relatively lower energy and floor space consumption.

## 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 the vendor's strengths and challenges. This IDC MarketScape for AFA should be considered on its own rather than in comparison with the previous IDC MarketScape for AFA, which was published in December 2015. Relatively few AFAs were

available in the market for the 2015 IDC MarketScape, but the market has matured considerably, and it was necessary to focus this document on a specific use case for AFAs (dense mixed enterprise workload consolidation) to keep the number of vendors manageable. Table 1 shows the specific vendor AFA platforms around which this evaluation was based.

## Pure Storage

Pure Storage is positioned in the Leaders category in this IDC MarketScape for AFA.

Pure Storage was among the first pure-play AFA vendors to ship a product back in 2H11. The company put an early focus on inline data reduction technology to drive down the cost per gigabyte of flash and has an extensive set of data reduction technologies that operate both in line and post-process. Its first product, the FlashArray, was initially targeted for high-performance databases and VDI environments, but over the years, it has evolved to have a full set of enterprise-class data services. Its systems are widely used in the enterprise and in cloud service and SaaS providers (where the company expects to generate roughly 30% of revenue for its latest fiscal year ending January 31, 2018). Initially, its storage platforms were built around commodity hardware, but over time, Pure Storage has moved to selectively integrate custom hardware where it felt that could provide the company with a performance, reliability, storage density, and/or cost advantage. The FlashArray is a block storage platform, but in 2016, Pure Storage introduced the FlashBlade, a scale-out system using CFMs (which the vendor calls DirectFlash Modules) that was targeted for use in unstructured and big data analytics environments. In April 2017, Pure Storage announced and shipped the next-generation NVMe-based FlashArray, the FlashArray//X.

With only two AFA platforms in its portfolio, Pure Storage can't really be said to be pursuing a portfolio strategy like competitors Dell EMC, IBM, and NetApp, but it can clearly cover block, file, and object storage needs with its arrays. Its FlashBlade product, while not evaluated for this study because it is targeted primarily at big data analytics, high-performance computing, and some secondary storage uses (e.g., backup), represents the industry's first truly flash-optimized big data analytics platform. The sole focus of Pure Storage on AFAs has been key in its ability to consistently grow its revenue over time. The company went public in late 2015.

The FlashArray//M family, Pure Storage's 12Gb SAS-based block storage platform, includes a range of options with different performance, capacity, and price points (M10, M20, M50, and M70). The FlashArray//X70 uses NVMe devices (DirectFlash CFMs), NVMe backplanes, and NVMe-based controllers but still attaches to hosts using either FC or iSCSI. A nondisruptive upgrade for the platform to NVMe over Fabric will be available in 2018. The use of NVMe significantly extends the vendor's high-end performance and scalability, enabling the array to deliver up to twice the IOPS and bandwidth of the M70 with lower latencies. It packs 333TB of raw storage capacity into 3U and uses only 3.97W/TB (i.e., NVMe, not SAS). 16/32Gb FC and 10/40GbE are supported for host connections. Pure Storage will continue to sell the SAS-based //M arrays while offering the //X70, letting customers choose which performance and cost profile are best for their individual environments.

There is a new version of Purity, the FlashArray storage operating environment, that includes a number of innovations to fully exploit NVMe but also draws on the functional strengths of that relatively mature software platform. With this NVMe-optimized design, the FlashArray//X handles all garbage collection, wear leveling, block allocation, and translation globally (in software) rather than at the

device level. The high efficiency of this implementation should be particularly important in maintaining predictively low latencies in highly consolidated environments.

## Strengths

Pure Storage has exhibited strong industry leadership in making inline data reduction technologies a key requirement for primary storage platforms; raising the bar on technology refresh requirements with a much more customer-friendly, lower-cost approach; supporting the move toward all enterprise-class data services being bundled with the array at initial purchase; leveraging cloud-based predictive analytics to drive value for its customers; and redefining how enterprises should be evaluating their customer satisfaction. All viable AFA competitors now really have to offer a full array of storage efficiency technologies (inline compression, deduplication, thin provisioning, pattern recognition, write minimization, and space-efficient snapshots and replication), and most also offer a data reduction guarantee that is workload independent (some vendor programs are workload dependent, whereas some are not). The Evergreen Storage program of Pure Storage is quickly making forklift upgrade requirements for technology refresh a reason not to buy competitive products, and most of its key competitors have tried to implement similar programs that offer trade-in credits for older technology, provide controller refreshes at designated points in the product life cycle, and guarantee stable maintenance costs for the life of the platform. For more information on Pure Storage's Evergreen Storage program, see *Evergreen Storage Is Changing Customer Experience Expectations in Enterprise Storage* (IDC #US41866916, October 2016).

The approach of Pure Storage to CX has clearly turned heads and has prompted some of its competitors to turn to similar programs that allow customers to compare (NPS) scores on an apples-to-apples basis. Ten years ago, no enterprise storage vendors mentioned NPS publicly, let alone published any numbers concerning their level of customer satisfaction. Driven by pioneers in this space like Pure Storage, Nutanix, and Nimble Storage, this is starting to change.

Outside of strengths Pure Storage exhibits around how its business is managed, it also offers strong product features. It offers several FlashArray platforms to provide appropriate capacity and price points for a wide range of customers. The company offers four platforms (an M10, M20, M50, and M70), which allow customers to start at under \$100,000 and span all the way to systems that can support 512TB of raw capacity in 7U. The X70, the new high-end platform, increases the infrastructure density of the company's highest end array. The average data reduction ratio, which is continuously measured and updated across Pure Storage's entire installed base through Pure1, the cloud-based predictive analytics platform, is above 5:1, so it's clear that Pure Storage's systems can support multiple petabytes of effective storage capacity. A customer can nondisruptively upgrade from the lowest end M10 platform to the highest end M70 or even the X70 platform. To help drive consistent performance across this entire range, Pure Storage was one of the first AFA vendors to use NVMe-based cache cards (which by the way were hot pluggable in its platform back in 2015). Its dual-controller architecture systems use fully redundant, hot-pluggable hardware components, and the measured availability of its entire installed base (again, from Pure1 metrics) surpasses six-nines in production use.

The FlashArray offers a full complement of APIs supporting integration into VMware and Hyper-V environments, management and monitoring tools, enterprise data protection products, and open source technologies like OpenStack, Docker, and Kubernetes and offers fully scriptable command line interfaces and REST APIs. The IP of Pure Storage is available through three different consumption models (appliances, converged infrastructure, and services based), and its pricing is all inclusive. One notable recent addition of note around bundled data services is that in 2H17, Pure Storage announced

portable snapshots that could be taken on a FlashArray but moved to certain other heterogeneous storage (as well as the cloud) and introduced support for synchronous replication and stretch clusters (included as part of the initial array purchase). While the company did not get credit for these features in this study for the purposes of the IDC MarketScape figure, this is worth mentioning since it is the only enterprise storage vendor that includes support for stretch clusters at no additional charge.

## **Challenges**

For use in mixed enterprise workloads that include at least some mission-critical workloads, a key functional area that was missing from Pure Storage's FlashArray was support for synchronous replication and derivative configurations like stretch clusters. Pure Storage did announce this capability later in the year but was marked down for not having this feature as of June 30, 2017. Given its recent release, the maturity of its technology in this space is below that of some of its main competitors. Pure Storage's reputation for extremely high customer satisfaction suggests that any issues would be rapidly resolved, but this is a new area for the company.

Pure Storage uses a dual-controller architecture. While additional controllers cannot be added to any of its platforms (the M10, the M20, etc.), customers can nondisruptively upgrade to higher-performance controllers while fully preserving hardware investments (through the use of trade-in credits on nondisruptive controller swaps). In fact, the multiple facets of its Evergreen Storage program go a long way toward addressing some of the classic complaints against dual-controller architecture systems like forklift upgrades, limited scalability, and loss of hardware investment when upgrading to next-generation platforms. Still, however, if customers specifically want maximum flexibility in creating the right balance between storage performance and capacity, AFAs based around web-scale architectures may be a better fit.

The FlashArray//X70 supports 333TB of usable storage capacity in 5U (67TB/U), and with data reduction, it easily supports multiple petabytes of effective capacity. The platform does offer relatively high storage density, although not the highest, but clearly supports sufficient performance and capacity for the most demanding mixed enterprise workloads. Still, storage density does have an impact on the cost per gigabyte of the system as well as energy and floor space consumption and must be evaluated in conjunction with actual data reduction ratios to determine the system's performance on these metrics. Data reduction ratios will, of course, vary based on workload types, a fact that is true for all vendors. Although the vendor does not release IOPS numbers for its systems, the //X70 delivers in the range of 1 million to 2 million IOPS (for those customers wanting to understand how to gauge the size of its failure domain).

Finally, Pure Storage has not, to date, been profitable. In its most recent quarterly financials call, it stated that it is on target to hit profitability with its fiscal 4Q18 results (which will likely be announced in February 2018). As a public company still growing revenue faster than the AFA market overall, Pure Storage has ready access to capital should it need it, but the company's ability to achieve and maintain profitability moving forward is at this time still unproven.

## **Consider Pure Storage When**

Functionally, the FlashArray//X70 clearly meets the requirements to handle mixed enterprise workloads that include at least some mission-critical workloads, including those that require stretch clusters. Customers comfortable with vendors' use of some proprietary hardware (i.e., DirectFlash Modules) will find much to like with the FlashArray. Note that the initial FlashArray purchase now includes the stretch cluster software bundled with the price of the system, the only vendor in this study

to do that. Like many of the other type 1 AFAs, the FlashArray will be easier to manage in most scenarios than the type 2 AFAs. But for customers that have managed enterprise storage products for years, the aspects of Pure Storage's offering that may be of most interest are the excellent cloud-based predictive analytics that drive a six-nines availability across its entire installed base and simplifies capacity planning, its consistently excellent ability to deliver excellent customer satisfaction, and the cost and ease-of-use advantages of its Evergreen Storage program.

## 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 tailor 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 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

Under IDC's AFA taxonomy, any external storage array that supports only all-flash media as persistent storage and is available under a unique SKU will be considered an AFA. System pedigrees are important, however, in helping customers understand differentiating functionality between platform types. Under the current taxonomy, there are three types of AFAs:

- **Type 1.** These are arrays that were originally "born" as AFAs, and they include products such as the Huawei OceanStor Dorado V3, the IBM FlashSystem A9000, the Kaminario K2, and the Pure Storage FlashArray//X.

- **Type 2.** These are arrays that originally began life as hybrid designs but have undergone significant flash optimization, do not support HDDs, and include at least some unique, high-performance hardware (typically controllers that are faster than those included in the vendor's HFAs) unique to the all-flash configuration. Examples are the Hitachi Vantara Virtual Storage Platform (VSP) F Series, the HPE 3PAR StoreServ 9450 and 20850, and the Tintri EC6000 Series.
- **Type 3.** These are arrays that originally began life as hybrid designs but have undergone significant flash optimization, do not support HDDs, and do not include hardware (other than the flash media) that is any different from the hardware the vendor ships on its HFA products. Examples are the Dell EMC VMAX All Flash and the NetApp All Flash FAS (AFF).

Note that some vendors have arrays that they will ship in either all-flash or mixed configurations. If an array is allowed by the vendor to support HDDs (regardless of whether it is shipped from the factory without them), then it is considered to be an HFA. Examples of this type of system are the Dell EMC SC4020 and the HPE 3PAR StoreServ 8440 and 20840. This study does not evaluate any of these types of systems.

Flash media options for AFAs include CFMs and SSDs, and some vendors offer customers the option to configure either in the same system (like Hitachi Vantara in the VSP G Series). Some vendors use CFMs for one platform and SSDs for another (e.g., IBM's CFM-based FlashSystem A9000 and the SSD-based DS8888), aiming these products at different target markets. CFM-based options use flash media that has not been packaged in an HDD form factor but is instead just accessed as a media store of a given capacity. These systems can use SCSI or NVMe (PCIe) as the storage interface, although versions based on NVMe (PCIe) offer significantly differentiated performance. Vendor products evaluated for this study that use CFMs include the IBM FlashSystem A9000 and the Hitachi Vantara VSP F Series, while those that use SSDs include the Dell EMC VMAX All Flash, the Kaminario K2, the NetApp AFF, the Pure Storage FlashArray//X, the Tegile IntelliFlash HD, and the Tintri EC6000 Series. Vendor products using SSDs but with modified firmware include the Huawei OceanStor Dorado and the HPE 3PAR StoreServ platforms. For more information on the differences between CFM- and SSD-based options and the pros and cons associated with each, see *Flash Media Packaging Decisions Secondary to System-Level Considerations in the All-Flash Array Market* (IDC #US41735716, September 2016).

## LEARN MORE

---

### Related Research

- *INFINIDAT Bucking External Storage Market Trends with Continued High Revenue Growth and Profitability* (IDC #US43193317, November 2017)
- *IDC FutureScape: Worldwide Enterprise Infrastructure 2018 Predictions* (IDC #US43137417, October 2017)
- *Rack-Scale Flash Vendor Excelerio Leverages Performance and Efficiency Advantages to Drive Compelling Customer Value* (IDC #US43168917, October 2017)
- *Western Digital Strongly Furthers Its Play in the Enterprise Space with the Acquisition of Tegile* (IDC #US43028617, September 2017)
- *An Enterprise Storage Consumption Model-Based Strategy Maximizes Vendor Return on Investment* (IDC #US42762617, June 2017)

- *A Framework for Evaluating Storage Efficiency Technologies in Enterprise-Class All-Flash Arrays* (IDC #US42464717, April 2017)
- *HPE Hits a Home Run with the Acquisition of Midmarket Enterprise Storage Vendor Nimble Storage* (IDC #US42389717, March 2017)
- *Worldwide All-Flash Array Market Shares, 3Q16: Top 5 Market Share Holders Consolidate Gains* (IDC #US42374817, March 2017)

## Synopsis

This IDC study provides an evaluation of 10 vendors that sell all-flash arrays (AFAs) for dense mixed enterprise workload consolidation that includes at least some mission-critical applications.

"All-flash arrays are dominating primary storage spend in the enterprise, driving over 80% of that revenue in 2017," said Eric Burgener, research director, Storage. "Today's leading AFAs offer all the performance, capacity scalability, enterprise-class functionality, and datacenter integration capabilities needed to support dense mixed enterprise workload consolidation. More and more IT shops are recognizing this and committing to 'all flash for primary storage' strategies."

## 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 and Trademark 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](http://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](http://www.idc.com/offices). Please contact the IDC Hotline at 800.343.4952, ext. 7988 (or +1.508.988.7988) or [sales@idc.com](mailto: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. IDC and IDC MarketScape are trademarks of International Data Group, Inc.

Copyright 2017 IDC. Reproduction is forbidden unless authorized. All rights reserved.

