White Paper

Architectural Design Decisions Directly Support a Better Customer Experience for Pure Storage FlashArray Users

Sponsored by: Pure Storage

Eric Burgener
September 2020

IDC OPINION

In all-flash arrays (AFAs), architectural design choices can have a major impact on the quality of the enterprise storage life-cycle experience. Pure Storage, a leading vendor in terms of enterprise-class AFA revenue, changed the industry forever with the introduction of its Evergreen Storage program back in 2015. The program gives customers a number of guarantees and investment protection offerings that enterprise storage providers had not provided to their customers before. These elements range from guaranteed data reduction ratios and overall customer satisfaction to nondisruptive upgrades and multigenerational technology refresh, capacity consolidation, included hardware upgrades, and on-demand trade-in credits. Evergreen Storage has prompted responses from Pure Storage's competitors, which are now offering programs of their own. In Pure Storage's case, many of the guarantees and offerings are not just marketing programs, they are directly supported by specific design choices and features in the company's FlashArray system architecture.

This white paper takes a closer look at FlashArray's architectural support for the Evergreen Storage program, delving into four very specific areas: software-driven architecture, stateless controllers, modularity of components, and simplicity of design. Real customer benefits — whether in terms of better performance, higher availability, easier deployment and management, improved efficiencies, nondisruptive multigenerational technology refresh, and/or lower costs — result from these architectural choices. These benefits are also a key factor in Pure Storage's industry-leading Net Promoter Score, an objective third-party metric that measures the quality of the overall customer experience (CX).

The typical legacy enterprise storage life cycle for a SAN array is 4-5 years, after which a need to move to newer technologies commonly drives an often painful "forklift" upgrade. Pure Storage claims to have doubled the enterprise storage life cycle for its FlashArray to 8-10 years, simplifying the enterprise storage life-cycle experience and significantly lowering the total cost of ownership. The flexibility to accommodate newer technologies nondisruptively is a key determinant of array longevity, and Pure Storage has a strong case to make for its FlashArray that is directly supported by architectural design decisions. As enterprises consider new storage platforms, they should seek a modernized version of the enterprise storage life cycle and expect vendors to validate their ability to deliver that.

IN THIS WHITE PAPER

Architectural decisions can make a huge difference in providing the performance, availability, efficiency, upgradability, and simplicity that effectively differentiate a storage system from those of competitors. In the case of Pure Storage's FlashArray, a number of architectural decisions were made that have allowed the storage platform to deliver a customer experience that is clearly different from those offered by legacy enterprise storage platforms. This white paper summarizes those architectural decisions, discussing implementation as well as the resulting impacts on CX. In addition, the paper explains the reasons for the architectural choice and summarizes the relevant impact on performance, availability, efficiency, upgradability, and/or cost. This discussion is specific to Pure Storage's low-latency block-based FlashArray storage platform. Readers should keep in mind that there are five FlashArray models currently being sold: //X10, //X20, //X50, //X70, and //X90. Each model uses the same architecture but differs in the power (the performance) of its controllers and the amount of raw storage capacity a system can support. All FlashArray systems are end-to-end NVMe, with optional NVMe over Fabrics host connectivity or SCSI-based host connections (using Fibre Channel or Ethernet).

SITUATION OVERVIEW

Information technology (IT) organizations that manage their own enterprise storage systems are very familiar with the legacy enterprise storage experience and life cycle. Over the years, storage managers have almost become inured to working with less flexible hardware-defined architectures, paying separately for various data services (snapshot integration, replication options, etc.) and managing systems through "nerd knobs" that require significant storage expertise. They have also endured maintenance price increases as a system ages, and then roughly every four or five years they are forced to go through a time-consuming and disruptive forklift upgrade process to move to newer technologies. Although not all of these aspects of enterprise storage ownership flow directly from system architecture decisions, many of them do. In a word, architecture matters.

When the first enterprise-class AFAs began to ship in the 2011-2012 time frame, the use of solid state storage as a persistent storage medium set these systems apart from their legacy brethren that were still architected for (and used) hard disk drives (HDDs). Some storage vendors just inserted NAND flash-based solid state disks (SSDs) into systems that had originally been designed around HDDs. Others made more sweeping architectural changes that made better, more efficient use of the new solid state storage medium. As one of the AFA start-ups that architected its systems from a "blank sheet of paper," Pure Storage made a number of design decisions intended to improve overall performance, availability, efficiency (of resource utilization), and simplicity that were directly linked to the decision to use only persistent flash media. One of the business objectives of Pure Storage's design intent was to at least double the useful life of an enterprise storage system, further enhancing the cost benefits associated with its more "flash optimized" design.

In considering Pure Storage's architectural decisions that have a direct impact on extending the storage life cycle, IDC notes that the vendor created and marketed an entire ownership program, called Evergreen Storage, to underline the advantages its design offers across a number of areas. The name "Evergreen" was chosen because the offering did at least double the useful life of a storage system while still allowing that system to nondisruptively accommodate newer storage technologies as they became available – in other words, the system stayed "evergreen" throughout its entire life cycle. Many of Pure Storage's competitors have characterized Evergreen Storage as "just a program," but a deeper look at the vendor's system architecture disputes that claim.

Software-Driven Architecture

Software-driven infrastructure has become a popular topic among vendors, particularly software-defined storage (SDS) vendors, in recent years. Compared with older, more hardware-defined designs, software-defined designs put a storage system's functionality and features (beyond just persistence) into software. With software-defined functionality, system capabilities in the areas of performance, availability, capacity, and storage management can be upgraded via software. This functionality makes it not only simpler but also less expensive to add new features to an existing system in the field because it does not require the purchase of any new hardware. This does not mean that system enhancements can never be tied to new hardware – for example, buying additional storage devices adds capacity – but software-defined functionality provides a much greater flexibility to upgrade system capabilities without requiring additional hardware. And even in the case of additional capacity, software upgrades that improve data reduction ratios can add effective capacity to a system without requiring any hardware purchases.

Limited support for different storage device geometries can be a problem with some legacy systems, precluding them from supporting newer storage technologies that may deliver faster performance, increased density, or other attractive capabilities. When storage device support is in some way limited by the storage controllers, controller upgrades may be required to support newer device geometries. Most systems do offer a "midlife kicker" that includes a "new" controller at some point in a system's life cycle that may introduce incremental improvements, but these are often not able to accommodate significant advancements like moving from SAS to NVMe or accommodating larger-capacity devices and/or new media types. In some systems, the only way to obtain a controller that can support desired new technologies is to perform a forklift system upgrade. When the software is the critical factor in storage device support (rather than the controller), there is significantly more flexibility to accommodate newer technologies.

Software-defined designs also enable more flexibility in handling metadata. If the metadata is extensible, a system can accommodate on-disk format changes that improve the performance, endurance, and/or efficiency of operation gains without requiring any data migration. If the metadata is static and tied to the media format, then there is little opportunity to change it to improve system capabilities without migrating data. More seasoned storage administrators may remember cases where they were forced to migrate all of the data on an array to obtain needed new "software" features because the on-disk format was changed with the new software release. Vendors are well aware of the pain this causes, but in some cases, the need to extend the capabilities (usually the capacity scalability) of the system is deemed to outweigh the pain.

The ability to take advantage of new technologies that offer significantly improved capabilities through just a software upgrade is one aspect of extending the storage system life cycle, and Pure Storage has an excellent track record here. Since 2012, new software releases of Purity FA, the FlashArray storage operating system, have improved performance, increased scalability through (among other things) support for new device geometries, accommodated new solid state media types (MLC to TLC to QLC), increased data reduction ratios, and enabled key new features, like cloud-based predictive analytics and support for native file systems, replication, and stretch clusters. The FlashArray's metadata is extensible, and over the years, the company has continuously evolved it to make its systems more performant, scalable, and efficient. All of these changes have been implemented in a manner that allows the installed base of systems to upgrade nondisruptively, helping keep the systems "evergreen."

The FlashArray's ability to simultaneously support multiple device geometries offers what Pure Storage refers to as "Capacity Consolidation." For an older FlashArray, a customer may at some point want to move from lower-capacity to higher-capacity (i.e., more dense) storage devices. This is an evolution that in competitive SAN systems often cannot be done and, even if it can be, often requires manual data migration. On the FlashArray, a customer can insert a new disk shelf (assuming that space is available) of the newer storage devices, and the system will, over time, migrate data from the older to the newer devices as a background process so as not to impact application performance. Once the data movement is complete, the administrator is notified, and the older shelf may be removed. In this way, a system that takes up the same amount of rack space can be expanded over time not only to higher-performance controllers (using the Evergreen Gold subscription's Upgrade Flex and Free Every Three features) but also to increased storage density (using the Capacity Consolidation feature).

Stateless Controllers

Many legacy storage arrays use a stateful controller design, an architecture that complicates controller replacements and upgrades. For example, if the memory that holds controller state information is held in the controller itself, removing a controller severely impacts array performance or even the ability to remain online. In the FlashArray, controller state information is kept in separate cache cards (which are NVMe based and can use either NAND flash or 3D XPoint media, depending on customer performance requirements and FlashArray model). When a controller is replaced, data on the NVMe-based cache cards is very rapidly reloaded onto the controller and the array is moved back into normal operation mode. This stateless controller design is a key element in enabling FlashArray to be upgraded nondisruptively via controller swap out. The fact that state information is kept separate from the controller also provides better data security than stateful controller designs.

FlashArray also uses a modified "active/passive" controller architecture where, during normal operation, both controllers accept I/O from the initiator side. The two controllers coordinate their operations such that data is moved to and from storage devices in a system through only a single controller. If a controller fails, the remaining controller continues to service all I/O with no change in performance. Storage administrators familiar with "active/active" controller designs may observe that having both controllers active during normal operation drives higher performance, and this may be true, depending on the power of the respective controllers being compared. Readers should note that in building each class of storage controller, Pure Storage designed each single controller to be able to handle the same amount of I/O as its competitors' dual controller designs (when operating normally). Another way to state this is that a single FlashArray controller is roughly twice as powerful as Pure Storage's competitors' offerings on a per-controller basis. The system was designed this way so that Pure Storage could provide the benefits of the "active/passive" design (a failed controller causes no performance impact) without the downside (a system runs only half as fast as during normal operation).

FlashArray controllers can be upgraded at any time during the life of an array to move to more powerful, latest-generation controllers within the same class or to move to a higher, more powerful class of system, also of the latest generation. (The "upgrade at any time" option under Evergreen Gold that includes full trade-in credit for old controllers is called Upgrade Flex. Evergreen Gold also includes a feature for no-cost controller upgrades to the latest generation of the same class of controller, called Free Every Three.) For more information and IDC commentary on the Evergreen Program, see *Pure Storage's Evergreen Storage Continues to Be the All-Flash Array Market's Customer Experience Program to Beat* (IDC #US45761920, January 2020).

Pure Storage's approach to controller design in FlashArray has several positive implications relative to "active/active" designs. First, upgrading a system to new controllers does not impact application services. The FlashArray architecture is such that a low-end //X20 system can be nondisruptively field upgraded to an //X50, an //X70, and even an //X90 without ever shutting the system down or impacting application performance. There are scenarios in which customers require more performance, not to mention higher capacity potential, from their arrays as workloads increase and requirements change over the life of an array. If that need arises, customers have the option to nondisruptively upgrade controllers to get to the next performance level (i.e., move from an //X20 to an //X50, from an //X50 to an //X70, etc.) without a forklift upgrade or a data migration. Readers should note that this upgrade capability also leverages the stateless controller design and modular components (discussed in the section that follows), both of which are featured in FlashArray. Second, this capability has a positive impact on the reliability of the controller pair. During normal operation, only one of the controllers is communicating with the back-end storage devices, not both. Because systems are operating normally instead of in failure mode most of the time, there is less overall stress on the controller pair (which combined run at a lower level of utilization) than with competitive "active/active" controllers in the same class (which are both handling initiator and target traffic during normal operation). Running components at lower levels of utilization has a positive effect on reliability.

Modularity of Components

FlashArray employs a very modular design in which major components that affect system performance, availability, and scalability can be hot-plug replaced quickly and easily, without tools. Controllers, cache cards, storage devices, disk shelves, power supplies, and fans are redundant where meaningful (the disk shelves aren't necessarily redundant) and provide the hardware foundation for the system's "six nines" availability track record. While the FlashArray backplane cannot be field replaced, starting in 2016, it could support both SAS and NVMe storage devices. After Pure Storage shipped its first NVMe AFAs (NAFAs) in 2017, existing customers began upgrading their installed base systems from SAS to NVMe, all completely nondisruptively. Although their competitors began shipping their first NAFAs in 2018 and through 2019, none of them ever allowed their installed base of SAS systems to be nondisruptively upgraded in place to NVMe – they all required a forklift upgrade to get to NVMe. FlashArray's nondisruptive SAS-to-NVMe upgrade path was unique in the industry, and no other vendor has offered this level of "evergreen" upgradability for in-place scale-up systems.

Systems that are more modular in design offer more opportunities to upgrade different storage components to newer technologies over the life of an array. Take the vendor's cache cards as an example. Like many enterprise storage vendors, Pure Storage uses solid state storage for its cache, sending all writes there first to take advantage of the extremely fast write acknowledgements that the medium can provide. Competitive caches in an array are often on the storage controllers, but Pure Storage's use of separate cards allowed the company to upgrade cache cards without upgrading controllers and work with higher-capacity caches that could handle all writes and also a good number of reads. In early systems, when the rest of the array was plumbed with SAS, the cache cards used NVMe to provide lower latencies. As the vendor shipped more NAFAs with larger capacities, Pure Storage introduced an optional storage-class memory-based cache card that lowered cache latencies yet again by an order of magnitude (available specifically on the //X70 and //X90 models). This new cache card, called the DirectMemory Cache, became available in September 2019, and FlashArray systems can support up to three of them. And in the spirit of its modular upgradeable design, previous-generation //X70 and //X90 systems can be retrofitted with DirectMemory Cache.

The name of the DirectMemory Cache plays off the name of Pure Storage's custom-built storage devices, the DirectFlash Modules (DFMs). In 2017, Pure Storage moved away from off-the-shelf SSDs, sourced from major storage component suppliers for the FlashArray, to the new custom DFMs. All along, the FlashArray's storage operating system had used some proprietary algorithms to manage the flash media directly for improved performance, higher endurance, and better capacity utilization. But the company could not always get the access it wanted to the underlying media when working with off-the-shelf SSDs. The SSD controller and flash translation layer provided by the storage component suppliers were often "in the way." Moving to the custom DFMs gave the company an unfettered ability to manage the underlying flash media directly (hence the "Direct" name), driving performance and efficiency improvements, while the custom form factor allowed Pure Storage to introduce higher-capacity devices for increased storage density than was generally available with off-the-shelf SSDs. Today, the largest enterprise-class SSD in common use is around 15TB, while Pure Storage sells an 18TB DFM for the FlashArray//X (as well as DFMs as low as 1TB in capacity). Both the DirectMemory Cache and the DirectFlash Modules are hot-pluggable.

The ability to manage the underlying solid state media directly in a media-agnostic manner is important for several reasons. First, it allows Pure Storage to implement I/O algorithms that are best suited for efficient use of the underlying media given the I/O patterns in its target markets and associated use cases. Pure Storage is not forced to work through SSD-resident controllers and flash translation layers that tend to be more attuned to the high-volume consumer rather than enterprise market requirements. This has implications not only for latencies but also for how garbage collection (i.e., free space allocation) and media endurance are managed – issues that are particularly important as Pure Storage moves to take advantage of other types of NAND flash-based media (TLC and QLC) that have lower performance and endurance. Second, by optimizing media endurance, the company can lower media overprovisioning ratios to cut costs out of the storage devices without sacrificing device-level longevity (the vendor overprovisions media at a lower rate than the established storage component suppliers but provides a lifetime-of-the-system guarantee on all solid state devices it sells). Third, the ability to manage underlying solid state media allows Pure Storage to accommodate new media types faster because it builds them into the DFMs and qualifies them for production use faster than competitors that must wait for storage component vendors to build and release a device before they can begin their own system-level qualifications.

Simplicity of Design

The FlashArray leverages a simpler design than the offerings of most of Pure Storage's competitors in the large enterprise systems market. First, the company has stayed with a dual controller design, while many of its competitors moved to a more complex design that allowed them to scale controllers beyond just two. Through the use of either a crossbar switch or a highly specialized ASIC, competitive systems could scale up to four and sometimes as many as eight controller pairs to extend the performance of their systems in place. In Pure Storage's case, customers needing more performance could just quickly and nondisruptively upgrade to a higher class of storage controllers. The crossbar switch and/or ASIC designs added significant complexity and even reliability issues to the competitive systems, both of which must be managed. Because of cache coherency overhead and other east-west traffic, these multicontroller systems did not scale performance linearly. The FlashArray uses existing system resources more efficiently with the much simpler dual controller design. While it is true that Pure Storage may not be able to build a single system that can drive as many IOPS as some competitive offerings at the very high end, many customers actually prefer to split such large workloads between the fault domains of multiple systems for availability reasons. For those organizations that prefer to consolidate tens of millions of IOPS onto a single system, Pure Storage may not be the best fit.

Traditionally, large enterprise storage systems have included administrative interfaces that offer very deep and granular configuration flexibility. These "nerd knobs" (as they are often referred to by sophisticated storage administrators) were originally needed because there was such a disparity between the performance of compute and storage (i.e., HDDs) resources that customers needed to continuously fiddle with settings to optimize performance as workloads evolved over time. When all-flash systems were introduced, one of the first things that storage administrators noticed was that, with the predictably consistent sub-millisecond latencies AFAs delivered under load, almost all storage performance problems just went away. Vendors that introduced AFAs by upgrading their existing HDD-based architectures kept these "nerd knobs" because they were already a part of their storage operating systems, but many AFA start-ups took the opportunity to think differently about what they demanded from storage administrators.

Pure Storage took this latter approach with its FlashArray. Instead of making everything configurable, the vendor chose a set of defaults that allowed a system to be unboxed and deployed very rapidly. Most of the enterprise-class data services, which in the past had to be individually configured, were now just "on" all the time. With HDD-based systems, there was a noticeable latency impact with many data services choices, but the performance of flash made that a nonissue at the application level for almost all workloads in AFAs. For example, instead of having to decide which storage device-based data protection method to use from a list, which included RAID 1 (mirroring), RAID 5 (single-parity RAID), RAID 6 (dual-parity RAID), or RAID 10 (mirrored stripes), and actually implement that method, Pure Storage assumes and implements a modified RAID 6 implementation for all data – automatically and without any operator intervention required. Features that were performance sapping in HDD-based and legacy-retrofitted AFA environments like thin provisioning and inline data reduction were instead written into the Purity OS to be "always-on" without impacting application performance in FlashArray.

While this is a great enabler for less experienced storage administrators (which by the way are becoming much more popular as the responsibilities for storage management in many IT shops are migrating more toward IT generalists like virtual, Linux, and Windows administrators), what about seasoned storage administrators who are comfortable with "nerd knobs"? Generally, administrators who may express a preference for them up front do so because, based on experiences with HDD-based arrays or legacy architecture AFAs, they do not trust that default configurations will give them all the performance they need for particular workloads and thus may keep them from meeting certain service-level agreements (SLAs).

Pure Storage's all-flash design effectively resolves this issue. Customer prospects concerned about this issue can talk to their colleagues who are using FlashArrays (there are tens of thousands of systems in the field). They will find that FlashArray systems can deliver sub-millisecond latencies at scale (i.e., while pushing millions of IOPS), and if they still have concerns about storage latencies, they can add optional DirectMemory Cache cards, which bring storage latencies down to the 150-microsecond range. Ultimately, even sophisticated storage administrators come to appreciate the simplicity of the FlashArray's default configurations once they see that they do not in most cases impact their ability to meet storage SLAs. And having features like inline data reduction on at all times, which can be guaranteed by Pure Storage and drives an average 5:1 data reduction ratio in mixed workload environments, produces very desirable economics. These are all critical parts of the FlashArray value proposition – that customers will not have to give up performance to get simplicity of management, efficient use of storage resources, or a compelling total cost of ownership (due in part to high data reduction ratios).

On a side note, when evaluating the total cost of ownership of any system, from Pure Storage or any other vendor, IT managers should focus on the effective capacity (the capacity available for use after formatting, data protection, and inline data reduction have been applied), not the raw capacity. Raw capacity advantages can quickly be overcome by higher data reduction ratios, particularly in larger configurations.

There are other simplicity aspects of the FlashArray CX that should be pointed out. Each independent system is one physical box with six cables that can be installed in 30 minutes. All array software, even advanced features like ActiveCluster (Pure Storage's stretched cluster option), are included with the base price of the array. Features like thin provisioning, inline data reduction, quality of service, and FIPS 140-2-compliant end-to-end encryption are always on. ActiveCluster uses a cloud-based mediator to resolve split-brain syndrome, removing the need to buy and maintain a separate quorum device for these configurations. Snapshots are "portable" in the sense that they are incremental, space efficient, and self-describing and can be shared between other FlashArrays, FlashBlades (the vendor's scale-out unstructured storage platform), third-party NFS storage, or the public cloud. Pure1, an AI/ML-driven cloud-based predictive analytics platform providing full-stack analytics, is included with the base price of each array. And FlashArray includes APIs that enable integration with VMware, OpenShift, and other automation/orchestration tools.

The Impact of FlashArray Life-Cycle Extension

IDC has commented at length about the features and benefits of Pure Storage's Evergreen Storage program under a separate cover. Aspects of that program, like all-inclusive software and guarantees for data reduction ratios and fixed maintenance pricing for the life of the array, clearly simplify the overall experience of owning and managing an enterprise storage array. Other features of that program extend the life cycle of a FlashArray while still enabling customers to take advantage of new technological advancements. The Evergreen Gold program provides an included controller upgrade every three years under the Free Every Three option, and the Upgrade Flex option allows customers to upgrade controllers at any time with full trade-in credit for old controllers. The Capacity Consolidation option lets customers move to denser and more modern storage devices in place to expand system capacity without increasing footprint. Unlike competitive programs, Pure Storage provides trade-in credits for up to the full amount paid for the components being retired toward the purchase of the new components, which means customers are not rebuying controllers and/or capacity at regular intervals — investment in these components is truly preserved.

It is the combination of these features that double the storage life cycle of a FlashArray into the 8- to 10-year range. In competitive systems, forklift upgrades are required when existing technologies run out of gas and can't be upgraded in place — a cycle that happens roughly every 4-5 years. Limitations around backplanes (SAS only), controllers, and supported storage device geometries are all factors that drive this legacy forklift upgrade, but on the FlashArray, these factors do not impede technology refresh. A SAS-based system purchased in 2016 from Pure Storage that used off-the-shelf SSDs could be nondisruptively upgraded to an end-to-end NVMe-based system that uses DFMs and supports newer device geometries (for increased storage density) and different media types (from MLC to TLC and/or 3D XPoint) — all without having to rebuy performance and capacity that are already paid for. Performance and capacity upgrades do generate some incremental costs, but those are always defrayed by the guaranteed trade-in credit amounts that customers know up front when they buy their FlashArray. These are the features that support the vendor's claim that a FlashArray has an 8- to 10-year life cycle relative to conventional SAN arrays, enabling the nondisruptive integration of the latest technologies to ward off the need to move to a completely new system.

The extended life cycle has very positive impacts. Storage administrators know the pain involved in data migration during technology refresh, and Pure Storage's approach means that, over the life cycle of a single array, a relationship with the vendor would cut the need for that at least in half. Customers are not rebuying entirely new storage systems and replacing storage capacity every 4-5 years. IDC would encourage prospects that understand the technology behind the Evergreen Storage program to compare the total cost of ownership and the management impact of this life-cycle extension to that of the traditional 4- to 5-year enterprise storage life cycle.

CHALLENGES/OPPORTUNITIES

For more seasoned IT personnel, it may be easy to dismiss Pure Storage's claims about the value driven by the Evergreen Storage program as "just another program." Because of the early popularity with customers of Evergreen Storage, competitors across the board have introduced their own versions that offer many similar guarantees, but a deeper look is required to realize just how and why what Pure Storage has done is different. That is a challenge that the vendor must overcome in the marketing of its architectural design choices (and the Evergreen Storage program). Existing customers get it, but if they never move beyond the prospect stage, other enterprises may not see the link between the vendor's system architecture decisions and the program specifics and how that drives differentiating value. Like many things in life, the Evergreen Storage program and its architectural tie-ins present both a challenge and an opportunity for Pure Storage as a vendor.

CONCLUSION

Many enterprise storage vendors have introduced "guarantee" programs that are incrementally changing the experience of enterprise SAN ownership for the better. But for most of these vendors, they are, in fact, just programs with no supporting architectural tie-ins in their storage systems. The guarantees that Pure Storage makes in its Evergreen Storage program are enhanced by specific features of its system design — namely, software-driven architecture, stateless controllers, modularity of components, and simplicity of design and operation. Taken together, these capabilities enable simple, nondisruptive technology refresh at reduced cost throughout the life of a Pure Storage FlashArray and can realistically extend the life of an array to 10 years without imposing the performance, availability, scalability, or functional limitations that typically drive the requirement for a forklift upgrade.

Related benefits, like avoiding the need to migrate data when moving to newer storage technologies, save time and reduce cost while inherently underlining the fairness of only having to pay for storage capacity once even as a system evolves to higher storage densities through the integration of newer media technologies. What Pure Storage has done with its approach has raised the bar in customer expectations around the experience of enterprise storage ownership, and it is one of the reasons why the company has consistently enjoyed a CX rating (as shown by its high validated Network Promoter Score) in the top 1% of all businesses across all industries.

The bottom line is that there are specific design features that support Pure Storage's Evergreen Storage claims and drive a differentiating CX around FlashArray ownership. These features all come together to define the FlashArray value proposition — that customers will not have to give up performance to get simplicity of management and growth, efficient usage of storage resources, or a compelling total cost of ownership for their enterprise storage platforms.

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