

MARKET NOTE

Pure Storage Doubles Down on All-Flash Unstructured Storage with FlashBlade//S Announcement at Pure//Accelerate 2022

Eric Burgener
Susan G. Middleton

Matthew Eastwood
Ashish Nadkarni

Phil Goodwin
Johnny Yu

EXECUTIVE SNAPSHOT

FIGURE 1

Executive Snapshot: Pure Storage Enters Second Phase of Its Evolution with FlashBlade//S Launch

The all-flash enterprise storage vendor Pure Storage held its Pure//Accelerate 2022 end user conference on June 8, 2022, in Los Angeles, California. The main thrust of the announcements was around the new FlashBlade//S platform for unstructured data storage (file/object). Executive presentations discussed the boom in unstructured data storage as well as the changing nature of storage requirements in a digital-first world. They also noted the increasing importance of environment, social, and governance (ESG) considerations when it comes to making information technology (IT) infrastructure decisions and how Pure Storage's product and engineering strategies provide unique differentiation in this area.

Key Takeaways

- FlashBlade//S features a new modular, disaggregated architecture that lets customers scale storage compute and capacity independently, enabling a significantly wider range of configuration flexibility. The vendor also discussed the power and cost efficiency of the new design, emphasizing how the company's focus on efficiencies across its entire line solidified the recent Meta win. Pure Storage expects its power and cost efficiencies to increasingly differentiate it from its competitors as ESG considerations become a more important determinant in storage infrastructure selection.
- The design changes on FlashBlade//S make it a peer of the FlashArray//X platform in terms of the vendor's Evergreen Storage, making program benefits consistent across both platforms. Several key changes were also made to Evergreen Storage, including the renaming of the Evergreen Gold option to Evergreen//Forever; the introduction of a new ownership option called Evergreen//Flex to provide pay-as-you-go pricing for on-premises, fully owned storage; and the renaming of the Pure as-a-Service managed infrastructure services to Evergreen//One.
- Pure Storage reviewed its most recent FY 1Q23 financials, which were quite impressive when compared with most of its United States-based enterprise storage competitors. The vendor turned in 50% year-over-year revenue growth, noted that its annual recurring revenue (ARR) hit a \$900 million run rate (which is larger than the ARR of several of its much larger competitors), and adjusted its FY23 outlook upward by over 8% to \$2.66 billion. Pure Storage also returned \$66 million to stockholders through share repurchases during FY 1Q23.

Source: IDC, 2022

IN THIS MARKET NOTE

Pure Storage held its annual user conference Pure//Accelerate on June 8, 2022, at L.A. LIVE in downtown Los Angeles. This Market Note reviews the announcements made at the event.

IDC'S POINT OF VIEW

Pure Storage – founded in 2009 by John "Coz" Colgrove, a serial inventor in the storage arena, and John Hayes, the former CTO of Yahoo! – has clearly transformed enterprise storage over the past decade. In his keynote, Pure Storage CEO Charles Giancarlo positioned the vendor's first decade of shipping product for revenue (which began in 2012) as the displacement of hard disk drives (HDDs) as the dominant media for performance-sensitive, mission-critical workloads. With all-flash array (AFA) revenue becoming the largest segment of the external storage systems market in 2021 and with virtually all of the vendor's competitors emphasizing solid state storage for primary workloads, a strong argument can be made that Pure Storage has achieved its goal. Along the way, the vendor has grown into a \$2.2 billion competitor (FY22 revenue) and garnered the number 3 AFA market share position in terms of revenue by gaining each year on AFA market share leaders Dell EMC and NetApp.

A true innovator cannot rest on their laurels, however, and Giancarlo talked about Pure Storage entering its second phase in which the company intends to make flash the dominant media type for most other workloads, both on premises and in the cloud, again replacing HDDs. Ultimately, the vendor envisions a world dominated by flash for "hot," "warm," and ultimately many "cold" workloads with tape for the "coldest" workloads. IDC does not see solid state media dropping below the price-per-gigabyte cost of raw HDD capacity by the end of the decade, but the argument that Pure Storage makes for the "all-flash datacenter" centers strongly around total cost of ownership (TCO), not the cost of raw capacity, as well as denser flash media types such quad-level cell (QLC), penta-level cell (PLC), and others that will become available later in this decade. With the vendor's recent multimillion-dollar win with Meta and the increasing importance of environmental, social, and governance (ESG) considerations to enterprises and cloud providers the world over, the vendor's efficiency story is resonating with more customers. (Meta's selection of Pure Storage to provide storage infrastructure for the social media giant's artificial intelligence [AI] Research SuperCluster [RSC] was based on the vendor's differentiating power and cost efficiencies relative to the other alternatives.) In FY 1Q23 alone, the company brought more than 300 new logos (not counting additional sales to other groups within existing customers) on board. At this point, Pure Storage counts 54% of the Fortune 500 as its customers.

In his keynote, Giancarlo also reviewed some financial and other data points about Pure Storage's recent performance. The vendor showed stronger-than-anticipated, year-over-year revenue growth of 50% in FY 1Q23 with \$620 million in overall sales. Annual recurring revenue (ARR) was up 29%, hitting a \$900 million annual run rate, which is a figure that puts them above several of their much larger enterprise storage competitors. Noting how their revenue exhibited a very broad-based performance across all products, geographies, and market segments, Giancarlo adjusted the company's FY23 outlook upward to \$2.66 billion (from the \$2.4 billion it was before). Pure Storage is now ranked on the Fortune 1000 list and has a Net Promoter Score (NPS) that puts it in the top 1% of B2B companies.

Major Changes in Pure Storage's Enterprise Storage Platform Portfolio

Pure Storage's product portfolio includes a scale-up architecture (FlashArray) for structured data storage and a scale-out architecture (FlashBlade) for unstructured data storage. While the vendor started out as a single product company, it now has an entire portfolio of all-flash offerings that include FlashArray//X for latency-sensitive, block- and file-based workloads; FlashArray//C for cost-sensitive, block- and file-based workloads; and FlashBlade for bandwidth-sensitive, scale-out, file- and object-based workloads. In December 2021, the vendor introduced FlashArray//XL, a new high-end offering that allows Pure Storage to bring its ease of use, unique technology refresh model, and differentiating customer experience (CX) to larger customers. FlashArray//XL has outperformed Pure Storage's initial expectations, garnering three times the forecast revenue over the past six months and significantly increasing the vendor's ability to win against the high-end platforms from more established legacy storage vendors.

At the event on June 8, the vendor introduced the new FlashBlade//S platform, which it expects to really begin to open up the more capacity- and cost-sensitive unstructured data storage opportunities over the next decade. As data continues to grow at 40-50% per year for most enterprises, 80% of this will be unstructured data, so FlashBlade//S is a particularly timely platform. While FlashBlade//S more than doubled the storage density, performance, and power efficiency of the prior generation FlashBlades, the real news around this release is its new modular, disaggregated architecture. The original FlashBlade design had blades that included both storage processors and storage capacity, and customers could nondisruptively scale a system by popping in additional blades. The new design still uses blades, but storage processors and storage capacity can now be independently and nondisruptively upgraded, providing a configuration flexibility that enables the system to be tailored based on workload requirements and higher-end scalability (in terms of both performance and capacity).

Enabling nondisruptive, multigenerational technology upgrades that support a 10-year life span for systems, all the while ensuring that they can support the latest technologies to keep the systems "evergreen," has been a key design tenet on both the FlashArray and the FlashBlade. The new 5U chassis has been future proofed in terms of backplane bandwidth and cooling so that the same chassis will be able to accommodate new processors and far denser storage capacities over the next decade as new technology emerges. The storage capacity is packaged in DirectFlash Modules (DFMs), which at initial release are available in 24TB and 48TB capacities. Each blade can accommodate up to 4 DFMs. Entry-level systems can start as small as 168TB and scale up to about 2PB raw in a single 5U chassis, which requires about 2,400W of power on average. The system will be able to scale to 10 chassis in a cluster. Along with the FlashBlade//S announcement, the next generation of the AI-ready infrastructure (AIRI) was also announced. AIRI//S has been developed jointly by Pure Storage in partnership with NVIDIA and uses the latest NVIDIA DGX systems, NVIDIA networking, and the new Pure Storage FlashBlade//S to simplify storage deployment and management for AI workloads.

While the focus of Pure//Accelerate 2022 was primarily on FlashBlade//S, there were some new features, such as hardware-assisted compression, that were announced on FlashArray. More notably, however, were the increasing engineering synergies between the two platforms. Pure Storage buys NAND flash media directly from suppliers and builds its own storage devices, a fact that the vendor trumpets as allowing it to deliver better and more consistent performance, higher storage density, lower energy consumption, improved endurance and reliability, and a lower cost of flash capacity relative to off-the-shelf solid state disks (SSDs). FlashArray uses DFMs, and with the new modular and disaggregated design of FlashBlade//S, it now uses the same DFMs. While many of the software capabilities in the two Purity storage operating system distributions (Purity//FA for FlashArray and

Purity//FB for FlashBlade) already leverage a common code base, the vendor continues to converge the two distributions for a more efficient use of software engineering resources. Pure Storage tells us to expect more of this over time.

Expanded Evergreen Portfolio: Evergreen//Flex and Evergreen//One

There's no doubt that enterprise customers love the simple, nondisruptive, multigenerational technology refresh model of Evergreen, but because of architectural differences between FlashArray and FlashBlade, the "evergreen" story has been better with FlashArray. With the release of FlashBlade's new disaggregated architecture, the two platforms are now on par with each other in terms of Evergreen Storage benefits. This brings meaningful benefits to FlashBlade//S owners. While the vendor promised that migration to FlashBlade//S for existing FlashBlade customers will be simple and nondisruptive, that upgrade path is not yet available. With what Pure Storage has already done with migration – from the original FA-200 and FA-400 FlashArray systems to the FlashArray//M, FlashArray//X, and FlashArray//XL systems – the vendor does have a track record of providing a cost-effective, simple, nondisruptive, multigenerational technology refresh, but the details of the migration path to FlashBlade//S are yet to be released.

Pure Storage also extended the scope of the Evergreen brand. Storage systems are still covered under the Evergreen program in terms of guarantees, as are capacity consolidation and options around multigenerational technology upgrades, and that program has been renamed Evergreen//Forever (formerly Evergreen Gold). Evergreen//Flex is a new option that allows customers to still own their on-premises storage infrastructure but to pay for it through a usage-based model. The older Pure as-a-Service offering, which allowed customers to purchase storage infrastructure from the vendor as a managed services offering, has now been renamed Evergreen//One. All of these programs do an outstanding job of delivering the cloud experience for storage infrastructure based on either FlashArray or FlashBlade, regardless of deployment location and consumption model.

Another key theme at the Pure//Accelerate event was sustainability with multiple executives discussing Pure Storage's focus on this topic. Pure Storage emphasized two areas: reducing power usage and offering extended product life cycles with Evergreen. Pure Storage has a strong track record of investment protection through its Evergreen model that enables customers to meet growth needs within the same frame and add capacity without significant disruptions. Pure Storage provided multiple examples of clients that continue to upgrade storage systems that have been in place for over six years. IDC observes that over the past 18 months, the importance of sustainability has reached new levels of awareness, including becoming a requirement for RFPs and an important factor in the vendor selection process for many organizations. Pure Storage's nondisruptive upgrade strategy and in-circle recycling program assures customers that they are working with a like-minded provider. This was underscored with Chief Product Officer Ajay Singh's statement that "Evergreen is an engineering philosophy, not a business model. We have a responsibility to reduce power consumption and improve the world around us."

Portworx Portfolio Expands to Include Additional As-a-Service Options

As storage vendors evolve, they inevitably increase their market value and reach by adding software capabilities. In this vein, Pure Storage acquired Portworx in September 2020. Portworx was positioned as one of the few "pure play" container data management and protection solutions in the market. And while Portworx continues to offer container data management, protection, and recovery, it is clear that Pure Storage has much greater plans. One must understand, however, that this is within the context of the Kubernetes (or cloud-native) ecosystem.

Pure Storage recently announced the general availability of Portworx Data Services for Kubernetes. The Portworx Kubernetes Data Services Platform includes database as a service, search as a service, and streaming as a service today, with the longer-term vision to bring an AI/ML as a service to customers. It is designed to deliver functionality such as single-click deployments in Kubernetes, scheduled backups, and single-pane-of-glass management for all your data services on one platform to provide disaster recovery and reduce cloud storage costs. The AI/ML functionality is targeted at assisting with provisions and hardware configurations, though it will certainly be expanded and enhanced over time.

Portworx Data Services (PDS) became generally available a few weeks before the conference. It supports six data services – Kafka, Cassandra, Redis, Zookeeper, Postgres, and RabbitMQ – with the promise of support for more databases in the near future. PDS targets platform engineers, database administrators, and application developers and essentially lowers the overhead of deploying, operating, and managing multiple different databases. Many organizations use different databases for different use cases, such as Kafka for streaming and Postgres for AI.

The drawback to this heterogeneity is that each database has its own unique methods for handling data protection, disaster recovery, security, life-cycle management, and other functions. This requires organizations to staff people with expertise in these different databases in order to deploy and manage them. In addition, an organization would have to coordinate with each database's vendor for support. PDS serves as a single platform from which multiple databases can be deployed as a service, reducing the need for unique skills for each database and lowering the barrier of entry for containerized database deployment. This is especially important to application developers' CI/CD practice, as it allows them to quickly access data services without any management baggage attached.

Pure Storage's Portworx portfolio also made Portworx-Backup as a Service (PX-BaaS) generally available recently. PX BaaS simplifies the day 0 and day 2 operations related to the data protection of Kubernetes applications in production environments.

LEARN MORE

Related Research

- *Enterprise Storage Vendor Pure Storage Turns Heads with Strong Financial Performance Over Its Fiscal 2022* (IDC #US48948322, March 2022)

Synopsis

This IDC Market Note reviews the announcements made at the Pure Storage annual user conference Pure//Accelerate 2022, which was held June 8, 2022, in Los Angeles, California. In his keynote, Pure Storage CEO Charles Giancarlo talked about how the vendor, having successfully led the evolution of primary storage workloads onto all-flash in datacenters across the world, is entering a new phase where it will be leading the charge to bring all-flash capabilities to secondary storage workloads (most of which are based on unstructured data). The core product announcements were around FlashBlade//S, Pure Storage's scale-out unstructured data storage platform, which features a new modular and disaggregated architecture for increased flexibility, impressive power and cost efficiency, and significantly increased performance and storage density. All in all, it was an impressive kickoff for the second phase of the vendor's evolution.

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

140 Kendrick Street
Building B
Needham, MA 02494
USA
508.872.8200
Twitter: @IDC
blogs.idc.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.

Copyright 2022 IDC. Reproduction is forbidden unless authorized. All rights reserved.

