White Paper

Pure Storage's QLC-Based FlashArray//C Provides a Compelling Alternative to HFAs for Capacity-Optimized Storage Workloads

Sponsored by: Pure Storage

Eric Burgener
May 2021

IDC OPINION

In enterprise storage, what's not to like about flash? Many vendors point to one answer: cost. But when the cost comparison focuses on cost per gigabyte at the storage device level, it can be misleading. Total cost of ownership (TCO) is a much better metric to compare the true cost of a storage system over time. The "cost" argument was used by hard disk drive (HDD)-based array vendors against the earliest shipping all-flash arrays (AFAs), but it was clear by 2016 that for many performance-sensitive primary storage workloads AFA TCO was better, and in many cases much better. And no one was denying the plethora of other benefits associated with AFAs (performance, storage density, ease of use, efficiency, energy and floorspace consumption, lower software licensing costs on servers, better reliability, etc.). Since then, the cost per gigabyte of flash media has continued to drop at a much faster rate than that of HDDs, and with the introduction of quad-level cell (QLC)-based flash media, some AFA vendors are turning their sights on less performance-sensitive capacity-optimized storage workloads. Pure Storage was the first established enterprise storage vendor to introduce an AFA targeted specifically at capacity-oriented storage workloads with its FlashArray//C product.

Vendors of hybrid flash and HDD-based arrays caught flat-footed by the introduction of QLC-based flash media are again using the "cost" argument. And with capacity-optimized workloads, the cost target is admittedly much lower with 7,200rpm large-capacity SATA HDDs than it was with 15,000rpm HDDs; at scale, systems based on these nearline HDDs often deliver a cost per gigabyte in the $0.02–0.04 range. While latency tends not to be an issue for capacity-optimized workloads, many of the other benefits of flash (e.g., throughput, bandwidth, density, efficiency, reliability) are issues. With the FlashArray//C, Pure Storage delivers these benefits in a package that, assuming data reduction ratios in the 4:1 to 5:1 range, boasts a cost per gigabyte in the $0.02–$0.04 range as well.

While there are clearly some colder storage-type workloads for which HDD-based systems deliver a better all-around value proposition, many virtual infrastructure, backup and disaster recovery, multicloud test/development, and hybrid cloud snapshot workloads can benefit significantly from these non-latency-related flash media characteristics and make excellent candidates for placement onto a FlashArray//C. This QLC flash media-based platform does not have to be cost effective for *every* capacity-oriented storage workload before it can be cost effective for many. Enterprises that are currently considering a capacity-optimized storage infrastructure refresh should take a look at FlashArray//C before they make another HDD-based storage system purchase.

May 2021, IDC #US47661320

IN THIS WHITE PAPER

At the Pure//Accelerate 2019 Conference, Pure Storage announced the FlashArray//C, an all-flash platform optimized for more capacity-oriented storage workloads. As cost per gigabyte has continued to narrow between NAND flash–based solid state disks (SSDs) and hard disk drives, IDC has been expecting vendors to begin to pursue capacity-optimized storage markets with these types of platforms. In this white paper, IDC discusses the benefits flash can offer for storage workloads that are today running on hybrid flash arrays (HFAs) and reviews Pure Storage's FlashArray//C product, the first real entry in this space from an established enterprise storage player.

SITUATION OVERVIEW

When Pure Storage announced its first AFAs in 2011, the firm touted its ability to deliver all-flash systems for primary storage workloads at near-HDD economics (with the point of comparison being the 15,000rpm HDDs that were routinely being used for latency-sensitive enterprise workloads at the time). Many enterprise storage systems at that time were performance constrained, causing storage administrators to often spend a lot of time tuning systems on a regular basis. Flash performance promised to resolve these issues, significantly reducing administrative overhead and enabling new applications, but at that time, the media was relatively unproven in enterprise settings and was much more expensive on a raw price-per-gigabyte basis than HDDs.

The fact that AFAs such as Pure Storage's FlashArray were aimed at performance-sensitive primary workloads was actually a critical factor in the total cost of ownership comparisons with HDD-based arrays. The ability to use data reduction, fueled by technologies such as compression, deduplication, thin provisioning, and space-efficient snapshots, in line with latency-sensitive workloads allowed Pure Storage to store four to five times the amount of data (on average in mixed workload environments) in a terabyte of raw capacity as HDDs. HDDs were just too slow to enable inline data reduction for latency-sensitive primary workloads that exhibited random I/O profiles (like most enterprise workloads).

In the subsequent years, flash media reliability and endurance have proven to meet enterprise requirements and flash media costs have decreased at a much more rapid rate than HDD costs, enabling flash to be cost-effectively used for less latency-sensitive workloads. By 2016, it was clear that, for most primary workloads, AFAs offered a better TCO proposition than most HDD-only solutions or HFAs. In 2016, IDC first published a study highlighting the six major areas that contributed to this lower overall TCO (relative to arrays still using HDDs):

- The need for far fewer devices to meet both performance and usable capacity requirements
- Lower energy and floor space consumption
- The need for fewer application servers (due to significantly reduced and much more consistent latencies, which led directly to much higher CPU utilization)
- Lower software licensing costs (due to needing fewer servers)
- Lower administration costs (time spent performance tuning systems frequently dropped to zero)
- Cost savings due to better device-level reliability (electronic flash media versus mechanical spinning disk media)

Today, AFAs dominate primary storage spend, driving over 80% of spending in the external storage systems market. Nagging questions in 2011 about media reliability and endurance and high flash media costs in primary storage environments have pretty much been put to bed, and there are a number of AFAs

from multibillion-dollar enterprise storage vendors that feature "six nines" availability and a comprehensive set of mature, proven storage management capabilities (RAID, snapshots, encryption, replication, stretch clusters, etc.). Pure Storage is today one of the top 5 enterprise AFA vendors by revenue.

Even outside of the performance advantages of flash media, higher reliability and increased storage density (at the device level) were key parts of the value proposition in the primary storage markets. Interestingly, aspects of these features are also of interest in the capacity-optimized storage markets that service the needs of less performance-sensitive virtualized infrastructure, disaster recovery, multicloud testing/development environments, hybrid cloud snapshot consolidation, and other colder storage workloads. While capacity-oriented storage environments typically don't need low latency, they can benefit from the higher throughput and bandwidth of flash to enable faster data movement and drive denser consolidation. And with many capacity-oriented storage environments tasked with supporting petabytes of data, or more, over potentially long time periods, flash media reliability and storage densities (and their reduced energy and floor space consumption requirements) represent very attractive features. Flash media costs have continued to drop, and the introduction of quad-level cell NAND flash will continue to drive those costs down over the next several years. IDC predicts that, from 2020 to 2024, enterprise flash raw cost per gigabyte will continue to drop at a compound annual rate of 16.6%.

From its beginnings selling an all-flash enterprise storage platform for structured (e.g., block-based) primary storage environments, Pure Storage has since expanded into unstructured (e.g., file- and object-based) data environments as well. In 2016, the vendor introduced FlashBlade, an all-flash storage platform targeted for use with scale-out file- and object-based workloads. Based on the 2019 acquisition of Compuverde, a software-defined file-based storage vendor, in 2020 Pure Storage added file support to its FlashArray products, including both the FlashArray//X and the FlashArray//C, turning these arrays into unified storage platforms that can simultaneously support both block- and file-based workloads. One of Pure Storage's original goals as a company was to bring enterprise-class all-flash technology to the masses in a cost-effective manner, and the introduction of its newest storage platform, the FlashArray//C, brings that value proposition to capacity-oriented storage workloads.

FlashArray//C leverages the proven FlashArray architecture and is based around NVMe, not SAS. It runs the same storage operating system as the FlashArray//X (called Purity), a design choice that bears significance for customers. Host connection options include both legacy Fibre Channel (FC) and iSCSI as well as newer NVMe over Fabrics (NVMe-oF) options. The FlashArray//C implementation does include some different capabilities that optimize the platform for more cost-sensitive, capacity-oriented storage environments:

- **Capacity-optimized DirectFlash Modules (DFMs) and DirectFlash Software (DFS).** Pure Storage does not use off-the-shelf SSDs in its arrays — it buys NAND flash media directly from suppliers in the United States, Japan, and Korea and builds its own custom flash modules (which Pure calls DirectFlash Modules). (Note that this means there is minimal impact on Pure's costs from tariffs between the United States and China, where many off-the-shelf SSDs are manufactured.) Pure Storage optimizes its flash translation layer (the DFS) to interact with the flash media directly (rather than through an off-the-shelf SSD controller and FTL), giving it the opportunity to make significant optimizations that are very specific to different FlashArray//X and FlashArray//C use cases. The combination of hardware and software components that deliver these optimizations in Pure Storage's AFAs are referred to by the vendor as DirectFlash, which includes four components: the DFS, the DFMs, the DirectFlash shelf (an external storage expansion shelf connected to the array using NVMe-oF), and the DirectFlash Fabric (NVMe-oF-based host connections).

On the FlashArray//X, the DFS is optimized for low latency, while on the FlashArray//C it is optimized for media reliability and endurance, capacity density, and low cost. This allows Pure Storage to use lower-cost QLC media in enterprise platforms earlier than the other vendors that are still waiting for the SSD manufacturers to produce cost-effective drives that can stand up to the media endurance requirements of write-intensive enterprise environments. The DFS on FlashArray//C manages media reliability in a manner requiring significantly less over-provisioned flash capacity on each DFM, a feature that drives meaningful cost savings relative to the overprovisioning required in off-the-shelf QLC-based SSDs. Ultimately, DFS is a major contributing factor to Pure Storage's ability to reliably leverage QLC technology to deliver an all-flash platform for capacity-optimized storage environments that is cost competitive with SATA HDD-based systems.

Although it is beyond the scope of this white paper, the ability to tune the DFS for different workload requirements is just one of many advantages that Pure Storage enjoys versus its competitors (all of which use off-the-shelf SSDs or are rapidly moving in that direction now that NVMe is here) because of the decision to create its own custom flash storage devices. Customers should ask about this because it will give Pure Storage the opportunity to discuss all the performance, density, reliability, media endurance, time-to-market, and cost advantages that flow from this strategic decision on both the FlashArray and FlashBlade storage platforms.

- **A performance and cost profile clearly targeted at killing HFAs, not replacing FlashArray//X.** The FlashArray//C is clearly differentiated from the latency-optimized FlashArray//X in two key areas: performance and cost. The FlashArray//X can consistently deliver sub-millisecond latencies at scale in mixed workload environments (and sub-150µs latencies when configured with the storage-class memory-based Direct Memory products), while the FlashArray//C delivers latencies consistently in the 2–4ms range. While this performance range is not attractive for most latency-sensitive primary workloads, it is very attractive for many capacity-oriented workloads and much more consistently delivered at scale relative to what most HFAs can deliver, particularly when you note that the FlashArray//C can pack 1.3PB of raw storage capacity (or over 5PB of effective capacity, assuming a 4:1 data reduction ratio) into just 3U of rack space. This is based on a new 49TB QLC-based DFM that just became available. A smaller 24TB QLC-based DFM is also available for FlashArray//C, and customers can mix and match different DFM types in the same system.

  Based as it is on QLC technology, the cost per gigabyte of FlashArray//C capacity is 30–40% less than that of FlashArray//X (based on street prices). While 7,200rpm SATA HDDs feature a lower cost per gigabyte for just the media, the flash media in the FlashArray//C offers much more on every other metric that is of interest for capacity-oriented storage workloads: throughput, bandwidth, storage density, energy and floorspace consumption, and reliability. Keep in mind that the cost per gigabyte at the system level includes storage controllers, cache and, in some cases, other components whose costs must be spread across the capacity of a given system, so a direct cost-per-gigabyte comparison at the storage device level can be misleading.

  But FlashArray//C doesn't have to cost exactly the same as HFAs to offer a better TCO, and the closer those cost-per-gigabyte numbers get, the more enterprises will find the advantages of "all flash" to outweigh any small acquisition price disadvantage for capacity-oriented workloads that have any performance sensitivity. FlashArray//C doesn't have to be cost effective for every capacity-oriented workload before it can be cost effective for many. Let's face it – there's nothing cheaper than tape if you're just going to write something once and then never access it, and tape performance will always be sufficient for some cold storage workloads. Even AWS still uses tape in Glacier. But FlashArray//C's economics and features make it attractive right away for many capacity-oriented workloads, and those economics and features will only get better over time as flash cost per gigabyte continues to decrease faster than HDD cost per gigabyte.

The FlashArray//C leverages a lot of other capabilities that have endeared the vendor to its customers. It runs the same Purity storage operating system as FlashArray//X and offers all the same enterprise-class data services. Other vendors with both AFAs and HFAs in their storage system portfolios have for years been highlighting the ability to replicate from AFAs to HFAs, letting them configure a lower-cost disaster recovery solution than vendors that only offered AFAs could deploy. With FlashArray//C at the same price point as HFAs and a common storage operating system running on both //X and //C, Pure Storage can now offer this same configuration (what Pure refers to as ActiveDR) with one important difference. The performance capabilities of FlashArray//C, in terms of latency, throughput, and bandwidth, are better than those of most HFAs. This means that on backups and recoveries, data can be moved faster, and when used as a disaster recovery platform in conjunction with backup solutions such as Veeam that support an "instant restore" capability that will run virtual machines (VMs) directly off of data in the backup platform, it offers better performance than most HFAs.

In addition to the ability to support Pure Storage's ActiveDR capability, the FlashArray//C is also a unified storage platform that can simultaneously run both block- and file-based workloads as needed. This support for multiple access methods provides additional flexibility for consolidating different capacity-oriented storage workloads that is not available on HFAs that support only block-based storage.

In addition to its use of Purity, the FlashArray//X storage operating system, the FlashArray//C uses the same hardware architecture as well. There have, however, been some changes to optimize the platform for more capacity- and cost-sensitive workloads. The FlashArray//C's storage controllers, though based on those of the FlashArray//X, have been engineered to maximize media endurance (a critical concern where QLC is used in enterprise environments) rather than low latency; the DFS is tuned differently; and the storage devices (DFMs) are designed for, and use, QLC flash media. A read cache on the DFMs minimizes the "read disturb" activity generated by enterprise workloads, improving media endurance. For error correction, FlashArray//C implements the low-density parity check (LDPC) soft decode method, an approach that provides better error correction than the hard decode method used by other vendors. With QLC flash media, there is an increased number of voltage thresholds that must be managed at the chip level, and the soft decode method does a better job of that. The multilevel compression in Purity operates slightly differently in the FlashArray//C because of the higher latency of the media, enabling more levels of compression to be completed inline (driving a slightly higher data reduction ratio) before the data is ever written to persistent storage. And to better optimize for cost efficiency in higher-capacity environments, Pure Storage increased the stripe width to reduce RAID capacity overhead without putting data durability at any greater risk.

The FlashArray//C is covered under both Pure1, the vendor's cloud-based predictive analytics platform, and the Evergreen Storage Program (along with its flash media endurance, nondisruptive upgrade, and other guarantees). Pure Meta, Service Orchestrator, and VM Analytics all work with the platform. Customers can purchase systems outright or under Pure Storage's Pure as-a-Service subscription model.

Use cases for FlashArray//C include capacity-oriented workloads running on virtualized infrastructure, backup and disaster recovery, multicloud test/development environments, and hybrid cloud snapshot consolidation. Pure Storage's cloud integration capabilities, which are the same for FlashArray//C as they are for the rest of the vendor's storage platform portfolio, enable easy data mobility between on-premises and cloud-based environments using CloudSnap (a portable snapshot capability) and/or replication. Pure Storage's cloud integration capabilities feature a unified subscription model that can span both on-premises and cloud-based infrastructure with Pure as-a-Service pay-per-use pricing, unified hybrid cloud management through Pure1, and unified cloud data technology (equivalent enterprise-class functionality in all locations) with Purity. These capabilities allow customers to move and run any application anywhere, anytime — the foundational capability underlying any effective hybrid cloud strategy.

CHALLENGES/OPPORTUNITIES

Despite all the advantages of all-flash storage in enterprise environments, anytime a vendor seeks to extend its use into less performance-intensive environments the issue of cost arises. To address this question head-on, Pure Storage has been very open about its target cost point. It has set its sights on HFAs that, based on nearline SAS drives, deliver a rough cost of $0.02–0.04/GB at the system level (assuming the blend of flash and spinning disk, the extra cache required to deliver performance, and the rest of the supporting system infrastructure [e.g., controllers]). With data reduction estimates between 4:1 and 5:1 assumed (which for many capacity-optimized workloads are in fact quite low), the FlashArray//C delivers a rough cost in that same range plus all the added benefits of all-flash configuration that were discussed previously. When Pure Storage first introduced the idea of all-flash storage for primary workloads, it had to overcome this same cost challenge, and it is hard to deny that the vendor has been extremely successful. The vendor faces this same challenge with the FlashArray//C and less performance-sensitive workloads.

If flash cost the same as spinning disk, the market for spinning disk would rapidly decline. By being honest and open about the cost-per-gigabyte point the FlashArray//C can deliver, Pure Storage is making it very easy for potential customers to examine the value proposition it offers and decide if the time is right to begin deploying AFAs for less performance-sensitive workloads. This is the opportunity.

CONCLUSION

In September 2019, Pure Storage introduced the FlashArray//C with QLC flash media-based pricing. The system at that time, however, still used triple-level cell (TLC) flash media as the vendor fine-tuned its ability to manage QLC media endurance and data durability to its high standards. In October 2020, Pure Storage released QLC-based DFMs in two capacities (24TB, 49TB) along with some enhanced software-based features (e.g., unified storage, ActiveDR) that open up new use cases and improve the efficiency of the FlashArray//C (a key consideration for its target workloads). Pure Storage is the first established enterprise storage vendor to introduce an AFA targeted at capacity-oriented workloads, and it offers a strong value proposition that is compelling for many of those workloads. Customers looking to refresh capacity-optimized storage platforms will want to evaluate what the FlashArray//C may bring to their own environments.

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
idc-community.com
www.idc.com