PURESTORAGE®
Uncomplicate Data Storage, Forever

BUSINESS WHITE PAPER

The Case for Replacing HDDs with All-flash Arrays

February 2023

Contents

Introduction

In 2012, the first all-flash storage arrays (AFAs) became available and competed directly with all-hard disk drive (HDD) arrays to handle enterprise storage workloads. Traditional storage vendors claimed AFAs were far too expensive for this purpose, but by 2016 all of them had introduced their own AFA offerings. By 2019, AFAs were driving over 80% of all primary (i.e., performance-sensitive, mission-critical) storage workload revenues across the industry. Enterprises had clearly come to understand how the total cost of ownership (TCO) of AFAs was far better for these workloads and how much easier they were to manage than older, HDD-based systems.

Flash won against HDDs not because of a lower $/GB raw capacity cost but due to a host of other reasons. AFAs offer a far lower TCO over the useful life of a system, as well as significantly better performance and higher storage densities. Flash storage devices are far more reliable than HDDs, reducing failures and minimizing the time spent replacing failed devices. Their lower latencies use server compute cycles much more efficiently, leading to a reduction in the number of cores required to meet any given performance requirement and lower software licensing costs. AFAs use media capacity much more efficiently, resulting in a more compact infrastructure that draws less power and uses less rack and floor space. Finally, AFAs are much easier to manage than HDDs because flash's high performance dispenses with all the manual performance tuning that HDDs with latency-sensitive applications and evolving workloads require.

The Case for All-flash for Secondary Storage Workloads

Secondary storage workloads do not require the same low latencies that primary storage workloads do and are much more capacity- and cost-sensitive, often requiring multiple petabytes (PBs) of data that must be retained for compliance or regulatory reasons over many years. Common secondary storage workloads include backup and disaster recovery, big data analytics, PACs in healthcare, electronic design automation, content repositories, and active archives.

Traditional secondary workload HDD storage vendors, anxious to protect their revenue streams from another all-flash incursion, have claimed that all-flash is too expensive. Just as they did back in 2012, these vendors focus on the $/GB cost of raw capacity at the storage device level. Pure Storage® led the charge that displaced HDDs in primary storage workloads over the last decade, and we firmly believe that all-flash will achieve the same thing with secondary workloads. Our belief in this outcome is based on a TCO argument, just like the primary storage rollover was, but includes a few new wrinkles as well.

Seasoned IT executives know that capital acquisition costs for storage infrastructure make up 35-40% of the overall TCO over the life of arrays, with the remaining 60-65% composed of operating costs. While having a lower acquisition cost is nice, to reduce TCO the real emphasis needs to be on lowering operational costs. With the FlashArray//C™, Pure Storage can equal the all-HDD platform acquisition costs for large scale secondary storage workloads (i.e., those growing to or already over 20PB in size) while driving significantly lower operational costs over the useful life of arrays. These huge operational cost savings result in a TCO that can be up to 70% lower than all-HDD systems, making a compelling economic case for enterprises to move these workloads to all-flash now.

Lowering operational costs is significantly more important than lowering capital costs in reducing the overall TCO of storage infrastructure.

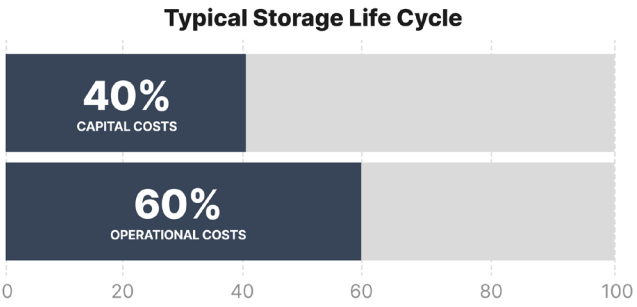

FIGURE 1 Capital vs operational costs during the typical storage life cycle

FlashArray//C is ideal for on-premises all-HDD system replacement with Evergreen//One™, a storage-as-a-service offering from Pure Storage that requires no upfront capital costs and has lower operational costs than the equivalent all-HDD storage system that it replaces. This acquisition model will allow enterprises to move to the benefits of an all-flash secondary storage system with no CAPEX while reducing OPEX. To understand how Pure can make this offer, the remainder of this white paper will review the capital (acquisition) and operating cost implications of the new all-flash storage infrastructure.

Comparing Pure FlashArray//C and All-HDD System Acquisition Costs

Flash costs have continued to decrease as flash media densities have increased. Flash $/GB costs are going down at roughly 20% per year while HDD $/GB costs are dropping at 2-3% per year. Newer media technologies, like quad-level cell (QLC) NAND flash, are driving major jumps in flash device density; meanwhile, HDD vendors are already struggling to increase device densities while keeping costs low. Even though decreasing flash media costs shift the cost of raw AFA capacity closer to that of all-HDD systems every year, there are compelling economic reasons to switch today:

- Flash and HDD storage devices are not comparable when it comes to performance. A flash device can easily deliver 10 to 100 times the performance of a single HDD depending on whether you're looking at latency, I/O operations per second (IOPS), or throughput.
- Flash devices have already outstripped the capacities of the biggest HDDs. HDD vendors are struggling to deliver 20TB devices in volume at reasonable prices, while flash devices that are 48TB in size are already shipping from Pure Storage on FlashArray//C. Terabytes per U of rack space (TB/U) is a common metric to compare storage densities between systems, and using its high density flash storage devices the FlashArray//C can have storage densities two to four times higher than all-HDD systems based on commodity HDDs.
- These first two aspects imply that to meet both performance and capacity requirements (whatever they may be) you will need far fewer Pure Storage flash devices than HDDs. In addition to storage devices, storage systems also need controllers, power supplies and fans, enclosures, and cabling. And the more storage devices a system needs, the more of these other components you'll need. FlashArray//C needs far fewer components and is much smaller in size. Smaller kits with fewer components are more reliable and require far less power and floor space.
- On-disk data protection algorithms employ data redundancy to ensure the data integrity and availability of data in the wake of storage device failures. HDD-based systems offer dual-parity RAID approaches that use small RAID group sizes. For example, a typical dual-parity RAID 6 implementation uses a 4(data) + 2 (parity) approach, resulting in a 33% capacity overhead due to data protection (2/6=33%). Pure uses RAID group sizes that can be as large as 8+2 for dual parity protection, resulting in a much lower 20% capacity overhead. Capacity overhead drives the need to purchase additional storage devices and capacity, and for a 20PB configuration that additional 13% of capacity overhead would require the purchase of an additional 2.6PB of capacity with the HDD-based system, increasing its cost accordingly.
- And finally, the relatively low performance characteristics of HDD will result in enterprises having to over-provision raw capacity just to meet their usable capacity requirements. HDD vendors typically caution customers against filling devices more than 70–80% full because of performance impacts, while Pure Storage flash devices are designed so customers can actually use 95% or more of their raw capacity without undue performance issues. This is another contributing factor, just like performance and on-disk data protection efficiencies, that minimizes the amount of raw all-flash capacity needed to meet customer requirements.

When configuring all-HDD, secondary storage systems, enterprises do not necessarily use the largest-capacity devices. Most HDD system vendors offer a number of options (which can range from smaller than 2TB to 20TB). Systems using smaller HDDs deliver a higher IOPS/TB ratio than those using larger HDDs, but this option can cost significantly more because more smaller HDDs must be purchased to hit the required capacity. The range in sizes is supposed to enable enterprises to configure the system that best meets their combination of performance and capacity requirements. In contrast, flash devices from Pure Storage deliver the same IOPS/TB regardless of their size, allowing customers to adjust system costs based entirely on capacity requirements. This is yet another factor that reduces the TCO of Pure Storage kits.

When comparing the acquisition costs of traditional multi-controller all-HDD systems in the 20PB range, the $/GB cost of usable capacity is around $0.15/GB (based on the use of 7200RPM HDDs that cost less than $0.02/GB each). The $/GB cost of usable capacity divides the total cost of a system, including not only the storage devices but also all the controllers, power supplies and fans, enclosures and cabling, and software licensing by the usable capacity (which takes into account the formatting, on-disk data protection overhead like RAID, and sparing). For HDD-based systems using traditional dual-parity RAID, usable capacity is generally at least 40% lower than raw capacity. For AFAs from Pure Storage, usable capacity is generally no more than 20% lower than raw capacity.

The acquisition cost of a Pure Storage FlashArray//C configured to provide 20PB of usable capacity is $0.20/GB. But this system supports a data reduction technology called data deduplication that is not supported on HDD-based configurations of competitive systems. While the data reduction ratios will vary based on the data types in each workload, Pure Storage assumes a very conservative 2:1 data reduction advantage on the FlashArray//C for mixed secondary storage workloads. With no data reduction advantage, purchasing a FlashArray//C costs roughly the same as an all-HDD system configured to meet a given performance and capacity requirement (due primarily to needing 20% less raw capacity and far less hardware). Any data reduction achievable with various workloads will drive that number even lower—in some cases even lower than an equivalent all-HDD system.

The takeaway from this analysis is that the $/GB acquisition cost of FlashArray//C is on par with all-HDD systems configured to meet a given performance and multi-PB capacity requirement. But remember, the acquisition cost is roughly only 35-40% of the overall TCO of storage infrastructure over its useful life. Let's take a look at an operational cost comparison between all-HDD systems and the Pure Storage FlashArray//C deployed at multi-PB scale.

Pure Storage Wins Hands-down on Lower Operational Costs

All-HDD systems have significantly higher operational costs. Because they require more components, they are more complex—a factor that makes them harder to deploy, manage, scale, and upgrade. With their lower performance at the device level, more time must be spent tuning systems as workload requirements evolve. Moving data to new locations takes longer. Features like snapshots that enable space-efficient data re-use are limited in their ability to share data between applications because of the high latencies and lower IOPS of HDDs. The larger kits required with HDDs drive higher energy and floor space costs and require more supporting infrastructure, like controllers, power supplies, fans, enclosures, and cables. Maintenance costs are higher because you're paying maintenance on a larger number of components. HDDs fail far more often than flash devices, and enterprises with deployments of hundreds to thousands of HDDs spend a significant amount of time finding and replacing failed drives and experiencing long rebuild times (during which performance can be impacted and data can potentially be at risk).

Two data points may be of interest in validating our claims in this area. Existing enterprise customers (of which Pure Storage has over 11,000) routinely testify that the administrative span of control (i.e., how much storage capacity can be managed by a single administrator) with Pure AFAs is 30% to 40% greater than for competing AFAs, leading to higher productivity that is particularly cost-effective with larger storage infrastructures. And it is significantly greater than that against all-HDD systems (which are much harder to manage than AFAs).

The other data point is with Meta, the parent company of Facebook. Early in 2022, Meta announced that they had awarded an exabyte scale storage contract for their Artificial Intelligence Research SuperCluster (AI RSC) to Pure. After having worked almost two years to create the needed system themselves using a software-defined and commodity server-based storage hardware approach, they concluded that they would not be able to build a system that could stay within their power budget for initial deployment (which was over 100PB), let alone accommodate the needed expansion over the life of the project. They considered storage systems from several enterprise storage vendors and selected Pure Storage, citing our ability to drive an 80% lower TCO than other vendors due primarily to our infrastructure efficiency advantages. These infrastructure efficiency advantages were driven by smaller kit, lower energy and floor space consumption, and lower administrative and maintenance costs.

And a Pure Storage Bonus: A Large TCO Impact from Significantly Extended Storage Life Cycles

We've noted that operational costs during the life of a storage system compose a significant portion of its overall TCO, and the costs incurred to migrate to next generation technology through a forklift upgrade every three to five years are the single largest contributor to these operational costs. While most storage systems allow enterprises to add both performance and storage capacity over time, efficiency considerations generally drive a comprehensive technology refresh requirement to move to more cost-effective systems that transcend both the performance and capacity densities of the prior generation.

The familiar forklift upgrade requires enterprises to not only buy new controllers and enclosures but also to re-license storage software and rebuy the storage capacity that they already paid for in their previous system. These costs, combined with the costs to migrate data and workloads to the new system, drive the significant expense (and potential risk) of technology refresh.

What drives the need for forklift upgrades is that customers cannot non-disruptively upgrade a multi-controller array in-place to next generation technology. While customers may be able to add additional controllers and storage devices to an existing system, these are based on older technologies that offer less performance and/or capacity and demand higher maintenance costs. Systems that enable the non-disruptive replacement of older components with next generation storage technologies over time result in more efficient systems and a longer storage life cycle—factors which keep systems simpler and costs down.

Our track record with our AFAs shows that our customers achieve at least ten-year storage life cycles. Ninety-seven percent of the arrays that Pure has shipped since 2012 are still in use. We have many customers that have upgraded their original FlashArray4xx systems all the way up to the current generation of FlashArray//X™ arrays. These customers never had to relicense storage software or rebuy existing capacity (although they have added additional capacity over time). They received access to all new software features at no additional charge, including capabilities that most vendors charge separately for, including AI-driven system (or fleet) monitoring and management, synchronous replication, and stretch clusters.This allows our customers to easily leverage new storage technologies as they become available to keep their systems continuously updated (or as continuously updated as they would like them to be).

Pure Storage's strategy to easily support non-disruptive multi-generational technology upgrades throughout the life of systems includes both a product architecture and program component. Dubbed Evergreen® Storage, this approach includes a storage architecture that was specifically developed to enable the non-disruptive, in-place integration of newer technologies over the life of an array, and it is unique in its ability to allow customers to non-disruptively upgrade components such as controllers, storage devices, backplanes, and enclosures over time to keep a system continuously updated and operating at peak efficiency without forklift upgrades. Program aspects of Evergreen Storage can include flash endurance, flat and fair maintenance, and 30-day money back guarantees, 7×24 worldwide support coverage, and next generation controller upgrades at no additional charge every three years (with the option to upgrade controllers more frequently with guaranteed trade-in credits for a nominal charge).

By not requiring a single forklift upgrade in over a decade, our customers save tens to hundreds of thousands of dollars (depending on the size of their storage infrastructure), as well as the aggravation and time that typically come with these upgrades. Evergreen Storage is not just a program; our product architecture comprehensively supports non-disruptive, multi-generational technology upgrades that make it unique in the industry in its ability to deliver large cost savings by significantly extending the storage life cycle. A 10-year life cycle likely removes the need for two all-HDD forklift upgrades, and as the Pure lifecycle track record stretches to 15 years that means three forklift upgrades that may not have to occur, driving much lower costs, reduced risks, and a much better customer experience.

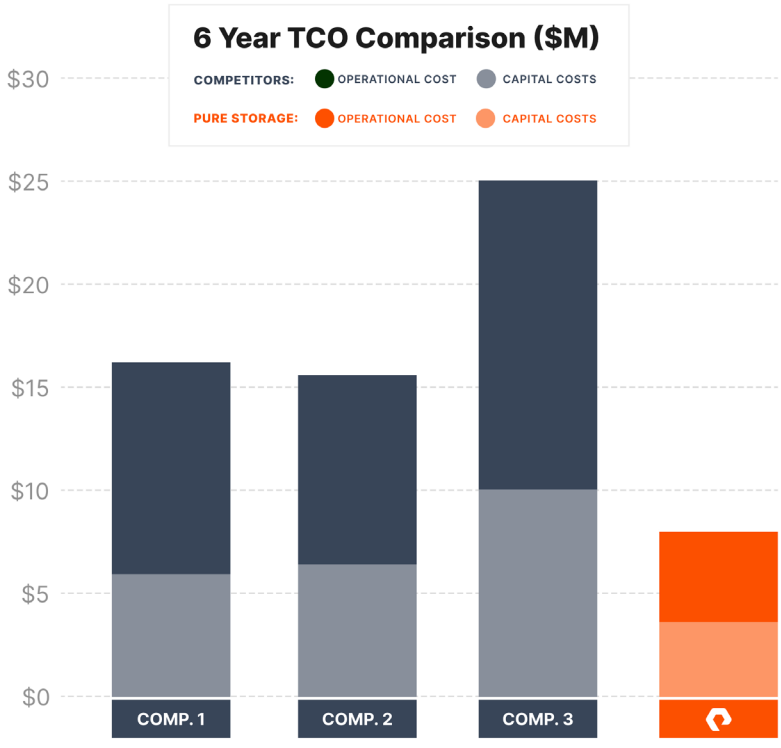

FIGURE 2 Six-year TCO comparison (assuming 2:1 data reduction with FlashArray//C)

Bringing It All Together

Figure 2 shows a six-year TCO comparison between Pure and three other well-known, multi-billion dollar enterprise storage vendors' systems. Note that Pure actually beats the acquisition cost of all three of these HDD-based systems with their FlashArray//C (which, by the way, is based on dense QLC NAND flash media) while exhibiting significantly lower TCO over a six-year period. Note also that this comparison leverages the use of deduplication to provide 2:1 data reduction, a feature not available on HDD-based systems, and assumes one forklift upgrade on the part of competitors.

The FlashArray//C drives a noticeably lower TCO relative to competitive all-HDD offerings after only six years, and this gap widens significantly as the Pure life cycle extends to 10 years and beyond.

How FlashArray//C Beats All-HDD Systems

We believe that a correctly designed AFA, leveraging QLC (and over time denser) flash media in a storage system designed specifically for solid-state storage (and not adapted from earlier HDD-based designs), will offer a significantly lower TCO over the life of the solution due to the following:

- Because of its larger storage device sizes (the largest of which is 48TB today), FlashArray//C needs roughly half (and likely many fewer) the storage devices to meet any given set of performance and multi-PB capacity requirements compared to HDD-based systems using 18TB HDDs.
- FlashArray//C systems are built from many fewer storage devices and thus need much less supporting infrastructure components like controllers, enclosures, fans, power supplies, and cables. FlashArray//C systems need far less rack and floor space, draw significantly less power, and experience increased reliability.
- HDD-based systems in the multi-PB range and beyond require at least hundreds (and potentially thousands) of individual HDDs, and with the higher failure rates of HDDs in practice enterprises with large configurations will likely be suffering a lot of device rebuilds and replacing at least several failed HDDs per month - this combines with much more hands-on administration and a much lower administrative span of control to drive higher ongoing maintenance costs.
- Because of its ability to support non-disruptive, in-place, multi-generational technology upgrades, FlashArray//C delivers at least a 10-year (and quite possibly longer) storage life cycle. FlashArray//C systems can be non-disruptively and easily upgraded to newer technologies as they become available while HDD-based systems will require at least two (and possibly three) very costly forklift upgrades over that same period.

In addition to these advantages, all-flash storage infrastructure based on FlashArray//C delivers the performance to ingest and recover data faster for data protection and disaster recovery purposes and enable space-efficient data sharing (using snapshots and clones) for workflows like analytics, e-discovery, and client/patient management to reduce the need for other separate storage silos. The denser workload consolidation that the FlashArray//C's performance and availability (six nines) enables will have a multiplicative effect on its TCO benefits as other storage silos can be retired as their data is consolidated onto the new infrastructure.

FlashArray//C delivers efficiencies that enable it to match HDD $/GB acquisition costs in multi-PB configurations, while the TCO of the system over its life cycle provides a compelling economic reason to move secondary storage workloads to all-flash today.

Learn More

If this discussion has piqued your desire to learn more about how FlashArray//C can offer a better storage infrastructure solution for your less latency-sensitive block- and file-based workloads than HDD-based systems, please contact us.

©2023 Pure Storage, Inc. All rights reserved. Pure Storage, the P logo mark, FlashArray, FlashBlade, Pure1, and Evergreen are trademarks or registered trademarks of Pure Storage, Inc. All other names may be trademarks of their respective owners.

PS2381-01 02/2023