

BUSINESS WHITE PAPER

The All-flash Data Center Is Imminent

The efficient storage infrastructure of the future.

September, 2023

Contents

Introduction	3
The Total Cost to Buy and Manage a Storage System over Its Useful Life	4
The Rise of the Cost per Effective Gigabyte Metric	4
Selecting Device Capacities When Building an All-HDD Storage System	5
Comparing Pure Storage DirectFlash Technology and HDDs	5
HDD Technology: The End is Near	5
Pure Storage DirectFlash vs COTS SSDs	6
So how is DirectFlash different from COTS SSDs?	7
Evergreen Storage: A Subscription to Innovation	8
Learn More	8



Introduction

In the enterprise storage market, flash media strongly dominates in performance-sensitive workloads. Although enterprise-class storage systems based entirely upon flash only entered the market in 2012, by 2019 up to 80% or more of revenues driven by primary storage workloads came from all-flash arrays (AFAs). By that time, hard disk drives (HDDs) had already been mostly relegated to secondary workloads that were more capacity- and cost-constrained than primary workloads, but they still accounted for roughly 90% of the stored enterprise data. Given the continued advance of flash semiconductor technology, a lively discussion is happening in enterprise storage about the ultimate fate of HDDs in the data center.

“By 2028, practically no new all-HDD storage systems will be sold for data center computing.”

Pure Storage® contends that by 2028, practically no new all-HDD storage systems will be sold for data center computing. While some HDDs may continue to be sold to expand existing systems that have not yet reached their end of life, HDD shipments have been significantly declining over the last decade. From a high of over 651 million units shipped in 2010, shipments have fallen almost 71% to 166 million in 2022.




FIGURE 1 HDD shipments have significantly declined since their high in 2010¹.

The market for “performance” HDDs (10K and 15KRPM) has basically disappeared in less than a decade (having been replaced by flash), and nearline (7200RPM) HDDs account for almost all the HDD spend going forward. This occurred despite the fact that the 10K and 15KRPM HDD cost per gigabyte for raw capacity was significantly lower than for flash over the last decade. The fact is, other cost variables outweighed that over the lifecycle of primary storage systems, and eventually got to the point that very few enterprises wanted to buy HDDs for performance-sensitive workloads.

It is that same dynamic that will ultimately result in the replacement of nearline HDDs by flash for almost all of the remaining enterprise workloads despite the fact that the raw cost of flash is not expected to ever drop below the raw cost of HDD capacity this decade. This may be a surprise for those that compare the cost per gigabyte of raw HDD capacity in the \$0.01—\$0.015 range with the cost of raw flash capacity that is ten to fifteen times higher today. How can Pure Storage be touting the all-flash data center and forecasting the end of HDDs? This white paper lays out the argument.



The Total Cost to Buy and Manage a Storage System over Its Useful Life

As AFAs came to dominate primary storage spend and push out “performance” HDDs, it was because the total cost of ownership (TCO) of AFAs was lower than the TCO of all-HDD systems. Interestingly though, the up front costs to buy an AFA with the same performance and capacity as a comparable all-HDD system were in fact still higher—the economic benefits come from the much lower operational costs over the life of the AFA. With the introduction of quad level cell (QLC) NAND, flash can now replace HDDs for capacity- and price-sensitive workloads at prices comparable to all-HDD systems (or even lower for larger systems).

Several developments have contributed to this reduction in up-front purchase costs for AFAs relative to all-HDD systems:

- Flash media cost per gigabyte has been declining at roughly 20% per year, while HDD cost per gigabyte has been dropping at only 9% per year.²
- Flash media density improvements have resulted in commodity off-the-shelf (COTS) solid-state disks (SSDs) that already have higher capacity than the latest high-capacity HDDs, and density projections for COTS SSDs show a two to three times advantage in device capacity by 2025.
- The maturation of Pure Storage DirectFlash® technology—which extends flash storage device density to an anticipated ten times advantage by 2025 (much higher than COTS SSDs) and improves effective capacity utilization at the device level to 70-80% (compared to 50-60% for HDDs)—is what allows Pure Storage to provide all-flash systems today with an upfront cost comparable to that of all-HDD systems configured to meet a certain performance and capacity target.

“Our contention that HDDs will comprise a niche market in the data center by 2028 is based on power, space and labor efficiencies that are only achievable with Pure Storage systems.”

Pure Storage systems built around our DirectFlash technology are far denser (smaller) and more efficient than storage systems built around HDDs (and even SSDs), so our contention that HDDs will comprise a niche market in the data center by 2028 is based on power, space and labor efficiencies that are only achievable with Pure Storage systems. Pure Storage’s flash storage devices are called DirectFlash Modules (DFMs), and today (2H, 2023) we ship 75TB DFMs. We will ship 150TB DFMs by the end of 2024 and 300TB DFMs by 2026. The industry, meanwhile, is currently shipping 2-20TB HDDs in volume and may get to 25-30TB HDDs in volume by 2026.

The Rise of the Cost per Effective Gigabyte Metric

Since its inception, the HDD industry has focused on the cost per gigabyte of raw capacity as a key metric. Experienced storage managers know, however, that it is not the raw capacity that determines how much data can be stored but the effective capacity. Factors like formatting, data reduction, on-disk data protection [RAID or erasure coding (EC)], and how much data can be stored in a device without unduly impacting performance together determine the effective capacity. Outside of data reduction (which increases effective capacity), these other factors all impose capacity overhead on all-HDD systems, lowering the effective or true capacity which can be used to store data. Storage devices that support a lower effective capacity force customers to buy more devices to hit a certain performance and capacity target.

“Pure Storage is shipping 75TB flash storage devices today and has a solid roadmap to 300TB flash storage devices with the highest effective capacity utilization rate in the industry.”

When enterprises measure and compare the cost of competing systems based on raw capacity, they will eventually have to address the effective capacity question (and associated costs). Why not address that up front? As data reduction, on-disk data protection, and effective capacity utilization at the system level become better understood, focusing on the cost per effective rather than raw gigabyte makes more sense.



Selecting Device Capacities When Building an All-HDD Storage System

While there are HDDs in the 20TB+ capacity range today, most enterprises are forced to purchase devices with lower capacities when configuring large, multi-PB scale systems. This is primarily due to two reasons:

- The number of input/output operations per second (IOPS) does not increase much with larger capacity devices, which means that the IOPS/TB goes down as the device size goes up. (This is due to the fact that HDDs are limited to one read/write head per disk, whose I/O speed is limited by the rotational speed of the disk itself.) As an HDD fills up, the ability to access all of its data with a minimum threshold of performance decreases, making capacity past a certain point not very usable. This is why disk vendors typically recommend not filling their devices more than 60-80% full.
- The limited throughput with which data can be written to or read from the device leads to longer and longer rebuild times as device capacity increases, which potentially puts both performance and data integrity at risk. HDDs are clearly less reliable than flash-based devices—they have an annualized failure rate of 1.54%³—and the failure rate increases with age. Smaller capacity HDDs are actually slightly more reliable and they can be rebuilt faster than larger devices, minimizing data risk. DFMs, on the other hand, deliver more than eight times the throughput at the device level and have an annualized failure rate up to ten times lower with twice the useful life.

These concerns cause many enterprises to decide not to leverage the largest capacity HDDs when building multi-PB storage systems. HDD storage system competitive cost comparisons, however, generally use the largest capacity HDDs that they support in calculations, a choice which leads to a lower cost per gigabyte but is often not a configuration that would actually be deployed in production.

Comparing Pure Storage DirectFlash Technology and HDDs

The Pure DFM roadmap enjoys a huge storage density growth advantage over HDDs, resulting in a much faster reduction in cost per gigabyte at the device level as NAND flash prices continue to decrease. And instead of managing the media indirectly through a flash translation layer (FTL) and through isolated controllers in each device that have no visibility to what else is happening in the storage system, our Purity storage operating system manages the media directly and globally (at a system level). This allows the system to deliver much better performance consistency, data availability, storage density, capacity utilization, endurance, reliability and device lifetimes.

“Pure Storage’s DFMs are already up to 10x more efficient in terms of energy and floor space consumption than HDDs today.”

Pure uses QLC NAND flash media in DFMs for the Pure//E™ family of storage systems. There are two platforms in the Pure//E family. The FlashArray//E™ is based around a scale-up architecture and supports unified block and file storage while the FlashBlade//E™ is based around a scale-out architecture and supports unified file and object storage.

HDD Technology: The End is Near

The only advantage HDDs offer against Pure Storage DFMs is that they cost less on a raw capacity basis. All-HDD systems have less performance, lower reliability, long rebuild times (putting data at risk), are harder and more expensive to manage, and require wholesale technology refreshes every three to five years. Enterprises may assume that the lower cost of raw capacity means that HDD-based systems cost less to buy. But do they really?



Specification	All-HDD Systems	All-flash SSD Systems	FlashArray//E
Device capacity	12TB	15TB	75TB
Target capacity	4096TB	4096TB	4096TB
Number of devices	342	274	55
Enclosures	1 base, 14 expansion (assumes 2U24)	1 base, 11 expansion (assumes 2U24)	1 base, 1 expansion
Rack space	34U	24U	6U

TABLE 1 Pure Storage FlashArray//E has an 84% lower storage device count than a comparably configured all-HDD system and an 80% lower storage device count than a comparably configured AFA configured with commodity off the shelf SSDs.

In looking at the storage device count for systems configured for the same capacity point, the storage density advantages begin to become apparent. Each 12TB HDD costs a lot less than a 75TB DFM, but you have to buy over 6 times as many of them. During normal read/write operations, a single 75TB DFM draws roughly 10 watts while six 12TB HDDs draw roughly 60 watts (6 × 10.1 watts). And the supporting infrastructure for those additional HDDs is much larger as well – more controllers, more enclosures, more power supplies, more fans, more cables, and maintenance charges paid on more devices. All of that additional storage infrastructure also costs more to buy, is more complex and harder to manage, draws more power, and takes up more rack and floor space.

Think about how this comparison will change as HDDs inch towards 30TB devices and Pure Storage will deliver 300TB DFMs by 2026. Keep in mind, however, that to be able to effectively use larger capacity COTS HDDs, the performance and rebuild time concerns must be overcome. The mere availability of a large capacity HDD does not address those challenges.

Keep in mind also that a single DFM can deliver thousands of times the IOPS of a single HDD. This performance density, when combined with the storage capacity density of DFMs, ensures that the much more compact Pure Storage infrastructure delivers significantly higher infrastructure density than all-HDD systems.

Pure Storage DirectFlash vs COTS SSDs

Solid-state disk (SSD) vendors today offer QLC NAND flash-based devices. Can't storage systems vendors building systems from these devices offer advantages similar to DirectFlash over HDD-based systems?

It is true that COTS SSDs offer better performance, storage density, and reliability than HDDs, but they aren't better enough in all those areas to pose an existential threat to nearline enterprise HDDs in the near term. They are, however, better enough to replace some all-HDD systems, depending on workload requirements. Pure Storage systems built with DirectFlash technology, however, deliver two to five times better cost efficiencies than COTS QLC SSDs today (and up to 10x better cost efficiencies than HDDs). The biggest COTS SSDs shipping in volume today are 15TB devices while Pure ships 75TB devices (2023), and our storage density advantage will increase rapidly over the next three years. This is enough of a difference to threaten HDD survival. Projecting out cost per gigabyte reductions with COTS SSDs does not suggest that they could cost-effectively replace all-HDD systems within this decade. But Pure Storage DFMs, with their 2x-5x better cost efficiencies than COTS SSDs, make that possible today. The end of disk is predicated on the spread of flash media technologies like Pure Storage DirectFlash, not the availability of storage systems based on COTS QLC SSDs.



So how is DirectFlash different from COTS SSDs?

The strategy of Pure Storage to develop DirectFlash technology in its Purity operating system was a momentous decision versus the legacy industry's use of COTS SSDs. The Purity Operating System manages the NAND flash chips system-wide both directly and globally, optimizing the storage semiconductor for performance, capacity, reliability and longevity. We create our own proprietary DFMs, and use Purity along with our high-capacity DFMs as the basis to build enterprise-class storage systems that are optimized for the full range of storage requirements.

Contrast this with COTS SSD-based systems. SSDs were designed by disk manufacturers to make it easier for HDD based systems to use flash technology. Given the wide range of systems and operating software which employed HDDs over the last several decades, the flash manufacturers came up with a brilliant idea – create a module that was hardware and software identical to an HDD. Thus the SSD was born. SSDs could thereby be used as general-purpose devices for consumer, commercial and enterprise systems that were originally designed for HDDs. But, in using a semiconductor to mimic a mechanical device, SSDs significantly reduce much of the native advantage of flash technology.

COTS SSDs have an FTL to make them emulate HDDs and need to contend with significant HDD baggage. While this makes it easier to integrate them into storage systems originally designed for HDDs, the FTL imposes performance, media reliability, endurance, capacity utilization, and cost disadvantages relative to flash storage devices that are managed directly as flash. The media in each COTS SSD is managed individually by a controller within that device without any visibility of what is happening elsewhere in the storage system, a design which results in less predictable performance (particularly as storage devices fill up), more write amplification and poorer capacity utilization.

Compared to the 15TB COTS QLC SSDs that are shipping in volume today, Pure Storage's 75TB DFMs as managed by Purity deliver more consistent performance as workloads scale, devote more board real estate to data, have five times higher storage density, demonstrate five times better reliability, squeeze out 15-30% higher capacity utilization, and have a lower cost per terabyte at the device level than COTS SSDs.

Pure Storage leads the industry in efficiency (as measured by watts/TB) and storage density (as measured by TB/U). This is no idle claim. With 75TB DFMs, Pure delivers system-level operating watts per TB well under 1W and a storage density of 1PB/U for systems in the 1–30PB range. Even if the SSD vendors could deliver larger device capacities today, storage systems vendors lack the ability to effectively address all that capacity. Larger capacity devices do not help lower the cost per gigabyte if that extra capacity cannot be effectively used. Pure Storage's DirectFlash technology enables systems to use and recover very large capacity devices – ask us about it and we'll tell you how we do that.

“Pure Storage Pure//E systems lead the industry in both efficiency (lowest watts/TB) and storage density (highest TB/U).”

In conclusion, it is the combination of our high density DFMs and the Purity storage operating system that makes us 2-5x better than COTS SSDs. DirectFlash technology forms the foundation of Pure Storage's industry-leading storage infrastructure efficiency and it cannot be replicated by competitors using COTS SSDs.



Evergreen Storage: A Subscription to Innovation

There is one additional feature of Pure Storage offerings that contribute to the impending demise of disk: how we handle technology refresh. From the beginning, Pure has designed its storage systems to support non-disruptive, in-service, multi-generational technology refresh with investment preservation, based on a system design we call Evergreen® architecture. This approach allows us to refresh systems to increase performance, availability, scalability and functionality non-disruptively without ever having to do a forklift upgrade, and without customers having to spend an additional dollar of capital. Pure has many customers who purchased original systems back in 2012 whose systems have been upgraded to the latest storage technologies over the years with no hardware re-buys, no software relicensing, no data migration over external networks, and no storage re-provisioning. A total of 97% of the systems we have ever sold are still in production use and are hardware- and software-identical to our latest generations of products.

“Evergreen architecture allows Pure customers to increase performance, availability, scalability and functionality non-disruptively without ever having to do a forklift upgrade.”

Contrast this with HDD-based systems that require costly, risky and time-consuming forklift upgrades to move to newer, faster storage systems generally every three to five years.

Explaining how Evergreen Storage works in detail is beyond the scope of this short paper, but we'll be happy to tell you how we do that while pointing out the other benefits of our Evergreen Storage program options (lifetime flash storage device warranty, free controller upgrades, flexible subscription pricing, access to all future Purity software enhancements, etc.). We also have Evergreen programs which deliver storage as a service to dispense with up-front capital equipment costs and guarantee performance and availability service levels.

Learn More

If you have all-HDD systems coming up for technology refresh and are interested in better understanding the cost savings available to you in a move to the Pure//E family of all-flash systems, we'd like to meet to discuss the business value we can deliver.

1 | Source: Statista, 2023 (<https://www.statista.com/statistics/398951/global-shipment-figures-for-hard-disk-drives/>)

2 | Source: BackBlaze, November 2022 (<https://www.backblaze.com/blog/hard-drive-cost-per-gigabyte/>)

3 | Source: BackBlaze HDD Annualized Failure Rates Study (<https://www.backblaze.com/blog/backblaze-drive-stats-for-q1-2023/>)

purestorage.com

800.379.PURE

