PURESTORAGE®
Uncomplicate Data Storage, Forever

BUSINESS WHITE PAPER

FlashBlade//E: The Cost-Effective Alternative for HDD-based Secondary Storage Workloads

Learn how Pure Storage matches all-HDD systems on acquisition costs and then drives an up to 80% lower TCO.

Contents

Introduction

As enterprises move through digital transformation, they are capturing, storing, protecting, and analyzing an increasing amount of data. Data growth rates for many enterprises exceed 30% per year, which means that many of them are already managing multi-petabyte (PB) data sets and considering how they will best manage tens of PBs in the coming years. The roughly 20% of all enterprise data that is used for latency-sensitive, mission-critical, primary workloads has largely already migrated to all-flash, but most of the data used for secondary storage workloads that are less latency-sensitive but much more cost- and capacity-sensitive still reside on hard disk drives (HDDs). These workloads are primarily made up of unstructured (i.e. file- and object-based) data.

Historically, large, secondary data sets have been hosted on scale-out architectures due to the need for easy scalability across a wide range. Due to HDD's historically lower cost of raw capacity (relative to flash), most enterprises have chosen to host these large, growing workloads on HDDs in the past. Recent industry developments, however, should prompt information technology (IT) organizations to question this strategy:

- Flash $/GB costs are going down at roughly 20% per year while HDD $/GB costs are dropping at only 2-3% per year. As this gap narrows, flash becomes cost-effective for more workloads.
- A Pure Storage flash device (called a "blade") can easily deliver thousands of times the performance of an HDD [depending on whether you're looking at latency, I/O operations per second (IOPS) or throughput]. Many secondary storage workloads don't need much performance, but flash capabilities deliver needed performance with far fewer devices.
- Flash device capacities have outstripped the capacities of the largest HDDs. While HDD vendors are struggling to cost-effectively deliver 20TB capacities, Pure Storage® is already shipping a 48TB flash device and will be roughly tripling that capacity within the next 18 months in the same footprint.
- Based on these performance and capacity trends, a single flash device can replace multiple HDDs in systems sized to meet a given performance and capacity requirement. These flash devices require far less supporting hardware infrastructure, use far less energy, and take up far less floor space in data centers. With large multi-PB data sets, rising energy prices, and looming power grid limitations, environmental, social, and governance (ESG) considerations are becoming important criteria in IT infrastructure purchases, an evolution which clearly favors flash.

Based on these trends, it is time to consider whether your secondary workloads can be more economically hosted on all-flash storage. Pure Storage helped lead the displacement of all-HDD storage for primary workloads over the last decade, and we believe that there are compelling reasons for moving many secondary storage workloads to all-flash now.

All-flash Secondary Storage from Pure Storage

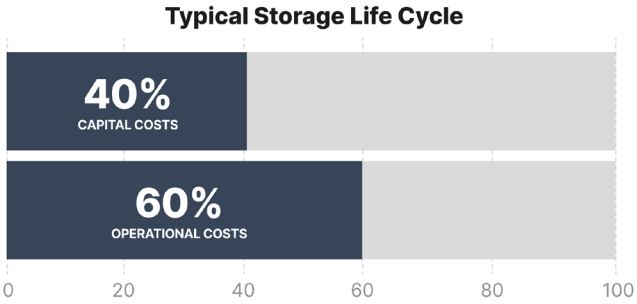

FIGURE 1 Capital vs operational costs during the typical storage life cycle

Pure Storage, a $3 billion vendor of enterprise storage solutions with over 11,000 customers, pioneered the concept of all-flash storage for enterprise workloads back in 2012. FlashBlade//E™, our new scale-out all-flash array (AFA), comes with industry-leading, storage infrastructure efficiency and ease of management. FlashBlade//E matches the acquisition cost of equivalent all-HDD systems configured for secondary storage workloads in the multi-PB range but comes with a reduced operational cost of as much as 80%. Enterprises can replace their all-HDD storage systems with FlashBlade//E on technology refresh, or they can buy FlashBlade//E as a service through Pure Storage Evergreen//One™ to get to that same low-cost operational model without any upfront capital expense.

Capital acquisition costs for storage infrastructure make up roughly 35-40% of the overall total cost of ownership (TCO) over the life of arrays, with operating costs making up the remaining 60-65%. Given this split, it is clear that focusing on lower operational costs has the biggest impact on overall TCO. But the value of FlashBlade//E for HDD-based secondary storage workloads does not depend solely on operational cost benefits to drive a lower TCO. It also depends on our high infrastructure efficiency to match the acquisition cost of all-HDD systems, allowing the economic benefits of switching to FlashBlade//E to begin immediately.

As the figure to the left demonstrates, lowering operational costs is significantly more important than lowering capital costs in reducing the overall TCO of storage infrastructure.

Validation of the 80% lower operational costs claim comes from Meta, the parent company of Facebook.In early 2022, Meta announced that they had awarded an exabyte scale storage contract for their Artificial Intelligence Research SuperCluster (AI RSC) to Pure Storage. Meta had worked for over two years to build a system from scratch using a software-defined and commodity server-based storage hardware approach. Meta concluded that they couldn't build a system that met their power budget for initial deployment (which was over 100PB), let alone accommodate the predicted expansion over the life of the project. After considering storage systems from several enterprise storage vendors, Meta selected Pure Storage, citing our ability to drive an 80% lower TCO than other vendors—due primarily to our infrastructure efficiency and ease of use advantages. These advantages were driven by smaller kit, lower energy and floor space consumption, and lower administrative and maintenance costs.

The Benefits of All-flash for HDD-based Secondary Storage Workloads

Few storage managers would quibble with the claim that if flash cost the same as HDDs on a $/GB basis, flash would quickly replace HDDs almost everywhere. With their higher performance and capacity, flash devices need far fewer components to meet a given performance and capacity requirement, require less supporting infrastructure (controllers, enclosures, power supplies, fans, network ports, switching infrastructure, cables), draw less power, and take up less floor space. Flash storage devices are far more reliable than HDDs, with fewer failures and less time spent replacing failed devices. Flash devices deliver more consistent performance under load, enabling denser consolidation of mixed workloads. They require far less effort to tune as workloads grow and evolve, lowering admin costs. And flash devices can handle data ingest and/or movement requirements with far fewer devices than all-HDD systems, another factor which helps to lower hardware costs and reduce energy consumption.

On the performance front, all-HDD systems handling secondary workloads will generally provide latencies in the tens to hundreds of milliseconds range, with throughput determined by data distribution across devices, workload access patterns, and the bandwidth of individual storage nodes. FlashBlade//E matches the fastest HDD response times, but across a narrower spread (10-20 milliseconds) and with an ability to support greater load. FlashBlade//E provides less volatile, more predictable performance while throughput is generally better than all-HDD systems, despite the fact that a FlashBlade//E system will be built from far fewer storage devices.

Higher flash performance was a key factor in leading to a much lower TCO when replacing HDDs for primary workloads. With FlashBlade//E, simpler scalability, easier management at scale, better infrastructure efficiencies and lower energy consumption drive a lower TCO for secondary workloads such as:

- Backup and disaster recovery
- Big data and log analytics
- PACs in healthcare
- Electronic design automation
- Content repositories and active archives

Let's take a closer look at the economics of FlashBlade//E.

Comparing TCO between FlashBlade//E and Scale-out All-HDD Systems

Conventional scale-out, all-HDD systems tend to be built around a software-defined design that uses commodity off-the-shelf server-based storage hardware and commodity off-the-shelf HDDs. Node-level building blocks, which include both storage compute and capacity, are clustered together using a dedicated Ethernet network, while storage software combines these individual nodes into what looks like a single system. To expand a system, administrators simply add more nodes. Commodity-based all-HDD systems have lower component costs, but need far more components, don't use raw storage capacity very efficiently, and have a high administrative overhead.

Usable not raw capacity is a key metric when configuring systems for production use. The efficiency with which raw storage capacity is converted into usable capacity is based on several factors. Before storage capacity is used to store data, an admin will format devices, deploy on-disk data protection algorithms, and configure spare capacity to minimize the impact of device failures.

On-disk data protection algorithms create redundant copies of data using schemes like RAID, erasure coding, and/or replication—each of which have different impacts in terms of capacity overhead. These algorithms employ data redundancy to ensure data integrity and availability if a storage device fails. That data redundancy, however, imposes a capacity overhead which adds cost. The efficient on-disk data protection algorithm in FlashBlade//E provides the same level of resiliency as all-HDD systems but requires significantly less capacity overhead, lowering costs by as much as 10-20%. On a 20PB configuration, that means 1-2PB less raw storage capacity that must be purchased, leading to large savings.

There are two additional considerations which lower the usable capacity of all-HDD systems: capacity over-provisioning and practical capacity utilization.

- **Capacity over-provisioning.** Due to their mechanical nature, HDDs exhibit very low IOPS/TB compared to flash storage devices, particularly with more random read-oriented workloads. An HDD's IOPS/TB yields decrease as its capacity increases. Because of this, system administrators need to very carefully take both performance and capacity requirements into account when determining the size and number of HDDs to purchase. In HDD-based systems, if capacity requirements are met but performance requirements are not, additional HDDs must be purchased to provide the needed IOPS. This results in capacity over-provisioning that directly increases the cost of a system.

  All-HDD system vendors offer a variety of HDD sizes (typically ranging from 1TB to 20TB), so customers can pick the size that best balances performance vs capacity. Keep in mind, however, that if you select smaller device sizes to better meet latency and/or data sharing requirements, you will end up having to buy more devices, again directly increasing your costs.
- **Practical capacity utilization**. Depending on the performance requirements, HDD vendors recommend only filling their devices 60-80% full. This is mainly because of performance concerns that kick in with access times when HDDs become too full. This means that a 2TB device will not actually contribute 2TB of raw capacity but somewhere between 1.2TB and 1.6TB, leading to the need to purchase additional devices to hit the needed usable capacity target.

The combination of these factors reduces the usable capacity relative to the raw capacity, and for scale-out HDD systems based on commodity hardware generally only 50% to 60% of the raw capacity of a system will be usable (assuming dual-parity RAID protection). It is far lower if replicas are used for data protection.

It's also important to note that the much larger storage device count required for all-HDD systems leads to needing additional infrastructure like controllers, enclosures, power supplies, fans, network ports, switching infrastructure, and cables. The additional network ports required by all-HDD systems drive higher CPU overhead to handle the larger network size and require more switching infrastructure, which in turn drives the need for more controller cores, DRAM, network cards, cabling, and potentially additional switches. The complexity of larger systems makes them harder to deploy, manage, scale, and upgrade.

FlashBlade//E: Cost-effective Scale-out All-flash Storage for Secondary Storage Workloads

FlashBlade//E is a scale-out, unified fast file and object storage system that users can scale by adding blades instead of full x86-based server nodes. Blades can add a configurable amount of storage compute and/or capacity depending on which blade type is purchased, giving users the opportunity to scale storage compute and capacity independently. Blades are designed to provide significantly more parallel access to media than HDDs, providing up to thousands of times as many IOPS at the device level. A FlashBlade//E chassis can hold up to 10 blades (or 1.96PB of raw capacity today). A fully configured FlashBlade//E system can support almost 8PB of raw capacity, taking up only 22U of rack space. FlashBlade//E is far more efficient in how it uses raw capacity relative to scale-out all-HDD systems based on the following characteristics:

- FlashBlade//E uses storage devices developed by Pure Storage that are far denser and support much higher parallel access to the media—which is based on quad-level cell (QLC) NAND flash—than HDDs. Because of the massive parallelism, even the densest blades (in terms of capacity) have IOPS/TB ratios that easily exceed HDDs by a thousand to one. This means that FlashBlade//E users generally don't need to consider performance as much as capacity when selecting device sizes. This allows users to select larger device sizes to achieve a lower $/GB without noticeably impacting performance for secondary storage workloads. This minimizes storage device count as well as supporting infrastructure requirements.
- These features drive very high infrastructure efficiency and reduce the raw capacity needed to meet a given performance and capacity requirement. The resulting smaller kit, which is reduced not only by number of devices but also by the need for far less supporting infrastructure, means a simpler, more reliable system that is easier to deploy, manage, scale, and upgrade and needs far less energy and floor space. In large multi-PB configurations using multi-parity data protection, FlashBlade//E can convert over 70% of its raw storage to usable capacity. That's 10-20% less raw capacity that has to be purchased relative to the most efficient scale-out all-HDD systems.
- The massive parallelism of FlashBlade//E also results in higher capacity utilization with short device rebuild times. While most HDDs will be capped at somewhere between 60% and 80% full, a FlashBlade//E blade contributes 95%+ of its raw capacity towards usable capacity without performance degradation.
- To provide on-disk data protection, FlashBlade//E uses a multi-parity protection scheme that can sustain up to three simultaneous storage device failures without impacting data integrity or availability. The data is distributed across a larger number of storage devices compared to HDD-based systems, minimizing the capacity overhead relative to HDD-based designs and reducing not only the number of storage devices required but the supporting infrastructure as well. Those factors all lower the cost of the system but also reduce its energy and floor space consumption.
- FlashBlade//E also supports compression which can provide up to a 2:1 data reduction ratio for some workloads. While this improves storage efficiency, most of the scale-out all-HDD systems support compression as well, so Pure Storage views it as a baseline capability rather than a differentiator for FlashBlade//E.

Due to availability considerations, some enterprises have established "blast radius" guidelines to minimize the impact of catastrophic failures that may occur at the enclosure, system, or data center level. Blast radius guidelines are often determined by how long it takes to restore data to a failure domain. Because of the much higher flash bandwidth on a per storage device basis, FlashBlade//E can restore data far faster than all-HDD based systems, minimizing both performance impacts and data risks in the wake of failures. Replication at both the file and object level is included, along with the base purchase of any FlashBlade//E system, providing resilience options that can address enclosure, system and/or data center-wide failures. The FlashBlade//E system is built on a platform with a history/legacy of 99.9999% uptime.

NOTE: *Customers with secondary storage workloads that are better suited to scale-up designs should talk to Pure Storage about the FlashArray//C, a scale-up QLC-based platform that supports block- and file-based access.*

Competitive vendors that use storage devices obtained from OEM suppliers rarely update the HDD firmware during the life of a drive, reducing opportunities to improve the efficiency of the devices over their life cycles. Because we manufacture our own storage devices, we frequently update the firmware over time to improve efficiency,reliability, and energy consumption.

NOTE: *Pure Storage can perform device firmware updates non-disruptively, a claim that not all of our competitors can even make.*

Purity, the storage operating system of FlashBlade//E, and our firmware operate far more efficiently than these components do in all-HDD systems because they are optimized to work with only our storage devices. This means that a FlashBlade//E system needs fewer compute and network (as well as capacity) resources to hit performance and capacity targets.

FlashBlade//E runs cooler than comparable scale-out all-HDD systems, due both to the lower heat radiation of flash relative to HDDs and the smaller kit size. Lower power consumption is not the only result of this; cooler running storage will also have less of a heat impact on other nearby IT infrastructure.

It's important to note that planned FlashBlade//E enhancements will noticeably lower the $/GB cost of usable capacity in deployed systems. Enhancements to the on-disk data protection algorithms will be available via software features accessible to all FlashBlade//E users at no extra charge. These enhancements will lower the capacity overhead without reducing resiliency, drive more aggressive data reduction ratios, and reduce energy consumption requirements per blade. Users can also expect an increase in the global namespace size and support for more expansion cabinets, extending the storage capacity under management for a single system. Denser flash media will more than triple the storage density (as measured in TB/U) within 18 months, reducing the energy and floor space consumption to hit a given performance and capacity target. And as always, all new capabilities can be non-disruptively integrated into deployed systems while preserving existing investment. What FlashBlade//E offers with the 48TB flash devices is compelling in many ways, but it will only get better with near-term enhancements, further outpacing the capabilities of scale-out, all-HDD systems for secondary storage.

We looked at several competitive, scale-out all-HDD systems configured to support 11PB of usable capacity. The $/GB of raw capacity, which takes into account not just the number of storage devices but also all the supporting infrastructure, for these competitor systems varies between $0.25 and $0.30 per GB (assuming a $/GB cost for nearline 7200 RPM HDDs of $0.015). The $/GB for usable capacity for the three competitive configurations ranges between $0.33/GB and $0.54/GB. FlashBlade//E weighs in at $0.20/GB for the raw capacity (and $0.29/GB usable) for the 11PB configuration. Figure 2 shows a comparison of the TCO of each of these options that shows the split between acquisition cost (CAPEX) and operational cost (OPEX) over a 6-year life cycle. FlashBlade//E drives a noticeably lower TCO relative to competitive all-HDD offerings over a 6-year life cycle.

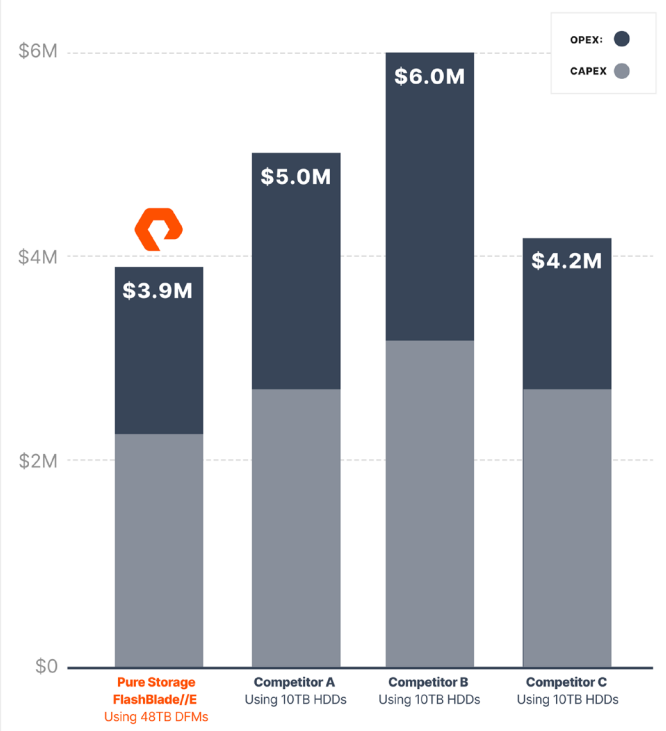

FIGURE 2 FlashBlade//E and Competitor All-HDD TCO Comparison

NOTE: *For the purposes of this comparison we are assuming no benefit from data reduction and a single forklift upgrade (complete hardware re-buy) on the part of the competitive systems in year 5. Due to its Evergreen® Storage architecture FlashBlade//E requires no forklift upgrade to move to later generations of storage technology but we are assuming two controller upgrades during the 6-year life cycle. The calculations also include 7×24 support during the 6-year life cycle for all the configurations.*

Validated Customer Experience

Pure Storage is the only vendor that tracks customer experience and publishes an independently validated metric (Net Promoter Score or NPS) that gauges it. We have an industry-leading NPS in the low 80s, while none of our competitors even publish an independently validated score. We have been tracking and publishing this number since 2015.

NOTE: *NPS is an independent metric, popularized by Bain and Company and used across 220 different industries, that ranks companies on their customers' response to a single question: "how likely are you to recommend the vendor to your colleagues?"*

Our NPS reflects the positive experience and strong loyalty we generate with our customers across all aspects of the storage life cycle:

- An all-inclusive base package that includes all existing software plus all future software enhancements at no additional charge. This includes features that other vendors either charge for like host multi-pathing, synchronous replication, and AIOps (delivered on FlashBlade//E through Pure1®), or do not have at all like non-disruptive storage device firmware upgrades
- System architecture that maximizes ease of management and enables non-disruptive, multi-generational technology upgrades in-place across a storage life cycle that can exceed 10 years.
- Guarantees on properly sized systems, customer satisfaction, flash media endurance, and flat and fair maintenance over the life of systems, with other programs that offer AI-driven monitoring and management, 7×24 worldwide support coverage, included controller upgrades across technology generations every three years, and options to implement more frequent multi-generational controller and media upgrades while preserving existing investment through guaranteed trade-in credits

Many scale-out storage systems can integrate newer technologies into an existing cluster by adding newer nodes, providing a form of multi-generational technology upgrade that extends the storage life cycle beyond the three to five years that is typical for most multi-controller arrays. Through our Evergreen Storage program, however, FlashBlade//E supports in-place multi-generational technology upgrades with guaranteed trade-in credits that keep systems operating at peak efficiency throughout their useful life. Our industry-leading NPS is based in large part on the positive customer experiences we generate through Evergreen Storage, a comprehensive effort that leverages architectural features and program aspects that pleasantly surprises seasoned customers exasperated with their storage experiences with other vendors.

Learn More

If this discussion has piqued your desire to learn more about how you can replace your HDD-based secondary storage systems with an all-flash solution from Pure Storage, please contact us.

PS2409-en-01 03/23