

WHITE PAPER

TRANSFORM HIGHER EDUCATION

SIX PROVEN USE CASES FOR UPGRADING
YOUR LEGACY STORAGE TO ALL-FLASH

TABLE OF CONTENTS

- INTRODUCTION** 3
- THE PROBLEM WITH TRADITIONAL STORAGE** 3
- A QUICK FLASH STORAGE PRIMER** 4
- THE PURE STORAGE DIFFERENCE** 4
- GETTING EDUCATED ON FLASH STORAGE** 5
 - Use Case #1: Driving Virtual Desktop Infrastructure 5
 - Use Case #2: Running Research Labs 6
 - Use Case #3: Managing Big Data and Analytics 7
 - Use Case #4: Fueling Online Learning 7
 - Use Case #5: Boosting Application Performance 8
 - Use Case #6: Powering Private and Hybrid Clouds 8
- EVALUATING ALL-FLASH OFFERINGS** 10
- CONCLUSION: A BREAKTHROUGH TECHNOLOGY FOR EDUCATION** 11

INTRODUCTION

Education in the 21st century is continuously undergoing transformation. Learning and instruction have escaped the boundaries of the lecture hall to become mobile, online, flipped, and active. Colleges and universities must compete for students, and are increasingly focused on student retention. Technology is so embedded into the daily activities of the institution that IT leaders are permanently seated alongside the leaders of academic and administrative units in strategic meetings.

Ubiquitous Wi-Fi, greater use of video, and an influx of data – among other technological advances – put intense pressure on underlying infrastructure. But as institutions work to modernize their infrastructures, one vital aspect is often overlooked: storage.

This paper looks at modern all-flash storage – outlining six use cases for its application in higher education and providing advice for evaluating flash storage solutions.

THE PROBLEM WITH TRADITIONAL STORAGE

While CPUs are measured in nanoseconds and servers move at gigahertz speed, mechanical hard drives are stuck in the realm of millisecond performance – akin to pitting the family Buick against Formula One racecars. There are three main issues when it comes to hard disk and server-based storage:

- **COMPLEXITY** Legacy storage technology requires hands-on management, tuning, and maintenance; ignoring this work can result in instability that creates disruptions for students, faculty, and staff.
- **LATENCY** Mechanical disk drives are slow, with input and output (I/O) latencies typically ranging from 10 to 20ms. Slow I/O means slow applications, which can negatively impact student experience.
- **WASTE** Hard disk-based storage is less efficient than all-flash. Hard disks are not suited for data reduction, which means you have to buy more raw storage for the same amount of data, and that takes up valuable data center space. Moreover, the number of application workloads that can be shared on a single storage system is limited due to performance bottlenecks.

Traditional storage has become the weak link in the data center. From the perspective of input/output-operations-per-second (IOPS), disk-based storage hasn't kept pace with compute or network IOPS – putting the entire data center stack out of balance. This will remain true until IT shifts its storage technology from traditional hard disk drives to the exponentially faster, low sub-millisecond latency solid-state disk (SSD), commonly known as flash memory.

All-flash storage delivers the same dramatic gains in service delivery previously brought about through virtualization and gigabit networks.

A QUICK FLASH STORAGE PRIMER

If you use a smartphone, you're using flash memory for storage. The same is true for your tablet, camera, and many laptops. Otherwise known as solid-state memory, flash is a medium that can be erased and reprogrammed electronically. Flash storage is becoming increasingly common in enterprise storage systems because it's incredibly fast. In the past, flash arrays were expensive and tended to be dedicated to specific, high-demand applications such as transactional databases. Organizations have migrated to all-flash arrays in mixed-load scenarios because:

- The performance of all-flash arrays exceeds that of hard disk drive (HDD)-based storage many times over
- There is better integration between server, fabric, and storage for sub-millisecond IO access
- They can increase efficiency, with built-in deduplication and compression capabilities that reduce the hardware and power footprint in the data center
- They can implement nondisruptive upgrades and achieve high availability
- There is enough capacity to handle petabytes of data

All-flash arrays are growing by nearly 30 percent each year, according to IDC.¹ Eventually, IDC predicts all-flash arrays will replace legacy HDD storage systems.

THE PURE STORAGE DIFFERENCE

Pure Storage's flagship product line, FlashArray, is an ideal solution to speed up application performance, ranging from the entry-level //M10, with a 30-terabyte effective capacity, to the //X70, with a 1.1-petabyte capacity. Each incorporates Purity, the Pure Storage operating environment architected from day one to manage flash.

For converged infrastructure and private cloud installations, Pure's FlashStack™ delivers pre-integration with popular private cloud platforms, including Cisco UCS Director, VMware vRealize, and the open source OpenStack.

Pure's FlashBlade™ storage system – intended for computing scenarios that are massively parallel, graphically intensive, or both – can handle gigabytes of bandwidth, petabytes of capacity, tens of billions of objects and files, and quick “scale-out” through the addition of more blades.

¹ http://www.purestorage.com/content/dam/purestorage/pdf/whitepapers/ps_g_flash-storage-buyers-guide_03.pdf

GETTING EDUCATED ON FLASH STORAGE

The following sections explore the impact of all-flash storage in six higher education use cases: virtual desktop infrastructure (VDI), research labs, big data and analytics, online learning, application performance, and private and hybrid cloud. It also highlights examples of successful implementations and discusses some of the features and functions institutions should consider as they evaluate all-flash solutions for their data center environment.

Use Case #1: Driving Virtual Desktop Infrastructure

Institutions have long recognized the advantages of using virtual desktop infrastructure, in which applications (and sometimes the desktop operating system) are hosted from a centralized server in a data center. Among the most popular scenarios:

- Running computer labs where there are numerous machines that need to be re-imaged and updated frequently
- Delivering professional- or industry-grade software programs to student computers that may not otherwise support such resource-intensive applications
- Setting up “pop-up” computer networks to address seasonal demand or special events, such as registration or orientation week

When the University of Texas at Dallas School of Engineering transitioned to a 10-gigabit Ethernet network to support virtual delivery of a growing number of essential engineering applications to students, IT leaders assumed they would just add more storage as needed. Unexpected latency problems stalled the initiative, however.

IT Manager John McConnell surveyed his computing infrastructure and concluded the storage model no longer served the school’s needs. Adding a flash cache card in each host improved speed but not scaling; when a new server was added, a flash card needed to be added too. It was only when the school used an all-flash storage array that the virtualized environment could meet growing needs.

“We literally have zero performance issues,” McConnell says.

The University of Portland in Oregon also upgraded to all-flash storage to sustain 150 desktop computers in student kiosks and computer labs. Infrastructure Manager Joey Houck says the move has given new life to older computing gear and reduced the university’s computer refresh budget by 20 percent. Being able to deliver software to students wherever they are — a dorm room or an overseas location — has allowed the campus to get by with a third as many computer labs. In some cases, space formerly allocated for computer labs has been repurposed for classroom usage.

Houck cites better end-user support as an outcome of the move to a flash memory system. “It just runs,” he says. In addition, while IT staff would previously visit sites for hardware troubleshooting, that can now be done over a web browser, allowing the IT team to “take care of more people in a shorter amount of time.”

“We needed both high performance and scalability in our storage infrastructure.”

—
AUSTIN NOTHAFT

GRADUATE RESEARCHER
RISELAB, UNIVERSITY OF
CALIFORNIA, BERKELEY

Use Case #2: Running Research Labs

The modern research laboratory, with its massive data sets, places major demands on high-performance computing resources. The better suited the computing infrastructure is to research needs, the faster the results arrive.

For storage, research labs have typically turned to massive clustered disk solutions, which tend to be slow, unstable and unpredictable. But an all-flash array uses compression algorithms that provide more usable storage in less space and minimize latency. This is critical for research projects that demand a lot of bandwidth and require rapid IO for faster access times and greater reliability, while also accommodating more capacity with non-disruptive upgrades and minimal tuning.

For example, at the University of California, Berkeley, genomic researchers are working to identify genetic markers that indicate the presence or absence of disease. Such research uses massive amounts of data storage capacity — a single sample of a person’s DNA is about 300GB of raw data, and the latest DNA sequencing research projects involve 50,000 to 100,000 participants, each of whose DNA will be sampled several times over the course of a research study. The database for a single project can grow into the petabytes.

At RISELab, the university’s new lab for tackling data-intensive research projects, a team of researchers is developing ADAM, an open-source, high-performance, distributed library for genomic analysis that is the basis for DNA sequencing work. Early implementations used spinning-disk storage technology and the Hadoop Distributed File System (HDFS), but that approach couldn’t keep pace with ever-increasing demands for high-performance storage.

“Our data storage needs are much greater than our compute needs,” says Austin Nothaft, a graduate student involved in the project. “We needed both high performance and scalability in our storage infrastructure.”

The lab adopted FlashBlade, a Pure Storage array built specifically for multi-petabyte, data-intensive visualization work in massively parallel environments.

“FlashBlade allows us to maintain high throughput while horizontally scaling bandwidth simply by adding blades,” Nothaft notes. “And, we can decouple the storage capacity from the number of compute nodes we have in our cluster.”

Use Case #3: Managing Big Data and Analytics

Research data isn't the only type of "big data" in play on campuses. Information is one of the most valuable assets for an institution, and the resulting push for data-driven decision-making, and the growing use of the Internet of Things (IoT), is putting pressure on campus IT resources. FlashBlade is the industry's first cloud-era platform purpose-built for modern analytics. It is engineered to be massively parallel at its core, delivering a big data platform that's easy to use and fast.

When databases run faster, users can obtain their results more quickly. As faculty and staff work to improve student outcomes, more responsive analytics can provide the data they need to intervene earlier. The result is improved student completion rates and better resource allocation.

Davenport University moved its mission-critical applications and data onto all-flash storage arrays. As a result, the Michigan institution's Blackboard Analytics and Banner enterprise resource planning (ERP) systems can keep pace with user expectations. The storage migration has enabled the program, which is underpinned by data streaming in from Oracle and Microsoft SQL Server databases, to deliver analytics performance in near-real time.

"I found that for less than the cost of the hardware maintenance on our legacy system, I could buy a FlashStack implementation and three years of maintenance."

—
RAY SMITH

ASSISTANT EXECUTIVE DIRECTOR FOR
TECHNOLOGY
MISSISSIPPI COMMUNITY COLLEGE BOARD

Use Case #4: Fueling Online Learning

Online learning continues to gain momentum. A 2017 report by Digital Learning Compass found the number of students at the undergraduate and graduate level who took online courses in 2015 totaled more than six million, representing a quarter of all higher education enrollments for that year.²

But online learning places huge demands on an institution's storage infrastructure for performance, capacity, and scale. Courses are more sophisticated, and the use of digital content – curriculum, video, and virtual labs – is growing exponentially.

A storage solution in an online learning environment must scale quickly and easily. The use of all-flash storage empowers IT to stay on top of service level guarantees and pivot as infrastructure requirements expand to accommodate program growth.

Mississippi Virtual Community College is a consortium of 15 state schools that offer online courses. Currently, the program supports 60,000 enrollments each year. As the program and other on-campus initiatives grew, the IT department at the Mississippi Community College Board (MCCB), which provides centralized IT services for the state's community colleges, looked for a solution to replace its existing converged infrastructure.

The IT department chose FlashStack – Pure Storage all-flash arrays combined with VMware virtualization software and Cisco UCS servers – the latter of which was already in place in MCCB's main data center and its disaster recovery site.

² <http://digitalllearningcompass.org/>

“I found that for less than the cost of the hardware maintenance on our legacy system, I could buy a FlashStack implementation and three years of maintenance,” says Ray Smith, assistant executive director for technology. “It was an easy decision for us.”

The result for students: faster performance. “It has made things a lot less stressful,” says Smith.

Use Case #5: Boosting Application Performance

IT infrastructure in higher education supports teaching and learning, but also management and operations.

Colleges often turn to hybrid storage solutions – spinning disks fronted by cache memory using flash – to boost the performance of certain applications. But that hierarchical philosophy doesn’t hold up well in a climate where student needs (and fees) drive IT priorities.

Institutions that move to all-flash storage can prioritize a wider breadth of applications to improve program performance for all users while also mitigating latency problems. Functionality built into the storage software stack brings other improvements as well. By exploiting built-in compression and deduplication, IT can increase data consolidation, which means the same storage capacity can fit more applications.

Up until just a few years ago, Davenport University’s IT department faced major hurdles trying to produce nightly reports from the Banner ERP software for the institution’s senior administrators. As one IT staffer points out, running the reports has “always been a headache because they take so long, and if something goes wrong, we miss the deadline.”

It was only after the university deployed an all-flash array that application performance finally improved. “When we switched over to the Pure Storage array, we watched user access time in the Banner environment fall to almost zero,” the staffer says.

Use Case #6: Powering Private and Hybrid Clouds

Institutional IT managers need systems that are simple, elastic, and available on-demand. They also need to accommodate continuous innovation. Cloud computing was designed to give users quick access to the computing resources they need, whether that’s for developing, testing, or rolling out new applications; delivering anytime, anywhere digital access; reducing costs; or simplifying administration and maintenance.

Yet fears of inadequate public cloud security and the desire to maintain oversight of critical applications, programs, and data are driving institutions to set up private and hybrid clouds for certain uses. A well-designed private cloud will often outperform a public cloud, if it’s built with best-of-breed equipment that is continually maintained. That’s a pricey proposition that often puts the concept of a private or hybrid cloud out of the reach of many IT budgets.

Here’s where Pure Storage’s unique Evergreen® Storage comes into play. This program was introduced in 2015 as a dramatic makeover of the standard storage company maintenance program. The idea: to allow a customer to buy

ISN'T FLASH STORAGE MORE EXPENSIVE?

With the total cost of ownership of all-flash now up to 50 percent less than the cost of high-performance spinning disk-based storage, it's an attractive solution to modernize storage infrastructure. In an economic impact study, analyst firm Forrester found an all-flash storage system could generate financial benefits in several areas, including: improved productivity, simplified deployment and management tasks, and cost avoidance for energy and legacy storage maintenance expenses. The bottom line: a return on investment (ROI) of 234 percent. As the report concluded: "A growing number of use cases have developed ... where the performance, resiliency, and ROI of flash make the decision to go all-flash easier."³ Eventually, IDC predicts all-flash arrays will replace legacy HDD storage systems.

storage once and keep it up to date forever with software and hardware upgrades that don't require re-buying TBs, migrating data, or disrupting the environment. This eliminates costly upgrade cycles and ensures optimal computing performance. The Evergreen Storage model has multiple aspects, but a key one is "Forever Flash," which has three provisions:

- **FLAT & FAIR** Once the institution purchases its first all-flash array from the company, maintenance pricing on a per-gigabyte basis is capped at that rate, which means it will stay flat or decrease over time for as long as the institution owns a Pure Storage system. Likewise, the maintenance and support pricing will never exceed the original contract price.
- **FREE EVERY THREE** Customers receive replacement controllers free of charge every three years; these span product generations and support the latest in storage technology.
- **FOREVER MAINTENANCE** If a component fails during the life of the array, including those facing wear-out conditions, Pure Storage will replace it. Oftentimes, the company will contact the customer proactively to replace the component before it fails.

The use of all-flash storage standalone or as part of a converged infrastructure – both of which are backed by the Evergreen Storage model – hits on the primary benefits of private and hybrid clouds. Institutions gain ease of maintenance and simpler management, reduce the cost of IT infrastructure, and can spin up new instances of applications and scale capacity as needed.

³ http://info.purestorage.com/HealthcareTEICaseStudy_TEI-Report.html

EVALUATING ALL-FLASH OFFERINGS

In developing the business case for the shift to an all-flash storage array, a big part of the argument is operational simplicity. Providing storage services is a complicated and time-consuming endeavor. A well-designed flash storage solution simplifies the computing environment by eliminating processes tied to defining storage pools, configuring RAID, and modifying settings for cache, tiering, and data reduction.

Simplified administration in Pure Storage products is backed by Pure1®, a web-based application (and mobile app) that provides a single dashboard to manage all of an institution's Pure Storage arrays. A new feature, Pure1 Meta™, facilitates "self-driving" storage. The arrays feed telemetry data back to Pure Storage, which continuously monitors cloud-connected arrays and applies machine learning and predictive intelligence to proactively resolve issues.

When evaluating all-flash offerings, an institution should also examine the data reduction technology built into the solution. As one analyst put it, "Data reduction may sound like a commodity but implementations are as diverse as the number of flash options within the market place."

When you're comparing solutions, consider how comprehensive the functionality is for data deduplication and compression. The first ensures that only unique data is stored on the array; the second reduces the capacity required to store unique data on the array.

In addition, evaluate the expansion path provided by the all-flash storage array. Enterprise storage is hardly ever a one-time purchase. Keeping pace with growing demands is a complicated equation. Pure's Evergreen Storage model cuts away much of the complexity as the demand for storage increases.

Likewise, resiliency is an aspect worth examining. Look for all-flash solutions with Tier 1 resiliency. That means six 9s of availability, nondisruptive upgrades, and the ability to tolerate disk or controller failure while maintaining full performance.

Finally, there's ease of installation. As numerous IT professionals have found, implementation of anything in the data center can be a multi-week, multi-person activity. That was the experience of the IT organization at the Mississippi Community College Board. But as Assistant Executive Director Smith attests, the entire installation of a Pure Storage FlashStack solution can be done by one individual in a fraction of the time the previous solution required. Deployment of FlashArray, Pure's converged solution, took only about an hour. As the technical director for one cloud consultancy expressed it, "I think of Pure Storage as the iPhone of storage, in the sense that its installation is almost plug-and-play."

CONCLUSION: A BREAKTHROUGH TECHNOLOGY FOR EDUCATION

Leading campuses have already discovered the benefits of implementing all-flash storage to support virtual desktop infrastructure, big data analytics, distance learning, research computing, and cloud usage (among other feats of technical prowess). By identifying and adopting the right solutions, these colleges and universities are getting the performance they want while simplifying operations, improving storage reliability, and reducing total cost of ownership. All-flash is a breakthrough technology in education, enabling institutions to bring their storage solutions into the 21st century to deliver the teaching, learning, and other services demanded by their community of students, faculty, staff, and administration.

Pure Storage (NYSE:PSTG) helps companies push the boundaries of what's possible. The company's all-flash based technology, combined with its customer-friendly business model, drives business and IT transformation with storage that is effortless, efficient, and evergreen. Pure Storage offers two flagship products – FlashArray//X, optimized for structured workloads, and FlashBlade, ideal as a platform for analytics – along with FlashStack, a converged infrastructure solution with components from Cisco and Pure Storage. With Pure's industry-leading Satmetrix-certified NPS score of 83.7, Pure customers are some of the happiest in the world, and include organizations of all sizes across an ever-expanding range of industries.

© 2018 Pure Storage, Inc. All rights reserved.

Pure Storage, Evergreen, FlashBlade, FlashStack, and the Pure Storage Logo are trademarks or registered trademarks of Pure Storage, Inc. in the U.S. and other countries. Other company, product, or service names may be trademarks or service marks of others.

The Pure Storage products described in this documentation are distributed under a license agreement restricting the use, copying, distribution, and decompilation/reverse engineering of the products. The Pure Storage products described in this documentation may only be used in accordance with the terms of the license agreement. No part of this documentation may be reproduced in any form by any means without prior written authorization from Pure Storage, Inc. and its licensors, if any. Pure Storage may make improvements and/or changes in the Pure Storage products and/or the programs described in this documentation at any time without notice.

THIS DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. PURE STORAGE SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

ps_wp11p_transform-higher-education-with-all-flash_ltr_01

SALES@PURESTORAGE.COM | 800-379-PURE | @PURESTORAGE