

Higher ed data storage just got smarter.

Discover the data engine for academic research.

Universities and research institutes have brilliant minds, talented data scientists, and powerful compute at the ready. What's often missing is an intelligent data storage platform capable of fully unleashing that potential.

We've designed the solution.

Small IT teams, multiple research groups, lean budgets, massive datasets, mixed workloads – these cease to be problems with Panasas as your data foundation. The PanFS® file system running on ActiveStor® systems powers some of higher education's most data-intensive research in bioinformatics, high energy physics, genomics, material sciences, machine learning (ML), and beyond.

Accelerate complex mixed workloads and AI initiatives



In addition to traditional HPC modeling and simulations, AI and ML are now being **deployed by researchers to drive advancements** in astrophysics, biology, climate science, medicine, and more.¹



The number of courses that teach students to build or deploy AI models on the undergraduate and graduate levels has increased in the last four years by **102.9% and 41.7%**, respectively.²

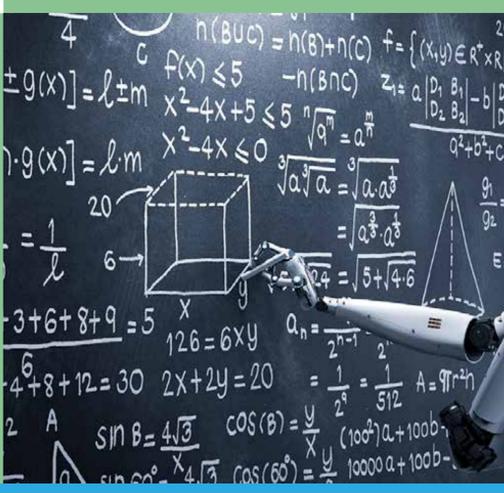
Challenge

Universities have seen a major increase in the quantity and complexity of their high performance computing (HPC) workloads, and researchers are adding cutting-edge artificial intelligence (AI) and ML technologies to advance their discovery pipelines. These types of mixed workloads are incredibly taxing – if not impossible – on legacy storage systems.

"Trying to do so many reads at the same time would bring other storage systems to their knees. There wasn't enough disk readability to handle that kind of workload in addition to the other tasks that execute on the cluster each day."

Solution

The PanFS file system has built-in Dynamic Data Acceleration (DDA), a proprietary software feature that automatically adapts to the changing and evolving small-file and mixed workloads in AI and HPC that other storage systems struggle to support. Your hard drives are never burdened by inefficient small-file and metadata writes, leaving them free to stream your large files, so you get consistently high performance, even for heavy concurrent and mixed workloads.



Consolidate with a scalable, intelligent, and cost-effective solution

Challenge

Research data volumes are increasing rapidly, but the academic funding pool certainly is not. Universities need adaptable and highly performant storage that can easily scale to support thousands of researchers working on hundreds of AI and HPC projects – without blowing the IT budget.

"When you're analyzing these quantities of unstructured data across this many teams, storage performance and capacity quickly become issues. We needed to centralize our infrastructure on an affordable solution that scaled efficiently."



Data stores can increase by **1 PB every 2 months** at top research universities.³



High-quality HPC resources **attract top talent** and position faculty members to win additional grants to further their research.⁴



Solution

Panasas ActiveStor systems utilize the most efficient combination of mixed media to optimize speed, capacity, and economics, so you never overpay for hardware that doesn't directly contribute to the performance levels you need. DDA cuts the complexity and costs of traditional tiered storage systems by providing a single all-hot tier with automatic capacity balancing for no hotspots.

Rest easy with multi-layer data security and protection



Security attacks against higher education more than doubled between 2019 and 2020, increasing from **6% to 15%**.⁵



Universities ranked data security as their **#2 liability concern** on a recent list of the top 10 risks affecting institutions.⁶

Challenge

Universities are particularly hot targets for cyberattacks, and data loss incidents can wipe out years of research progress in an instant. Aging data infrastructures, untrained user networks, and understaffed IT departments can be a recipe for disaster. With so many research teams needing to manage their own environments, security has become a top concern.

"We deal with mountains of sensitive data, and our chief worry is keeping them safe. We also foster a culture of openness in our research, so we need to strike a balance between authorized access and strong security frameworks."

Solution

To safeguard your data, we take a defense-in-depth approach: ACLs, SELinux support, and hardware-based encryption at rest prevent unauthorized access while the realm is both online and offline. The PanFS solution also uses patented per-file object encryption coding to ensure maximum protection, as well as built-in prevention and automated rapid failure recovery logic engineered to avoid corruption issues common in open-source solutions. Translation: We've got you covered.



Eliminate time-consuming storage complexity and downtime

Challenge

Open-source storage solutions require considerable staffing effort to set up and maintain. To make matters worse, they lack the high-grade reliability needed for 24/7/365 operations at most research institutes. As the landscape in academic research becomes increasingly competitive, unplanned downtime and frequent outages wreak havoc on project timetables.

"We don't have a huge IT team or weeks of time to integrate discrete systems, and we don't have a full-time storage admin that can keep things running 24/7. We needed a turnkey platform that just worked."



In a 2020 survey, nearly **50%** of the respondents reported the storage system failures once a month, with users coming to expect downtime as the norm in HPC storage.⁷



The most common reported challenges for HPC storage operations are **recruiting and hiring qualified staff**, followed by the time and cost needed to tune and optimize the systems.⁸



Solution

ActiveStor appliances roll in preconfigured, taking you from dock-to-data in less than a day. Our low-touch solutions require no deep technical skills or large dedicated staff for ongoing administration – a single part-time admin can easily manage up to 1,500 Panasas storage nodes. One Panasas customer has gone eight years without a single minute of unplanned downtime.

Panasas has a long history of partnering with the world's leading research institutes and universities to support their groundbreaking projects. We know that data storage is a bedrock for research, and we're committed to delivering the versatile, simple, and economical solutions that academic researchers need. You bring the vision and the talent – we'll handle the data storage.

Do your research

For more information, visit www.panasas.com/industries/academic-research

Request a meeting <https://www.panasas.com/book-a-free-consultation/>

Send an email to Richio Aikawa, Sr. Director, Strategic Marketing at raikawa@panasas.com.

¹ Nvidia. 2019. "How AI, Machine Learning Are Advancing Academic Research." <https://blogs.nvidia.com/blog/2019/03/27/how-ai-machine-learning-are-advancing-academic-research/>
² Stanford University Human-Centered Artificial Intelligence. 2021. "Artificial Intelligence Index Report." https://aiindex.stanford.edu/wp-content/uploads/2021/11/2021-AI-Index-Report_Master.pdf
³ EdTech Magazine. 2022. "The Technology Behind America's Top Research Universities." <https://edtechmagazine.com/higher/k12/article/2022/08/technology-behind-americas-top-research-universities-perform>
⁴ EdTech Magazine. 2019. "Universities Leverage High-Performance Computing for Multiple Returns on Investment." <https://edtechmagazine.com/higher/article/2019/04/universities-leverage-high-performance-computing-multiple-returns-investment>
⁵ Education Technology. 2020. "Ransomware Attacks on Education Sector More Than Doubled Since 2019." <https://edtechnology.co.uk/cybersecurity/ransomware-attacks-education-sector-doubled-since-2019/>
⁶ Collegis Education. 2021. "10 Concerning Stats About Cybersecurity in Higher Ed." <https://collegiseducation.com/news/technology/10-concerning-stats-about-cybersecurity-in-higher-ed/>
⁷ Hyperion Research (commissioned by Panasas). 2020. "New Study Details Importance of TCO for HPC Storage Buyers." <https://www.panasas.com/resources/hyperion-research-new-study-details-importance-of-tco-for-hpc-storage-buyers/>
⁸ Ibid.