

PANASAS®

# Converge, Consolidate, Control: Solving the Enterprise HPC and AI/ML Challenge



# Let's maximize your time to value. How would you like to proceed?

Start at Intro →

Go to Chapter 1 →

Go to Chapter 2 →

Go to Chapter 3 →

Skip to Finale →

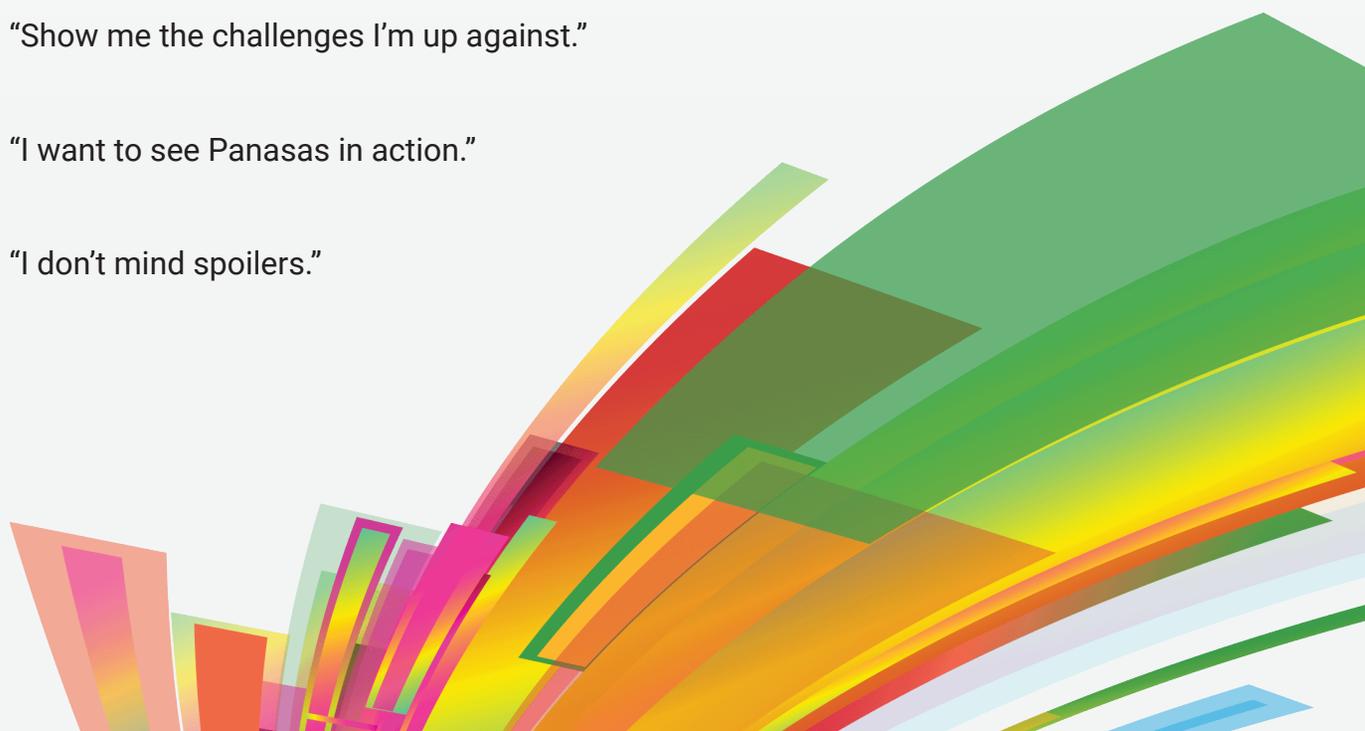
"I'm like [PanMove](#) and can move large volumes of data easily."

"Let's deepen my knowledge about the convergence at hand."

"Show me the challenges I'm up against."

"I want to see Panasas in action."

"I don't mind spoilers."



# Introduction

by Jeff Whitaker

If you're like me, you wish you could know ahead of time if reading an e-book will yield a meaningful return on investment. To this end, I hope you won't mind if I skip the traditional introduction and jump straight to the value of having reliable and simple-to-manage storage that automatically adapts to today's converging high-performance computing (HPC), artificial intelligence, and machine learning (AI/ML) workloads.

## Three critical points that not even senior executives can ignore

1

To summarize the convergence at hand and the thoughts of HPC IT teams everywhere: you cannot deploy HPC and AI/ML at scale without the right high-performance storage. Indeed, IDC recently reported that, while AI is growing in importance for enterprises, most AI projects fail because of inadequate or non-purpose-built infrastructure.<sup>1</sup>

2

Most HPC parallel file systems aren't known for their reliability. And in the face of converging workloads and rapidly growing volumes of data, compute requirements are increasing, and data storage failures cost organizations more than ever before. Given this, enterprises need a scalable, flexible, and reliable high-performance storage solution.

3

It's no secret that HPC left the "back room" and barged into the front room of enterprise IT in nearly all top industry verticals. This convergence of HPC and AI/ML in the enterprise presents a series of evolving challenges—categorized and covered ahead—that can result in disparate levels of performance, manageability, and security, as well as a lot of attrition.

Put (1) to (3) together and it's clear: The value of a simple alternative to unreliable HPC storage has turned the corner from "nice to have" to "need to thrive."

<sup>1</sup> <https://www.idc.com/getdoc.jsp?containerId=prUS48870422>

## Why is Panasas uniquely positioned to solve your enterprise HPC challenges?

We know better than anyone else what you *don't want*.

You don't want to pay a hefty price to meet the performance demands for running multiple applications with different I/O patterns and file sizes, or mixed workloads. You don't want to keep bringing in experts to manually tune multiple storage environments and deal with the management burdens associated with those systems. And, if possible, you don't want to worry about whether your systems meet data compliance and control requirements.



How do you get what you need? This e-book provides the full answer. Here's the abridged version:



At Panasas, we continue to add and enhance enterprise features in our storage solutions that rise to best-in-class needs, right where HPC and the enterprise converge. This is the intersection where HPC scalability and performance meet enterprise-grade manageability, reliability, security, and support. This is also where HPC, high-performance data analytics (HPDA), and AI/ML workloads converge for multiple business groups and diverse research groups at academic institutions and research labs. The common thread across all verticals is the need for a reliable, performant, and consolidated storage infrastructure.



## Your key takeaway from this e-book

You do not need to compromise simplicity, reliability, or flexibility for high performance. These terms get thrown around a lot in the storage space, so you may be thinking: “Sounds nice, but what does that mean for me and my team?” I’ll tell you:



**Simplicity** means easy to install, configure, use, monitor, and manage, starting as a design mandate (not an afterthought).



**Reliability** means both the storage system and the data on it are reliable through fault tolerance and self-healing, maximizing uptime.



**Flexibility** means running multiple-file-size workloads—from scratch storage to project files to home directories—concurrently, with thousands of users.

We’ve found that staying ahead of the curve in this highly competitive space is easy. It means having a portfolio of data storage solutions built on the world’s leading parallel file system: PanFS®—the data engine for innovation.

By the way, that’s not just a marketing slogan. It’s the design philosophy of the PanFS parallel file system, where the most innovative reliability algorithms—such as per-file erasure coding, quadruple redundant directory copies, checks on parity, and automatic background capacity balancing—keep the system tuned to the workloads automatically.

Ready to see how we make storage systems run like a fine-tuned race car? Then buckle up and enjoy the view of your next-generation high-performance storage platform.

### Jeff Whitaker

Vice President of Marketing and Product  
at Panasas

P.S. If you’re pressed for time, do yourself a favor : save this read for later and [check out our portfolio of simple, reliable, high-performance storage platforms now.](#)

# Clashing Designs and Conflicting Paradigms



## The evolution of HPC to Enterprise HPC

Discovering patterns of meaning in your data and extracting valuable insights from it in real time—that sums up the real-world value of emerging technologies. It's not a question of whether this value is being harnessed by enterprises.<sup>2</sup> The applications of HPC-driven innovations in AI and ML are visible in every market and valued by every enterprise.<sup>3</sup>

What's equally apparent is that the HPC processing technologies needed to run compute- and data-intensive applications and workloads have opened up a huge catalogue of pain points. These range from surprisingly complicated I/O patterns to security concerns—all of which we'll discuss ahead and can be remedied by [storage consolidation](#).

But let's not get ahead of ourselves and overlook the all-important question: "How did the enterprise converge with HPC and AI/ML and create the need for storage consolidation?" No response would be complete without mentioning that most large enterprises (and likely yours)

didn't grow up in the HPC space. The only reason they (and you) are reading about HPC is the same reason you've adopted the applications it supports. Namely, it is a *necessary means* to achieve a *desired end*.

The means are what you're doing now: continuously updating your infrastructure with "bleeding-edge" hardware and software. Using your updated infrastructure to create groundbreaking innovations that thrive is the desired end.

These innovations include seismic processing in the energy sector, enhancing the eloquence of chatbots in retail, smart farming in agriculture and livestock production, and the list goes on.

Yet, no matter the use case, capitalizing on the investment necessitated the formation of an HPC ecosystem. This wouldn't be so problematic were it not for the fact that IT teams—not C-Suites—are the ones tasked with delivering on the promise of HPC and AI/ML.



<sup>2</sup> Contrary to constantly repeated belief, the notion of HPC entering the enterprise is old enough to drive. Just pick up (and dust off) your January 2006 edition of InfoWorld and flip to page 17: "High-Performance Computing: Supercharging the Enterprise." The subtitle informs us: "Thanks to lower barriers to entry, compute clusters and grids are moving out of the labs and into the mainstream." It's safe to say the novelty of what we're discussing is about more than the consequences of ignoring the prominence of storage in the mainstreaming of HPC.

<sup>3</sup> The applications include everything from precision agriculture to chatbots to enhanced diagnostics.

In other words, you've inherited a complex infrastructure complete with unpredictable and dynamically evolving computing systems. Complex, unpredictable—the fact that HPC systems were traditionally deployed by researchers and developers is not an insignificant piece of IT trivia. The HPC space reflects the mentality of its original dedicated user group, and their concerns didn't include incongruent policies and SLAs or business continuity risk reduction. What made sense for those traditional users—the ones that weren't using shared HPC systems that could be accessed by various teams within an organization—doesn't make sense for you.

You need shared HPC systems, and you are very concerned about incongruent policies and SLAs, among other challenges that traditional HPC users never could have predicted. Let's take a closer look at them.

The HPC space reflects the mentality of its original dedicated user group, and their concerns didn't include incongruent policies and SLAs or business continuity risk reduction.



## The Enterprise HPC market at a glance

The HPC market is growing at a CAGR of 5.5% and is projected to exceed 64.66B (USD) by 2030.<sup>4</sup> The convergence of enterprise and HPC is driven by industries where complex problem solving in real time is both a necessity and a competitive advantage. Other uses of HPC systems include shortening product development cycles and running data-intensive industrial applications for specific workloads and projects. In every case, simplifying complex HPC and AI data environments with a simple, reliable, and flexible HPC storage solution represents a unique opportunity for these enterprises to stay ahead of the curve.



<sup>4</sup> <https://www.globenewswire.com/en/news-release/2022/04/04/2415844/0/en/High-Performance-Computing-Market-Size-to-Surpass-USD-64-65-Bn-by-2030.html>

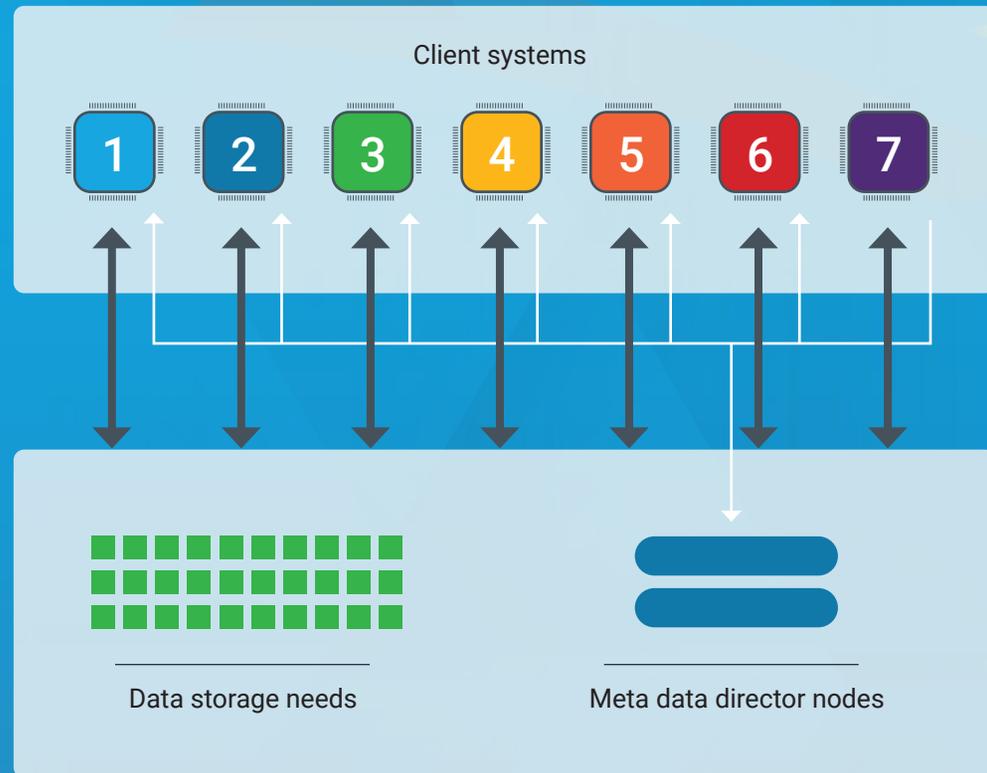
**HPC Storage:  
Afterthought?  
Think Again.**



## This doesn't compute:

### A "consolidated" list of Enterprise HPC and AI/ML challenges

You are now in the position of having to retrofit a traditional HPC environment to fit the demands of your not-so-traditional enterprise environment. And, in the quest to achieve the performance needed for your HPC and AI/ML workloads, you may have missed the importance of HPC storage. No need to get defensive about it: more often than not, storage is viewed as an afterthought. We know this because A) we're Panasas, and B) we have seen enough storage infrastructures to know why they're literally not up to speed. This gives us an intimate knowledge of the unique challenges you're facing. We've gathered the most pressing below. As you review them, keep this in mind: all the obstacles slowing down your time to insights and jacking up your TCO are temporary. We'll clear them up with our portfolio of tailored storage solutions for your HPC and AI/ML workloads.



## PanFS Parallel File System



1

## Sprawl at the center of it all

HPC systems are needed for certain business-critical, compute-intensive applications. You're running these applications to harness the power of HPC resources, but the clusters involved are tuned only for a few workloads. Given the variability in different HPC workloads, compute and storage silos begin to form in the infrastructure since custom-tuned environments are thought to be a necessary solution to this variability. And before long, you've got underused islands of HPC compute and storage clusters. The outcome is degraded performance, slower data access times, and a ton of administrative overhead.

2

## Compatibility issues

Parallelism is what allows organizations to analyze large datasets and turn mountains of data into actionable insights faster than ever before. Parallel frameworks allow processing across multiple nodes, but they also introduce compatibility issues. In lieu of a long (and unavoidably tedious) discussion about workload managers, it's sufficient to say that your IT teams are likely deploying a wide variety of separate clusters in an irregular and disparate fashion. They're also using exotic hardware and *ad hoc* solutions to "just make things work." Unfortunately, this usually worsens the problem and won't deliver the data sharing, scalability, or throughput you need.

3

## Massive volumes of data

Given the price and performance, the presence of cluster sprawl is unsurprising. Combining the computational horsepower of multiple computers makes it possible for multiple people to work on multiple projects using multiple data-intensive applications. However, the vast amount of data your teams are creating isn't always used to develop better products or to plant the seeds of the next breakthrough. More often, it's a reminder of how difficult it can be to scale storage capacity to handle an unprecedented growth in data (most of it unstructured).



4

## Data growing pains

It can be said without the slightest hint of hyperbole that it's never been so easy to generate staggering amounts of data. The problems organizations face in trying to manage exponential rates of [data growth aren't lacking in coverage](#). What's lacking is an efficient means of consolidating all the unstructured data that's accumulating in your business on a daily operational basis. Because when data grows, you're storing more data. And when you're storing complex HPC data, [the cost of storage outages becomes staggering](#).

5

## I/O bottlenecks

Next to the compute and networking aspects of HPC, storage is often left playing second fiddle in the design of an HPC system. Putting HPC storage in the back seat is the main reason so many storage systems are woefully mismatched to performance needs. New workloads introduce new I/O patterns. And when this happens, storage takes the blame for being a bottleneck. Of course, storage often isn't designed for the scale and performance that compute requires. But that's only because the right solution wasn't selected up front. As a result, the performance—along with the price tag—is much higher than it needs to be.



6

## Accelerated TCO

As data piles up unseen and high-priced CPU and GPU-accelerated compute nodes sit idly by, waiting to be fed reams of data, the question naturally arises: “Do we need to keep this mountain of data online at all times?” Likely not. But it’s hard to predict existing and future I/O requirements of HPC workloads that run the gamut from exceptionally large files to medium-sized files to countless small files. This unpredictability of I/O demands—coupled with the inability to efficiently handle mixed workloads—goes some ways toward explaining the prevalence of severely bloated CPU/GPU compute budgets.

7

## HPC insecurities

Shared HPC systems are exactly what they sound like. And it’s precisely their inherently “shared” nature that makes it difficult to enforce security requirements for application usage on these systems. Questions arise: “Does everybody who needs it have access? Do I have all the necessary firewalls and is everything separated?”<sup>5</sup> Figuring this out on the application and server side ignores what’s needed at the back end (storage). Meanwhile, meeting strict security requirements for enterprise-level HPC environments is hard enough. Thwarting increasingly sophisticated cyberattacks without degrading performance is even harder.

<sup>5</sup> This speaks to the security risks that are inherent to the containerization of HPC. Containers themselves cannot be isolated from the host operating system. They must use the underlying host OS kernel for security policies, provisioning control, user communications, and so on. So, if there’s a vulnerability in the host OS kernel, it will render all containers susceptible to security risks. Note the domino effect. A security flaw in a single container can (and will) compromise all other containers.

## Spoiler for the chapter ahead: Panasas is the panacea

Against the backdrop of these challenges, the ongoing struggle of life sciences firms presents itself as a teachable moment. They are suffering from a “lack of scalability in existing HPC systems, too much customisation [*sic*] and complexity, undesirable pricing models, as well as support issues and high support costs”—not to mention an “inability to scale storage capacity.”<sup>6</sup> Life sciences is a highly regulated industry, so we can add “regulatory compliance issues” to that long and growing list of innovation inhibitors.

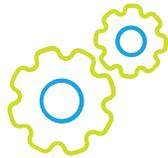
By now, we hope it's obvious that life sciences firms aren't the only ones facing an uphill battle in the attempt to harness the power of HPC and AI/ML at scale. Complex HPC environments transcend the lines that divide large enterprises into verticals. The pressure to innovate is universally felt.

That's why organizations across the globe are investing in the design and maintenance of hugely expensive HPC systems. Their intention is to seize the power and promise of the latest technologies. But, as it so often happens, that forward-looking investment ends up taking the shape of a bottleneck. All because the storage infrastructure that underlies these HPC-level aspirations is inflexible, expensive, unreliable, and hard to manage.

<sup>6</sup> <https://www.computerweekly.com/news/252521243/APAC-life-sciences-firms-face-HPC-challenges>

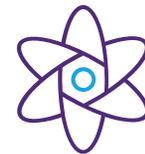


## Unique high-performance storage for industries



### Manufacturing

Manufacturers need simulation tools for multidisciplinary design exploration and engineering development of products in automotive, aerospace, consumer products, and semiconductors.



### Life Sciences

Innovations in genomics, bioinformatics, and other life sciences disciplines demand speed to store, retrieve, and analyze unprecedented volumes of information.



### Financial Services

Investment and trading firms demand fast and accurate quantitative analysis for exceptional algorithmic trading, risk management, compliance checking, pricing, and portfolio management.



### Energy

Oil and gas enterprises rely on the right tools to manage and analyze large data sets, quickly and accurately locate new natural resources, improve user productivity, and reduce overall project time.



### Academic Research

Institutions of higher learning must support a broad range of high-performance applications—and their massive datasets—across scientific, engineering, and humanities disciplines.



### Government

Government agencies and laboratories need the right tools to work with enormous datasets for initiatives in energy research, biomedical advances, climate science, economic modeling, and defense and homeland security.



### Media & Entertainment

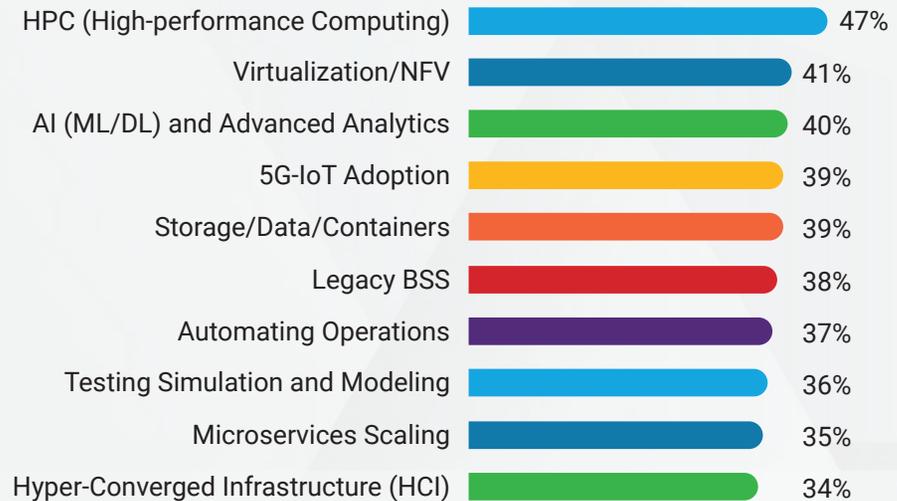
Media and entertainment organizations that focus on creative, data-intensive content need powerful processing horsepower and scalable storage to supercharge productivity and meet aggressive deadlines.

Whether this sounds painfully familiar or reads like a cautionary tale, it doesn't change the facts. What you need are data storage solutions that deliver optimized performance, enterprise-grade reliability, and infinite—yes, infinite—scalability for your HPC and AI/ML workloads. As for maximizing the value of those terabytes and petabytes generated by your high-intensity workloads, the answer is that a scale-out storage system delivers consistently high performance for whatever type of workload you can throw at it.

What happens when you don't have to trade manageability, usability, and reliability for performance? Glad you asked. You're about to see what Panasas makes possible with next-generation storage technologies.

## Overcoming the adoption barriers

Which applications are creating the most significant adoption challenges for your network or your users today?



Source: [Futurum research](#)

# Contributors to the Enterprise HPC and AI/ML challenges



## Treating storage as an afterthought

Infrastructure team determines the compute needed to solve an application requirement, and only asks “what kind of storage do we need?” after the fact.



## Performance bottlenecks

When the performance of the storage is dictated by the choice of infrastructure, the performance needed from the environment may be unachievable.



## Storage option overload

Traditional storage, storage inside servers, cloud-based storage—there are plenty of options. Yet the choice will determine the outcome and dictate the results (or lack thereof). Is there a right one?



## Various teams with various needs

There’s the application side of a team, the IT team running certain apps, and the research team doing image processing. How do you build a storage platform to support various teams with various demands?



## Security and compliance requirements

The infrastructure team needs to solve security and compliance requirements. There’s also the challenge that HPC systems need safeguards that don’t compromise performance.

# Laying the Storage Foundation for the Future of Innovation



Do all the challenges with enterprise HPC and AI/ML “converge” into ancient history once Panasas enters the picture? That’s our claim. As for the proof, that can be found in how the Minnesota Supercomputing Institute (MSI) at the University of Minnesota (UVM) is accelerating its research initiatives. **Note well:** Even if you find it hard to identify with researchers that sequence zebra mussel genomes, you’ll immediately recognize the business value of MSI’s results. In other words: with slight variations, this could be you.

## Case study in brief



### Customer

The Minnesota Supercomputing Institute at the University of Minnesota



### Scope

The new **\$7 million** Agate supercomputer cluster provides advanced research computing infrastructure to **900 lead researchers** and over **4,500 users**.



### Challenge

System selection, including balancing increasing performance and capacity needs, GPU use, and access demands: technology availability, funding, floor space, power, and cooling.



### Solution

The Agate heterogeneous high-performance Linux cluster of **770** AMD Milan 7763 64-core processors and **264** NVIDIA A100 Tensor Core GPUs, all connected with Mellanox HDR-100 InfiniBand network, with **10 PB of Panasas storage** to support various data-intensive research initiatives.

### Results

MSI is positioned at the frontier of advanced academic research as one of the top ten University-funded U.S. academic systems with **7 times the performance** of its preceding cluster.



## Advancing the most advanced research: **The Minnesota Supercomputing Institute**

Whether engineering autonomous solar-powered mowers to manage weeds or sequencing zebra mussel genomes to combat climate change, research teams at MSI depend on computing and data storage technologies to support their research projects. And there are a lot of them. Today, MSI supports over 900 principal investigator groups and 4,500 active users.

Because the research at UMN and MSI covers so many fields, their systems must support myriad use cases. And as data generation rates grow, so do research needs. Given this scope and variety, they need storage solutions that can handle distinctive high-performance data loads with overlapping I/O patterns (i.e., mixed workloads).

That's why MSI is committed to building out an integrated data infrastructure that not only keeps up with its performance and capacity needs but can also adapt and scale with it. But "adapt" and "scale" have taken on new meanings. What is needed now is nothing short of next-generation computing and storage technologies. And that's why UMN and MSI got to know Panasas.

## Meeting high-performance needs with a high-value partnership

As MSI upgraded its supercomputers over the years, one element of its system selection process has remained constant for over a decade: an enduring partnership with Panasas for multipurpose, reliable, and high-performance data storage.



# Panasas and MSI Partnership timeline



Panasas **ActiveStor**<sup>®</sup> became the MSI storage standard for the institute's "Itasca" cluster. With **1.28 petabytes (PB)** of usable ActiveStor capacity, this deployment supported up to **12,000** simultaneous clients.

ActiveStor<sup>®</sup>

2012



MSI upgrades its Panasas storage once again; this time for its "**Mangi**" cluster (an expansion of Mesabi).

"Mangi" cluster

2019



"Mesabi" cluster

MSI upgraded its ActiveStor deployment for its new "**Mesabi**" cluster. This increased its total capacity and throughput to serve growing volumes of workloads and applications.

2015



ActiveStor<sup>®</sup> Ultra

MSI upgrades to the next-generation Panasas **ActiveStor<sup>®</sup> Ultra** hardware, adding approximately **10 PB of new capacity** to drive new research initiatives.

2021-present day

"Agate"

This latest storage upgrade empowered MSI to support new research projects on its latest supercomputer, which is generating massive quantities of data at unprecedented rates. MSI calls it "**Agate.**"

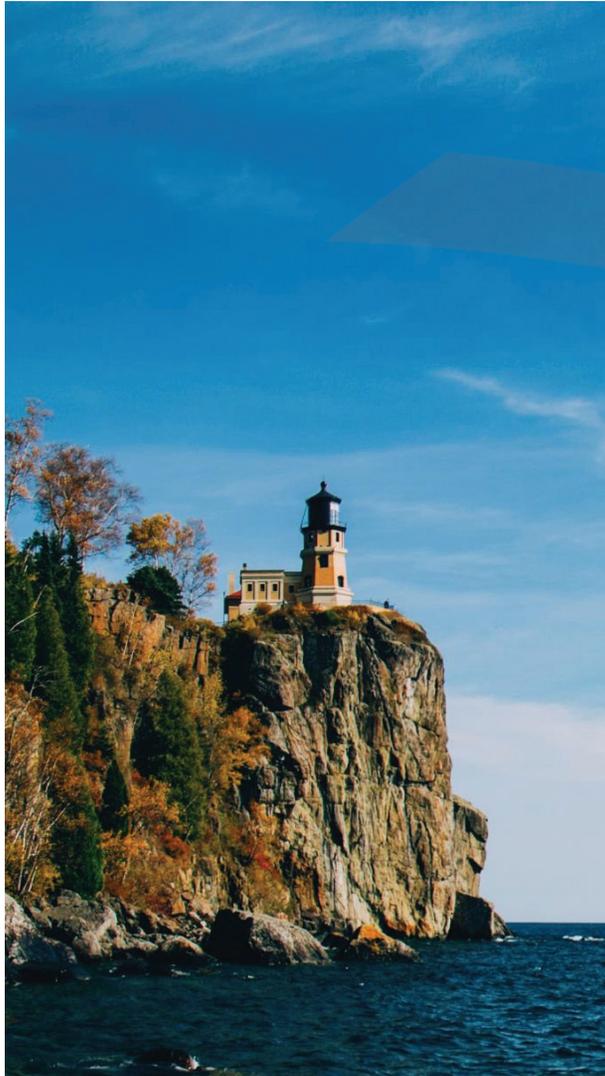
## Unlocking Agate of possibilities

Named after a Lake Superior rock formation and Minnesota's state gemstone, Agate started out as a vision of a beyond-cutting-edge supercomputer in 2018. At the time, MSI was still uncertain about two things: the build and the timelines surrounding CPUs.

This uncertainty prompted it to delay a major system upgrade. MSI opted instead for the smaller 2019 Mangi deployment. If nothing else, pushing Agate to 2021 allowed MSI to harness newer and more powerful CPUs.

But then came the concerns over acquisition and requirements—not to mention funding. Mangi's deployment alone consumed 20–30% of the budget. On the bright side, floor space had become less of an issue with systems becoming denser. Yet power and cooling costs took pride of place.

Other changes became prominent as well. Researchers were using larger datasets. New use cases drove up memory needs significantly. Emerging ML and molecular dynamics applications increased compute demand and so did interactive access to generate simulation input and view simulation output. GPU usage had nowhere to go but up.



**“Because our datasets are so large and technically complex, we face challenges on a daily basis in terms of how fast we can process and analyze our data, and in terms of how much space we have to store our analyses.”**

**Steven Friedenberg**

Assistant Professor in the College of Veterinary Medicine



Despite these developments and ever-increasing requirements, the Agate cluster came online at the end of 2021. For now, it is the definition of an “impressive system,” sporting 15 racks from HPE with 16 tons of Direct Liquid Cooling. Agate includes 44,032 CPU cores and 322 TB of memory. The cluster boasts a whopping 8 quadrillion floating-point operations per second (that’s 8 petaflops).

With these enhanced capacities and capabilities, Agate is something more than a “game changer” for researchers. It’s a new beginning for the future of innovations across all modes of human inquiry.

**“The expanded capabilities of MSI will give my lab a competitive advantage in genomic analyses. High-performance storage space and computing power, especially RAM, are absolutely critical to our analyses, and we will work much more efficiently with the additions provided by this new resource.”**

**Suzanne McGaugh**  
Associate Professor in the College of Biological Services

None of this would’ve been possible if data storage was unable to keep up with Agate. So, how does it keep up?

## We call it “ActiveStor Ultra” for a reason

When it comes to the high-performance data storage needed to keep pace with a hyper-sophisticated cluster and a plethora of disparate use cases, researchers and system admins normally expect frequent downtime, performance bottlenecks, and complex, time-intensive data management. Fortunately, Panasas ActiveStor Ultra storage recalibrated their expectations.

The ActiveStor Ultra appliance runs the Panasas PanFS parallel file system. This storage solution’s single-tier architecture is uniquely suited to the mixed workloads that multiple research groups at large institutes carry out every day on their clusters. Automation features—optimal data placement, capacity balancing, tuning and retuning, and

recovery logic—deliver consistent high performance that actively adapts to different I/O patterns that other storage systems would noticeably struggle to support simultaneously.

The 10 PB of Panasas storage underlying the Agate system gives MSI researchers the exceptional reliability needed to support 24/7 operation. And that reliability only increases with scale. With unmatched uptimes and an easy-to-use single management console, Panasas increases productivity because MSI researchers can focus on their work, not their storage.

We could end by saying that Panasas plays an integral role within the Agate system, but Dr. Jim

Wilgenbusch, the Director of Research Computing at UMN, already said it for us: “The Panasas partnership gives us access to the kind of high-performance computing technology that is critical to tackle an increasing variety of sophisticated use cases with MSI’s Agate cluster. This long-term partnership has provided our researchers with the top-flight technology and technical support they need.”



If we can take the complexity and cost out of a system running 138,612 GPU-accelerated jobs, imagine what ActiveStor storage can do for you and your business. Better yet, don't imagine. Come see for yourself.

[Let's talk today](#)