

HPC Without Compromise

Don't Let Storage Get in the Way of Science



The performance, availability, and affordability of storage systems can have a direct impact on the speed of scientific research and analyses. Unfortunately, most scale-out file systems and network-attached storage (NAS) systems force researchers to make unwanted tradeoffs.

The highest performing storage systems can be prohibitively expensive. More affordable systems often struggle with the volume and variability of genomics, CryoEM, drug discovery, bioinformatics, personalized medicine, and other scientific workloads. And regardless of performance and cost, most HPC storage systems require ongoing management, fine-tuning, and downtime – all of which can get in the way of scientific progress.

Panasas PanFS® is different.

A true parallel file system with Dynamic Data Acceleration, PanFS delivers consistently high performance at any capacity. It automatically adapts to evolving workloads to minimize manual intervention and tuning. It can be updated and scaled without downtime. And it leverages commodity hardware to give researchers the very best performance at the most affordable price.

Ultra-high performance

With the explosion in genomic and structural information and the emergence of data-intensive molecular diagnostics, researchers need high-performance storage systems that keep up with the pace of change. PanFS delivers ultra-high performance and capacity that scale linearly without limitation. And the system dynamically adapts to workloads, regardless of complexity, to continually optimize performance.

| 2 PB usable system performance | Read/write performance ratio (approx.) | Performance loss as system fills | High availability | NVDIMM (speeds application writes) |
|--------------------------------|--|----------------------------------|--|------------------------------------|
| 30+ GB/s | 1 to 1 | No | N+2 included, increases reliability at scale | Yes |

Easy to manage

PanFS implements a single global namespace that reduces storage complexity and simplifies management, and dynamic workload optimization minimizes manual intervention and tuning. Only one part-time administrator is typically required to manage a PanFS storage solution, regardless of scale.

| Data scrubbing | Asynchronous replication | Minimal administration effort required | Estimated FTE requirement | In addition to scratch, reliable enough to be used to store home directories, applications, and general-purpose storage | 99.99% uptime | Operational headaches |
|----------------|--------------------------|--|---------------------------|---|---------------|-----------------------|
| Yes | Yes | Yes | <0.25 | Yes | Yes | None |

Highly reliable

PanFS delivers high availability and protects data assets with fully automated online failure recovery. Network-distributed erasure coding and fully parallel rebuilds recover nodes quickly. And background data scrubbing ensures data integrity.

| Erasure coding | Data protection overhead | Snapshots | Automatic SSD <-> HDD tiering | Tuning/re-tuning required when changing workload to optimize performance | Declustered RAID rebuild | End-to-end checksum |
|----------------|--------------------------|-----------|--|--|--------------------------|---------------------|
| Yes | 20% | Yes | Yes, Adaptive Small Component Acceleration (DDA) | No | Yes | Yes |

Conclusion

More than ever, scientists and researchers need a true parallel file system. One that maximizes the efficiency of hundreds or even thousands of compute nodes without undue cost. One that can handle the quantity and complexity of scientific data without manual intervention or continuous fine-tuning. And one that delivers exceptional performance, availability, and affordability – without compromise.

PanFS delivers the highest performance, most dynamic optimization, and easiest to manage file system for scientific workloads – at the lowest total cost of ownership of any HPC storage solution.

About Panasas



Panasas builds a portfolio of data solutions that deliver exceptional performance, unlimited scalability, and unparalleled reliability – all at the best total cost of ownership and lowest administrative overhead. The Panasas data engine accelerates AI and high performance applications in manufacturing, life sciences, energy, media, financial services, and government. The company's flagship PanFS® data engine and ActiveStor® storage solutions uniquely combine extreme performance, scalability, and security with the reliability and simplicity of a self-managed, self-healing architecture. The Panasas data engine solves the world's most challenging problems: curing diseases, designing the next jetliner, creating mind-blowing visual effects, and using AI to predict new possibilities.

Worldwide Office
1-888-PANASAS
info@panasas.com

Panasas Headquarters
San Jose, CA, USA
Panasas Research & Development
Pittsburgh, PA, USA

Panasas EMEA
Oxford, United Kingdom
emeainfo@panasas.com

Panasas APAC
Sydney, Australia
apacinfo@panasas.com

Panasas China
Shanghai, China
chinainfo@panasas.com