


# 5 Reasons You Need to Prioritize HPC Reliability

## 1 You're dealing with more HPC data every day

The amount of data processed in HPC and AI environments is



increasing aggressively.<sup>1</sup>

HPC and AI/ML data drive key processes for critical workloads, so preventing data loss and corruption while ensuring uptime has become just as crucial as peak performance.




## 2 But more data means more problems

As data volumes expand, HPC storage systems must grow to keep pace. Scaling necessitates the use of additional hardware, which inevitably decreases system reliability, resulting in longer rebuild and restoration times.




AI projects can quickly scale from terabytes to petabytes, with



300-400PB capacities becoming increasingly common.<sup>2</sup>

## 3 System failures become commonplace

HPC users experience approximately 10 storage system failures per year (nearly once a month)



with an average recovery time of 1.7 days.<sup>3</sup>

As rebuild times increase, so does the chance of widespread, secondary failures requiring significant rebuild and restoration times. Effective recovery typically requires multiple storage experts whose services are costly and may be in short supply.



## 4 And the cost of downtime and lost data is substantial

In addition to the loss of valuable data, failures can lead to potentially serious implications for project completion and business continuity.



Typical HPC downtime costs approach

**\$127,000 per day,**



including the lost productivity of an expensive compute cluster sitting idle.<sup>4</sup>

## 5 Unlike other HPC solutions, only Panasas is built on a superior reliability architecture to maintain performance and reduce downtime

At Panasas, reliability is in our DNA. We develop HPC and AI/ML solutions that deliver the reliable storage and data mobility you need to support critical operational priorities. With our parallel file system, PanFS®, data reliability *increases* with scale and your HPC system is protected in 5 ways:

- 1 Built-in reliability; no extras required**

Unlike legacy RAID, where data is protected at the drive level, PanFS protects across storage nodes or object storage devices (OSDs). The result is much higher reliability with protection against drive and node failures. Additionally, all OSDs contribute to reconstruction performance, decreasing rebuild time and increasing overall reliability.
- 2 Eliminated tiering with Dynamic data acceleration**

DDA intelligently determines the best media type for each file: large files on high-bandwidth, cost-effective HDDs, small files on high-IOPs SSDs, and metadata on low-latency NVMe. In contrast, old-school tiering is based on access times, reducing performance while increasing complexity and cost.
- 3 Patented per-file object erasure coding**

Each file is individually erasure coded and striped across many OSDs, high data integrity, and parallel rebuilds at the file level.
- 4 Uniquely stable architecture**

Extended File System Availability, continuous background data scrubbing, a "shared nothing" architecture, fault tolerance, and deep automatic failure recovery that detects and corrects issues without interrupting operations.
- 5 Data visibility and mobility**

Panasas offers a data protection suite that defends across storage systems. Enabling easy data movement and visibility across environments, Panasas ensures safe backup and archiving of HPC and AI/ML datasets across data centers and to the cloud.

*The system architecture in PanFS has been shown to reduce rebuild complexity, time, and expense.<sup>5</sup>*

<sup>1</sup> [https://hyperionresearch.com/wp-content/uploads/2021/11/Hyperion-Research-SC21-Market-Update-Briefing\\_AI-HPDA-Growth\\_Norton.pdf](https://hyperionresearch.com/wp-content/uploads/2021/11/Hyperion-Research-SC21-Market-Update-Briefing_AI-HPDA-Growth_Norton.pdf)  
<sup>2</sup> <https://www.techtarget.com/searchstorage/opinion/Why-the-future-of-AI-storage-may-have-to-exclude-flash>  
<sup>3</sup> [https://www.panasas.com/wp-content/uploads/2020/04/Hyperion\\_Importance-of-TCO-for-HPC-Storage-Buyers\\_Q1-20\\_FINAL\\_2020-04-22.pdf](https://www.panasas.com/wp-content/uploads/2020/04/Hyperion_Importance-of-TCO-for-HPC-Storage-Buyers_Q1-20_FINAL_2020-04-22.pdf)  
<sup>4</sup> Ibid.  
<sup>5</sup> "Characterizing and Modeling Reliability of De-clustered RAID for HPC Storage Systems" Z. Qiao, S. Laing, S. Fu, H.-B. Chen, B. Settlemeyer – 2019

**Panasas Headquarters**  
San Jose, CA, USA

**Panasas Research & Development**  
Pittsburgh, PA, USA

**Worldwide Office**  
1-888-PANASAS  
[info@panasas.com](mailto:info@panasas.com)

**Panasas APAC**  
Sydney, Australia  
[apacinfo@panasas.com](mailto:apacinfo@panasas.com)

**Panasas EMEA**  
Oxford, United Kingdom  
[emeainfo@panasas.com](mailto:emeainfo@panasas.com)

**Panasas China**  
Shanghai, China  
[chinainfo@panasas.com](mailto:chinainfo@panasas.com)