

PanFS[®] High Performance Parallel File System

Built for Price/Performance, Engineered for Manageability

Panasas[®] PanFS eliminates the historical trade-off between HPC storage performance and ease of use for research and simulation workflows in life sciences, energy, computer aided engineering, media & entertainment and government. PanFS on Panasas ActiveStor[®] Ultra leverages 20 years of experience in the most demanding HPC environments to deliver a powerful combination of high performance, management simplicity, workload-and-configuration flexibility, rock-solid reliability, and lowest TCO. PanFS provides a versatile and scalable foundation for emerging applications such as deep learning, precision medicine, and autonomous vehicles that is far simpler to manage than other parallel file systems.


Key features of the PanFS parallel file system include:

- A portable operating system interface (POSIX) parallel file system on an object back-end delivering highly scalable, massively parallel performance.
- A balanced architecture that delivers consistent performance, regardless of scale or workload complexity
- Intelligent data placement with separate, parallel, bottleneck-free data paths to metadata and data.
- Data reliability and availability that increases with scale.
- Simple, automated management that requires only a single, part-time administrator, regardless of scale.

Consistently High Performance, Regardless of Workload

Requiring little need for labor- and skill-intensive tuning/retuning to maintain peak performance, PanFS balances and optimizes every part of the storage system – CPU, NIC, DRAM and storage media – ensuring consistent, predictable performance across varying workloads, regardless of complexity, and with no need for manual tuning or reconfiguration. To eliminate hotspots and bottlenecks, PanFS divides file system workloads between director nodes that process metadata (information about files) and storage nodes that process data (files). Client systems first obtain access permission and data locations from director nodes, then read or write directly to multiple storage nodes in parallel.

Storage can be optimized for different workloads by configuring the number and ratio of director and storage



PanFS Parallel File System

nodes. Capacity scaling adds compute, memory, network bandwidth, and raw storage as balanced building blocks, eliminating the need to overprovision any one resource. This results in delivering storage that scales linearly and delivers consistent I/O performance across large numbers of concurrently executing applications, exercising a complex mix of large and small unstructured data sets. And unlike most file systems, PanFS maintains consistently high performance as capacity is expanded.

Surprising Simplicity

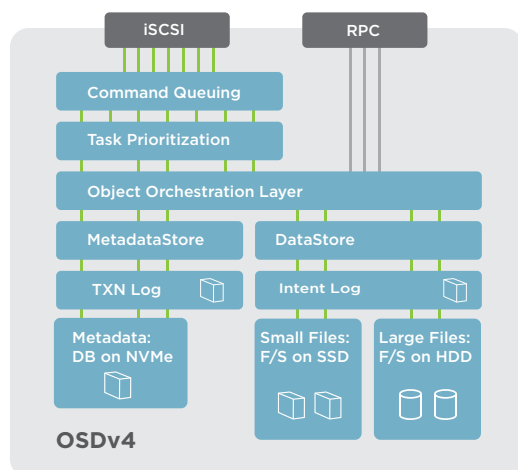
PanFS implements a single global namespace that reduces storage complexity and simplifies management. Only one part-time administrator is typically required to manage a PanFS storage solution, regardless of scale. Storage nodes can be added easily with the click of a mouse. PanFS automates key workflows such as new storage discovery, load balancing to streamline performance, and enterprise data services such as reporting, snapshots, and user quota enforcement. PanFS automatically recovers from node and drive failures and tuning/retuning to accommodate workload transitions is not required.

Intelligent Data Placement Accelerates Performance

PanFS features a multi-tier intelligent data placement architecture that matches the right type of storage media to each type of data to deliver the highest performance at the lowest cost:

- Small files are stored on high IOPs flash SSDs
- Large files are stored on low-cost, high-capacity, high-bandwidth HDDs
- Metadata is stored in a database on low-latency NVMe SSDs
- An NVDIMM-based intent-log protects both inflight data & metadata operations
- Unmodified data and metadata are cached in DRAM

Because PanFS protects newly written data in NVDIMM, it allows the other drives to write data fully asynchronously, coalescing writes and placing data on the right storage media in the most efficient manner for best performance. In addition, ActiveStor Ultra features an extremely balanced design that deftly optimizes the right amount of CPU, memory, storage media, and networking — with no hardware bottlenecks or hotspots from the NIC down to the storage media — to deliver best price/performance.



PanFS Intelligent Multi-tier Architecture

Enterprise-Grade Data Protection and Availability

PanFS uses software-based, network-distributed erasure codes to separately protect individual files rather than traditional RAID groups that protect whole drives. This eliminates the multiple risks, costs, and performance penalties of traditional RAID architectures.

- N+2 encoding protects against up to two simultaneous failures of either drives or whole nodes

- Affected files are rebuilt on distributed free space, eliminating the need for spare drives
- Storage nodes across the entire cluster are used to rebuild affected files in parallel, greatly accelerating recovery
- Reliability increases as the storage cluster grows
- Background data scrubbing continuously checks for errors and corrects them
- Single point in time immutable file system snapshots protect against data loss
- Extended File System Availability ensures users can access the namespace, even in the event of three simultaneous failures NVMe SSDs
- An NVDIMM-based intent-log protects both inflight data & metadata operations
- Unmodified data and metadata are cached in DRAM

Multi-protocol Access

Clients requiring the highest performance access to the ActiveStor Ultra solution use the Panasas DirectFlow data protocol. The DirectFlow parallel data access protocol supports direct access to Linux clients. The ActiveStor Director also provides scalable access for client systems via the standard NFS or SMB protocols via “gateway” services. User authentication is managed via a variety of options including Active Directory and Lightweight Directory Access Protocol (LDAP).

Flexible Network Connectivity

ActiveStor Ultra provides flexible Ethernet and InfiniBand networking options to match cluster connectivity and performance requirements.

Timely High-Quality Service and Support

Unlike open-source solutions and even commercial alternatives from broad portfolio vendors, Panasas offers timely world-class L1-L4 support with just one phone call to resolve issues within minutes to hours rather than days and weeks.

- Single-vendor problem resolution, with centralized responsibility for addressing your issues
- Focused attention on your issues the moment you need it
- Support expertise from highly qualified technicians at your fingertips
- Rapid time to problem resolution