PANASAS PANFS FILE SYSTEM

Perfecting Parallel Storage for High-Performance Computing

Panasas® PanFS®, the intelligent parallel file system, delivers massive scale-out performance, exceptional reliability, and simple manageability for tomorrow's high-performance computing (HPC) workloads.

Storage File Systems Can't Keep Up with HPC

Back in the late 1990s, research labs and enterprises were taking advantage of the huge potential of HPC applications and while compute power quickly increased to keep up with the demands, storage systems could not.

Today, many file systems still haven't kept pace with the rapid evolution of HPC systems, where hundreds of computationally intensive applications may concurrently read and write very large data sets to a single, shared storage pool. These systems require an unprecedented combination of extreme input/output (I/O) performance and unlimited scalability.

While legacy parallel file systems such as Lustre and Spectrum Scale are powerful, they can also be extremely complex and expensive to manage. They introduce significant administrative overhead for tuning, optimizing, and maintaining storage performance for different HPC workloads, driving higher total cost of ownership (TCO). IT organizations urgently need parallel storage solutions that combine extreme performance and linear scalability with plug-and-play management simplicity and low TCO.

Panasas ActiveStor® HPC storage solutions don't have the performance constraints of traditional scale-up network-attached storage (NAS) systems or the management complexity of legacy parallel file systems. Panasas high-performance, plug-and-play ActiveStor parallel storage appliance is optimized to accelerate workflows and simplify data management and supports many industry and research innovations around the world.

PanFS: The Parallel File System Built for HPC

Panasas PanFS eliminates every trade-off between HPC storage performance and ease of use. The parallel file system software that powers Panasas ActiveStor® Ultra solutions 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.

Key features of the PanFS parallel file system include:

- A portable operating system interface (POSIX) compliant, 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.
- A Linux operating system (OS) platform that enables hardware and software flexibility and leverages the latest technology advances.

Balanced, Parallel Storage Architecture

PanFS balances and optimizes every part of the storage system — compute, memory, network, 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, it 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 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.

Intelligent Placement Accelerates Data Access

PanFS efficiently leverages the unique performance characteristics of SSDs and HDDs to deliver the highest performance at the lowest cost. Tiered storage architecture separates metadata files into a database on NVMe SSDs, small files on SATA SSDs to optimize for IOPS and latency, and large data files on HDDs to optimize for streaming bandwidth. This allows storage and metadata services to scale independently and incrementally.

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 because:

- N+2 encoding protects against up to two simultaneous failures of either drives or whole nodes.
- Extended File System Availability ensures users can access the namespace, even if some of the files have lost data.
- 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.

PanFS also provides inflight protection for data and transactions by caching newly written data in power-protected NVDIMMs.

Exceptional Ease of Use and Management

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

Today's Storage Solution for Tomorrow's HPC

With Panasas PanFS parallel file system, you have a storage solution that is powerful, scalable, reliable, and easy to manage, with the industry's lowest TCO. It's the only storage file system that is ready today for tomorrow's extreme HPC workloads.

For more information about Panasas PanFS, visit www.panasas.com.