PANASAS®

Comparing the Performance, Reliability and Administration Requirements of Different Parallel File Systems

COMPETITIVE COMPARISON

Comparing the Performance, Reliability and Administration Requirements of Different Parallel File Systems

Comparing commercially available solutions that integrate parallel file systems can be challenging. Every user's application mix and use-cases are different and specific actual targeted benchmarking would be required to see how each system would perform against an organization's specific workload. However, there are some straightforward ways to compare parallel file systems that will give you a very good sense of whether or not the system meets your expectations for an effective HPC storage deployment.

The comparison below evaluates the BeeGFS, GPFS, Lustre and Panasas® parallel file systems on the basis of performance, reliability, administration requirements and support. We looked at throughput per hard disk drive (HDD) as a comparable figure of merit, since the number of disks has the primary impact on the system footprint, compared read/write performance, data availability, tiering, tuning requirements, uptime, rebuild times, and what it takes to run and support the system.

The PanFS® parallel file system stands out for being 2X faster than BeeGFS, GPFS and Lustre, without any of their notorious complexity and brittleness. With PanFS, performance and capacity scale linearly without limitation, and performance adapts to dynamically changing workloads and is consistently fast, regardless of workload complexity. In addition, PanFS delivered on the ActiveStor® appliance combines the industry's leading price-performance with easy manageability at the lowest total cost of ownership of any HPC storage solution. HPC organizations no longer need to trade off performance and price-performance for simplicity, uptime and great support. With PanFS on ActiveStor Ultra, HPC organizations can have it all.

Performance

Features	PanFS on Panasas ActiveStor Ultra	Lustre DDN Exascaler	BeeGFS	IBM ESS GPFS	Notes
2 PB system estimated performance	20 GB/S	10.4 GB/S	9.4 GB/s	8.9 GB/S	All assuming approx. 120 drives at 16 TB each, DDN Lustre, BeeGFS and GPFS would need to use 2X more drives to achieve similar performance compared to Panasas.
Read/write performance ratio (approx.)	1 to 1	1 to 1	2 to 1	2 to 1	BeeGFS writes half speed of reads with Buddy Mirroring.
Performance loss as system fills	No	Yes	Yes	No	See Wiki.Lustre.org report titled: The Effects of Fragmentation and Capacity on Lustre File System Performance.
High availability	N+2 included, increases reliability at scale.	N+1 included, but complex and unreliable, gets more brittle at scale.	Optional Buddy Mirroring.	N+1 included, but complex and unreliable, gets more brittle at scale.	Panasas N+2 can survive any two component failures. DDN Lustre, GPFS and BeeGFS N+1 can survive any one component failure. Buddy Mirroring HA doubles BeeGFS HW requirements and cost.
NVDIMM (speeds application writes)	Yes	No	No	No	Panasas NVDIMM accelerates application performance.

Reliability/Tuning/Rebuild

Features	PanFS on Panasas ActiveStor Ultra	Lustre DDN Exascaler	BeeGFS	IBM ESS GPFS	Notes
Erasure coding	Yes	Yes	No	Optional	
Data protection overhead	20%	20%	140%[1]	20%	
Snapshots	Yes	No	No	Yes	
Automatic SSD <-> HDD tiering	Yes Adaptive Small Component Acceleration (ASCA).	Yes SSD <-> HDD	No	Yes SSD <-> HDD	For DDN and IBM, data movement latency from HDD to SSD. For BeeGFS, customer would have to move data between tiers manually.
Tuning/ re-tuning required when changing workload to optimize performance	No	Yes	Yes	Yes	Panasas excels at mixed workloads.
Declustered RAID rebuild	Yes	Yes	No	Yes	BeeGFS RAID rebuilds take significantly longer.
End-to-end checksum	Yes	Yes	No	Yes	

[1]BeeGFS = Protection by Buddy Mirroring + RAID overhead.

Data Management/Administration

Features	PanFS on Panasas ActiveStor Ultra	Lustre DDN Exascaler	BeeGFS	IBM ESS GPFS	Notes
Data scrubbing	Yes	No	No	No	
Automatic data balancing	Yes	No	No	No	
Asynchronous replication	Yes	Yes	No	Yes	
Minimal administration effort required	Yes	No	No	No	
Estimated FTE requirement	<0.25	1+	1+	1+	
In addition to scratch, reliable enough to be used to store home directories, applications and general-purpose storage	Yes	No	No	No	
99.99% uptime	Yes	No	No	No	
Operational headaches[2]	None	Unreliable failover. Requires complex variable striping, mostly not done. PFL unstable, loses data.	Poor support from Germany.	Complex, spotty support.	

[2] Based on customer survey data and disinterested 3rd party input.

Support

Features	PanFS on Panasas ActiveStor Ultra	Lustre DDN Exascaler	BeeGFS	IBM ESS GPFS	Notes
Support quality[2]	Excellent	Spotty	Poor	Spotty	
US-based development organization	Yes	Yes	No Germany-based.	Yes	
US-based support organization	Yes	Yes	No L3 from Germany.	Yes	

[2] Based on customer survey data and disinterested 3rd party input.

For more information about Panasas visit **www.panasas.com**.