

SMASH THROUGH DATA PROCESSING BOTTLENECKS WITH FUNGIBLE STORAGE CLUSTER AND IBM SPECTRUM SCALE



Blazing Performance with Secure Scale-Out Storage for High-Performance AI/ML and HPC Workloads

IBM's Spectrum Scale (a.k.a. GPFS) is a well-known highly scalable, highly performant parallel file system used in the most demanding HPC and Artificial Intelligence (AI)/Machine Learning (ML) applications. The Fungible Storage Cluster, powered by the Fungible Data Processing Unit™ (Fungible DPU™), is the highest performance, scale-out disaggregated All-Flash storage platform in the market. The combination of the two offers a powerful solution that eliminates the data bottleneck that feeds the compute cluster, delivering a much shorter time to results and at a lower cost.

THE DATA PROCESSING BOTTLENECK

HPC and AI/ML application architectures often have a large number of compute nodes and massive compute capabilities. These applications need to ingest and process large amounts of data. However, traditional storage solutions in the market are bottlenecked in the amount of data they can deliver to these applications. This lack of storage throughput is often the bottleneck that constrains the overall cluster performance.

Let's take GPU servers used in AI/ML. They are constrained by the data ingestion bottleneck. The dataset required is most often larger than the server's local storage so remote storage is required, but today's remote storage is not fast enough to provide the throughput needed by the GPUs. For example, in a genomics application with many small files, a highly performant shared storage infrastructure with very high IOPS can significantly speed up the processing. Imagine 30 million IOPS delivered by shared storage to the Spectrum Scale cluster in a durable 6 node configuration (12RU) that is linearly scalable to hundreds of nodes.

The traditional active-active dual controller All-Flash array does not scale and is not performant enough to feed the compute servers or the GPUs. Software-Defined Storage (SDS) on commodity hardware has the same issues with scalability and performance. Direct Attached Storage is inefficient, does not share data well and requires significant computational power for data protection, data reduction, and data security. These storage options force customers to significantly over-provision storage to get to the desired performance.

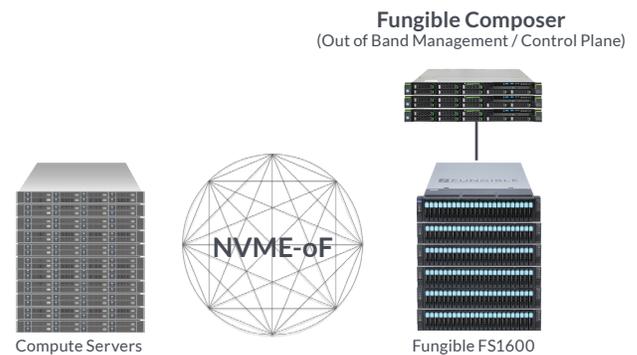
These suboptimal approaches result in expensive GPU servers or compute nodes sitting idle, waiting for data. By eliminating the data I/O bottleneck, the efficiency of the GPUs and compute nodes will be higher and time to results will be much faster. In addition, if the data is shared, the infrastructure would be much more agile and lower cost. Pooling the storage increases utilization, provides agility and enables independent scaling of compute and storage. With storage services such as erasure coding, compression and encryption offloaded to the Fungible Storage Cluster, the GPFS nodes have more cores to run GPFS and so, fewer servers are required.

The combination of the Fungible Storage Cluster and IBM Spectrum Scale smashes through the data processing bottleneck.

THE FUNGIBLE STORAGE CLUSTER

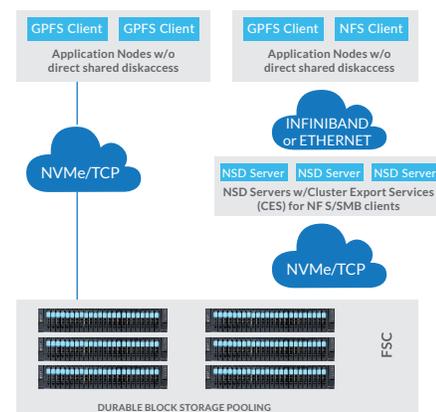
Powered by the Fungible DPU, the Fungible Storage Cluster (FSC) is the highest performance, secure, scale-out disaggregated all-flash storage platform in the market today. The FSC comprises a cluster of two or more Fungible FS1600 storage target nodes and three Fungible

Composer nodes. The control plane is managed by the Fungible Composer software, a centralized management solution that configures, manages, orchestrates, controls and deploys the Fungible Storage Cluster.



THE SOLUTION - BETTER TOGETHER

Fungible Storage Cluster is connected to GPFS Client or the NSD servers via NVMe/TCP. The Fungible Storage Cluster forms a pool of replicated or erasure-code protected, compressed and encrypted shared NVMe storage that are shared by all the application servers. Control plane functions such as volume creation are performed through the Fungible Composer. Storage services are offloaded from the GPFS clients or NSD servers to the FS1600 nodes. Data durability, through replication or network EC, ensures there is no single point of failure. true scalability of performance and capacity can be achieved by adding additional FS1600 nodes.



UNPARALLELED PERFORMANCE

Want results faster? Can't wait for days or hours for the job to finish? Expensive GPUs sitting idle waiting for data? Fungible Storage Cluster with IBM GPFS is an unbeatable combination for applications that are starved for data.

With replication, a cluster of 6 storage nodes is targeted to deliver up to 30 million IOPS¹, with NVMe/TCP, as measured by GPFSperf. The performance scales linearly in 2RU increments.

SUMMARY OF BENEFITS

- NVMe All-Flash array powered by the Fungible Data Processing Unit for the highest performance
- Data durability schemes protect against drive, processor, system and rack failures and eliminates any single point of failure
- Network EC provides improved reliability and media efficiency compared to replication
- Line-rate data reduction techniques applied to hot and warm data offer more than 5X media, achieved without performance impact
- Line-rate encryption secures data with no impact to cluster performance
- Offloading of storage services from GPFS clients or NSD servers to the Fungible Storage Cluster reduces the number of servers required
- Standard IPoE network compliance (NVMe/TCP) obviates need for expensive InfiniBand or complex RDMA that does not scale
- Disaggregation of compute and storage nodes enable improved infrastructure agility through independent scaling of servers
- Shared pool of storage removes copies and increases utilization. For example, if multiple GPU servers require the same dataset for training, only one copy of that dataset is needed.
- Per-volume level granularity of features for data durability, reduction, security makes the FSC ideal for multi-tenant environments

¹ Projected metrics based on actual 4K random read measurements with a 2 node cluster that delivered 10 M IOPS in replication mode. With 4+2 EC, a 3 node cluster delivered 8.5M IOPS.

CONCLUSION

The combination of ultra-high performance Fungible Storage Cluster with the highly scalable IBM Spectrum Scale creates the market-fastest distributed file system solution, while eliminating storage bottlenecks and achieving unprecedented IOPS/PB for AI/ML, HPC and other IO-intensive workloads. Faster time to insights creates greater business value - higher Return on Investment (ROI) and lower Total Cost of Ownership (TCO).

NEXT STEPS

For additional information and demo, contact sales@fungible.com. See more detailed performance results in this [whitepaper](#).

ABOUT FUNGIBLE

Silicon Valley-based Fungible is reimagining the performance, economics, reliability, security and agility of today's data centers.

CONTACT US

sales@fungible.com

FUNGIBLE, INC.

3201 Scott Blvd., Santa Clara, CA 95054, USA
669-292-5522

www.fungible.com | [in](#) [▶](#) [🐦](#) [✉](#)