

Qumulo Distributed File Storage for High Performance Computing (HPC)

Many organizations today, including research and academic institutions, national labs, energy companies, financial institutions, manufacturing and government organizations, use high performance computing (HPC) to solve complex problems. For example, to help research and find cures for disease, gain a better understanding of the universe, build better machines that help drive innovation, or even predict the weather.

Today's HPC workloads require file system solutions that can manage billions of small and large files, while providing consistent high performance and extreme scalability to keep up with massive data growth. They also need to provide a seamless way to leverage industry standard hardware and employ hybrid and cloud environments, as and when needed for burst and elastic computing and to leverage cloud applications, such as AI engines, not available in their data center.

Qumulo Meets the Demands of Today and Tomorrow

Qumulo delivers high performance cloud-native distributed file system that meets the performance and capacity demands of HPC environments, in the data center and in the cloud.

Qumulo's file storage efficiently manages mixed IO performance for billions of small and large files at petabyte scale to accelerate productivity for scientific research, genomics, machine learning, and artificial intelligence across multiple industries. In addition, Qumulo's real-time analytics eliminate data blindness for organizations, helping to save time and money by providing visibility across billions of files.

Performance to Support HPC Workloads

Qumulo file storage delivers extreme low latency, IOPS and throughput performance to meet the demands of today's HPC workloads. Qumulo's software is uniquely optimized for both fast reads and writes through built-in hybrid intelligent predictive caching and proactive prefetch which enables fast reads and identifies read IO patterns to proactively move data to the fastest media.

FEATURES AND BENEFITS

Real-Time Visibility

- Eliminate data blindness
- Instant visibility across billions of files

Unify & Consolidate

- Single namespace for NAS, home directories, instruments, processing and AI/ML
- Consolidate storage silos

Scale Across on-prem and the Cloud

- Cloud-native distributed file storage
- Burst compute, collaborate, archive

Performance to support HPC workloads

- Low latency performance with all-NVMe nodes
- Mixed I/O performance for billions of large and small files

Workflows

Genomic sequencing, microscopy analysis, seismic processing, reservoir modeling, research & engineering simulations, data/image processing and analysis, molecular dynamics, risk analysis, fraud prevention, home directories and file share, archive file storage

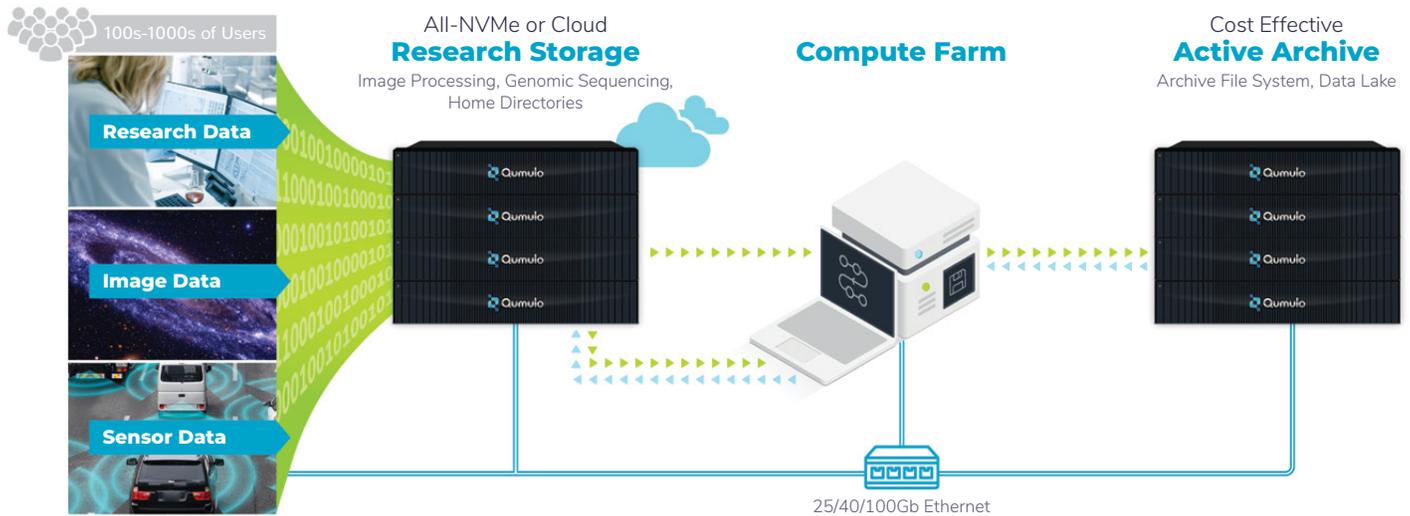
"It's impossible to make smart decisions about data when your storage is a figurative black box. Qumulo lets me know in an instant how the data is used, who touches it and how often – so storage is no longer a technical issue, but a management decision."

— SCI Institute



Distributed File System for HPC Workloads

NVMe. SSD. Disk. Active Archive. Cloud



Linear Scalability Across On-Prem and Cloud Environments

As HPC performance and capacities expand and grow, organizations can benefit from Qumulo's simple, modular, scale-out file storage. Qumulo's software-defined architecture utilizes clusters of nodes made up of Qumulo hardware, or pre-qualified industry-standard hardware from HPE or Dell. Managing growth is simple. Simply add nodes to the existing infrastructure to increase performance and capacity levels uniformly, when required, with no disruption or downtime.

In addition, due to the growing size of data sets and the compute-intensive nature of AI and ML organizations are taking advantage of the cloud. By using Qumulo's single platform, organizations can seamlessly scale workloads to AWS or GCP cloud environments, as and when needed, for data storage, multi-site collaboration, and compute performance processing.

Unify and Consolidate to Expand Collaboration

With Qumulo's single namespace, organizations can consolidate legacy NAS and parallel system storage silos, to improve data accessibility, as well as simplify and reduce the cost of data management. With a single SMB share or NFS mount for NAS for home directories, instruments, processing and AI/ML, Qumulo makes it easy for regionally dispersed teams to share resources and expand collaboration.

Simplicity and Flexibility

Unlike inherently complex parallel file system environments, Qumulo is simple to install and manage. Offering a wide variety of protocols, including NFS, SMB, FTP, HTTP and S3 (through partnership with Minio), Qumulo's software can interface with existing applications easily. It integrates with Active Directory and LDAP for user security, to manage permissions, controls, and access restrictions for file applications, to create the smoothest workflow possible for HPC environments.

Real-Time Visibility Across Billions of Files

Qumulo provides built-in real-time analytics, to provide visibility to the entire file system. These actionable insights identify storage usage and capacity trends, as well as user-based performance trends. With this superior visibility, organizations can proactively manage ongoing requirements, and plan for future performance and capacity growth.

ABOUT QUMULO

Qumulo is the leader in enterprise-proven hybrid cloud file storage, providing real-time visibility, scale and control of your data across on-prem and cloud.

For more information, visit www.qumulo.com