


Extend Real-Time Data Access Anywhere – Edge, Cloud, and Hybrid Environments

Qumulo Accelerators are lightweight, software-defined data engines that extend the power of the Qumulo Cloud Data Platform to edge locations, remote sites, and multi-cloud environments. They empower organizations to project massive datasets exactly where they are needed — in real time and without costly, complex, or time-consuming data duplication.



Modern enterprises increasingly need consistent, high-performance access to large file datasets across distributed environments — from branch offices and rugged remote sites to cloud regions spanning multiple providers. Other hybrid-cloud architectures often require data replication, involve long transfer times, and create management silos that impede collaboration and slow workflows.

Qumulo Accelerators Solution

Qumulo Accelerators solve these challenges by bringing fast, consistent, low-latency access to shared data anywhere it's needed, while still centralizing governance and control through a unified global namespace powered by Qumulo Cloud Data Fabric.

Core Benefits

- Instant Data Access:** Qumulo Accelerators eliminate long waits for data movement by delivering reads and writes from the nearest endpoint, while still enforcing strict global consistency across all locations.
- Intelligent Performance:** NeuralCache anticipates access patterns and brings data where it's needed most — reducing latency and accelerating application workflows regardless of data location.
- Cost Efficiency:** By avoiding unnecessary replication and minimizing idle data copies, organizations reduce storage costs while maximizing performance.
- Simple, Flexible Deployment:** Accelerators can be deployed on lightweight hardware, virtual machines, or cloud instances — enabling frictionless adoption across diverse environments.



A Modern File Data Experience

Qumulo Run Anywhere

Unlike legacy storage systems, Qumulo is 100% software-defined, meaning it runs consistently across on-premises, cloud, or hybrid environments—on any x86/ARM server. This lets modern enterprises avoid vendor lock-in, simplify infrastructure, and deploy anywhere while maintaining a single, unified file system.

Any Data, Any Location, Total Control
Qumulo's Cloud Data Fabric removes capacity and distance limits, making any data instantly available to any workload—anywhere—with strict consistency while ensuring sovereignty and security.

Intelligent Performance with NeuralCache

Qumulo NeuralCache uses machine learning to optimize reads and writes, improving performance and drastically reducing cloud I/O costs by up to 99%.

Cloud-Native, But Not Cloud-Limited
Qumulo is cloud-native by design, using object storage with real-time file semantics to enable secure, predictable scaling. It eliminates the need for data replication and app rewrites, ensuring cost control and compliance with data sovereignty requirements.

Customer Success

Qumulo zero-latency support enables fast issue resolution through direct access to experienced Customer Success Managers or file system engineers—no tickets required. Qumulo support has earned an NPS score of 95.

Key Capabilities

Extend Data to the Edge — Edge Accelerator

- **Fast, Local Access:** Enable secure, low-latency reads and writes from lightweight VMs or portable physical appliances — ideal for remote teams, satellite offices, field operations, and mobile professionals.
- **Consistent Collaboration:** Users at any location can create, view, modify, and share data in real time without waiting for replication or resolving version conflicts.
- **Intelligent Caching:** Powered by NeuralCache, Qumulo's machine-learning prefetch engine, Edge Accelerator anticipates needed data and brings it locally in advance.
- **Flexible Deployment:** Run on lightweight hardware (e.g., Cisco Unified Edge) or as a virtual appliance — easily deployable in minutes on your choice of hypervisor.

Edge Accelerator Use Cases:

- ✓ Small office/home office data access
- ✓ Mobile and remote workflows with large media or forensic data sets
- ✓ Autonomous outposts (e.g., seismic, weather monitors) requiring fast local writes and ingestion

Optimize Cloud Usage — Cloud Accelerator

- **Cloud Access Without Replication:** Deliver secure, low-latency access to on-premises datasets directly from cloud compute environments — without costly data copies.
- **Resource Hunting and Burst Compute:** Seamlessly connect to additional cloud infrastructure for on-demand burst compute (e.g., GPU clusters, rendering farms) and take advantage of cross-region pricing and availability.
- **Real-Time Collaboration:** Collaborate across cloud and on-premises environments with strict data consistency and immediate visibility.
- **Workload-Specific Accelerators:** Customize Cloud Accelerator for specific workflows such as AI training, rendering, or analytics.

Cloud Accelerator Use Cases:

- ✓ Cloud burst compute for AI and analytics
- ✓ Multi-cloud collaboration and cross-region workflows
- ✓ Pay-per-use access to cloud resources without data egress penalties

Ideal for These Organizations

Enterprise IT & DevOps seeking to unify distributed workloads, Media & Entertainment needing remote editing and cloud burst rendering, Life Sciences & Research workflows with large datasets across hybrid platforms, AI/ML Teams requiring immediate access to training data across locations, and edge-centric operations where remote ingestion and low-latency access are critical

Licensing & Deployment Considerations

Edge Accelerator

- Available in two tiers: SoHo (up to 32 TB) and RoBo (up to 100 TB local cache).
- Used as a local cache endpoint with at least one Cloud Data Fabric instance.

Cloud Accelerator

- Charged based on throughput/IOPS consumed.
- Not sold as a standalone product; pairs with Qumulo Cloud Data Fabric.

Note: Accelerators are designed for caching and read-write access, not as the primary data store — full data durability and protection remain within the core Qumulo Cloud Data Platform.



Conclusion

With Qumulo Accelerators, organizations no longer need to choose between performance, consistency, and flexibility. By extending the Cloud Data Fabric intelligently to edge and cloud environments, Qumulo enables real-time collaboration, efficient data access, and workload acceleration — while simplifying operations and reducing costs.

Get started today — project data where you need it, collaborate without friction, and unlock new levels of hybrid data performance.

About Qumulo

Qumulo is the leading provider of cloud file data platforms, offering unrivaled performance, scale, and data management solutions. Qumulo's platform is trusted by Fortune 500 companies and global enterprises to manage petabytes of data, enabling them to unlock the value of their data and drive innovation. For more information, visit www.qumulo.com