Qumulo

Solution Brief

Qumulo Cloud Data Fabric

When businesses look to enhance global collaboration across their various sites or leverage the public cloud, they often find multiple technical challenges in their way:

- Centralizing data or moving it to the cloud can lead to high latency, which affects performance for users and applications that remain on-premises and for those accessing on-premises data from cloud-based applications.
- Distributed organizations often replicate and migrate data to where needed to mitigate performance challenges. This results in data duplication, poor data governance, data security challenges, and increased costs.
- Existing on-premise storage vendors lack a public cloud counterpart that offers the performance and functionality of on-premise storage at a reasonable price point.

These challenges drive many organizations to defer modernization projects, or execute them with unsatisfactory results.

Breaking the Barriers to Business Transformation

Qumulo Cloud Data Fabric is an enterprise solution for managing unstructured data, addressing the challenges of traditional network-attached storage appliances. It facilitates data portability, allowing organizations to take advantage of the economic and operational benefits of private, public, and hybrid clouds. This platform enables the rapid acquisition of talent and partnerships regardless of geographic location and enables organizations to utilize public cloud services based on capability, availability, and cost, ultimately driving market competitiveness and advantage.

Why Cloud Data Fabric Stands Above The Rest

Qumulo Cloud Data Fabric removes the technical and economic barriers to making data available where and when it is needed. It consists of five core features that are directly integrated into the Qumulo Cloud Data platform. This delivers a unique solution that seamlessly integrates edge, data center, and cloud into a single cohesive data platform, where data is accessible where and when it is required, instantly.

- Seamlessly connects cloud with on-premises storage clusters to enable hybrid cloud use cases such as burst compute, cost-effective data archival and consumption of innovative cloud application services.
- Facilitates AI initiatives by expressing data directly adjacent to Cloud AI training, services, and GPU compute, based on availability and cost.
- Facilitates access to data anywhere, streamlining skills and partnership integration, regardless of geographic location.
- Enables organizations to consolidate their data footprint into the public cloud and provide access to it globally across all branch and home offices.

Run Anywhere

Qumulo ensures feature parity across all platforms, giving you access to enterprise-class data services on AWS, Azure, on-premises data centers, or edge computing. It operates seamlessly on major hardware brands like Dell, Cisco, and HPE, as well as various hypervisors.

Cloud Native

Qumulo's cloud-native architecture enables scalable capacity, I/O performance, and throughput, allowing you to pay only for what you use. This approach simplifies the economics of cloud storage, offering a total cost of ownership (TCO) comparable to on-premises solutions.

Cloud Data Fabric

Today’s organizations have workloads literally everywhere – not just on-premises, but in the cloud, at the edge, in data centers, and across the globe. Qumulo is the only data platform and file system that enables instant data access to power your workloads no matter where they run.

Cloud Data Platform

Qumulo allows you to unify file data from all Qumulo instances globally into a single logical volume. Our Global Namespace ensures that every workflow can continue using the same data as it transitions from on-premises to cloud to archive.

Qumulo Cloud Data Fabric

Global Namespace

At its core, the Cloud Data Fabric is powered by Global Namespace which provides visibility and access to all files combined with a strictly consistent, geographically distributed, high-performance cache to enable data to be accessed anywhere. Strict consistency across the fabric eliminates the risk of accidental corruption and maintains application integrity, simplifying integration with existing applications. In addition, file permissions and metadata are immediately consistent across the entire Cloud Data Fabric. Changing permissions on a file will immediately update permissions across all different locations where that file is visible.

Predictive Caching

Predictive caching combines heat map technology with patented statistical models to analyze data access patterns and intelligently prefetch data that can be cached before it is requested based on client activity. As users or applications access files in a Global Namespace, the data is persistently cached locally allowing subsequent requests for the same data to be served from the local cache, significantly reducing WAN round-trips and providing sub-millisecond performance for users and applications. The cache is automatically managed, maintaining active data and metadata while evicting stale data as it fills up.

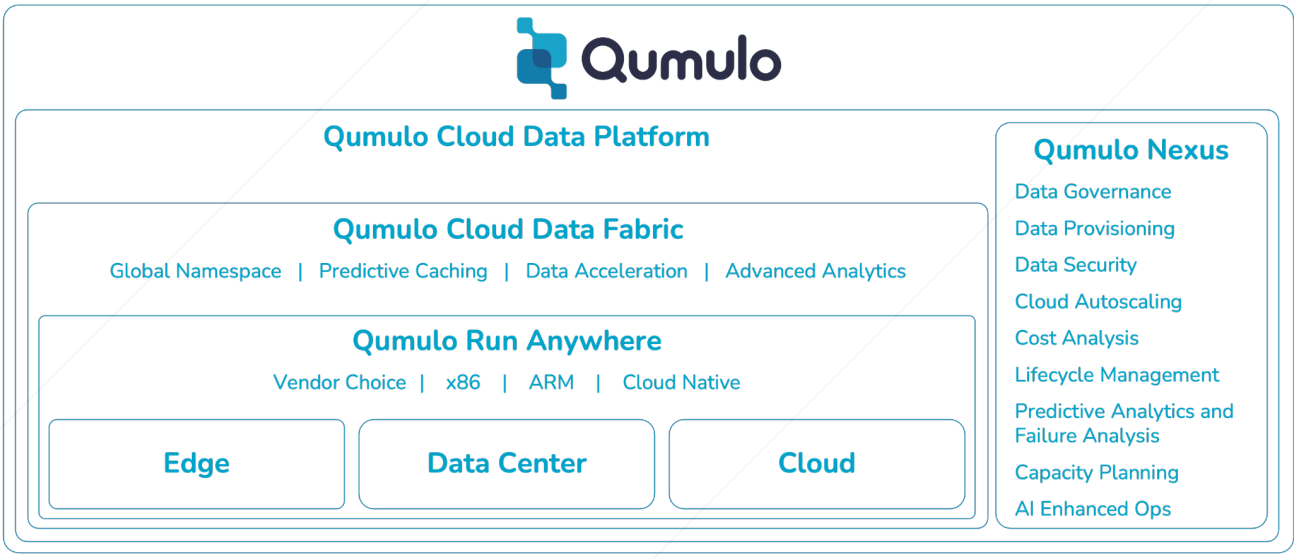

Data Portals

Data Portals are a core technology of Cloud Data Fabric that facilitates the transfer of data between Qumulo instances, providing real-time access to data without users needing to know that it's stored remotely in another data center or in the cloud. Organizations can create a Data Portal instantly anywhere within the Qumulo Cloud Data Fabric. This capability allows them to scale from a single Edge node in branch offices to thousands of cores, accommodating the most demanding high-performance burst compute and AI workloads, and to scale down again once the demand decreases.

Qumulo Cloud Data Fabric

The Data Portal controls access and manages the remote cache, evicting the least recently accessed data. If the Data Portal is terminated, all cached data and metadata are cleared from the cache, ensuring no further access is possible. This feature allows organizations to share all of their data or specific subsets, thereby improving data governance, security, and compliance with regulatory requirements.

Data Acceleration

All data transferred across a Data Portal is transferred in blocks independent of the file access protocol, accelerating application responsiveness. When a user accesses a file, SMB and NFS are terminated locally, and the Data Portal transports metadata and data blocks between a hub and spokes. When a user writes to a file, only the blocks that have been written need to be updated across the fabric, not the entire file, which reduces data transferred and improves performance. This combination of data block with local protocol termination and acknowledgment improves application responsiveness and accelerates data transfers when compared to running native file access protocols across WAN links.

Advanced Data Analytics

Qumulo's Cloud Data Fabric provides extensive metadata through a RESTful API, enabling advanced data analytics for informed business decisions. It allows for quick data discovery, helping users locate files by criteria like creation date, size (e.g., over 256GB), and access patterns. Users can search for specific keywords to analyze file content and track access and modifications for better data governance. With the Qumulo Copilot Connector, customers can use generative AI to uncover insights in unstructured data. The RESTful API also facilitates custom workflows and automation to enhance operational efficiency.

Summary

Qumulo Cloud Data Fabric marks the dawn of a new era of data potential by breaking through the technical and economic barriers of today's unstructured data platforms. Organizations can now extract insights from their unstructured data, no matter where it is stored. This shifts the focus from managing data limitations to unlocking business potential. It gives organizations the agility to grow, adapt, and innovate by connecting global teams, unifying data across locations, and accelerating insights through cloud integration.

For More Information

Contact your local sales representative today for more information, or reach us at https://qumulo.com/contact to speak with a Qumulo technical expert.

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.