SOLUTION BRIEF

Cloud Native Qumulo on Google Cloud

Cloud Native Qumulo on Google Cloud for any File-based Workload

Introduction

Enterprises today are managing more than 100 exabytes of unstructured data in their data centers—a testament to how vital and fast-growing this information has become. But with growth comes opportunity: the chance to unlock the value of your data with the flexibility, scalability, and innovation of the cloud. With Qumulo, your data isn't stuck—it's free to power the next generation of insights, innovation, and growth

Cloud Native Qumulo (CNQ) makes that transition seamless. With exabyte-plus scalability and full support for any file or object-based application, Qumulo empowers you to accelerate your journey to the cloud by bringing every workload—from massive archives to the most demanding HPC applications—out of legacy infrastructure and into the cloud without extensive replatforming. Instead of being constrained by compatibility gaps, limited features, or runaway costs, you gain a platform designed to scale effortlessly and economically.

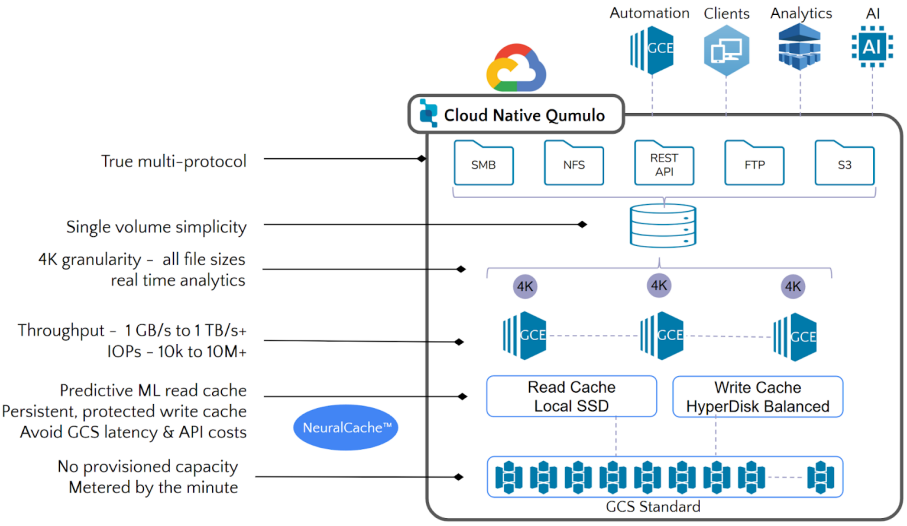

Cloud Native Qumulo (CNQ) on Google Cloud is explicitly engineered to leverage native cloud compute, object storage, and networking resources, creating the only fully scalable, fully elastic file and object data service available on Google Cloud today. CNQ's fully customizable architecture can be configured for the specific capacity, throughput, and IOPS requirements of any file or object-based workload. It also means that CNQ can deliver comparable performance and services to on-premises file storage, and at a similar TCO. In fact, a CNQ instance can also save you up to 80% on costs compared to other cloud-based file services!

Dynamic scalability
CNQ on Google Cloud can scale automatically to exabytes in a single namespace simply by adding data. Not only that, CNQ on Google Cloud can be easily reconfigured to dynamically scale throughput and capacity independently to support your workload needs. Plus, you pay only for the capacity and performance you actually use.

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—while ensuring sovereignty and security.

Deployed in minutes
CNQ runs in your own Google Cloud VPC, deployed via Terraform, enabling you to build a complete file data platform on Google Cloud in less than 15 minutes! You choose the specific GCE instance type you need to satisfy your workload's performance requirements.

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%.

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.

Qumulo

 | info@qumulo.com

SOLUTION BRIEF

Cloud Native Qumulo on Google Cloud

No limits, no compromises

Beyond the inherent scalability and dynamic elasticity in every deployment, CNQ supports enterprise-class features such as snapshots, replication, quotas, and multi-protocol support – NFS, SMB, S3, and REST APIs – for all your data needs.

By enabling you to share the same data via both file and object protocols, CNQ not only enables collaborative and mixed-use workloads, but it also eliminates the need to import file data into object storage. It delivers consistent time-to-first-byte latencies of 1-2ms – up to 100x faster than native Google Cloud Storage – to satisfy your most performance-intensive AI and HPC workloads.

Disruptive Features at a Disruptive Price

CNQ is also the only file service on Google Cloud that can run in any Google Distributed Cloud (GDC) location as easily as in any Google Cloud region worldwide, giving even your on-premises data centers the fastest possible access to your data.

CNQ can also be dynamically reconfigured without taking services offline, enabling you can boost performance temporarily as needed, or permanently if your workloads change. As an example, a CNQ instance initially deployed as a disaster recovery or archive target can be converted from a low-cost, high-capacity service to a high-performance data platform in seconds, improving business resilience.

Use Cases

- Video editing, rendering, burst processing
- PACS and Vendor Neutral Archive
- Genomic sequencing
- Long-term data archives
- Enterprise backup storage
- Disaster recovery from primary / on-premises storage
- Ransomware protection
- ArcGIS maps
- ADAS pipelines
- Energy subsurface modeling

Cross-site data mobility

If you already use Qumulo storage on-premises or on other cloud platforms, Qumulo's Cloud Data Fabric (CDF) services enable seamless movement of data between your on-premises, edge, and Google Cloud-based deployments, creating a single, globally consistent file system namespace. It enables you to easily build a follow-the-sun

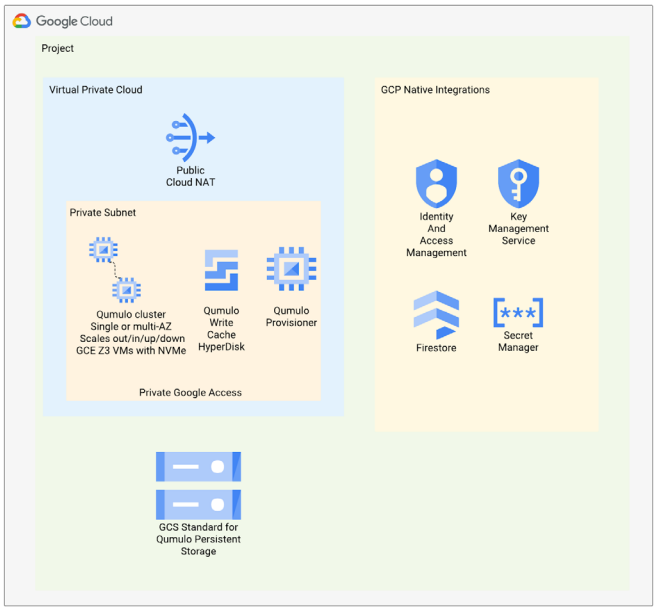

workflow or move files through large-scale data pipelines. CDF also simplifies data governance and security by unifying access controls and data visibility across sites and platforms.

Additional information for Google Cloud users

If transacted via Google Cloud Marketplace, any compute, storage, and licensing fees may count toward your Google Cloud EA.

Use custom connector modules to enable direct integration between data on your CNQ instance and Google Gemini's powerful AI tools to gain new insights into and uses for your data.

About Qumulo

Qumulo is the leading provider of cloud file data services, providing real-time visibility, massive scale, and API control of data across the data center, private, and public clouds. Organizations trust Qumulo to solve their most complex and mission-critical data challenges quickly, efficiently, and securely. For more information, visit www.qumulo.com.

Qumulo

 | cloud@qumulo.com