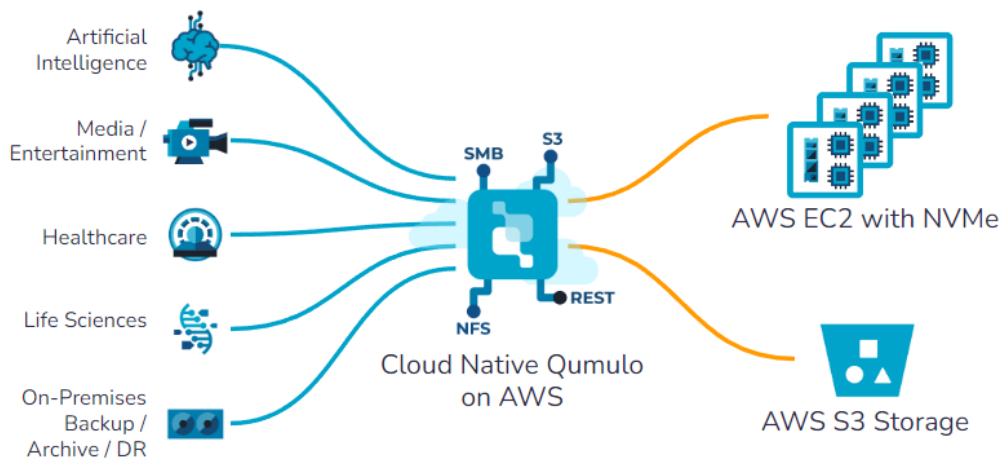


Cloud Native Qumulo on AWS for any file-based workload

Today’s enterprises are collectively hosting over 100 exabytes of unstructured data in on-premises data centers that are maxed out and can’t support the ever-expanding footprint of the modern organization. You want to move those workloads to the cloud, but you’re limited by the lack of compatibility in cloud storage, or by the lack of available features on today’s cloud-file service platforms, or by the high cost of running those workloads at scale in the cloud.

If your enterprise is looking to escape the confines of the on-premises data center, Cloud Native Qumulo’s exabyte-plus scalability and support for virtually any file- and object-based application let you migrate virtually any workflow, from your largest archive to your most performance-intensive HPC applications out of your legacy data center and into the cloud.



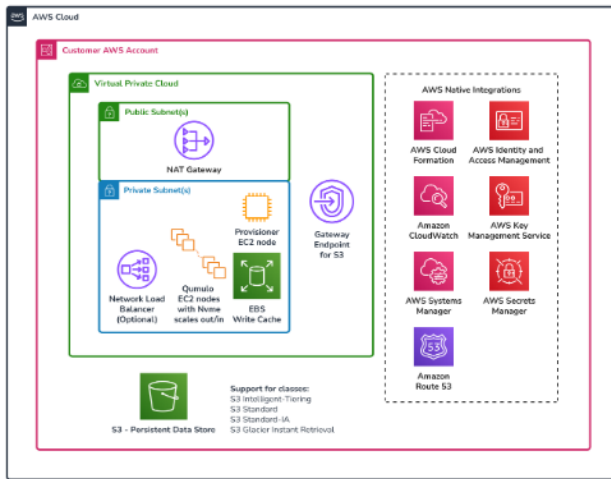
Any workload at any scale on AWS

Cloud Native Qumulo (CNQ) was engineered specifically to leverage native cloud compute, object storage, and networking resources to create the only fully scalable, fully elastic file and object data service available on AWS today. CNQ’s fully customizable architecture can be configured for the specific capacity, throughput, and IOPS requirements of virtually 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 save you up to 80% on costs compared to other AWS-based cloud file services!

Dynamic scalability
Every Cloud Native Qumulo (CNQ) instance on AWS can scale automatically to an exabyte in a single namespace simply by adding data. Not only that, but CNQ on AWS can be easily reconfigured to add or reduce performance simply by adding or removing compute instances as needed for your workload demands to increase throughput or IOPS. You pay only for the capacity and performance you actually use.

Deployed in minutes
CNQ runs in your own AWS VPC, deployed via either Terraform or CloudFormation, letting you choose the specific EC2 instance type you need to satisfy your workload’s performance requirements and build a complete file data platform on AWS in less than 15 minutes.

Manage your own TCO
Since CNQ can use any AWS S3 storage tier for the persistent data layer, you can choose the optimal S3 configuration for your specific workload, balancing the cost of storage used against the cost for data access.



Cloud Native Qumulo on AWS architecture

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, REST and S3 – for all your data.

Additionally, by letting you share the same data via both file and S3 protocols, CNQ not only enables collaborative and mixed-use workloads and eliminates the need to import file data into object storage, but also consistently provides time-to-first-byte latencies of 1-2ms – up to 100x faster than native AWS S3 – to satisfy your most performance-intensive AI and HPC workloads!

Disruptive features at a disruptive price

CNQ is also the only file service on AWS that can run in any AWS Local Zone as easily as in any AWS 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, so you can boost performance temporarily as needed, or permanently if your workloads change. A CNQ instance originally 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.

Suggested use cases

- Video editing, rendering, burst processing
- PACS and Vendor Neutral Archive
- Genomic sequencing
- Long-term data archives
- Enterprise backup storage (S3 with object locking recommended)
- Disaster recovery from primary / on-premises storage
- Ransomware protection for primary-site data.

Additional information for AWS users

- If deployed via AWS Marketplace, any compute, storage and licensing fees may count toward your EDP.
- Use S3 Intelligent Tiering to optimize costs for performance-intensive workflows, and S3 Glacier Instant Retrieval class storage for archive and other cold-data use cases.
- Use custom connector modules to enable direct integration between data on your CNQ instance and Amazon Q's powerful analytic tools to gain new insights into and uses for your data.

Cross-site data mobility

If you already use Qumulo storage on-premises or in other cloud platforms, Qumulo's native replication services enable seamless movement of data between your on-premises, edge, and AWS-based deployments, letting you build a follow-the-sun workflow or to move files through a large-scale data pipeline.

Use continuous replication for disaster recovery scenarios, or combine replication with Qumulo's snapshot feature to protect older versions of critical data from loss or ransomware.

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.