Qumulo Stratus - An Overview

Technical White Paper

January 2026

This document described Qumulo's fully disaggregated, share-nothing, product "Stratus," which is a cryptographically isolated, multi-tenant, exabyte-scale data platform. Engineered for a wide range of hybrid- and multi-cloud use cases, including the most demanding high-performance computing and artificial intelligence applications to cost-efficient cold archive storage, Qumulo Stratus empowers large enterprises and customers building secure, sovereign clouds to optimize their file and object workloads for capacity, price, and performance.

Introduction to Qumulo Core

Prior to outlining the motivations that led to Stratus, Qumulo's fully disaggregated, shared-nothing cryptographically isolated multitenant exabyte scale dataplatform, this section provides a refresher on the software architecture of Qumulo's Data OS.

Qumulo Core is a software data platform that can be deployed on a wide variety of storage server hardware and in various locations ranging from traditional on-premises data centers, mobile edge locations and all major public clouds. It supports file & object based workflows. Additionally, a robust public API set enables automated management and real-time visibility into system and data usage.

Fundamentals of Qumulo Core

The following foundational characteristics of Qumulo Core are important to enumerate:

1. Qumulo Core provides a 100% software-defined, exabyte-scale, distributed file system that presents a single namespace. An on-premises Qumulo cluster consists of an aggregation of independent storage server nodes, each node contributing to the cluster's overall capacity and performance. Individual nodes stay in constant coordination with each other. Any client can connect to any node and read and write across the entire namespace.

2. In the cloud, Cloud Native Qumulo uses object storage (AWS S3, Microsoft Azure Blob Google Cloud Storage, OCI object storage and their variants) for the data layer, in which the blocks associated with any given file are abstracted and distributed across a logical collection of discrete objects. The details of Cloud Native Qumulo are covered in later sections.

3. Qumulo Core is optimized for scale. Any on-premises or Cloud Native Qumulo instance can comfortably support tens of petabytes to exabytes of data, hundreds of billions of files, millions of operations, and thousands of users, all in a single namespace.

4. Qumulo Core is self-optimizing for maximum performance. Every Qumulo instance tracks data access using a heat map to identify frequently-accessed data blocks. These blocks are proactively moved by an internal prefetch algorithm:

    ○ Data blocks on hard disk media or cloud object storage are moved to flash storage as their heat score increases

    - If a given block’s heat score continues to rise, data that is already on flash storage is proactively moved to system memory for even faster access
    - At a global level, across all Qumulo instances for all Qumulo customers, the cache hit rate is ~95%

5. Qumulo Core is highly available and immediately consistent, built to withstand component failures in the infrastructure while still providing reliable service to clients. This is accomplished through the use of software abstraction, erasure coding, advanced networking technologies, and rigorous testing. When data is written to Qumulo’s file system, the write operation is not confirmed to the service, user, or client until the data has been written to persistent storage. Thus any subsequent read request will result in a coherent view of the data (as opposed to eventually consistent models).

6. Qumulo Core delivers platform-agnostic file and data services for public-, private-, and hybrid-cloud workflows. Qumulo’s software abstracts the underlying physical or virtual hardware resources thereby delivering feature parity across both public and private cloud infrastructure. This enables Qumulo to keep pace with ongoing innovations in compute, networking, and storage technologies driven by the cloud providers and the ecosystem of component manufacturers.

7. The Qumulo management model is API-first. Every capability built by Qumulo is first developed as an API endpoint. We then present a curated set of those endpoints in our command line interface (CLI) and the WebUI, our visual interface. This includes system creation, data management, performance and capacity analytics, authentication, and data accessibility.

8. Qumulo ships new software according to a routine monthly update schedule, which limits administrative burden while ensuring access to the newest features quickly. New updates are released every few weeks, enabling a rapid response to customer feedback while driving constant Qumulo Core improvements.

9. Qumulo’s container-based architecture enables a unique upgrade process that minimizes disruption to users and workflows. On a rolling, node-by-node basis, the new operating software is deployed in a parallel container to the old version. Once the newer container has successfully initialized, the older container environment is non-disruptively shut down and the upgrade proceeds to the next node until the entire cluster has been upgraded. Regardless of cluster size, even a major platform upgrade typically completes in minutes with minimal impact to active workloads.

Details about Qumulo Core's software architecture are outlined in the whitepaper Cloud Native Qumulo: an Architectural Overview.

Qumulo's Customer Success team is highly responsive, connected, and agile. Each Qumulo deployment has the ability to connect to remote monitoring via our Mission Qontrol cloud-based monitoring service while maintaining complete data privacy. Our customer success team uses that data to help customers through incidents, provide insight into product usage, and to alert customers when their systems are experiencing component failures. This combination of intelligent support and rapid product innovation powers an industry-leading NPS score of 95.

The case for a disaggregated shared-nothing architecture

Enterprise unstructured data estates double every two to three years. As recently as a decade ago the maximum density of state-of-the-art storage platforms was between 120 TB per rack unit for all-flash platforms and 180 TB for disk-based systems. In comparison, today's storage platforms are capable of hosting between 750 TB (SSD) per rack unit and 500TB for HDD deployments.

Modern enterprise file and object workloads have expanded in scope, with even modest datasets requiring dozens of terabytes of capacity and hundreds of thousands of IOPS. Today's AI / ML pipelines can require multi-petabyte datasets be served simultaneously to hundreds of GPUs and tens of thousands of CPU cores. While they may not require the intense real-time compute horsepower, today's archive datasets can scale to hundreds of petabytes of additional capacity.

Additionally, modern enterprise workloads now span geographies and teams, often crossing both logical organizational boundaries and physical locations and infrastructures.

With data centers reaching limits and prevailing power constraints, enterprises are increasingly prevented from deploying new hardware on-premises, effectively negating the benefits of today's data-dense storage servers. As a result, many data-heavy workloads have found their way into public clouds. Public clouds offered compelling economics and a level of elasticity not available on-premises; however, public cloud architectures also posed drawbacks, such as: (i) a lack of support for multi-protocol file data services in general, (ii) limited features compared to on-premises NAS platforms, (iii) feature disparities between cloud vendors, and (iv) a lack of interoperability between cloud and on-premises data services.

Cloud Native Qumulo (CNQ) was engineered to address these limits: it delivers the same enterprise-class data management features – exabyte scalability, multiprotocol support, snapshots, replication, quotas – on all public cloud platforms. Cloud Native Qumulo is covered in more detail in a later section of this document.

While multi-cloud hybrid has been compelling for a vast majority of customers, many larger enterprise organizations, with a broad range of large-scale workloads, have identified a need for public cloud-like resource elasticity, a consumption-based economic model, workload isolation, and strict tenant separation – but contained entirely within their own private data centers.

There are a variety of factors, including regulatory compliance, data sovereignty, ultra-low latency requirements, or simply a strategic decision to maintain direct physical control over sensitive infrastructure and data. Effectively, these enterprises require public cloud capabilities from on-premises infrastructure.

Faced with these strategic decisions, IT organizations must resolve new technical challenges within the same constraints. For example, data center operations are not configured to support workload elasticity, tenant isolation, or pay-as-you-go resource management, but neither are the underlying infrastructure platforms that power most of these same on-premises workloads.

The next section of this document will explore those limits in more detail, as well as outline how Qumulo Stratus addresses these requirements.

Legacy storage architecture limitations

Scale-out architectures designed in the 1990s were a compelling remedy to the limitations of scale-up storage designed in the 1980s.

Scale-up storage, in use for over 30 years, relies on centralized controllers that handle all storage operations. Storage capacity consists of one or more physical disk arrays, which can be increased by adding more physical disks, in the form of additional SSD devices, additional HDD devices, or both. To increase performance, customers need to either add compute capability to the existing controllers, or replace them with higher-powered hardware.

Every scale-up storage instance can reach a point at which no further expansion is possible because the central controller will reach capacity, regardless of how robust. This means no further disk arrays can be added, or the sheer volume of client traffic exceeds what the controller pair's network configuration can sustain. The only options are a forklift replacement of the existing storage, or the addition of another scale-up instance and subsequent rebalancing of data and workloads. Both of these require significant time and expertise while risking operational downtime.

In contrast, scale-out storage consists of multiple independent storage servers (each with its own complement of CPU, DRAM, network interfaces, along with integrated storage capacity in the form of HDDs and/or SSD/NVMe). Scale-out storage comes with several advantages over scale-up architecture:

- Performance and capacity scale linearly simply by adding more storage servers
- Each additional storage server adds one or more network paths to the whole, eliminating the inherent network bottleneck of scale-up controllers
    - *I.e., a single scale-out storage instance may have 100 or more unique paths to the shared network, whereas a scale-up storage array is often limited to two per controller*
- A hardware failure on any one storage server will have minimal effect on overall data availability since the workload can fail over to any of the remaining active nodes

The difference between scale-up and scale-out storage architectures are well illustrated by Figure 1 below:

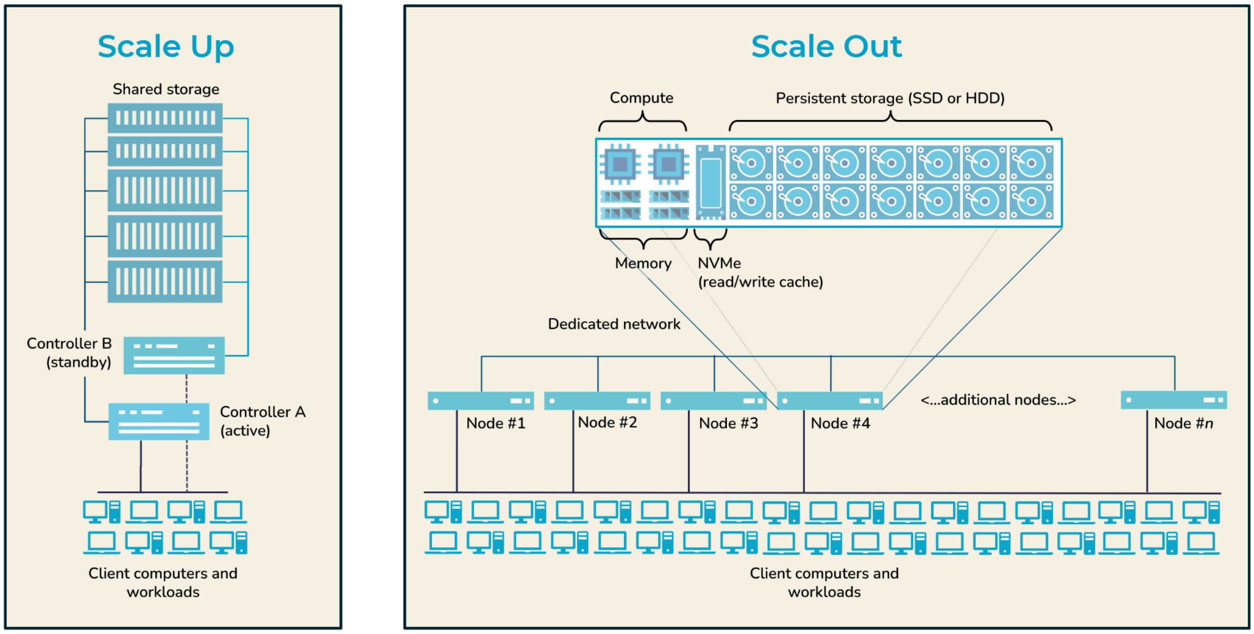

Figure 1: Scale-Up and Scale-Out Architectural Differences

As workloads become more varied and sensitive to change, even legacy scale-out storage architectures exhibit limitations:

Inability to independently scale performance and capacity

Adding a storage server to an existing scale-out instance linearly increases both the total amount of capacity available and the aggregate performance; however, this feature becomes less useful for high-performance workloads that need more IOPS, or capacity-heavy workloads (e.g. large-scale archive datasets) that have minimal performance needs.

IOPS intensive workloads have modest datasets (dozens of terabytes) and need a significant compute to deliver the needed hundreds of thousands to millions of IOPS. Examples of workloads that fit this profile are: OLAP databases, virtualized infrastructures, IOT data ingestion pipelines, messaging and queuing services. When using other scale-out storage systems, customers find themselves with “wasted” storage capacity when they size their systems to support the workload IOPS requirements.

Large scale, long term data archives (dozens of petabytes to exabytes) and related workloads require minimal performance. Examples of these workloads include: compliance retention, backup, reference datasets, etc. As such, customers with archive scale-out systems tend to have excess CPU, DRAM and network capabilities which, at scale, is economically inefficient.

Where scale-out storage forces some customers to purchase unnecessary capacity in order to achieve a given performance target, other customers find themselves having to buy more compute than they need in order to satisfy their capacity objectives.

A related limitation is the inability to mix high performance (high throughput and and/or high IOPS) workloads with cost-sensitive archival workloads, reducing the economy-of-scale that scale-out storage initially promised.

Further, the right mix of CPU, DRAM, network interfaces and storage media for the underlying storage server for these two workloads are different. A high performance cluster is typically composed of storage servers with more CPU, DRAM, network cards and NVMe storage, while an archive cluster is built from less expensive storage servers with less CPU, DRAM, NICs and mostly HDD storage.

Forced trade off: Security and QoS vs. workload consolidation and economies-of-scale

The inherent architectural performance and scaling limitations of scale-up storage systems result in siloed data for each application. The initial benefits of dedicated services with (i) customized Quality of Service (QoS) agreements and (ii) maximal security isolation are sidelined by the cost of operational complexity and stranded capacity due to inefficient resource utilization as business requirements evolve.

Similarly, scale-out systems offer benefits in the form of (i) the ability to aggregate multiple workloads and (ii) groups of users. Nevertheless, the consolidation of high-performance and cost-sensitive workloads onto a single shared storage infrastructure may compromise performance, economics, or both.

It is inherently impossible to deliver true QoS or guaranteed security isolation between workloads in any storage infrastructure where all physical resources – CPUs, DRAM, network capacity and storage media – are shared. Additionally, strict QoS guarantees across disparate workloads are not available with stateful protocols like NFS and SMB.

Logical security isolation via network segmentation and legacy multi-tenancy functionality (RBAC, per tenant admin, etc.) may satisfy the due-diligence requirements in some heavily-regulated industries; however, security-conscious organizations ranging from global financials to national intelligence agencies and sovereign cloud providers, find the limitations of these features unacceptable.

Lessons from Public Cloud Providers

Public cloud providers are able to deliver what was once considered impossible:

1. Maximum utilization of underlying physical infrastructure resulting in lower CapEx and OpEx
2. Virtually guaranteed QoS on a per service basis
3. Total security isolation between tenants (customers)

They do this by strictly adhering to a few core architectural principles:

- Elastic scaling of all services
- Single purpose services
- Standardized physical infrastructure that is operationalized via code (Infrastructure-as-Code)
- Strictly isolated logical infrastructure on top of shared physical infrastructure
- Hard boundaries between layers and services

The Qumulo team was inspired by the scale and reliability of cloud infrastructure while designing Qumulo Stratus. Qumulo set out to deliver a comparable architecture in an on-premises platform.

The first phase involved delivering Cloud Native Qumulo (CNQ), which can elastically scale both performance and capacity, and is limited only by the underlying cloud provider’s own resource availability.

CNQ’s architectural disaggregation (which is enumerated in more detail below) allows it to meet the specific requirements of all workloads, whether they need high IOPS, high throughput, or simply large scale, minimal data access.

The second phase is the launch of Qumulo Stratus, a fully-disaggregated, shared-nothing, cryptographically isolated multi-tenant storage platform. Together, CNQ and Stratus enables deployment across private data centers, neocloud facilities, and hybrid environments, delivering the same physical flexibility as any public cloud while maintaining logical isolation between resources and data.

The architecture of Qumulo Stratus is laid out in a later section of this document.

Cloud Native Qumulo Architecture

One of the biggest challenges to large-scale cloud adoption is the absence of cost-effective, scalable file services among public cloud providers. Most native file services in the cloud lack the rich data management features that enterprises require. Moreover, the few enterprise-class file services available in the cloud offer limited scalability at a high cost.

Most cloud-based file services, even those that include some enterprise-level features and functionality, lack an elastic consumption model, forcing customers to provision multiple storage silos for different workload classes, or overprovision capacity to meet performance objectives. Despite the pay-as-you-go model that cloud computing inherently suggests, cloud file-services customers are still billed for the capacity they provision, not the capacity they actually use.

Qumulo completely re-engineered the data persistence layer in the file system for the cloud. CNQ leverages the cloud provider's own cost-effective native object storage services (e.g. AWS S3, Azure Blob Storage, Google Cloud Storage, or Oracle Object Storage) for the durable data store. Low-latency, durable write-back caching is provided by the cloud provider’s own block storage services, and local NVMe delivers high-performance read caching.

Figure 2 below shows how Qumulo Core has been engineered to leverage cloud-object storage to provide the underlying durable data layer.

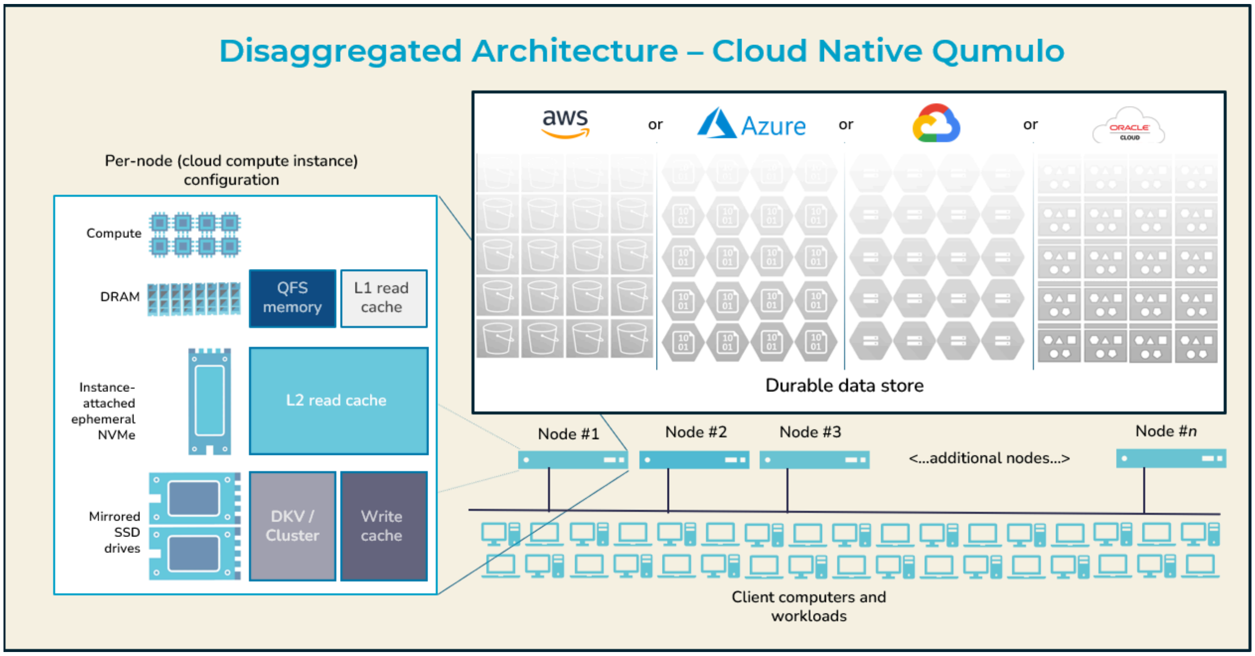

Figure 2: Cloud Native Qumulo Architecture

This approach provides built-in performance, cost-effectiveness, high availability, resiliency, elasticity, and security. Combined with CNQ's architecture, performance can now be dynamically adjusted up or down by adding, changing, or removing nodes without disruption. This process takes only minutes, as there is no need to restripe or reprotect data. CNQ can deliver any performance at any capacity.

As a result, CNQ has eliminated the need for performance tiering altogether. Unlike other cloud file storage solutions, CNQ customers pay only for the capacity and performance they actually use. It delivers low-latency file access due to architectural components like its Intelligent Cache Manager, and allows customers to elastically change the file system's performance capabilities in real-time and without disrupting active workloads.

Unlike other cloud storage file systems that are architecturally locked into a performance-per-capacity model, CNQ allows adjustments across performance and capacity independently. Customers have the flexibility to change underlying components, such as the compute instance type, compute instance count, and cache disk capacity, allowing for rapid and non-disruptive performance adjustments. This architecture, which includes Qumulo's NeuralCache, enables exceptional elasticity and virtually unlimited capacity. CNQ delivers an adaptive storage platform that enables efficient management and scaling data storage business needs evolve. CNQ does not compromise on performance or reliability. For variable workloads, customers can even scale back down again during low-demand periods to minimize costs.

CNQ retains all the core Qumulo functionalities, including real-time analytics, robust data protection, security, and real-time global collaboration. Its architecture is fully integrated into the cloud's elastic resource model, providing exceptional flexibility and efficiency, making CNQ ideal for both hybrid-cloud and multi-cloud enterprises.

Cloud Native Qumulo's architecture is outlined in a technical white paper that can be accessed on the Qumulo Website (Cloud Native Qumulo - An Architectural Overview).

The explanation of Qumulo Stratus builds upon the innovations provided by the CNQ and public cloud relationship by decoupling the public cloud dependency and supporting private- and neo-cloud architectures.

Qumulo Stratus Architecture

Qumulo Stratus is essentially the on-premises version of Cloud Native Qumulo: a fully disaggregated, shared-nothing, cryptographically isolated, multi-tenant, exabyte scale data platform for a wide range of use cases. Stratus supports the most demanding high-performance computing and AI applications as well as the cost efficient cold archive storage, all engineered and optimized to run in a private data center or private cloud.

Qumulo Stratus employs a hybrid deployment of Qumulo clusters spanning on-premises and (optionally) public-cloud environments, creating a single, scalable DataCore (shown in Figure 3 below) that hosts compressed, cryptographically-isolated data from numerous tenants. Each tenant data workload is supported by a dedicated Qumulo Core instance that can be deployed on physical or virtual compute infrastructure, with any single deployment able to scale to 100s of compute nodes. Critically, each tenant instance is fully independent of all others, ensuring dedicated QoS and secure tenant isolation for every deployment.

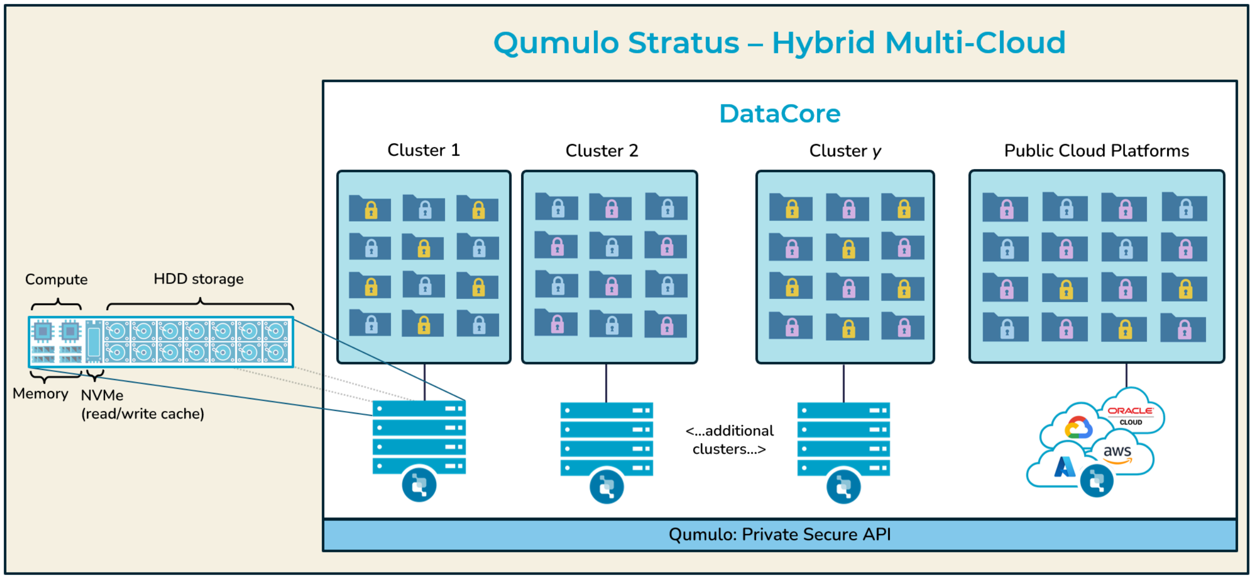

Figure 3: Hybrid-cloud Qumulo Stratus deployment

The core architectural principles of Qumulo Stratus are:

1. **Qumulo Stratus DataCore** - A Shared Backing Store

   The Qumulo DataCore provides a unified, scale-out storage tier. The DataCore can span multiple datacenters and optionally extend into public cloud providers (AWS, GCP, Azure, OCI). Powered by Qumulo's long established Data OS (See Qumulo Software Architecture Overview for technical details), Qumulo Stratus DataCore houses compressed and encrypted data for multiple tenants. Just as in public clouds, object storage services (e.g., S3) provides a giant, durable and scalable storage tier for a wide variety of users and applications, the Qumulo Stratus DataCore functions as a scalable, durable storage substrate for the modern hybrid enterprise.

2. **Qumulo Stratus Accelerators** - Shared Nothing, Scalable Compute & Cache

   The Qumulo Stratus Accelerators provide a scalable computing environment for an instance of Qumulo's Data OS dedicated to a single tenant. An unlimited Qumulo file system namespace, inherently multi-protocol (NFS, SMB, S3, FTP, etc.) can be elastically deployed on one or more (up to 265) Qumulo Stratus Accelerators delivering the performance (throughput and IOPS) required by the tenant application.

3. **Per-tenant Cryptographic Data Isolation**

   Each tenant possesses a unique cryptographic keyset, managed independently and never exposed to other tenants or even the Administrator of the Qumulo DataCore.

Tenant data at rest and in transit is encrypted using the tenant specific keys that are stored in tenant-specific Key Management Services (KMS).

4. **Per-tenant Network Authentication / Authorization Infrastructure**

   Per-tenant isolation of network authentication / authorization infrastructure (DNS, AD, LDAP, etc.) ensures that teams and workloads spanning different trust domains can utilize the same data infrastructure.

5. **Per-tenant Isolation of Administrator Control**

   The principle of least privilege is a core architectural principle of Qumulo Stratus. Qumulo Stratus DataCore provides a scalable shared data storage substrate for one or more tenant deployments, which are each dedicated to one tenant on a distinct Stratus cluster. Therefore, it is possible to guarantee that no administrator has access to any privileges more than minimally required and cannot compromise the data infrastructure as a whole (shown in Figure 4 below).

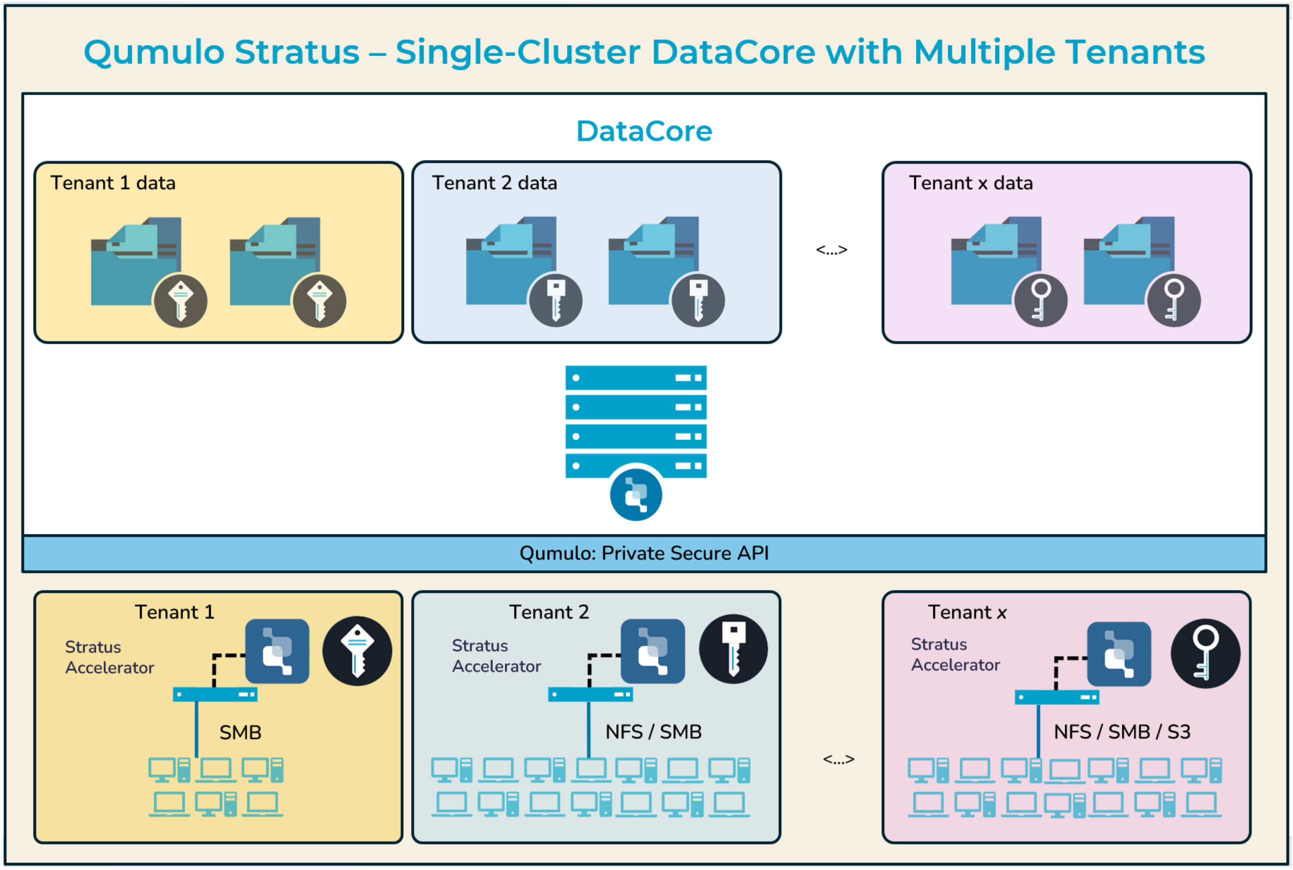

Figure 4: Qumulo Stratus - Single core multi-tenant deployment

Stratus DataCore

The Stratus DataCore provides storage capacity unmoored from the performance requirements of various workloads. Stratus DataCore is the shared but cryptographically isolated storage layer analogous to the object storage services in various public clouds.

Even a root administrator of the Stratus DataCore will only have visibility to the object stores that contain the obfuscated tenant data. Attempting to view the actual tenant data will yield a bit stream that is both compressed and encrypted with AES-256 keys that are inaccessible to the DataCore admin by design, whereas each tenant administrator is able to see all Qumulo-hosted resources and services throughout the organization, regardless of where it's physically located, as shown in Figure 5.

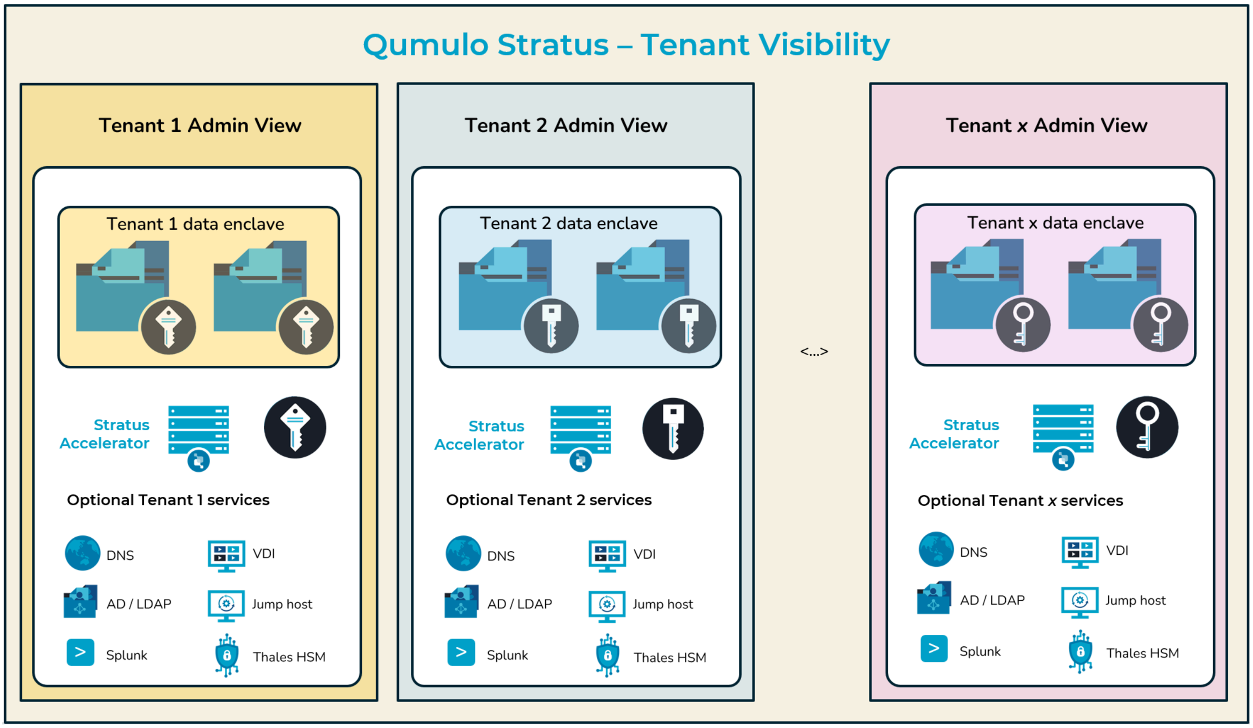

Figure 5: Qumulo Stratus - Per-tenant data isolation

DataCore Physical Infrastructure

1. A Qumulo Stratus DataCore is powered by Qumulo's long established Data OS (See (Qumulo Software Architecture Overview for technical details).
2. Data written to the DataCore is erasure-encoded across the entire cluster to ensure a Minimum Time to Data Loss (MTTDL) of 10,000 years (4x9s of data durability).
3. A DataCore cluster can consist of any supported storage server. Qumulo strongly recommends the use of cost-efficient, data-dense hybrid storage servers, since the

DataCore is optimized for capacity only – performance is managed exclusively at the per-tenant level.

4. Qumulo recommends that DataCore nodes be deployed with dedicated back-end 100GbE networking for intra-cluster (east-west) traffic, with a separate, dedicated front end 100GbE network for client traffic originating from the Stratus Accelerator layer.

DataCore Security

1. This layer does not terminate any client (NFS, SMB, FTP, etc.) traffic directly and is accessible exclusively by a Stratus cluster deployed on one or more Stratus Accelerators using a secure, S3-like private API that is not published or directly accessible by customers.
2. This layer does not need to be part of any AD / LDAP trust domain.
3. Data from multiple tenants residing in the DataCore is both compressed and encrypted (by Qumulo Core running on Stratus Accelerators) using AES 256 keys that are only accessible to the tenant administrator.
4. Data is segregated on a per-tenant basis into numerous logical folders / buckets each containing files that vary in size from 8MB to 64MB with compressed & encrypted tenant data. As such, files on the Stratus DataCore are logically comparable to sectors on a physical disk.
5. What the DataCore root Admin can do:
    a. Activities related to the monitoring and management of the DataCore cluster itself, such as:
        i. The addition of new nodes to increase cluster capacity
        ii. Replacing failed components
        iii. Network configuration
        iv. Software upgrades to Qumulo Core environment
    b. Control security policies for API access to the DataCore storage platform by tenant instances running on Stratus Accelerators.
    c. Implement and manage per-tenant quotas
    d. See the persistent object stores which host obfuscated tenant data
6. What the DataCore root Admin cannot do:
    a. Access tenant data in their original uncompressed, unencrypted format
    b. Access, manage or monitor any Accelerator
    c. A side-by-side comparison of a tenant administrator’s view, along with a Stratus administrator’s view of the same data, is shown in Figure 6.

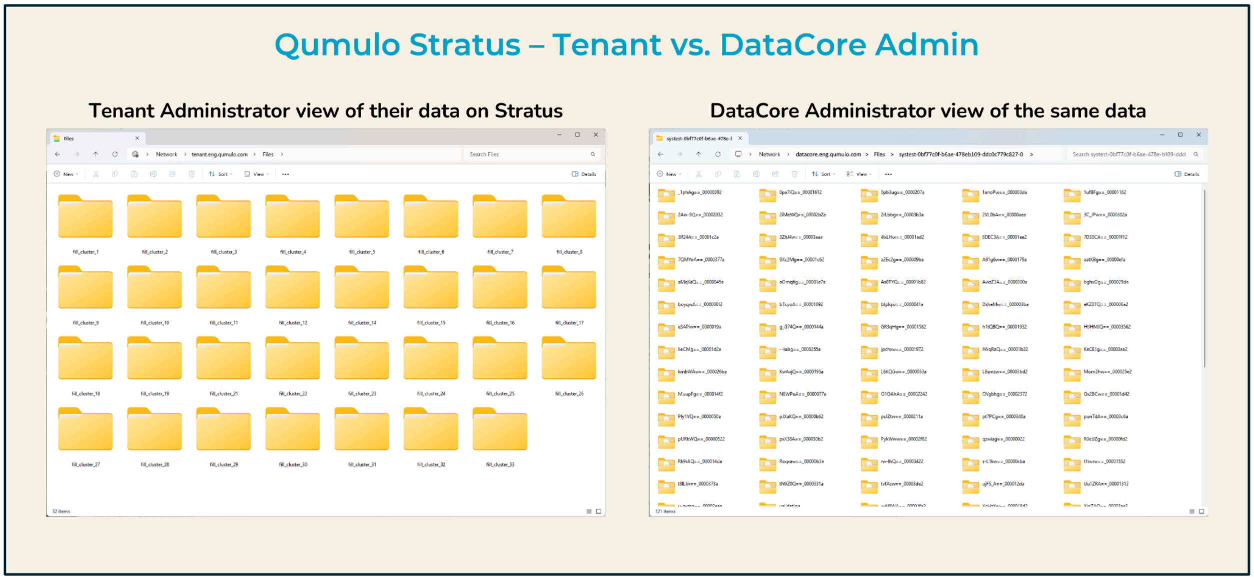

Qumulo Stratus – Tenant vs. DataCore Admin

Tenant Administrator view of their data on Stratus

DataCore Administrator view of the same data

Figure 6: Qumulo Stratus - role-based visibility to the same data

DataCore Scalability

1. A single DataCore can be as small as 10 nodes (or 10PB of raw capacity) and as large as 265 nodes (190PB - 356PB raw based on current data densities). Qumulo recommends that each DataCore be restricted to a single physical rack (containing all east-west, intra-cluster traffic within a single top-of-rack switch). Based on current drive densities, a Qumulo DataCore cluster that fully occupies a single rack would have a raw capacity of 28PB to 40PB+.
2. Multiple DataCore clusters can be deployed within the same data center or across data centers.
3. Since a tenant file system can span multiple DataCore clusters, the one-cluster-per-rack recommendation does not limit per-tenant file system size in any way.
4. Data in a DataCore is exclusively accessed via a private and secure RESTful API.
5. Customers can architect a storage platform similar to cloud storage infrastructure, mimicking cloud availability zones and regions, within their own data center.
6. The DataCore can also be expanded into public-cloud providers (AWS, GCP, Azure, OCI), enabling customers to leverage public clouds for burst storage capacity. Figure 7 depicts a Qumulo Stratus DataCore spanning multiple datacenters and public clouds.

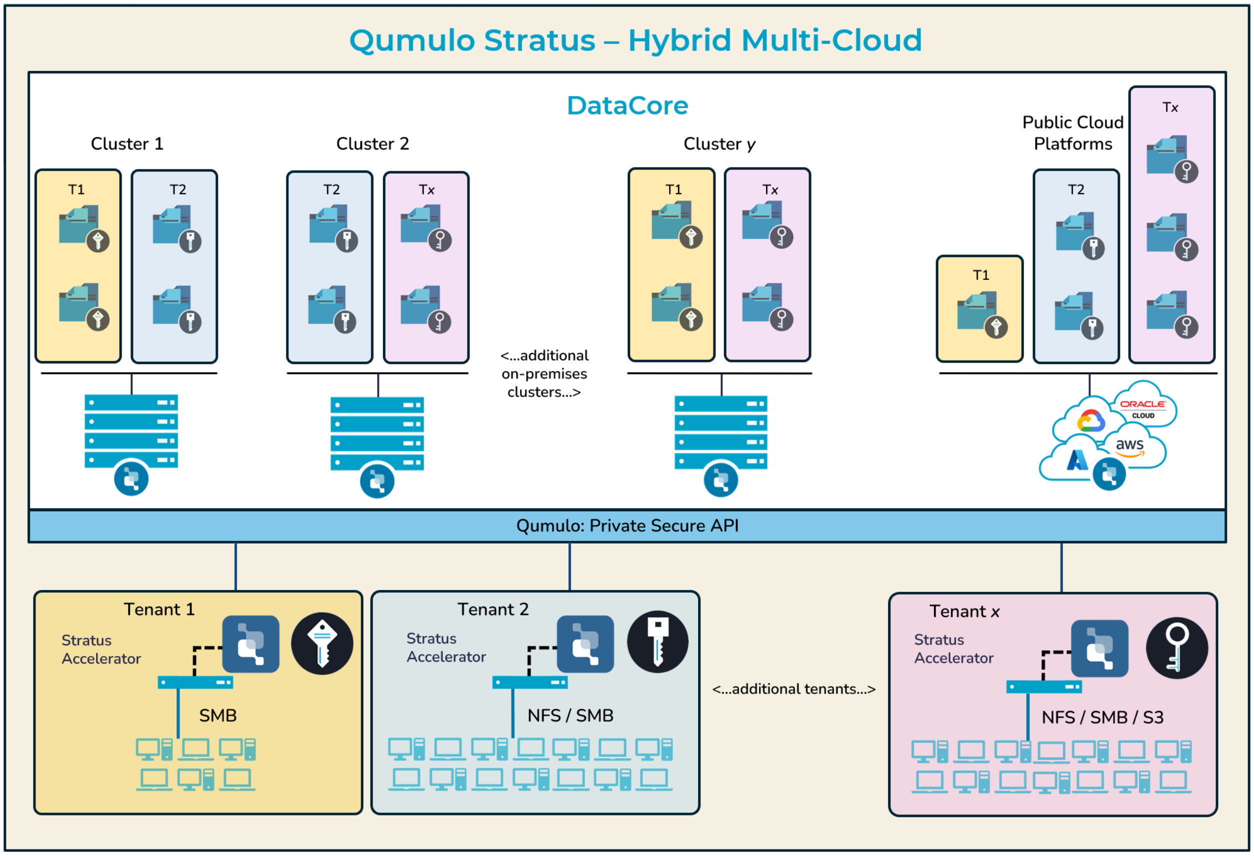

Figure 7: Qumulo Stratus - Hybrid Configuration with tenant data spanning DataCores

Stratus Accelerator

Stratus Accelerators can be configured to deliver unrestricted performance independent of storage capacity, with each Accelerator hosting data for a different tenant. Figure 8 below depicts this flexibility.

The software that is deployed on Stratus Accelerators is logically identical to Cloud Native Qumulo (see Cloud Native Qumulo - An Architectural Overview for a detailed technical description of this software) deployed in cloud compute instances. As such, the Stratus Accelerators can be thought of as similar to cloud compute instances (AWS EC2, Azure VM, GCP GCE).

From an administrative and user / application point of view, a Stratus cluster is similar to a single instance of Cloud Native Qumulo, with all the associated file, protocol and data services.

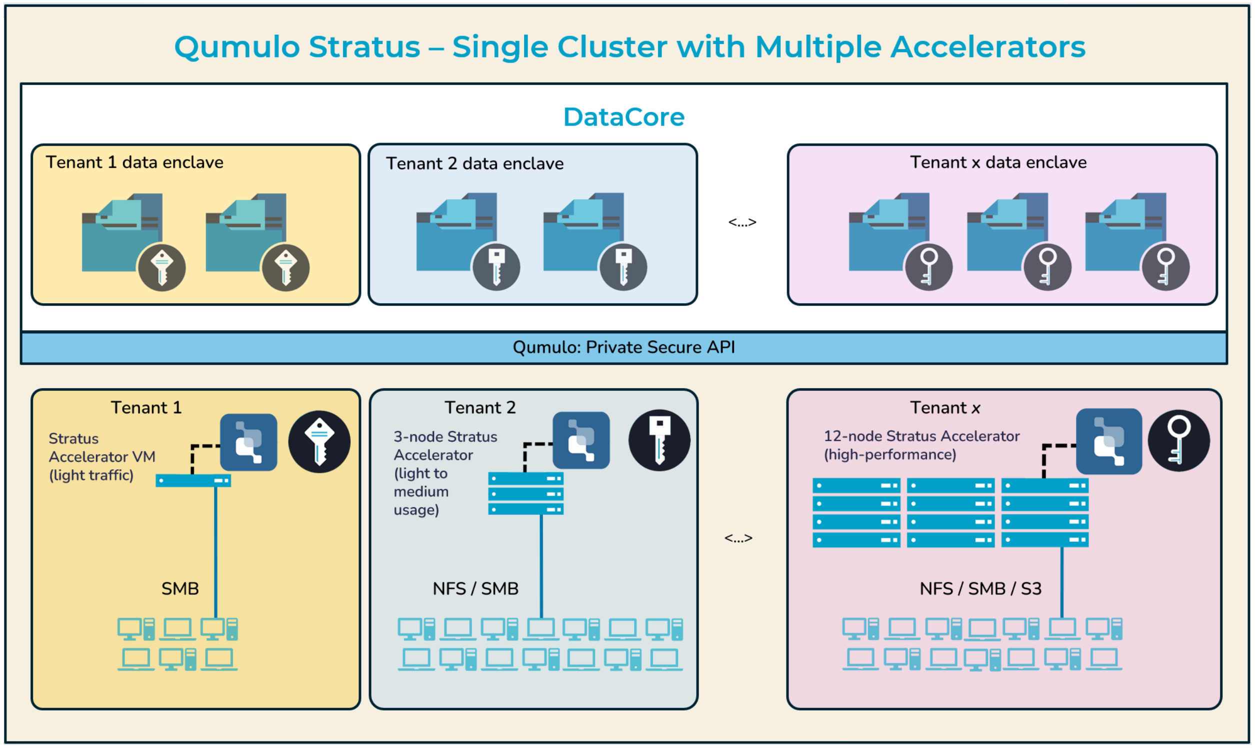
Qumulo Stratus – Single Cluster with Multiple Accelerators
DataCore
Tenant 1 data enclave
Tenant 2 data enclave
<...>
Tenant x data enclave
Qumulo: Private Secure API
Tenant 1
Stratus Accelerator VM (light traffic)
SMB
Tenant 2
3-node Stratus Accelerator (light to medium usage)
NFS / SMB
<...>
Tenant x
12-node Stratus Accelerator (high-performance)
NFS / SMB / S3

Figure 8: Qumulo Stratus Accelerators - Various Configurations

Stratus Accelerator Physical Infrastructure

1. Stratus Accelerators host the per tenant facing Qumulo file system (described in detail in Qumulo Software Architecture Overview and Cloud Native Qumulo - An Architectural Overview), protocol services (NFS, SMB, S3, FTP, etc.) as well as data services (Replication, Quotas, Real-time analytics, Cloud Data Fabric).
2. A Stratus Accelerator can be built up using a very wide variety of compute servers. The key requirements for a Stratus server are:
    a. CPU – x86-based, with more cores equating to better performance
    b. Memory – a minimum of 32GB of DRAM; performance increases with more memory
    c. Networking – Qumulo recommends the use 100GbE NICs that support RDMA
    d. Ephemeral Memory (non-boot NVMe media for caching) – Qumulo recommends at least four independent NVMe devices, adding up to a total of at least 12TB, per Accelerator instance
3. A Stratus Accelerator can be made up of virtual infrastructure as long as each of its constituent nodes adheres to the requirement laid out in (2) above.

4. A Stratus Accelerator does not contain any persistent user data. The entire cluster acts as a high-performance read and write-back cache and is powered by Qumulo NeuralCache.

Stratus Cluster Security

1. Qumulo recommends that one Accelerator be deployed per tenant.
2. Each Stratus Accelerator compresses and encrypts all associated tenant data (e.g., user / application file data, user / application file metadata, cluster metadata, etc.) using its own dedicated, private AES 256 keys prior to writing data to the Stratus DataCore, which serves as the long term persistence layer. Software encryption and related security features (including AD / LDAP integration) of Qumulo's Data OS are described in detail in the Qumulo Security Architecture and Practices whitepaper.
3. Each Stratus Accelerator has its own dedicated network and security services, e.g., DNS, DHCP, AD / LDAP, Logging Infrastructure, KMS, etc. It should be noted that such rigorous separation of tenant infrastructure is optional. Nothing prevents a customer from deploying multiple (or all tenants) with identical network and security services.
4. What a Stratus Accelerator root administrator can do:
    a. Cluster monitoring and management activities, including:
        i. Add or remove constituent nodes to increase or decrease cluster performance
        ii. Replace failed cluster components
        iii. Cluster network configuration
        iv. Cluster software upgrades
    b. Cluster-level data services, such as Cloud Data Fabric portals, snapshots, replication, share and export management, audit and syslog management, etc.
    c. Security management, such as RBAC roles, POSIX and ACL permissions settings
    d. Full access to the file system hosted by this cluster
5. What a Stratus cluster root administrator cannot do:
    a. Access Qumulo Stratus DataCore or any data that is stored in the DataCore, including the data belonging to the cluster they have root access to, except across normal file system operations
    b. Access any Stratus Accelerator without explicit entitlements/role-authorizations

Stratus Cluster Scalability

1. A Stratus Accelerator can comprise a single node (although Qumulo recommends a minimum of three nodes for availability) and up to 265 nodes depending on the performance needs of the tenant workflow
2. Tenant data is stored on one or more Stratus DataCore clusters (or in supported public clouds) that is exclusively accessed via a private and secure RESTful API
3. A customer can deploy a single-node Stratus Accelerator to support applications with infrequent multi-protocol file access
4. A customer can also deploy a Stratus Accelerator using high-performance compute nodes supporting performance-intensive applications (requiring high throughput and/or high IOPS)
5. A Stratus Accelerator can be scaled up (by adding nodes to increase performance) or scaled down (removing nodes and reclaiming unused capacity) non-disruptively in under 90 seconds

Advantages of Qumulo Stratus

With Stratus, Qumulo delivers a fully disaggregated, shared-nothing data platform with public cloud-like elasticity and security for all data workloads everywhere, including those hosted exclusively on-premises.

Cloud Economics

By consolidating all storage needs (both performance-sensitive workloads and large-scale, long-term archive data) into a single data substrate (Stratus DataCore) customers can leverage the most economical storage server hardware for all workloads, reducing capital costs and minimizing operational expenses. With Qumulo’s usage-based metering, Stratus enables customers to accurately bill each tenant based on their actual consumption, thus explicitly aligning end user costs incurred to the value of the service delivered.

A fully hybrid-cloud Qumulo Stratus environment is shown in Figure 9 below.

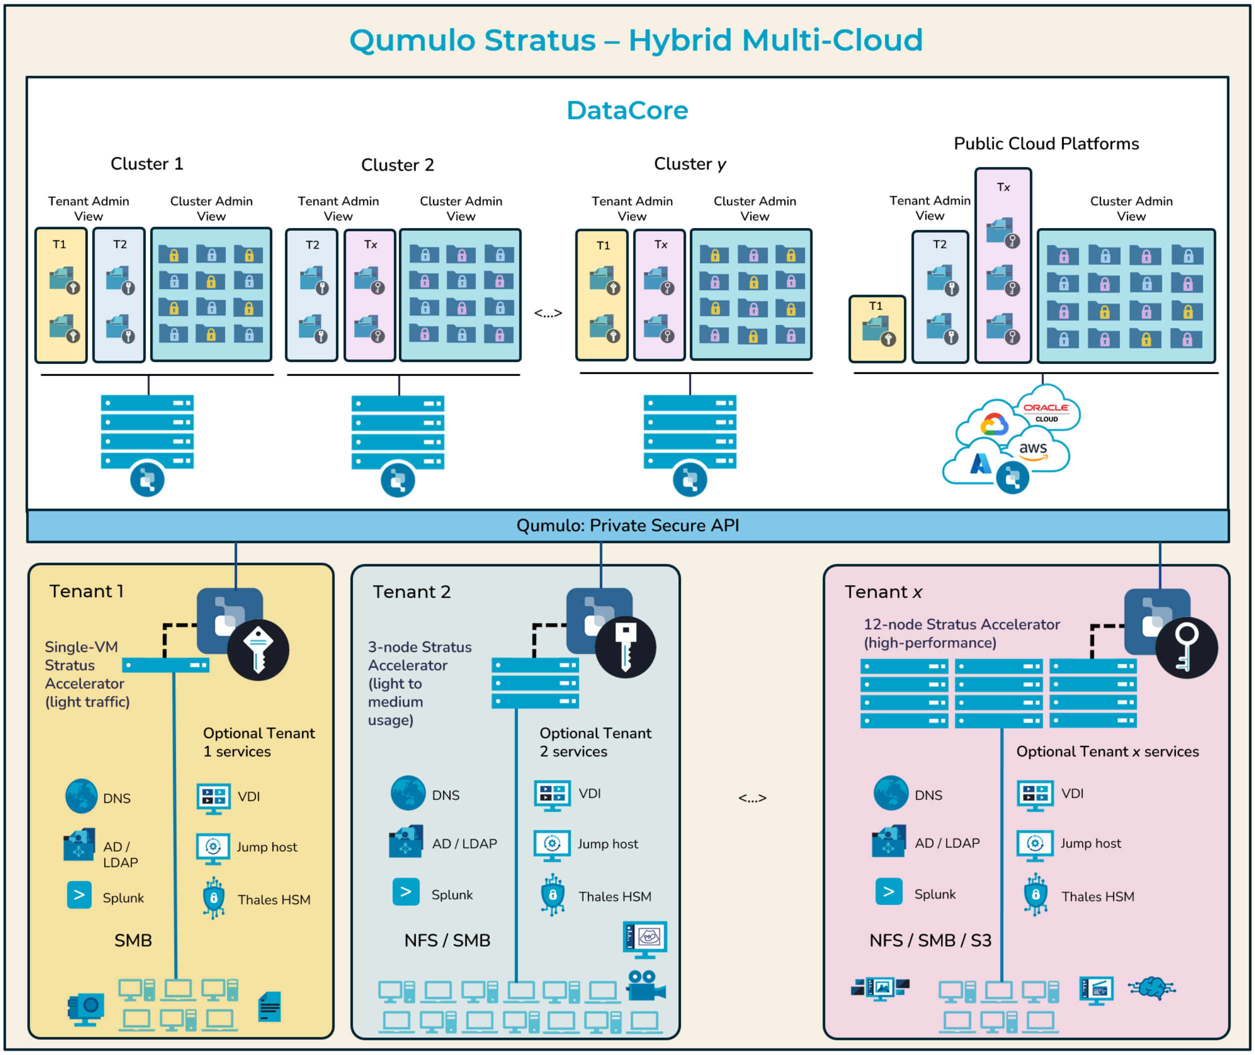

Figure 9: Hybrid-cloud Qumulo Stratus deployment showing per-tenant services and Accelerators, along with side-by-side Tenant and DataCore administrative viewable data

Per-tenant Storage QoS

By fully disaggregating storage capacity from performance, Qumulo Stratus enables complete isolation of one workload (Tenant-1) from another (Tenant-2), effectively resolving the noisy neighbor problem without having to deploy dedicated (and captive) storage infrastructure per workload.

Each tenant's workload is served by its dedicated Stratus Accelerator, whose Qumulo NeuralCache-powered performance is not impacted by the workloads of any other tenants whose data also resides on the same Stratus Datacore substrate.

Cryptographically Assured Multi-tenancy

The hard separation of the data-persistence layer (Stratus DataCore) from the data-access layer (Stratus Accelerator Clusters) creates a new gold standard for multi-tenancy. The innovative architecture of Qumulo Stratus delivers a level of QoS and security comparable to deploying dedicated and captive storage infrastructure per tenant, but at a fraction of the cost.

Any given tenant's persisted data is inaccessible to any other tenant or any outside user / administrator. With support for per-tenant network authorization and authentication services, Qumulo Stratus can support tenants that do not share a trust boundary while avoiding the expenses associated with per-tenant storage silos.

Conclusion

Qumulo Stratus delivers a complete solution for any large, multi-tenant environment demanding robust security, privacy assurance, and QoS guarantees, while simultaneously driving down the capital costs and operational expenses endemic to large data footprints. This contrasts with "shared everything" systems, which compromise security, data isolation and storage QoS.

Qumulo Stratus leverages both on-premises infrastructure and cloud storage forming a unified data substrate even in a hybrid- or multi-cloud enterprise. By providing each tenant with dedicated cryptographic keys, network, authentication and authorization services as well as compute and caching resource, Qumulo Stratus delivers complete isolation while harnessing the efficiencies of a unified storage platform

Contributors

This article is maintained by Qumulo. It was originally written by the following contributors.

Principal authors:

Kiran Bhageshpur | Chief Technology Officer at Qumulo

James Walkenhorst | Sr. Technical Marketing Engineer at Qumulo

Related Resources

- Qumulo Technical Overview
- Qumulo Resource Library
- Qumulo Scale Anywhere
- Qumulo Nexus
- Qumulo Stratus Press Release
- Qumulo Stratus Announcement Blog Post
- Qumulo Software Architecture Overview
- Cloud Native Qumulo - An Architectural Overview
- Cloud Data Fabric - A Technical Overview
- Qumulo NeuralCache Technical Webinar
- Qumulo Security Architecture Whitepaper

Appendix: Terminology

1. **Qumulo Core:** Refers to Qumulo's software defined storage and data management platform.
2. **Stratus DataCore**: Refers to the unified, storage tier that is completely disaggregated from all tenant services.
3. **Stratus Accelerators**: Refers to the per tenant scalable performance and caching layer fully disaggregated from data persistence / durability constraints.
4. **Stratus Cluster**: A cluster (1- 265) of Stratus Accelerators that deliver file, protocol and data services to a single tenant
5. **Stratus FS**: Same as a Stratus Cluster
6. **Stratus Tenant**: A distinct application / users (or group of related applications / users with similar storage needs and performance characteristics) within a distinct and possibly unique trust boundary
7. **Qumulo Filesystem**: Qumulo's scalable filesystem (`qfsd`) deployed across thousands of customers both on-premises and in various public clouds. Also abbreviated to Filesystem. See Qumulo Software Architecture Technical Overview for details
8. **Protocol Services**: Various network protocols supported natively by the Qumulo filesystem. Includes NFS v3, NFS v4.1, SMB 2.2, SMB 3, *FTP*, S3
9. **Data Services**: Refers to various enterprise features of the Qumulo Filesystem including Replication, Snapshots, Storage Quotas, Real-time storage analytics, Cloud Data Fabric, etc.