

Qumulo and Commvault — A Faster Data Protection Solution Built for Today's Massive Unstructured File Datasets

Even with the most reliable file data platforms like Qumulo, backing up data is an enterprise best practice, and in fact may be mandated in many organizations. Qumulo delivers data protection features at its core with **Qumulo Protect**, which includes efficient erasure coding, snapshots, replication of snapshots, and replication of an entire cluster to support disaster recovery failover configurations. For off-cluster data protection, **Commvault Backup & Recovery** integrates directly with Qumulo for a modern, comprehensive solution.

Backup to the Future

It's hard to beat an offsite backup for a simple, cost-effective data protection strategy. The challenge is managing backups in this era of massive data volumes and unprecedented file data growth—the sheer volume of data created in a given day can overwhelm the capabilities of many traditional backup solutions. More data silos solve the storage and backup problems in the short term, but create more management burden in the long run.

Together, Qumulo and Commvault offer a modern solution to this challenge. Both built for scale, we partnered to deliver an integrated solution to meet today's challenges and the challenges of the future, where data growth continues to increase year-over-year and datasets become increasingly unwieldy to manage and protect.

Why Traditional Solutions Can't Keep Up

Traditionally, once an initial backup is complete, incremental backups rely on 'tree walks' to discover file changes across the file system to ensure that changed files are copied but unmodified files are skipped. The problem is, these time-consuming tree walks can dominate the backup window, consuming massive amounts of compute and network resources, and leaving little time to actually back up files.

The result can be a system that is constantly struggling to keep up, leaving critical data at risk of not being backed within backup windows—which leaves IT departments out of compliance with their stated backup SLAs.

Qumulo + Commvault = Tightly Integrated Data Protection

Qumulo and Commvault worked to radically simplify and accelerate data protection management for large datasets. Commvault is tightly integrated using the robust Qumulo REST API framework to access Qumulo snapshot engine data. This data service provides real-time intelligence on which files have changed, eliminating the need for long delta tree walks—allowing incremental backup scans to complete in a few minutes start-to-finish, rather than hours. Since Commvault provides an “incremental forever/virtual full” backup model, consistently completed incrementals will provide customers with the consistency and reliability they are looking for.




COMMVAULT

Features & Benefits

- Built-in, end-to-end encryption, including at-rest and in-flight
- Intuitive, web-based user interface with real-time Qumulo file systems status
- Policy-driven automation, monitoring, and reporting; no cumbersome scripts required for efficient backups
- Quick, incremental backups ensure Data Protection SLAs can be met
- Modern data protection solution based without reliance on legacy protocols like NDMP
- Recovery of entire system, instance, application, or granular single file

How It Works

- Qumulo is a highly scalable, high-performance file data platform, providing enterprise clients with file data services across multiple protocols
- Qumulo offers robust REST API frameworks to access all its features, making it simple for Commvault to integrate tightly with Qumulo, gaining incredible performance and data awareness benefits without relying on legacy protocols like NDMP
- The result is a truly modern file data platform and integrated data protection solution that is standards-compliant, easy to scale or automate, and extremely reliable
- Data awareness makes incremental backups quick to start and finish reliably. Management of entire data platform is simple, with integrated analytics on a single pane of glass



First, the companies tackled simplifying management by integrating Qumulo v3 and higher into Commvault's Command Center, first appearing in Commvault Backup & Restore R11.21. This release significantly improved manageability by giving administrators a single view for managing backup and restore jobs and enabled them to gather information about the file system and backup job status.

Backup the Clock

It turns out, a good portion of the time consumed in running incremental backups can be the process to determine which files have changed, particularly on large datasets with petabytes of data and billions of files. Knowing just which files have changed ensures only the modified files are copied, saving time, network bandwidth, and backup media. Qumulo and Commvault tackled the issue of accelerating backup of large unstructured data sets by dramatically reducing the time it takes to get a list of changed files.

By leveraging Qumulo snapshot engine data accessible via Qumulo REST APIs, Commvault R11.22 and later can generate a delta list in a fraction of the time required by traditional solutions, massively accelerating incremental backup jobs. Scan times drop from what were once hours to mere minutes even at million or billion-file scale!

Incremental Backup Over SMB	Earlier versions	Commvault R11.22
1 million modified files	58 minutes	30 minutes (2x faster)
1 million new files in a directory	1 hour 20 minutes	3 minutes (16x faster)

Incremental Backup Over NFS	Earlier versions	Commvault R11.22
1 million modified files	7 hours 2 minutes	37 minutes (11x faster)
1 million new files in 1,000 directories	7 hours 13 minutes	15 minutes (28x faster)

The time saved scanning the system during an incremental backup means customers can dramatically reduce backup windows and spend less time monitoring them. This ensures critical business data is protected and available to recover within the allowance required by the company or departmental level SLAs.

Together, Qumulo and Commvault provide a fully protected solution for storage and backup of massive amounts of unstructured data, making incremental-forever backups of Qumulo systems with petabytes of data and billions of files fast and simple. To the delight of customers, this process takes place with no manual intervention and no dependency on legacy protocols like NDMP.

Qumulo and Commvault are both well known for scale and performance. Now our joint customers can backup their massive file data environments with ease and confidence!

About Qumulo

Qumulo hybrid cloud file storage delivers real-time visibility, scale, and control of data across on-prem and cloud. Qumulo customers understand storage at a granular level; programmatically configure and manage usage, capacity, and performance; and are continuously delighted with new capabilities, 100 percent usable capacity, and direct access to experts. More: www.qumulo.com

About Commvault

Commvault empowers you to recover intelligently through holistic risk-mitigation and prevention strategies, software, and support. The result is you can quickly recover data and applications from data-loss events — like ransomware — or anything else that could threaten your operations, systems, users, or customers.

