



Qumulo and Splunk: Modern Scale-Out Storage for Splunk Environments

Scalable Storage Solution for Splunk Big Data Analytics

Qumulo, a leader in scale-out storage, and Splunk, a provider of award-winning software Big data Analytics, have delivered a solution for providing scalable, efficient storage for Splunk Data, as well as API integration directly with Splunk.

Splunk is the global industry standard for Big Data Analytics. Splunk software transforms machine-generated data into valuable insights that can help make your business more productive, profitable and secure.

Utilizing Splunk, customers can search, explore, browse, navigate, analyze and visualize petabytes of data from one place.

Real-Time Visibility and Control at Multi-Petabyte Scale: Qumulo Cluster Management Application for Splunk

By design, when Splunk aggregates incoming data from many discrete sources, it stores that data in a storage pool called a data source. This architecture allows Splunk to ingest and process massive amounts of data from one central location rather than having to manage data from individual sources spread throughout the data center. The data source can grow exponentially based on which clients are pushing data to Splunk, how often, type of data, etc.

Qumulo's Analytics REST API provides real-time performance and aggregate information about the state of your Qumulo cluster. Qumulo Cluster Management for Splunk is a Splunk-certified application that uses Qumulo's REST API to provide real-time visibility and control for your Qumulo cluster. Our application defines a data source for Qumulo data that can be queried and cross-referenced with other enterprise machine data indexed in Splunk to help administrators reason about capacity and performance in their environment.

Qumulo's built-in dashboard UI provides great, at-a-glance information about the state of your Qumulo cluster. Using Splunk's advanced query language and auto-generated visualization, you can also create custom visualizations based on Qumulo cluster stats to augment the data visualization capabilities you already have with Qumulo. The Qumulo Splunk application ships with a built-in dashboard showing capacity and performance info, and it can be easily customized by enterprise admins to show different information as needed.

Lastly, the Qumulo Splunk application provides real-time statistics about your Qumulo cluster directly in the same management UI (Splunk) that is used for other enterprise assets... making it easier for administrators to manage Qumulo assets using the same tools they use for other assets in their environment.

Manage Your Splunk Data, Not the Storage

Splunk generates many petabytes of data, which needs to be indexed and made searchable for the entire enterprise. The massive volumes of data processed by Splunk can create challenging demand on legacy storage infrastructures.

Benefits for Splunk from a Modern Scale-Out Storage Solution

Multi-Petabyte Scale

Modern scale-out storage designed from the ground up for the new era of multi-petabyte data scale

Linear Scale-Out

As ingest of Splunk data increases both Capacity and Performance must grow in linear, predictable fashion.

Real-Time Visibility & Control

The ability to quickly and accurately understand how Splunk data is being written and utilized is critical to a successful an enterprise Splunk deployment.

Direct API Integration

The storage solution should have direct API integration with the Splunk application to be effective.

Qumulo Core is a modern scale-out storage system designed from the ground up for the new era of multi-petabyte data scale on premises and in the cloud.

Qumulo Core stores tens of billions of files with scalable throughput and is the only product that provides real-time visibility and control for file systems at petabyte scale. Storage administrators can instantly see usage, activity and throughput at any level of the unified directory structure, no matter how many files in their Splunk data footprint, allowing them to pinpoint problems and effectively manage how storage is used.

Qumulo Core and Splunk Implementation

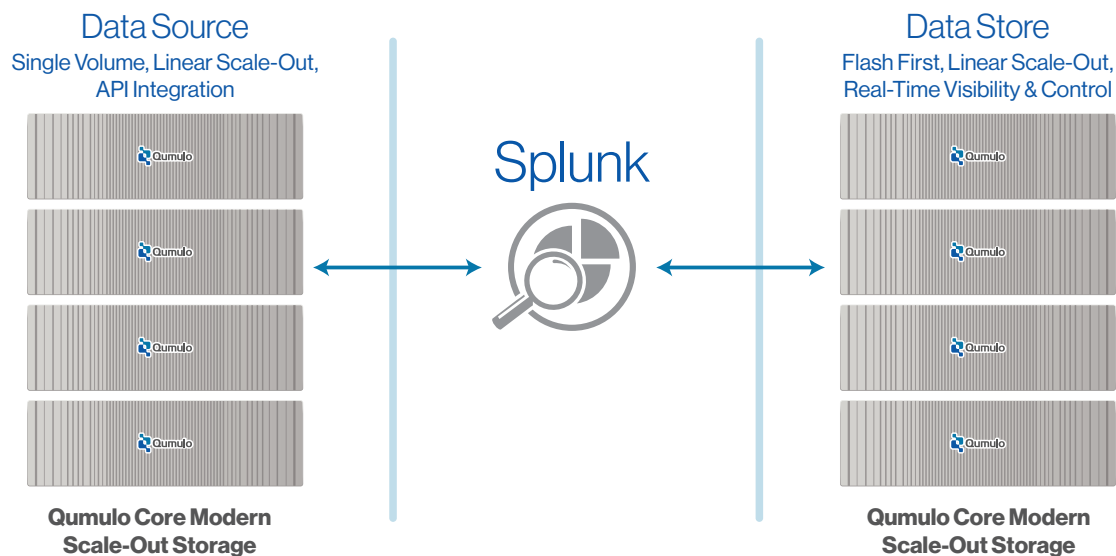
Qumulo Core was designed to address massive volumes of unstructured data in a linear and predictable fashion. Qumulo can be integrated with Splunk as a single volume Data Source to allow collections from a centralized mount point, or as the Splunk Data Store, providing the performance and capacity needed to address Splunk workloads.

Qumulo allows customers to scale Splunk Data Sources and Data Stores seamlessly and transparently with its modern scale-out storage architecture.

Qumulo provides built-in data analytics to provide detailed information on the efficiency and usage of the Splunk deployment. Qumulo can provide real-time and historical visibility and control into:

- IOPS
- Throughput/MBPS
- Capacity Usage
- Capacity Trends
- Predictive Bottleneck Detection
- Client Activity
- Top Storage Consumers
- Largest Files/Directories
- Most Active Clients
- Chargeback/Showback Models

Qumulo Core in a Splunk Architecture



Contact Qumulo to learn more about our Splunk solutions:
1-855-478-6856 • info@qumulo.com

