



SOLUTION BRIEF

Quantum ActiveScale with Splunk SmartStore Provides Superior Scalability

Highlights

- **Decouple storage and compute.** Splunk’s SmartStore architecture allows users to scale their storage capacity independently from the indexers.
- **Improve data retention.** ActiveScale provides extremely efficient and highly-scalable storage capacity, enabling users to keep their data for much longer.
- **Improve performance.** SmartStore uses locally attached SSDs as a smart caching layer to store data dynamically based on user activity.
- **Increase data protection.** ActiveScale provides up to 19 9’s of data durability and integrity, even at petabyte scale.

Challenge

- Companies are generating terabytes of machine data every day, rapidly consuming disk space and forcing users to scale out.
- Increasing data retention policies require data to be kept for years in order to maintain compliance with regulatory standards.
- Protecting business critical data with traditional replication requires three full copies of their data, further increasing storage issues.
- IT budgets aren’t growing as fast as the data, making it difficult to keep pace with the rapid growth of data without more efficient storage techniques.

ActiveScale™ Provides Superior Scalability and Cost Savings for Splunk Environments

Machine generated data is everywhere. More than 44 exabytes of data are generated every day, flooding out of billions of laptops, smart phones, and internet connected devices. For those who know how to take advantage of it, this information can be extremely useful. With the right tools, it can provide valuable insight into how systems are functioning, improve customer experiences and give you a real-time understanding of what’s happening in your business.

The best-known tool for analyzing machine data is Splunk. It allows users to quickly make sense of large amounts of machine data from any source, regardless of the data type. It can reveal long term trends and patterns of activity that you’d never know about otherwise.

The huge volume of data being created is often more than Splunk environments based on traditional storage infrastructures can handle. Large companies can easily generate terabytes of machine data every day, and at this rate, it can be difficult to maintain a year or even 6 months’ worth of data before running out of space.

The Quantum-based solution leverages the unique and powerful capabilities of ActiveScale object storage with Splunk SmartStore that enables the storage infrastructure to scale independently from compute. With ActiveScale and SmartStore, petabytes of storage can be added to any Splunk environment, without adding a single indexer!

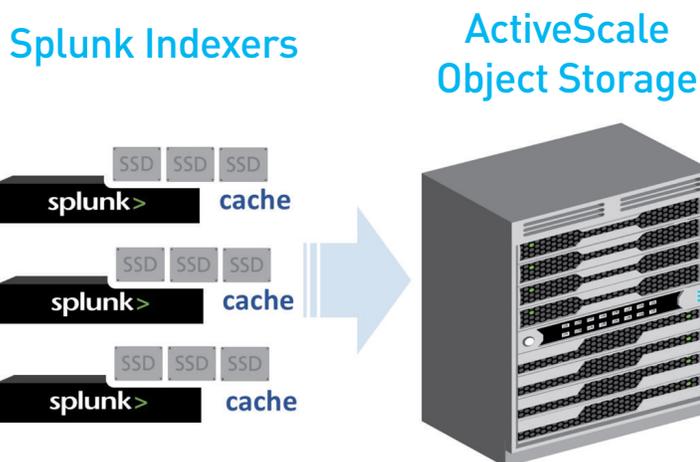


Fig. 1. Splunk Indexers to ActiveScale Object Storage Workflow Diagram

Solution

The combination of Quantum ActiveScale object storage with and Splunk SmartStore is a radical departure from the classic storage approach which allows Splunk users to scale their storage independently from the indexers.

This powerful solution changes the economics of Splunk storage infrastructures by providing petabyte scalability without imposing excessive scale-out costs. By combining Splunk with ActiveScale, users can grow their Splunk indexes well into the petabyte range without adding a single indexer. Performance is greatly improved, thanks to the new smart caching layer, which caches data dynamically, based on search criteria.

With ActiveScale and SmartStore, users can easily meet their data retention requirements while still remaining within budget.

The Traditional Approach is Good but Scaling Can Be an Issue

In Splunk environments based on the traditional architecture, data is stored on the indexers, using locally attached drives and files are organized into a series of directories called buckets. To prevent data loss, three copies of each bucket are created and distributed across multiple servers. This does a good job of protecting your data but consumes a lot of disk space. Maintaining three full copies of your dataset works for small environments but can be difficult and expensive to scale.

The Modern Approach is Scalability and Cost Effective with SmartStore

With a modern architecture that leverages ActiveScale and SmartStore, the storage is decoupled from the indexers so that you can grow them independently as needed. With indexes and raw data stored on ActiveScale object storage, scalability is not an issue. The locally attached drives, which previously held the hot, warm, and cold buckets, can be re-purposed as cache to increase performance and accelerate searches. This is a radical departure from the classic architecture and offers several compelling benefits.

Greatly Increased Capacity

ActiveScale object store uses erasure coding technology to protect the data. Erasure coded volumes can survive more disk failures than RAID or replication with only 35% additional disk space. This approach not only increases data durability but delivers more than DOUBLE the usable disk capacity compared to the classic architecture!

Better Search Performance

Search performance is largely dependent on the speed of the underlying storage devices. Splunk SmartStore provides a significant boost in search performance thanks to a new smart caching layer that caches data dynamically, based on the user's search criteria. If you're searching within a particular time period, it can read in all of the data from that time period and cache it, ensuring that searches are always run from the fastest possible storage tier.

Improved Data Retention

The addition storage capacity and easy scalability of ActiveScale makes it much easier for users to meet their data retention requirements. And since the data now spans a much longer time period, it is possible to study long term trends and uncover patterns of activity that were previously unexposed.

ActiveScale Object Storage System with Splunk SmartStore

The SmartStore with ActiveScale object storage solution handles today's big data storage challenges and helps forward-thinking businesses transform themselves with modern approaches to data storage.

The Quantum logo is displayed in a white, sans-serif font against a dark blue background. To the right of the logo, there is a decorative graphic consisting of several overlapping squares in various shades of blue, arranged in a stepped, ascending pattern.

Quantum technology and services help customers capture, create, and share digital content—and preserve and protect it for decades at the lowest cost. Quantum's platforms provide the fastest performance for high-resolution video, images, and industrial IoT, with solutions built for every stage of the data lifecycle, from high-performance ingest to real-time collaboration and analysis and low-cost archiving. Every day the world's leading entertainment companies, sports franchises, research scientists, government agencies, enterprises, and cloud providers are making the world happier, safer, and smarter on Quantum. See how at www.quantum.com.

www.quantum.com