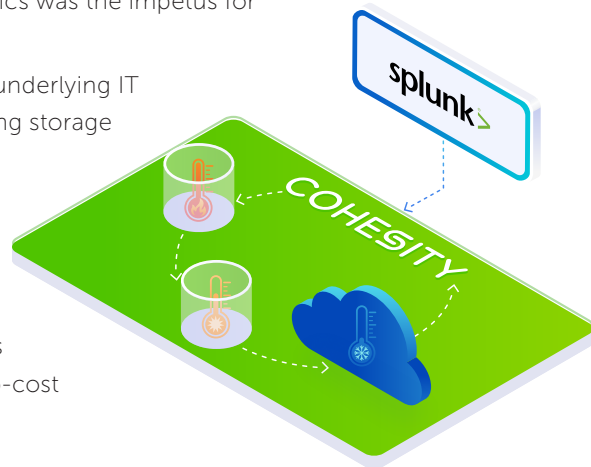


Boost Agility, Lower Costs with Splunk SmartStore and Cohesity SmartFiles

Splunk empowers organizations to gain insights from their machine data. Yet as data volumes have continued to grow, enterprises have struggled to cost-effectively search, manage and retain all of it for business advantage and compliance needs. Even distributed horizontal scale-out models, designed to process sizeable data volumes by co-locating compute and storage, fall short because demand for storage outpaces compute. Eliminating the need for organizations to continue to incur the performance impact and cost of always adding both storage and compute to harness their data for analytics was the impetus for Splunk SmartStore.

Software-defined Splunk SmartStore relieves pressure on underlying IT infrastructure. So does Cohesity SmartFiles by consolidating storage infrastructure and transforming enterprise file and object services. These complementary technologies maximize data management flexibility while maintaining search performance. As Splunk SmartStore allows compute and data storage to be independently scaled up or down based on business demands, Cohesity SmartFiles provides boundless capacity with limitless frequent snapshots, zero-cost clones, and more—on-premises and in the cloud.



Splunk SmartStore Advantages

Splunk SmartStore automatically evaluates users' data access patterns via an application-aware cache and places actively accessed data in local storage for real-time analytics and unused/inactive data in lower-cost, remote storage. This boosts scalability and data availability, expanding organizations' data retention capabilities at significantly lower TCO and with greater flexibility from built-in applications and an index-aware cache. Splunk SmartStore also simplifies indexer maintenance as all data is pushed to remote storage to assist with patching, upgrading, or replacing indexes without impacting data integrity.

Splunk SmartStore allows enterprises to use AWS S3—or other S3 API compliant object stores such as Cohesity SmartFiles—as volumes for storing indexed data in Splunk. This presents a fundamental shift in how organizations can store data in Splunk. Before, Splunk relied on dedicated storage for each indexer while trying to keep costs down by enabling various classes of storage. Now, Splunk SmartStore replaces warm, and cold storage buckets with a cache manager and an object store, enabling historical data to be stored longer at a lower cost than traditional on-premises storage options.

Splunk SmartStore's architectural goal is twofold: Optimize the use of local storage and ensure the fast indexing and search capabilities of Splunk Enterprise. That's why it has a remote storage tier and cache manager for data to reside either locally on indexers or on a remote storage tier, allowing data movement between the indexer and the remote storage tier to be managed by the cache manager on the indexer. Unlike regular deployments, Splunk SmartStore's remote object store becomes the location for master copies of warm buckets, while the

indexer's local storage retains newly ingested hot buckets and caches copies of warm buckets frequently used by search queries. Splunk SmartStore lets enterprises manage their indexer storage and compute resources more cost effectively by scaling these resources separately.

Why Add Cohesity SmartFiles to Your Splunk SmartStore Deployment?

How organizations deploy Splunk SmartStore has recently become more critical because the efficiency of operations directly impacts data processing speed, storage performance, and resource management. Organizations deploying Splunk SmartStore with Cohesity SmartFiles gain agility, performance, and cost advantages.

Unlike alternative solutions which have not been well-suited for managing large scale volumes and are hosted on a more expensive tier than direct attached storage, Cohesity SmartFiles provides web-scale and flexible deployment on-premises, or moving cold/frozen data to the cloud. In all scenarios, data can still be easily searched and analyzed.

Splunk SmartStore delivers performance at scale in a decoupled architecture, dynamically bringing active datasets closer to compute on-demand while processing data without impacting the user search experience. This allows organizations to process data independent of data age or storage placement - keeping all data searchable all the time.

Together, Splunk SmartStore and Cohesity SmartFiles deliver the industry's leading software-defined, data-centric, object solution for the enterprise - delivering data management for unstructured data.

Better Together

In combination, Splunk and Cohesity technologies deliver a robust data management solution from core to edge to cloud that is simpler and more cost effective to deploy. The combination provides a scalable object storage based solution for hosting Splunk's warm and cold data —with advanced caching and fast search of backup data.

For all organizations choosing Splunk SmartStore, Cohesity SmartFiles uniquely provides:

- **Freedom of hardware choice** – Workloads are free from proprietary hardware, running in the data center with easy on-ramping to cloud and seamless connectivity within hybrid environments.
- **Global space efficiency** – Data optimization with global dedupe, compression, and small file optimization.
- **Broad compatibility** – Cohesity SmartFiles supports seamless data interoperability in Windows, Linux, and cloud/object environments with support for NFS, SMB, and S3 protocols.
- **Stress-free storage management** – A single, intuitive UI, automation, and policy settings keeps IT in control of data.
- **Global actionable search** – To quickly help you find what you need across all data under management.
- **Integrated security** – A multilayered security approach helps reduce risk through threat prevention, detection, analysis, and remediation

Learn more about how your organization can take advantage of Splunk SmartStore and Cohesity SmartFiles at www.cohesity.com/solution/scale-out-storage.