

Reimagining Splunk at Enterprise Scale

Pure Storage® FlashBlade™ and Splunk SmartStore deliver simpler and faster analysis of cybersecurity threats.

The exponential increase in machine data provides opportunities to address business challenges by analyzing data to identify cybersecurity threats, application performance issues, or even new products. Combined with the complexities and cost of IT infrastructure, the traditional approach of analyzing data across “hot,” “warm,” and “slow, cold” tiers using distributed scale-out architectures is inadequate to address log analytics needs. You need a flexible architecture that provides cloud-like agility and all-flash performance. Enter [Splunk® SmartStore](#) powered by the [Pure Storage FlashBlade solution](#).

The Challenges of Distributed Splunk Scale-Out Models

Splunk helps address data management and network security challenges. It enables the search, analysis, and visualization of machine data from IT infrastructure or business applications, and delivery of insights and business value to customers. Splunk had used the distributed scale-out model that evolved over a decade ago by co-locating storage and compute. However, with the exponential growth of data, customers are facing challenges for multiple reasons.

Increasing cost of infrastructure: The distributed scale-out model provides high availability by replicating the data, which eliminates any benefits gained by Splunk compressing the data and increasing storage requirements. Co-locating storage and compute means that when you need more storage, you have to add both compute and storage. To further increase total cost of ownership (TCO), Splunk indexers with a distributed scale-out architecture usually have more servers with less storage each to minimize the amount of data and time associated with server maintenance and failures.



Accelerate searches

Get all-flash performance for data operations and searches in object. Optimize indexer bursting.



Reduce TCO

Use a single copy of warm data. Use compression to reduce storage needs by 30% to 40%.



Enhance availability

N+2 erasure coding protects data. Reduce node offline time by 94% during failures.



Simplify at scale

Add index nodes and rebalance data 99.7% faster. Expand capacity of both indexer cluster and FlashBlade object store.

purestorage.com

800.379.PURE




Search performance degradation: Splunk searches with a distributed scale-out architecture are associated with significant performance degradation as data ages. As it ages, data is tiered to cheaper and lower-performance storage tiers in warm and cold buckets. Cold buckets reduce storage required via Splunk’s TSIDX reduction feature, but this significantly impacts search performance. This storage approach is impractical when responding to search requests related to regulatory or compliance requirements, cybersecurity, and legal discovery—all of which demand information beyond the most recent data.

Infrastructure complexity and operational overhead: High availability by replication in distributed scale-out architecture also requires that all replicas are online all the time. This means server maintenance and software updates of index cluster servers have to be serialized and involve evacuation of data, performing the update and rehydration of the data. Secondly, indexer cluster expansion and hardware refresh requires (at times multiple) data rebalances. Lastly, in the event of a failure of an indexer node, you need to reconstruct the data in the indexer node. Such architecture limitations add significant complexity, time, and cost while reducing the search and ingest performance due to the loss of computing resources and/or data migration or reconstruction process.

Splunk SmartStore with FlashBlade

A Splunk SmartStore solution powered by the Pure FlashBlade captures the best of Splunk capabilities while addressing the limitations of a traditional distributed scale-out architecture. Splunk SmartStore disaggregates compute from storage and consists of stateless indexer servers, S3 object store, and an index-aware cache.

		
<p>Splunk SmartStore Bring the agility and simplicity of cloud to on-premises Splunk deployments.</p>		<p>Pure Storage FlashBlade Ultra-fast, S3-compatible storage for on-premises deployments, architected to deliver scale-out capabilities with all-flash performance.</p>

A SmartStore solution with Pure Storage FlashBlade as the high-performance S3 object store delivers multiple benefits:

Faster insights from searches irrespective of the age of the data: SmartStore’s architecture eliminates the cold buckets and includes a new local cache, which is ideal for searching through recent and/or recently searched data. However, using typical low-cost and low-performance object storage solutions with SmartStore limits search performance. A SmartStore solution with FlashBlade provides all-flash performance with high bandwidth and parallelism for data operations and searches outside of the SmartStore cache. It also ensures that you can efficiently complete critical, non-repetitive tasks related to regulatory or compliance requirements, cybersecurity breaches, and legal discovery. Additionally, FlashBlade’s bandwidth is ideal for supporting the bursting of SmartStore indexers.

purestorage.com

800.379.PURE



SOLUTION BRIEF

Lower cost of ownership of overall Splunk deployment: Using SmartStore with FlashBlade optimizes infrastructure utilization and lowers the storage and compute requirements when compared to Splunk’s classic direct attached storage architecture. With this approach, you can size indexers based on ingest rates and concurrent search volumes instead of worrying about storage. Additionally, SmartStore requires storage of a single copy of the warm data and leverages the data resiliency of the object storage solution. FlashBlade further reduces storage requirements for the object tier by 30% to 40% by providing data compression.

Increased availability: FlashBlade delivers highly efficient [N+2 Erasure Coding](#) data protection. Data-at-rest encryption enables SmartStore to provide a highly available and secure solution. Splunk SmartStore doesn’t require you to reconstruct warm data if an indexer goes down. To meet Splunk’s replication- and search-factor parameters, you only need to replicate metadata between the nodes in the cluster. Pure’s internal testing revealed that—in the event of a node failure with SmartStore solution with FlashBlade—the node was offline for 94% less time compared to a classic Splunk architecture with a similar data set.

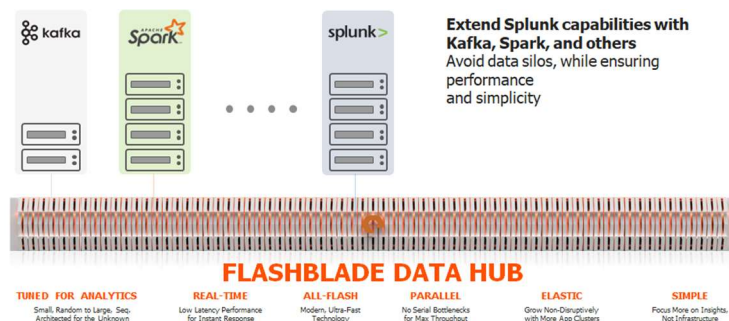
Simplified management of Splunk clusters at any scale: SmartStore simplifies the data management of the Splunk infrastructure to address the growing needs by enabling indexer node additions, data rebalance, and indexer removals. It only requires metadata replication between the nodes in the cluster instead of requiring data migration. Based on Pure’s internal tests, node addition and data rebalancing are almost 99.7% faster on SmartStore with FlashBlade compared to classic Splunk. Indexer upgrades are performed in parallel. A SmartStore solution with FlashBlade can seamlessly expand both the indexer cluster computing and the capacity and performance of FlashBlade object-store. This reduced operational complexity enables Splunk administrators to focus on monitoring ingestion and indexing performance without spending time managing infrastructure.

Additional Advantages

Splunk SmartStore with FlashBlade delivers additional benefits that are possible only with a customer-centric partner like Pure Storage.

Pure’s simplicity and ease of use: Simplicity and ease of use are central to any Pure solution. Splunk SmartStore with FlashBlade is no different. [With Pure1® AI-driven management](#), [full-stack analytics](#), and [predictive support](#), management and planning of the Pure FlashBlade environment is simple and efficient.

Share analytics data with a data hub: A data hub is a modern, data-centric architecture for storage. Data hubs power a wide range of workloads like Splunk and other modern analytics and data-intensive workloads by enabling enterprises to consolidate data silos and share data. A data hub allows you to share log data across various applications and teams to maximize insights and drive innovation.



purestorage.com

800.379.PURE



Future-proof architecture offers investment protection: FlashBlade is an ideal future-proof infrastructure investment if you're not currently considering SmartStore but using Classic Splunk Architecture with the option to deploy SmartStore in future. A Pure solution with FlashBlade for cold tier is optimal in Classic Splunk environments. You can repurpose it as a high-performance all-flash object storage tier with SmartStore in future. With [Pure Evergreen™](#), you never have to rebuy your storage or worry about it becoming obsolete.

Accelerate Machine Data Analytics with SmartStore and Flashblade

Using Splunk SmartStore with Pure accelerates searches with an [all-flash, cloud-native architecture](#) while reducing overhead costs, increasing availability, and improving operational efficiencies. Take the next step to modernize your Splunk environment and build a foundation for the next generation of search and analytics. Don't let a legacy log analytics architecture slow your business.

purestorage.com

800.379.PURE

