

Splunk Infinite Cache by VAST Data


Break the performance vs. capacity tradeoff
in Splunk Enterprise



Search On Demand, Redefine TCO

Splunk Infinite Cache on VAST gives observability and security teams the freedom to search all their data on demand without the performance impact of searching remote buckets. VAST Increases search performance, accelerates historical queries, and enables more ad-hoc searches, so organizations can identify and act on threats and performance issues sooner. Teams gain these advantages while cutting total Splunk storage needs by up to two-thirds, reducing cost, power, cooling, and rack space.

VAST Data Provides 3:1 Data Reduction for All Splunk Data (Not Just S3)



Performance

Keep hot data fast and make months of warm searches feel just as responsive.



Consistency

Replace eventually consistent object storage with strict consistency for data everywhere.



Simplicity

Consolidate block cache and SmartStore S3 on one universal, all-flash data platform.



Efficiency

Thin, flashless indexers shift storage to VAST, reducing cost, space, power, cooling.



Economics

Cut total storage in half; VAST innovations deliver unparalleled data reduction savings.



Reliability

Eliminate duplicate replicas across tiers; fewer PBs to protect, back up, and heal.

Introducing the VAST Data Infinite Cache Architecture

Splunk SmartStore has two storage components in its architecture: local block storage for the cache and remote S3 storage for data resiliency. This results in the deployment and management of multiple storage architectures – commonly local NVMe SSDs or an enterprise SAN for the local cache and any flavor of S3 offering for the remote storage.

Splunk search execution relies on buckets residing in local cache. If a remote bucket is required for a search, the Splunk cache manager retrieves it to local cache. The downside of this function is the search execution must wait for the download to complete before the bucket can be read, which slows search execution time. Larger local caches store buckets longer, reducing the volume of remote downloads, but at the cost of increased storage infrastructure.

The VAST Data platform is unique in supporting SmartStore as it provisions both block and S3 storage on a single platform. Going further, VAST reduces all data between block and S3, resulting in the most efficient infrastructure on the market. This is Infinite Cache, a cache that scales to store all Splunk buckets, eliminates the need to download remote buckets, accelerates all searches, and essentially costs nothing!

Key Capabilities

Highly resilient infrastructure that's also highly efficient

Splunk is a mission critical application. The VAST Data Platform is architected for six-nines availability (99.9999%). VAST's Disaggregated Shared Everything (DASE) architecture is designed to be highly resilient and highly efficient. VAST clusters have zero single points of failure and can be built with rack level resiliency. VAST's locally decodable erasure codes bring data protection overhead down to as low as 2.7%, further extending savings beyond Data Reduction!

Eliminating incomplete/partial search results with Splunk multi-site HA clusters

Splunk's S3 partner ecosystem is comprised of eventually consistent S3 replication implementations, which, by design, have a lag in the data that is available between each site (replication lag). These architectures prioritize local response time over data availability. VAST implements a strictly consistent S3 synchronous replication architecture. This results in no out of sync data between two sites and thus no incomplete or partial search results. VAST's synchronous replication does incur higher S3 upload latencies, but that's meaningless, as these latencies do not impact Splunk indexing or search performance

Simplified sizing and lifecycle management

With VAST Infinite cache, customers are no longer locked into a rigid configuration with the local cache; block storage capacity can be increased on demand to meet search needs. By implementing VAST vs local NVMe SSDs, Splunk deployments can scale on demand – whether the Indexers are bare metal, virtual machines or containers.

Support for Splunk's DB Connect with the VAST DataBase

The VAST Data platform provides a multitude of data services above block and S3 storage, including a Kafka compliant event bus, and an exabyte-scale transactional data warehouse, which supports multiple query engines. Splunk Search head Clusters can directly query data in the VAST DataBase via the Splunk DB Connect, expanding the data types available to Splunk data engineers and observability teams.

Why VAST Data for Splunk Smart Store

VAST Data helps Splunk customers ask bigger questions across more data without slowing day-to-day operations, while reducing infrastructure costs and eliminating search-time trade-offs.

About VAST Data

VAST Data is the leading data platform software company bringing businesses into the AI era. The VAST Data Platform delivers scalable performance, simplifies data management, and enhances productivity by accelerating time-to-insight for workload intensive applications. Launched in 2019, VAST has become history's fastest-growing data infrastructure startup. Looking to simplify your data center and discover insights in all of your data? Contact us at hello@vastdata.com