

SOLUTION BRIEF

Splunk SmartStore powered by NetApp StorageGRID

Enterprise-class object storage
that enables performance at
scale with lower TCO for
Splunk SmartStore



Machine-generated information is complex, growing, and largely becoming a key element to operational and business assessments. Handling these large datasets a huge challenge, because there is rapid development in machines generating vast amounts of data to be accessed and processed. Splunk plays a vital role in dealing with the assessment and presentation of data analytics, and NetApp provides storage at scale efficiently in multicloud environments.

Splunk Enterprise is a full-featured, powerful software platform for collecting, searching, monitoring, and analyzing machine data. Splunk rapidly turns machine data into visibility, insight, and intelligence. As your business harnesses the power and ease of Splunk data analytics, you naturally want to index an ever-growing amount of data. And as this data grows, so does the compute and storage infrastructure required to service it. When older data is referenced less frequently, you don't want to commit the same number of compute resources to it and or let it consume expensive primary storage. To operate at scale, organizations need a way to move less frequently searched data to more efficient tiers, freeing compute and primary storage for the data that is used the most.

The combination of Splunk SmartStore and the NetApp® StorageGRID® object-based storage solution provides storage at scale and performance at lower cost. The solution can be deployed on single sites or multiple sites, and it offers decoupled compute and storage for optimal flexibility.

Why StorageGRID?

NetApp StorageGRID is a software-defined solution that supports industry-standard object APIs like the Amazon Simple Storage Service (S3) API. It allows you to build a single namespace across up to 16 data centers worldwide, with customizable service levels for metadata-driven object lifecycle policies. Integrated lifecycle management policies optimize where your data lives throughout its lifecycle.

StorageGRID enables ubiquitous data access by providing a global object namespace formed by a distributed “grid” of storage nodes. These nodes can have multiple storage tenants and use standard S3 API calls to service both cloud and enterprise unstructured data applications. The solution's object lifecycle policies let you create various service levels, optimizing durability, protection, performance, and location across geographies. You can adjust policies and realign your data landscape as requirements change.

StorageGRID optimizes data durability and availability across multiple geographies. It enables hybrid cloud workflows—whether data is on premises or in a public cloud—to fit your business demands with access to Amazon Simple Notification Service (SNS), Microsoft Azure Blob, Amazon Glacier, and Splunk SmartStore.

Intelligent tiering with SmartStore and StorageGRID

Splunk SmartStore takes advantage of NetApp StorageGRID as the remote storage. By using the layered erasure coding (EC) of StorageGRID, you can reduce costs without sacrificing durability. With node-level EC, you can protect against failed drives and rapidly rebuild lost data segments; with geo-distributed EC, you can protect against site-level disasters. You can also combine replication and geo-distributed EC to balance performance needs and cost savings between different sets of data or during an object's lifecycle in Splunk SmartStore deployments.

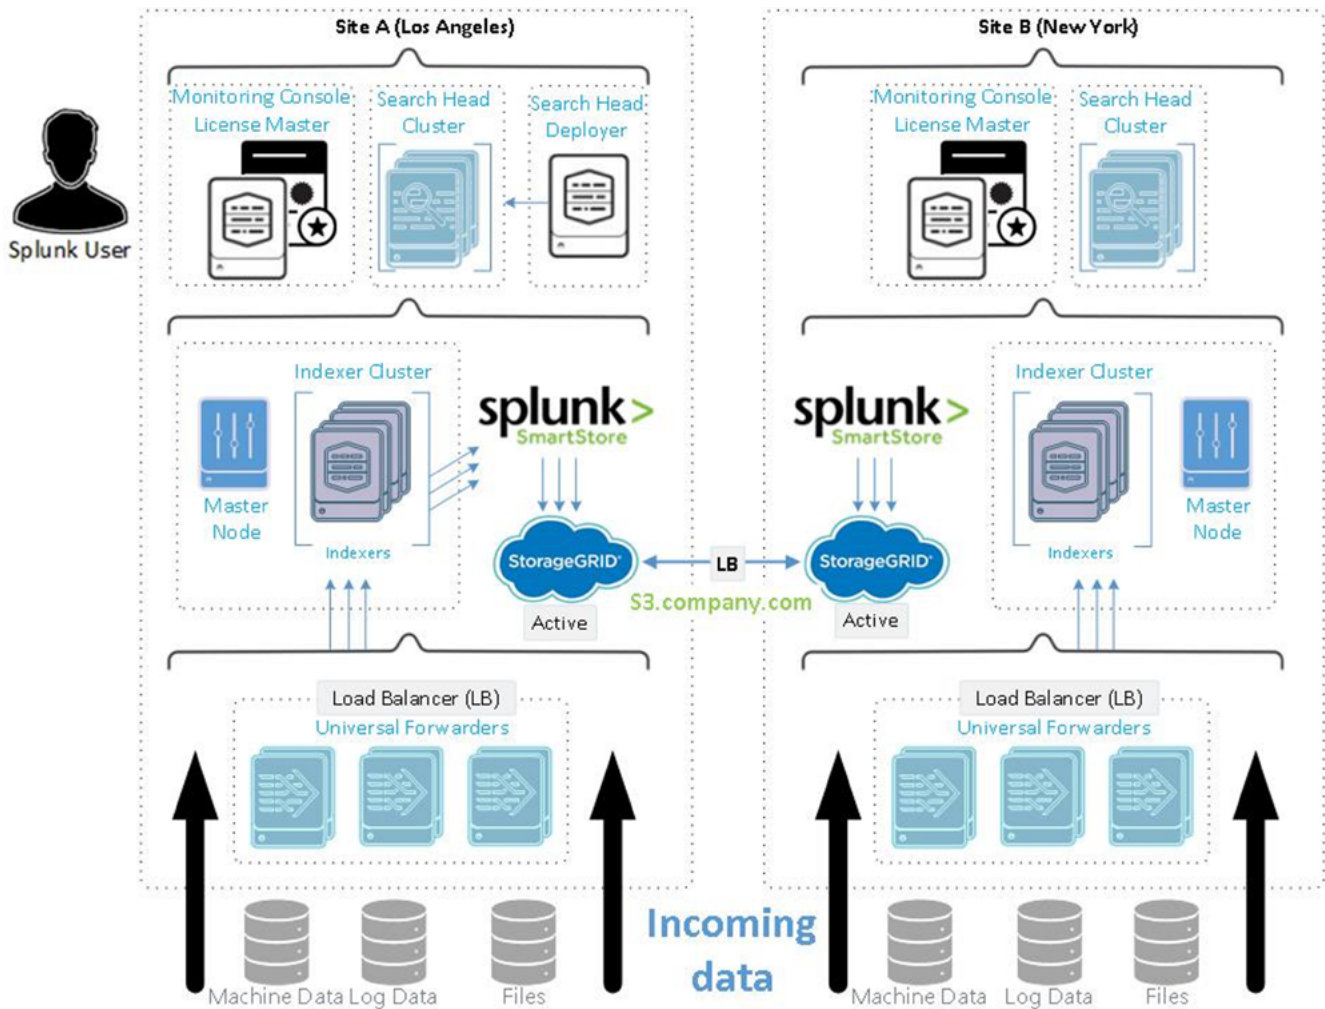


Figure 1) Splunk SmartStore multisite deployment with NetApp StorageGRID.

Proven software designed for nondisruptive operations

StorageGRID is an 11th-generation object store with nearly two decades of production deployments in the most demanding industries. NetApp has demonstrated more than 20 years of product hardening with over 1 million systems shipped. With advanced features such as the NetApp Active IQ® digital advisor tool for proactive, immediate response backed by NetApp's world-class support organization, StorageGRID is a solution that you can trust with your critical data assets.

Flexible deployments

Because every Splunk SmartStore deployment is unique, StorageGRID aligns with your environment, whether that includes nodes as virtual machines (VMs), as optimized hardware-based appliances, as bare-metal servers with Docker containers, or combined across virtual and physical environments. In all cases, designing and deploying StorageGRID is a centrally managed, streamlined process that allows rapid deployment of petabytes of storage. And the range of StorageGRID appliances makes the process even simpler.

Speed:

StorageGRID SGF6024 appliance

Accelerate your object workloads with the industry's first scale-out all-flash object appliance. You can add it to existing grids to provide pools of fast object storage, or you can use it to build all-flash grids with at least three appliances. With various densities that range from 800GB SSDs to 15.3TB SSDs, you have options for speed not previously seen with object workloads.

Scale:

StorageGRID SG6060 appliance

The StorageGRID SG6060 appliance is already a thriving member of the StorageGRID family, and now it can scale to new heights with multishelf deployments. That means you can buy a single SG6060 appliance and add one or two expansion shelves to scale a single namespace to over 400PB. The SG6060 appliance balances cost and speed with the most cost-effective performance-to-scalability ratio in the StorageGRID portfolio.

Simplicity:

StorageGRID SG1000 appliance

The SG1000 is the first services appliance in the StorageGRID family, streamlining large deployments by providing both high-availability (HA) load balancing and improved grid administration on a single set of nodes. When deployed in pairs, the services appliances can provide high availability across sites while still offering the flexibility to operate as a load balancer, an administrative node, or both simultaneously. The SG1000 appliance now makes it possible to deploy an all-appliance grid.

Deploying Splunk SmartStore on NetApp

StorageGRID object-based storage creates an enterprise-grade turnkey solution that is easy to implement. Each appliance is built to solve specific performance or capacity needs. You can also deploy StorageGRID in a software-only model, as containers, or on physical or virtual servers.

Enable the hybrid cloud

Cloud-to-cloud data management can also mean cost savings. StorageGRID can manage and store objects in its own globally distributed infrastructure, in Amazon S3 or S3-compatible object stores, or in public clouds. Depending on your location or cost needs, you can use Cloud Storage Pools to add a cloud copy as a cloud or backup tier for Splunk SmartStore to Amazon Glacier or to Microsoft Azure Blob.

About NetApp

In a world full of generalists, NetApp is a specialist. We're focused on one thing, helping your business get the most out of your data. NetApp brings the enterprise-grade data services you rely on into the cloud, and the simple flexibility of cloud into the data center. Our industry-leading solutions work across diverse customer environments and the world's biggest public clouds.

As a cloud-led, data-centric software company, only NetApp can help build your unique data fabric, simplify and connect your cloud, and securely deliver the right data, services and applications to the right people—anytime, anywhere. www.netapp.com

