HPE STORAGE SOLUTIONS FOR SPLUNK

Get more with HPE Storage solutions for Splunk

Hewlett Packard Enterprise provides compute and storage for a complete Splunk solution stack based on industry-leading HPE ProLiant servers or HPE Synergy solutions. HPE Storage offers an array of choices to best fulfill the performance and scalability requirements at an optimal cost structure. Through HPE GreenLake, Splunk infrastructure can be acquired in a cloud-like, consumption-based fashion.

Splunk was recently ranked #2 in Gartner's "Named ITOM Performance Analysis Software Vendors, Ranked by 2018 Software Revenue, 2017-2018".

Splunk software helps simplify and modernize IT and expands into business analytics, security, and IoT use cases. HPE Storage provides a range of differentiated solution architectures for Splunk to fulfill the most demanding Splunk requirements at an optimal price-point.

MACHINE DATA IS BIG DATA

Every IT application, system, and piece of infrastructure in a company, right down to the light switches, generates data at millisecond intervals. This machine-generated data is as complex as it is invaluable. It is also one of the fastest growing areas of Big Data.

Machine-generated data contains a detailed record of all user transactions and individual customer and component behavior. It further captures sensor activity, machine performance, security threats, and fraudulent activity, to name a few. The data holds valuable insights, critical to both the security and the profitability of the enterprise.

SPLUNK AND THE ROLE OF STORAGE

With 13,000 installs across 120 countries, Splunk is a leader in Big Data analytics. Splunk core capabilities include the universal collection and indexing of machine data from nearly any source. Splunk provides real-time monitoring for patterns, thresholds, and alerts in addition to search and analyze capabilities for historical purposes.

Splunk is commonly used for enterprise security. Splunk also has use cases extending to areas such as online retail order tracking, insurance claims processing, and marketing analytics.

Storage is a critical enabler for a successful Splunk deployment.

- **Performance:** Incoming events (ingest) must be able to be stored while data is accessed at the same time for reporting and historical analysis.

- **Scalability:** The storage architecture must provide near-limitless scalability as ingest devices are being added (ingest increase) or data retention is extended for new use cases.

- **Economics:** IoT and operational analytics environments tend to see data grow at dramatic rates; hence, the cost aspect of the underlying storage layer is critical.

HPE STORAGE SOLUTIONS FOR SPLUNK

HPE Storage offers a range of solutions to provide the ideal performance, price, and scalability point while reducing complexity. These solutions are tailored to the specific environment parameters in order to obtain the perfect performance/scalability/ cost ratio.

DIRECT-ATTACH STORAGE (DAS) WITH AN OPTIONAL AMAZON S3 OBJECT TIER

HPE Apollo 4x00 or HPE ProLiant servers with S3 objects based on HPE Scalable Object Storage with Scality RING (or Ceph)

An architecture based on a scale-out approach leveraging x86-based servers with an optional Splunk SmartStore tier on S3-compatible object storage caters to massive-scale deployments and environments with an existing object store or interest in building a multi-use case object store. HPE Apollo 4000 servers provide unmatched capacity-density as well as a combination of Persistent Memory, NVMe, SSDs, and HDDs.

RESOURCE

Evaluator Group Technical Insight Paper: "Enterprise Storage for Splunk"

splunk>

Check if the document is available in the language of your choice.

Make the right purchase decision. Contact our presales specialists.

Chat   **Email**   **Call**

Share now

Get updates

HIGH-PERFORMANCE, DATA-SERVICES-RICH SHARED STORAGE ARRAYS

HPE Primera, HPE 3PAR, HPE Nimble Storage

Shared-array based architectures provide several advantages in addition to the ability to independently scale compute from storage:

- High performance through the use of all-flash and hybrid deployments

- Multi-workload and multi-tenant consolidation onto same platform

- Mission-critical reliability

- Efficient use of storage capacity through compression, deduplication, and system-inherent redundancies

- Advanced data protection options through the use of snapshots and data mobility through built-in HPE Recovery Manager Central

- Pain-free management of infrastructure through HPE InfoSight that offers predictive analytics

"Supporting the growing demand for creating business intelligence and value out of machine data by providing server and storage solutions that deliver high performance, reliability, and massive scalability, while also answering today's DevOps challenge for containerization and on-demand instance creation."

– Theresa Melvin, HPE Big Data Chief Architect

ARCHITECTURAL AND FINANCIAL ADD-ONS FOR ULTIMATE FLEXIBILITY

High-performance shared storage with S3 object tier

As retained data capacity increases, a second tier of storage based on S3 objects assists in keeping costs under control. HPE provides architectural choices where arrays are combined with an S3 object tier leveraging the Splunk bucket tier capability to transparently move data to the most appropriate storage tier. The private cloud on-premises S3 object capabilities that HPE provides are through HPE Scalable Object Storage with Scality RING.

Cloud-like, consumption-based acquisition

HPE GreenLake allows cloud-like procurement through a pay-as-you-go, consumption basis model. HPE GreenLake allows Splunk users to scale storage and compute capacity up and down while paying only for the capacity consumed.

LEARN MORE AT

hpe.com/storage
hpe.com/ai

Hewlett Packard Enterprise