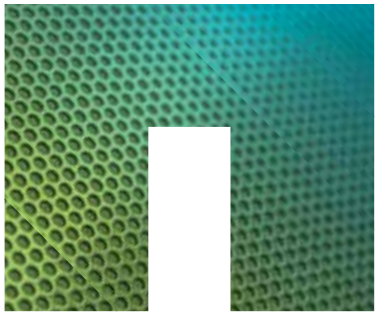

Solution Brief

NetApp Solution for Splunk

Enterprise-class building blocks for machine analytics

Key Benefits

- Increase search performance by up to 111% on average compared to commodity servers with internal disks.[1]
- Proven reference architecture solutions make configuration and ordering easy.
- Optimize performance and capacity buckets for Splunk's hot, warm, and cold data tiers.
- Scale compute and storage independently to better match application workloads.
- Reduce storage requirements and maintain availability with fewer copies of data.
- Get single-interface management across your storage environment.
- Realize better performance than commodity servers with internal storage during data rebuilds.
- Protect your data with encryption.

Operational Intelligence: Getting from Data to Insights

An organization's data is a definitive source of intelligence, because it is a categorical record of activity and behavior, including user transactions, customer behavior, machine behavior, security threats, and fraudulent activity. Splunk is operational intelligence software that enables you to monitor, report, and analyze live streaming and historical machine-generated data. Splunk helps you distill, sift, and understand this machine data to improve service levels, reduce IT operational costs, mitigate security risks, enable compliance, and create new product and service offerings.

Splunk Enterprise Deserves an Enterprise Storage Infrastructure

As the use of Splunk for operational intelligence grows, its operational integrity becomes even more critical. Splunk deserves a storage infrastructure that promotes optimal and consistent performance with minimal maintenance and expense. The NetApp® E-Series and EF-Series storage systems provide better performance, data availability, scalability, data protection, and single-interface storage management compared with Splunk workloads that run on commodity servers with internal drives.

The NetApp all-flash E2800 and EF570 storage solutions combine robust, full-featured storage management software, a rugged array chassis, and the most recent solid-state drive (SSD) innovations to offer superior technological and business value. The NetApp all-flash E2800, EF570, and E5700 use the same chassis, which is deployed in thousands of installations that demand high-performance, dense, cost-effective storage. Together, these storage systems have a proven record of five 9s or greater reliability across a million deployed systems.

These building blocks are configured to support the Splunk hot, warm, and cold data tier model, as shown in Figure 1. Operations can effectively accelerate data indexing and searching with flash and minimize cost and space for colder data with high-capacity near-line SAS (NL-SAS) drives. Preconfigured solutions are available to accelerate Splunk deployments and optimize performance.

For Splunk's hot data tier, the NetApp EF570 delivers sub-100 microsecond latency and up to 1M IOPS, greatly accelerating search performance. It scales to 1.8PB and can deliver 21GBps of throughput. Also, because of the implementation of NetApp Dynamic Disk Pools (DDP) technology, this exceptional performance is only negligibly affected during disk failures. DDP technology also performs much faster data recovery from disk failures, and there is no need to immediately replace failed disk drives.

1. NetApp E-Series and Splunk Technical Report: TR-4460.

NetApp®

Figure 1) One architecture for Splunk's hot, warm, and cold data tiers.

Superior Performance

For many operations, the most compelling reason to power your Splunk environment with NetApp storage is the performance advantage that you will realize compared to the performance of commodity servers with internal storage. Recent testing that closely simulated real-world Splunk search performance showed conclusively that operations have much to gain from the NetApp storage approach.

Searching was significantly faster with NetApp compared to commodity servers with internal storage—on average up to 111% faster. Dense searching was up to 22% faster, and rare search performance was up to 200% faster on average, as shown in Figures 2 and 3. NetApp installations have seen even faster search run times.

NetApp E-Series storage systems also support advanced data security features, including media erasure and full hardware disk encryption. The NetApp SANtricity® System Manager software not only configures these features, it also manages encryption keys for each disk in the entire pool of NetApp storage systems.

"With all the data tiers in a single Shared Nothing DAS (Direct Attached Storage) such as E-Series, network latency is reduced as the buckets roll from one data tier to the next."

—NetApp E-Series for Splunk Enterprise, Function1

Splunk Apps and Technology Add-Ons

The alliance partnership between NetApp and Splunk includes developing apps for the NetApp storage platform portfolio. These apps and add-ons make the NetApp storage part of the infrastructure landscape that Splunk is surveying, facilitating better use of resources and supporting a more secure overall operation. Platforms that are currently supported with Splunk apps include SANtricity (EF-Series and E-Series), as shown in Figure 4, NetApp StorageGRID® object stores, and NetApp ONTAP® data management software. These apps are available on the Splunk base portal, apps.splunk.com.

"With the significant amount of machine-generated data captured every day, we rely on NetApp E-Series to deploy Splunk solutions for monitoring and troubleshooting the multiple platforms in our environment."

—Roy Shiladitya, head of Information Technology at ING DIRECT Australia

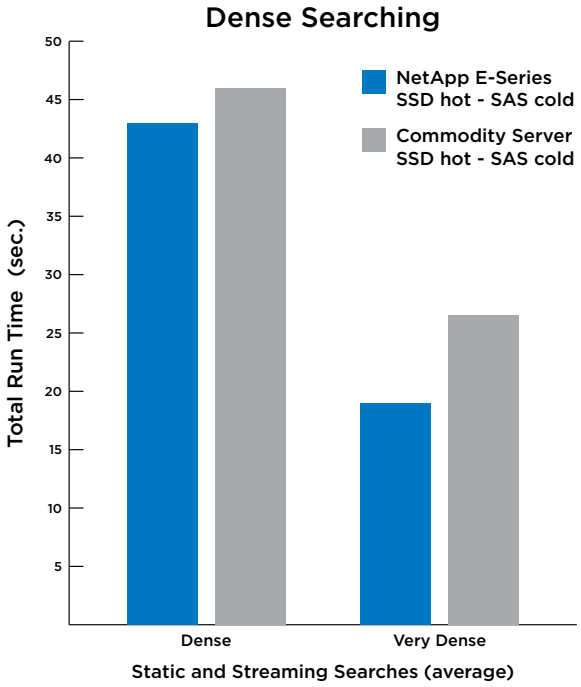

Figure 2) E-Series performance compared to commodity servers with internal storage for dense searching.

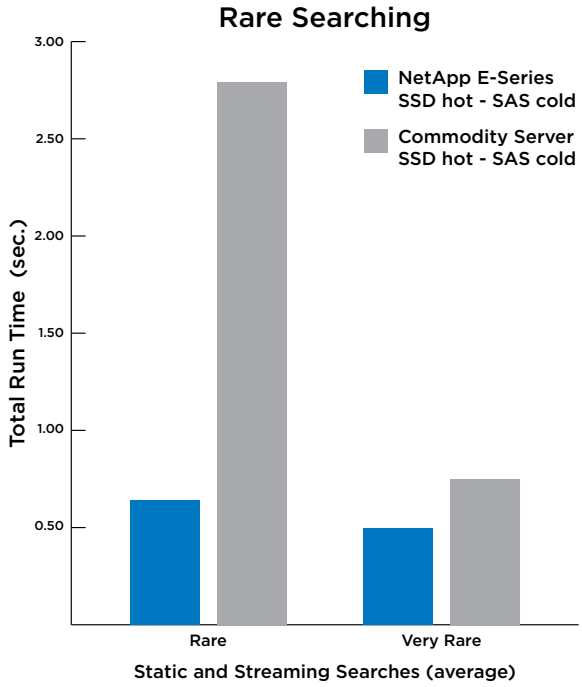

Figure 3) E-Series performance compared to commodity servers with internal storage for rare searching.

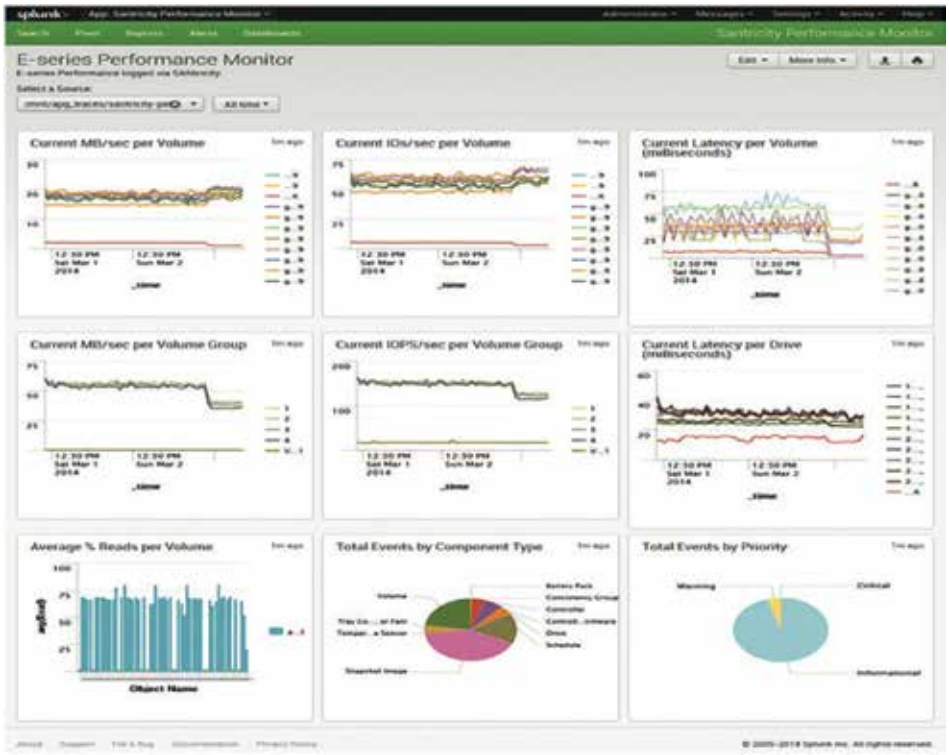

Figure 4) NetApp SANtricity performance app for Splunk Enterprise.

Validated Reference Architecture Solutions

Preconfigured and proven solutions for converged infrastructure, and different ingest rates, make sizing, ordering, and deploying solutions easy and cost effective.

Conclusion

When performance, reliability, cost, and convenience are taken into account, NetApp storage proves its business value in Splunk implementations.

About Splunk

Splunk Inc. (NASDAQ: SPLK) is the market leader in analyzing machine data to deliver Operational Intelligence for security, IT and the business. Splunk® software provides the enterprise machine data fabric that drives digital transformation. More than 13,000 customers in over 110 countries use Splunk solutions in the cloud and on-premises. Join millions of passionate users by trying Splunk software for free: http://www.splunk.com/free-trials.

About NetApp

NetApp is the data authority for hybrid cloud. We provide a full range of hybrid cloud data services that simplify management of applications and data across cloud and on-premises environments to accelerate digital transformation. Together with our partners, we empower global organizations to unleash the full potential of their data to expand customer touchpoints, foster greater innovation and optimize their operations. For more information, visit www.netapp.com. #DataDriven