SOLUTION BRIEF

Hyperfast Data-to-Everything with Splunk and Diamanti

KEY BENEFITS

TURNKEY SOLUTION FOR SPLUNK

The Diamanti platform is fully compatible with SVAs which helps run any Splunk topology without modification, enabling dynamic scaling.

LOWEST INFRASTRUCTURE TCO

Achieve more with less. Reduce infrastructure footprint dramatically and save up to 80% on TCO.

EXTREME PERFORMANCE

Utilize the full power of bare metal. Achieve 24 times faster indexing rates than those offered by legacy infrastructure.

Introduction

Every aspect of an organization is now generating data: IT Operations, Marketing, Supply Chain, Security, Development, and more. IT infrastructure and IoT devices have contributed to the exponential growth of machine-generated data over the last decade. This data contains valuable information that can drive efficiency, productivity, and visibility for the business. Gaining insights into this data in real-time would help organizations proactively monitor infrastructure, gain visibility into end-to-end business processes by analyzing data streams to identify patterns, outliers and trends and investigate potential root causes of problems to drive continuous improvement. Data is not only restricted to on-premises environments but also spans across the cloud leading to complications and complexities in the analysis. Thus, there is an increasing need for platforms to access, examine, process, and analyze this data in real-time to produce useful insights while reducing the total cost of ownership (TCO) for an enterprise.

Splunk: The Data-to-Everything Platform

The Splunk platform uses machine data—the digital exhaust created by the systems, technologies, and infrastructure powering modern businesses—to address big data, IT operations, security, and analytics use cases. The insights gained from machine data can support any number of use cases across an organization and can also be enriched with data from other sources. The enterprise machine data fabric shares and provides access to machine data across an organization to facilitate these insights.

Diamanti Platform: Purpose-built for Modern Applications

The Diamanti platform is the first and only Kubernetes solution integrated with a patented I/O-optimized architecture, delivering transformational application performance. With Diamanti, Kubernetes becomes an out-of-the-box solution, allowing organizations to focus on deploying modern applications across on-premises and hybrid cloud infrastructure. The key components of the Diamanti platform are – Diamanti Spektra, Diamanti Ultima and Diamanti D20 Series that together make Diamanti the infrastructure purpose-built for modern applications.

Diamanti Spektra is the prevalidated, pre-packaged and fully-featured software stack including Kubernetes, container runtime, operating system, enterprise-class DP/DR features, access controls and Management UI. Diamanti Spektra eliminates unnecessary layers of abstraction and deploys containers and virtual machines on bare metal, resulting in efficient resource utilization for actual application workloads. Diamanti Ultima is a pair of second generation PCIe based I/O acceleration cards that offload networking and storage traffic freeing up compute resources to power modern applications and deliver dramatically improved performance. The Diamanti D20 family of modern hyperconverged platforms consists of D20, D20X, G20T and G20P, and each includes ultra-fast NVMe storage.

Running Splunk on the Diamanti Platform

Splunk deployments provide optimal levels of service as long as the environment's underlying infrastructure resources complement the topology. Splunk Validated Architectures (SVAs)[1] are proven reference architectures for stable, efficient, and repeatable deployments. The Diamanti platform is fully compatible with SVAs, enabling any Splunk topology to run without modification. Dynamic scaling of the underlying infrastructure and application performance become as simple as a click of a button, eliminating lengthy infrastructure provisioning and resource tuning processes. An example of a compatible topology is shown in Figure 1.

Distributed Clustered Deployment - Single Site (C1/C11)

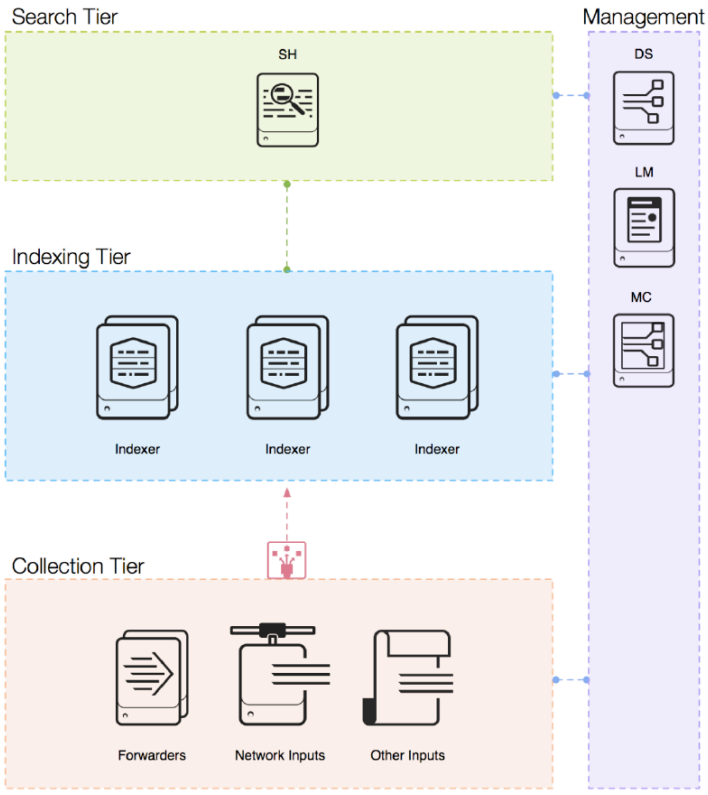

Figure 1: Splunk Validated Architecture (SVA) for Distributed Clustered Deployment - Single Site

1 https://www.splunk.com/pdfs/technical-briefs/splunk-validated-architectures.pdf

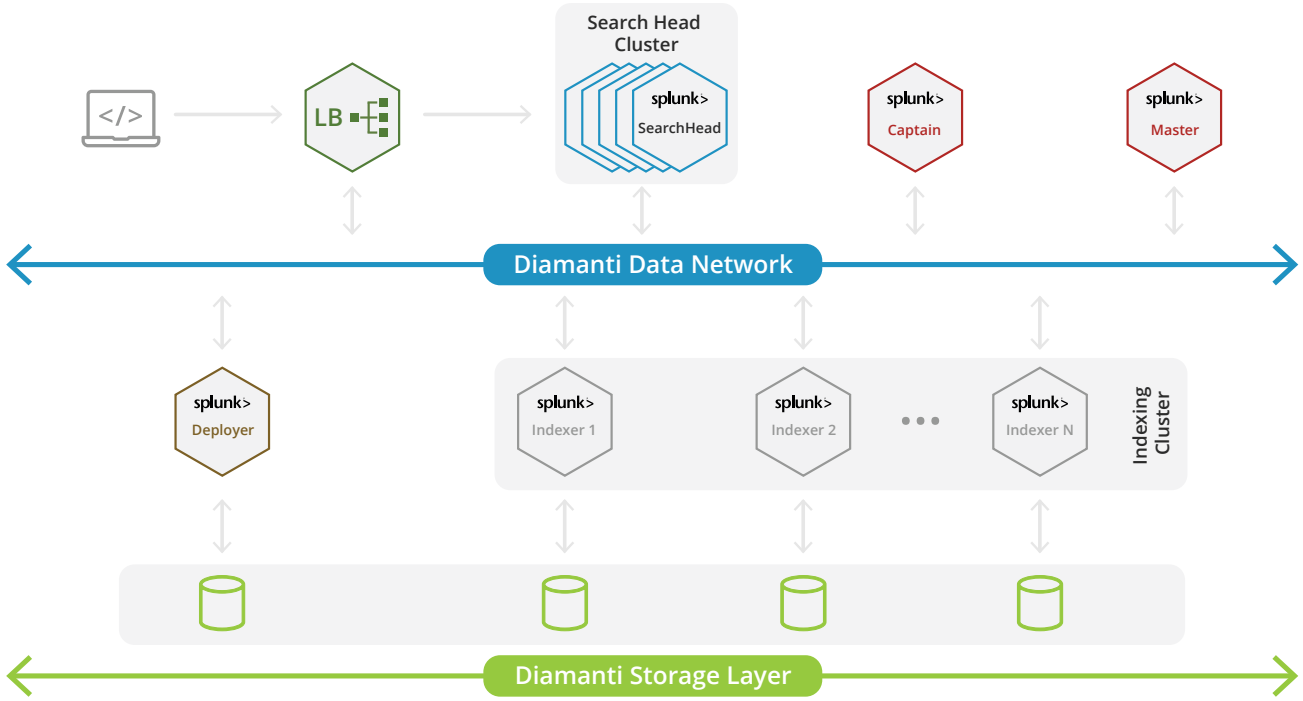

Figure 2: Splunk Deployment on the Diamanti Platform

Figure 2 illustrates a Splunk deployment on the Diamanti platform with multiple indexers. Using Diamanti's intelligent storage architecture and NVMe based storage, Splunk generates extremely fast indexing with low latency and consistent high performance. Full enterprise-grade storage services provided by Diamanti Ultima include backups, volume mirroring, snapshot-based replication, and volume resizing. Cluster-level features such as High Availability (HA) and Disaster Recovery (DR) required for system fault tolerance are supported by the storage management layer which ensures disaster recovery for the indexing tier as well as the search tier and enables seamless movement of Splunk deployments.

Case Study: NBCUniversal

NBCUniversal, a Fortune500 entertainment giant, used Splunk to process event data and produce insights with the aim of driving actionable measures on a real-time basis. With the existing infrastructure, the Splunk instance processed about 1 terabyte (TB) of data per day. This resulted in a huge backlog of unprocessed event data due to Splunk's slower data ingestion and indexing rates. Hence, NBCUniversal was not able to achieve its objective of driving actionable measures in real-time.

NBCUniversal then resorted to the Diamanti platform. They deployed a 32-node Diamanti cluster with a total usable storage capacity of 200 TB across the cluster. As a latency-sensitive application, Splunk was able to benefit from the hardware offload capabilities of Diamanti Ultima that physically isolated network and storage traffic to eliminate noisy neighbors and guarantee quality of service (QoS). The Diamanti platform maximized the hardware utilization since it eliminated unnecessary layers of abstraction and inefficiencies, and delivered exceptional Splunk performance without expensive and intensive overprovisioning. By migrating off legacy infrastructure that required a hypervisor and leveraging a bare metal Kubernetes solution, NBCUniversal was able to process approximately 24 TB of data per day, a 24 times increase, driving actionable measures on a real-time basis. Additionally, NBCUniversal reduced the infrastructure footprint by six times to achieve an 80% savings in total cost of ownership.

24x
Faster Indexing

80%
TCO Savings

1M
IOPS per Node

Based on 32 node Diamanti cluster with 200 TB storage and TCO over 5 years.

Figure 3: Splunk Deployment on Diamanti at NBCUniversal

Summary

With the exponential growth of data, infrastructure scalability, agility, and performance are crucial to a successful deployment of the Splunk application. As organizations contemplate the shift from deploying Splunk using traditional, legacy data center infrastructure to modern, cloud native architectures, a Kubernetes-ready infrastructure has become a requirement. The Diamanti platform natively supports SVAs to empower organizations to deploy Splunk with confidence and performance. As Splunk deployments grow within an organization, the underlying infrastructure scales in lockstep offering both stability and overall ease of management. Maximum hardware resource utilization reduces an organization's IT costs, consolidates platforms and legacy infrastructure, and attains high container density for the application. The Splunk platform running on Diamanti allows companies to be truly successful in achieving a real-time data-to-everything engine.