

Splunk on VMware vSAN

Hyperconverged Infrastructure is Ideal for Big Data

VMware and Splunk Address Key Challenges in the Data Center

Organizations have placed unprecedented demands on IT; as the business value of IT continues to grow, so does data center complexity, from the applications deployed to the infrastructure that supports them, both on-prem and in the cloud. As IT becomes more integral to the business, it needs insights to rapidly resolve incidents or avoid them all together. They also need highly available infrastructure that's easy to use and can stretch to the edge and the cloud for operational efficiency.

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.

VMware vSAN is the industry-leading storage virtualization software powering VMware's hyperconverged infrastructure solution. vSAN helps customers seamlessly evolve, as it is native to vSphere and requires no new tools. vSAN's industry leading ecosystem empowers customers to run HCI on certified solutions with their preferred vendor, and its multi-cloud capabilities provide customers consistent operations from edge to core to cloud, with intrinsic security throughout.

“We chose vSAN over traditional SAN because we could scale our storage performance linearly, which resulted in a predictable and performant experience. Also, we were able to reduce our cost and minimize our real-estate footprint.”

— CHANH CHI SR. DIRECTOR, IT
VMWARE

Splunk Needs the Right Infrastructure

As data centers have grown in complexity, it has become much more difficult for IT personnel to identify the root cause of serious incidents. Splunk software helps to create value from ever growing machine data. It continuously ingests data from a wide variety of sources, including infrastructure, applications and IoT devices. Splunk turns unstructured data into real-time insights, enabling IT to reduce the length of serious incidents by up to 90%, and reducing the overall number of serious incidents by up to 45%. In order to provide value with minimal operational impact, Splunk needs the right infrastructure. Hyperconverged infrastructure (HCI) can provide the performance, architecture, scalability and cost required to run Splunk optimally.

As a big data application, Splunk needs high-performance storage, massive capacity, and the infrastructure supporting it needs to be flexible, scale-out and provide linear performance. As Splunk needs to run continuously, it needs a highly-available, resilient storage platform. While traditional storage can provide some of these benefits, hyperconverged infrastructure is a much better fit.

The Benefits of HCI Powered by vSAN

Organizations need to quickly procure, install and deploy infrastructure cost effectively for all applications, including Splunk. Hyperconverged infrastructure powered by vSAN, can be procured in just weeks, can be deployed quickly, and can provision infrastructure in just minutes. Hyperconverged infrastructure scales incrementally, similar to Splunk, and can scale one node as a time, adding resources only when they're needed. In addition, HCI reduces data center complexity and accelerates decision making, enabling quick provisioning by IT generalists without the need for deep storage expertise.

Splunk also scales linearly, and HCI is well suited to support Splunk. For larger Splunk deployments, admins can distribute instances across multiple virtual machines to help scale the deployment, with each instance performing a specialized task. VMware HCI enables admins to quickly and easily provision new virtual machines and granularly control the resources of each VM or VMDK with policy-based management. Admins can add capacity or change service levels on the fly.

Why vSAN

vSAN primes businesses for growth through

- **Seamless Evolution to Full-Stack HCI**, as it is vSphere native, requires no new tools and integrates with the VMware SDDC
- **Industry-leading deployment flexibility**, with over 500+ ReadyNode configurations from 19 OEM vendors and a turnkey appliance, Dell EMCvRail
- **Multi-Cloud Support**, vSAN supports the most hybrid cloud use cases and offers native services with leading cloud providers

For many organizations, Splunk is a mission-critical application that needs the highest level of availability and resiliency. vSAN enables admins to specify the number of failures to tolerate at the VM or VMDK level and supports RAID 5/6 erasure coding for double failure of disks, hosts or domains. Disaster recovery of Splunk solution is provided by combination of vSphere Replication and VMware SRM and Splunk level replication with RPOs as low as five minutes.

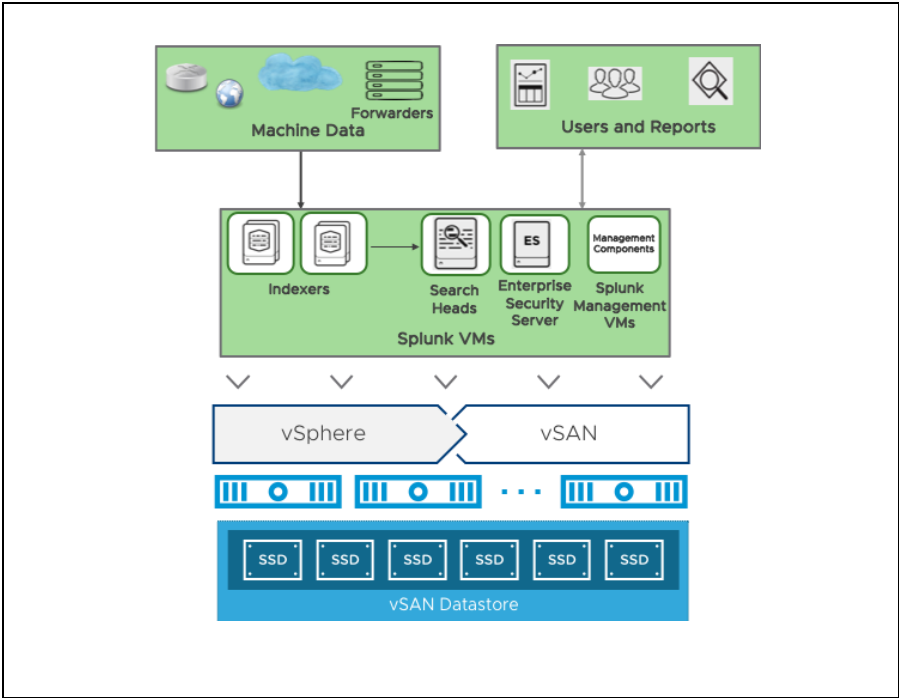


Figure 1: Splunk on vSAN Architecture

The Business Value of Splunk and vSAN

In addition to the operational benefits of vSAN: fast procurement, deployment and provisioning, linear scaling and high availability, organizations can also lower their CAPEX and OPEX when deploying Splunk and vSAN.

Splunk reduces operational costs. By proactively surfacing insights and providing visibility into the infrastructure, Splunk enables organizations spend less time monitoring, troubleshooting and preventing fires. Admins and app developers are in fewer war room meetings and freed up to spend time on higher level tasks.

VMware vSAN lowers CAPEX by leveraging industry-standard x86 servers and pay-as-you-grow-economics to lower capital costs; some customers have reported up to 75% lower CAPEX over external, all-flash arrays with fibre channel networking. In addition, vSAN highly automates processes, reducing management OPEX by up to 58%.

Learn More

Interested in learning how to deploy Splunk on vSAN? Check out the reference architecture.

Learn more about vSAN on our [website](#) or take a technical deep dive at [StorageHub](#).

