



PERFORMANCE, SCALE, AND SAVINGS

A Modern Approach to Splunk Architecture Utilizing Dell VxRail

A hyper-converged infrastructure reference design leveraging the power of Dell VxRail, virtualization, and world-class Splunk engineering expertise.

Copyright © 2022 Kinney Group, Inc., and Dell Inc. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA March 2022.

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Executive Summary

In today's world of mission-critical applications, having real-time access to analytics supporting security, IT operations, and application performance is a core requirement. Since its founding in 2003, Splunk Inc. has been recognized as the industry leader for collecting application and infrastructure machine data, and turning that data into actionable intelligence.

Splunk's flagship platform, Splunk Enterprise, provides IT professionals with a flexible and powerful capability for taking any form of machine data (structured or unstructured), indexing it, and then providing analytical insights for supporting security, operations, and application performance use cases. For eight consecutive years, until the latest 2021 report, Splunk Enterprise has been recognized by Gartner as a leader in the Security Information and Event Management (SIEM) Magic Quadrant. Today, over 22,000 Commercial and Public Sector organizations and 95% of the Fortune 100™ trust Splunk as the platform of choice for turning machine data (i.e. log data) into analytical insights.

The Splunk Enterprise platform is typically deployed in three primary ways for customers:

1. Splunk Enterprise deployed and operated “on-prem” within an organization's own data center;
2. Deployed via Splunk Inc.'s SaaS offering — Splunk Cloud;
3. Splunk Enterprise deployed via a “bring-your-own-license” approach where Splunk is installed and managed within a cloud provider (e.g. Amazon Web Services, Microsoft Azure, Google Cloud Platform) or within a dedicated hosting environment (e.g. Rackspace, Kyndryl)

For many customers, options 2 and 3 above are not viable for security, operational, total cost, or policy reasons. Organizations that deploy Splunk Enterprise within their own data centers take on the requirement for ensuring that the infrastructure elements supporting the Splunk platform — compute, storage, network — are properly sized and configured to ensure acceptable performance and ease of operation. Further, given the nature of the wide array of use cases Splunk supports, the on-prem architectures supporting Splunk Enterprise must be designed in a way that provides for effective scaling of compute and storage resources as use case count and data ingest grows.

Kinney Group, Inc. (KGI) is a longstanding, certified Splunk Services Delivery Partner and authorized Splunk Technology Alliance Partner. Since 2013, KGI has delivered over 600 Splunk engineering services engagements to Commercial and Public Sector organizations in North America and Europe. Over the course of the past nine years, KGI has had first-person experience with organizations operating Splunk Enterprise “on-prem,” and are doing so in ways that are expensive to scale, expensive to operate, and not designed for Splunk to perform optimally at-scale. When Splunk Enterprise is expensive and does not perform, organizations cannot realize the rich returns Splunk can provide.

Based on experience, KGI believes customers required to operate Splunk Enterprise on-prem in their own data centers can do so very effectively with the Dell VxRail hyperconverged infrastructure (HCI) system, which combines the power of PowerEdge server hardware integrated with VMware vSphere, VMware vSAN, and the VxRail HCI system software.

KGI sees the following key advantages associated with running Splunk Enterprise on the Dell VxRail HCI system infrastructure:

- **VxRail HCI provides the ease and effectiveness of scale-out of compute nodes supporting the unique requirements of Splunk Enterprise indexers and search-head clusters.** Within the Splunk platform, indexers perform the tasks that enable Splunk Enterprise to ingest terabytes of data per day. Search heads, combined with Splunk’s innovative “schema-on-read” architecture, provide end-users with access to rich analytics provided via Splunk reports, dashboards, and knowledge objects. Optimal performance of indexers and search heads is a requirement for a high-performing Splunk environment. The VxRail architecture, combined with VxRail HCI software, allows organizations to easily scale out their Splunk systems with minimal disruption.
- **Virtualizing Splunk Enterprise within a hyperconverged infrastructure provides flexibility and maximizes processing output for each hardware node.** Historically, Splunk has recommended that “on-prem” customers deploy Splunk on bare metal versus virtualized infrastructures. (Note: KGI has always found this recommendation interesting given that Splunk Cloud operates within a multi-tenant, virtualized environment in AWS). Deploying Splunk Enterprise on bare-metal server hardware is expensive and inefficient. Using VMware vSphere combined with the VxRail HCI system can more effectively provide the appropriate compute and memory resources required for Splunk Enterprise configurations compared to a bare-metal deployment approach.
- **VxRail HCI enables a common architecture that is suitable for Splunk Enterprise environments of all sizes.** This is an important consideration for organizations that may be running multiple instances of Splunk Enterprise in physically disparate data center environments. Having a consistent architecture approach that for environments of all sizes provides simplicities and efficiencies for customers operating Splunk Enterprise at multiple locations.
- **VxRail HCI combined with KGI’s Atlas platform dramatically simplifies and optimizes the operation of Splunk Enterprise.** Splunk environments are powerful and are also complex. The

VxRail HCI system approach simplifies the deployment and scaling of compute, network, and storage infrastructure needed to operate Splunk Enterprise. KGI's Atlas platform provides Splunk administrators with a comprehensive platform of utilities, automations, and tooling needed to optimize the operation and sustainment of Splunk Enterprise. The combination of VxRail and Atlas create a new paradigm of simplicity, performance, and efficiency for operating Splunk Enterprise environments of all sizes.

Dell EMC VxRail Using Splunk SmartStore and ECS Enterprise Object Storage

Traditional Splunk Enterprise deployments use a data life cycle management scheme designating data as being “hot,” “warm,” or “cold,” storing the data in “buckets” that align with these designations. Newer data is stored in hot and warm buckets requiring high-performance storage. Once data is considered archival in nature, it is moved into cold or “frozen” buckets for long-term storage. This data is stored and replicated on Splunk indexers, meaning that as data volume increases, additional indexers must be added or expanded.

Splunk SmartStore takes a different approach. SmartStore is a Splunk indexer capability that makes use of remote object storage — such as Dell EMC ECS — to store indexed data. As a deployment's data volume increases, demand for storage typically outpaces demand for compute resources. SmartStore provides the ability to manage indexer storage and compute resources in a more cost-effective manner by scaling these resources separately.

With SmartStore, the indexer storage footprint can be minimized, and I/O optimized compute resources can be chosen. Most data resides on remote storage, while the indexer maintains a local cache containing a minimal amount of data, including hot buckets and copies of warm buckets participating in active or recent searches.

Combining VxRail with ECS Enterprise Object Storage and Splunk SmartStore offers the following key advantages:

- **Lower Total Cost of Ownership (TCO).** Compute and storage are effectively decoupled, providing the ability to scale storage for longer retention and scale Splunk indexers to meet ingest or other performance demands. Indexer footprint is also reduced, as warm/cold storage is replaced with the SmartStore cache.
- **Performance at scale.** Searched data is brought to high-performance cache storage in VMware vSAN as needed, reducing ever-growing cold storage requirements. Cached data is based on age, priority, and access patterns.
- **Faster failure recovery.** Because direct storage for Splunk indexers is minimized and the majority of data is decoupled from indexers, this allows for significantly faster indexer recovery and rebalancing.

To harness the power of the VxRail HCI system for supporting Splunk Enterprise, KGI has created this VxRail Reference Design for Splunk. This design makes optimal use of the VxRail HCI system architecture and VMware vSphere virtualization, and does so in a manner that aligns with published Splunk best practices and infrastructure requirements.

The designs created by KGI have been tested and validated in Dell's Customer Solution Center (CSC) environment. KGI executed performance testing on two primary design sizes that make use of VxRail HCI nodes in combination with Dell EMC ECS storage. The results of KGI's efforts were conclusive – the Dell VxRail HCI system is a superb choice for organizations seeking to operate Splunk Enterprise within their data centers, and do so in a way that is performant, efficient, and flexible.

Solution Design

Overview

The following reference architecture describes using Dell EMC VxRail hyper-converged infrastructure combined with Dell EMC ECS enterprise object storage for a virtualized Splunk Enterprise environment using Splunk SmartStore for data retention.

VxRail is a fully integrated, preconfigured, and pre-tested hyper-converged infrastructure solution. Powered by industry-leading vSAN and vSphere software, VxRail is the easiest and fastest way to streamline and extend a VMware environment while dramatically simplifying IT operations.

Figure 1 (following page) represents a reference architecture utilizing Splunk instances as virtual machines on a VMware vSphere 7.0 cluster following Splunk's documented virtualization best practices.

In the storage layer, VxRail leverages VMware vSAN technology to build a vSAN on groups of locally attached disks. This configuration provides rapid read and write disk I/O and low latency using an all-flash array. This Virtual SAN is used to hold all virtual machines and the Splunk SmartStore cache. For longer-term data retention, Dell EMC ECS is used as an S3-compatible object storage solution for Splunk SmartStore.

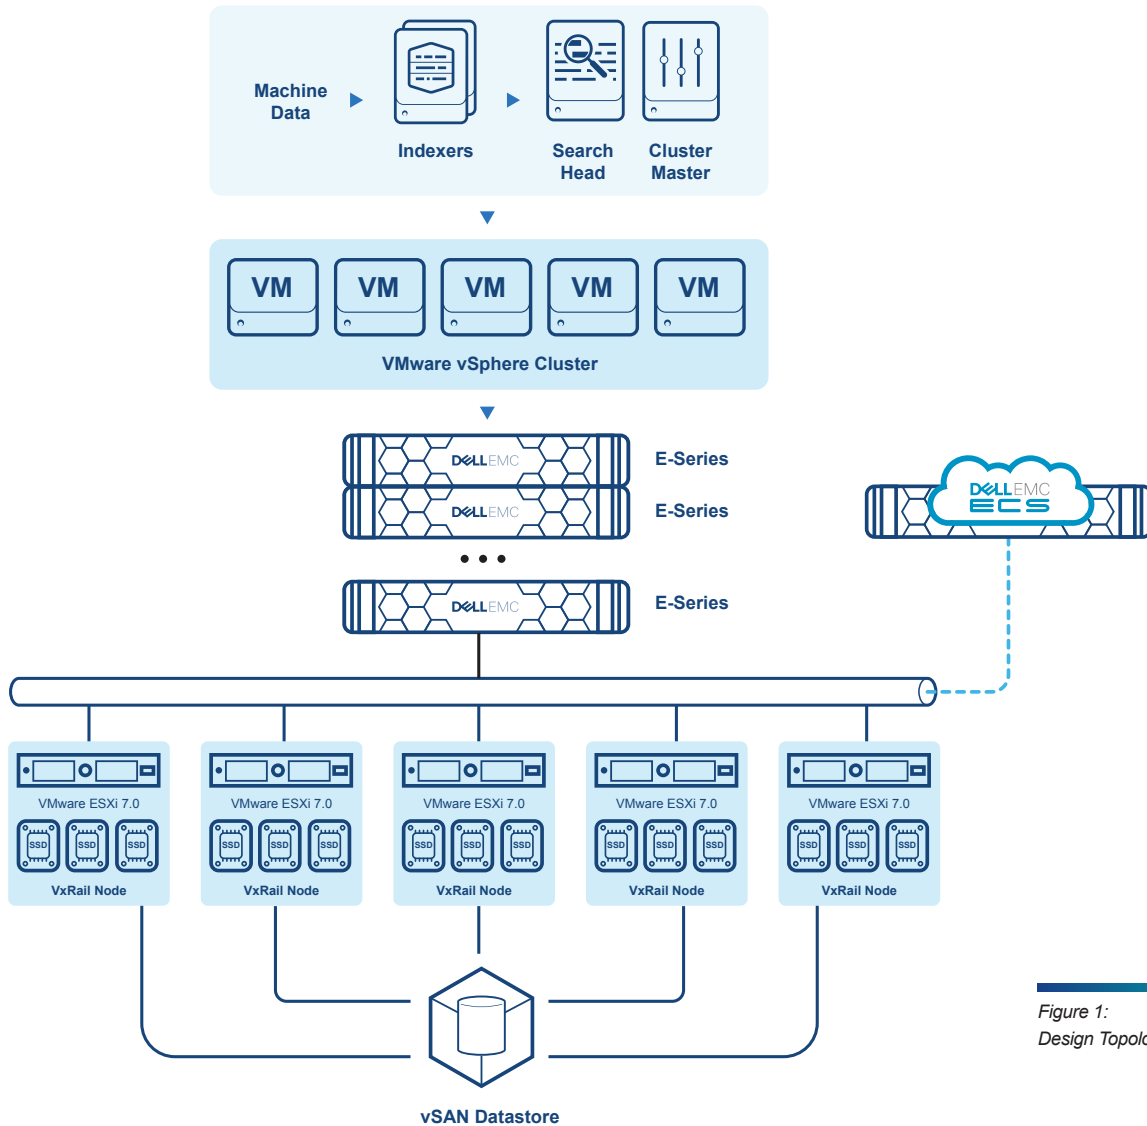


Figure 1:
Design Topology

Representative Use Cases

This solution design outlines the requirements for a Splunk Enterprise deployment utilizing Splunk Enterprise Security at three daily ingest rates (500GB, up to 2TB, and up to 5TB). Of course, if Enterprise Security is not required, ingest needs differ, or longer-term retention requirements change, the solution can be adjusted due to the scalable nature of VxRail and ECS. In all scenarios a SmartStore cache equivalent to approximately 30 days of ingest was chosen.

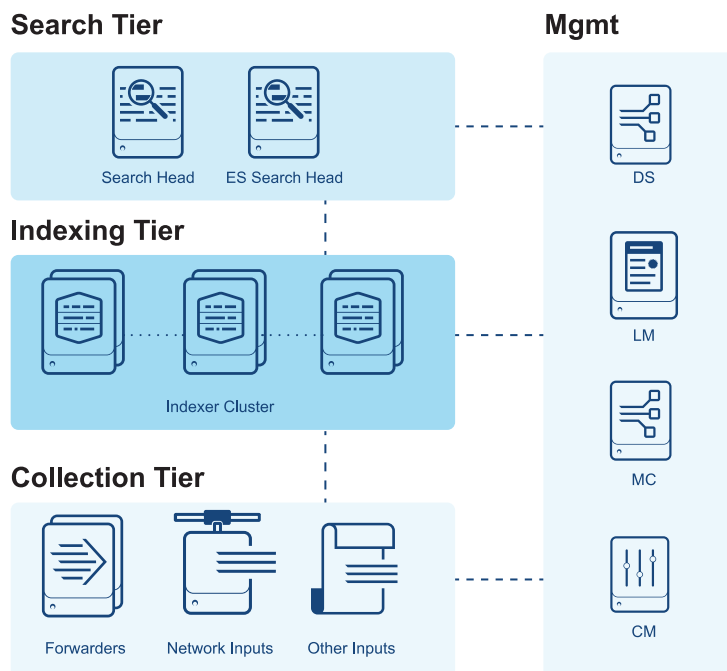
Utilizing Splunk SmartStore and Dell EMC ECS, long-term data retention can scale independently of the size and number of Splunk indexers. Examples of ECS storage requirements based on daily ingest vs. desired retention periods are shown in the Table 1 below:

Daily Ingest	120 Day Retention	1 Year Retention	2 Years Retention
Up to 500GB	30 TB	91 TB	183 TB
Up to 2TB	120 TB	365 TB	730 TB
Up to 5 TB	300 TB	913 TB	1,825 TB

Splunk Architecture

Figure 2 shows the Splunk architecture chosen for this solution, which comes from Splunk’s Validated Architectures. For more information about the architecture chosen, please see <https://www.splunk.com/pdfs/technical-briefs/splunk-validated-architectures.pdf>

The design makes use of a single search head to run Splunk Enterprise and a separate search head for Splunk Enterprise Security, and multiple indexers in a cluster. For larger environments where more concurrent users are expected, a search head cluster is specified.



Splunk Virtual Machine Requirements

Table 2 below shows the virtual machine requirements used for this design, chosen to meet or exceed Splunk recommendations for an Enterprise Security deployment. Splunk Deployment Server(s) were not included as part of these requirements as their usage depends on customer needs.

Daily Ingest	vCPU	RAM	Storage (vSAN)
Cluster Manager	12	12 GB	300 GB
Administrative Server	12	12 GB	300 GB
Search Head	32	16 GB	300 GB
Search Head (ES)	64	32 GB	300 GB
Indexer ≤ 500 GB/day	24	24 GB	1.85 TB
Indexer > 500 GB/day	48	64 GB	

The Administrative Server noted in the table above includes Splunk server roles such as the license manager, monitoring console, and (for larger deployments), a search head deployer.

Storage requirements for each indexer includes allocation for the SmartStore cache and assumes an overall cache size approximately equivalent to 30 days of ingest. It also assumes the default Splunk settings of indexer cluster replication factor of 3 and search factor of 3.

VxRail Nodes and Virtual Machines

For this reference design, VxRail E series servers were chosen for their balance of compute vs. storage and ability to support a vSAN configuration comprised entirely of NVMe storage. Given the disk I/O requirements of Splunk, Intel Optane for cache and NVMe for storage is specified as part of this design and is strongly recommended to ensure the highest possible vSAN performance.

Table 3 below shows the specification for VxRail E560N servers chosen for this design:

Processor	2 x Intel 6258R (2.7GHz, 28 cores each)
Memory	192 GB
Cache	2 x 375 GB Intel Optane (1 per disk group, 2 disk groups/node)
Storage	6 x 1.92 TB NVMe (3 per disk group, 2 disk groups/node); Leaves 2 free capacity slots for additional disks.
Network	2 x 25 GbE Ports

Using the individual node specifications noted above, the resulting VxRail HCI cluster sizes for the three chosen daily ingest rates are shown below in Table 4:

Daily Ingest:	500 GB / day	2 TB / Day	5 TB / Day
E560N Nodes	5	11	24
Total CPU cores	280	616	1,344
Total RAM (GB)	960	2,112	4,608
Total raw storage (TB)	57.6	126.7	276.5
Total usable storage (TB)	17.7	45.5	105.9

Note: The usable storage shown above assumes a vSAN "Failures to Tolerate" (FTT) policy of 1.

Following Splunk recommended guidelines, Table 5 below shows the number of virtual machines required for the daily ingest rates specified in this design:

Daily Ingest:	500 GB / day	2 TB / Day	5 TB / Day
Cluster Manager	1	1	1
Administrative Server	1	1	1
Search Heads	1	3 (SH Cluster)	3 (SH Cluster)
Search Heads (ES)	1	1	1
Indexers	5	20	50

Testing and Validation

Overview

This solution design was tested and validated at the Dell Customer Solutions Center (CSC) using multiple daily Splunk ingest rates from 500 GB / day up to and including 2 TB / day.

Hardware Used

For compute, a cluster of VxRail E560 servers was used, configured to meet the specifications noted in Table 4 above. For Splunk SmartStore storage, a Dell EMC EX300 cluster was targeted.

Virtualization Settings

For each scenario, virtual machines comprising the Splunk Enterprise cluster were configured using the CPU and memory requirements noted in Table 3 and followed Splunk recommendations for deploying in a virtualized environment. All CPU and RAM resources were fully reserved for each virtual machine.

Indexer and SmartStore Configuration

A single 1.85 TB storage volume was configured for each Splunk indexer. While this amount of storage can handle the Splunk Enterprise installation and a SmartStore cache size equivalent to approximately 30 days of ingestion, the SmartStore cache size was kept low during testing to force cache eviction, causing searches to fetch data from SmartStore more frequently. To reduce cache thrash, the SmartStore cache can and should be adjusted to a higher amount in a production environment. Splunk Enterprise was configured with an indexer cluster replication factor of 3 and a search factor of 2.

The Dell EMC ECS storage cluster was configured with an S3-compatible bucket dedicated to Splunk SmartStore storage. Splunk was then configured with the relevant URI and authentication information to access this bucket.

Software Used

The following Splunk related software packages were used in the testing of this design:

- Splunk Enterprise 8.2.4
- Splunk Enterprise Security (latest version)
- Splunk_TA_cisco-asa (latest version)
- Splunk_TA_cisco-esa (latest version)
- Splunk_TA_cisco-wsa (latest version)
- Splunk_TA_isc-bind (latest version)

Note: Ingest rates above 2 TB/day were not tested due to limitations in the resources available at the Customer Solution Center where testing and validation of this design took place. Design calculations, however, indicate 5 TB/day is quite achievable using the requisite number of VxRail server nodes indicated in Table 4.

- Splunk_TA_mcafee (latest version)
- TA-crowdstrike (latest version)
- splunk_app_gogen (version 0.5)

Testing Procedure

Testing was designed to mimic a real-world customer environment running Splunk Enterprise Security (ES) on a single indexer cluster. One search head was dedicated to ES, and additional search head(s) were configured depending on the size of the environment. For ingestion rates at or below 500 GB / day, a single additional search head was deployed, while larger environments were configured with three search heads in a search head cluster configuration.

For sample data generation, the synthetic log tool Gogen was used, which is an application capable of reading samples and generating events to simulate various data sources for ingest, parsing, and load testing of various systems. The application also includes splunk_app_gogen, a modular input wrapper for streaming Gogen events into Splunk. For testing, Gogen configuration files were generated to simulate twenty (20) sourcetypes matching the various Splunk TAs mentioned above. Gogen and the related splunk_app_gogen were installed on multiple Splunk universal forwarders external to the Splunk cluster, the number of which was determined by the daily ingest amount being targeted.

Enterprise Security Tuning

Out of the box, Splunk Enterprise Security (ES) contains many inefficiencies in the search configuration. Enterprise Security was tuned to avoid skipped searches as much as possible while attempting to maintain the number of scheduled searches in the environment. In addition to enabling Splunk ES correlation searches, all common ES related data models were accelerated.

As a reminder, this tuning should be performed in any production Splunk ES environment to avoid skipped searches. While results were seen in the testing environment, the tuning should be continued for additional gains in scheduled search capacity in production.

Indexing and Search Performance

To illustrate the indexing and search performance of the representative use-cases for this design, the 500 GB/day and 2 TB/day ingestion rate scenarios were chosen.

500 GB / Day Ingestion Rate

Figures 3 through 5 below show Splunk monitoring console output for cluster indexing performance, scheduled search performance, and SmartStore cache hits/misses at an ingestion rate of 500 GB/day. These indicate that 5 indexers were ingesting at or above the stated amount and that skipped searches remained at or near zero even with a high SmartStore cache miss ratio since the SmartStore cache settings were kept low to force more frequent downloads from SmartStore remote storage for testing.

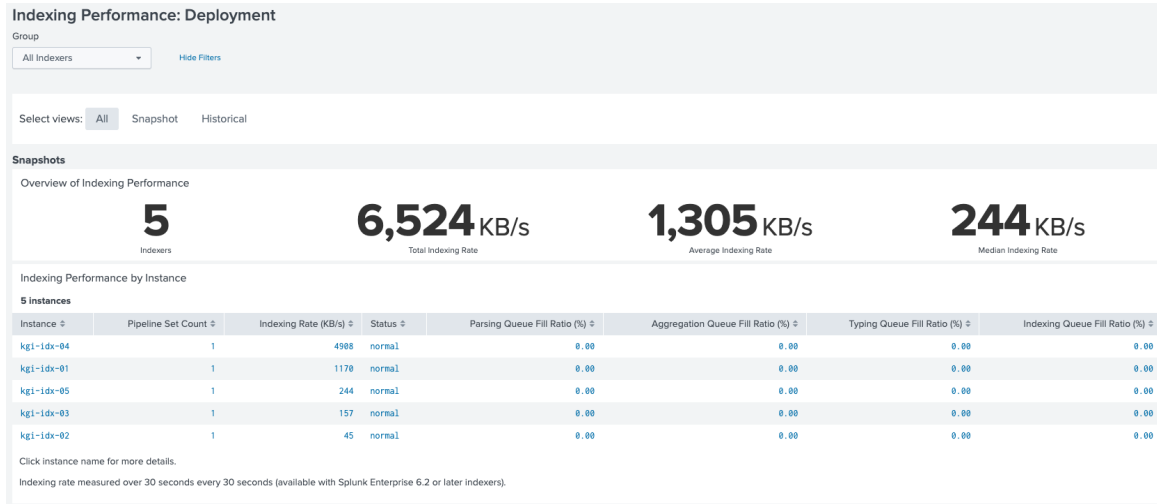


Figure 3: Indexing performance at 500 GB/day

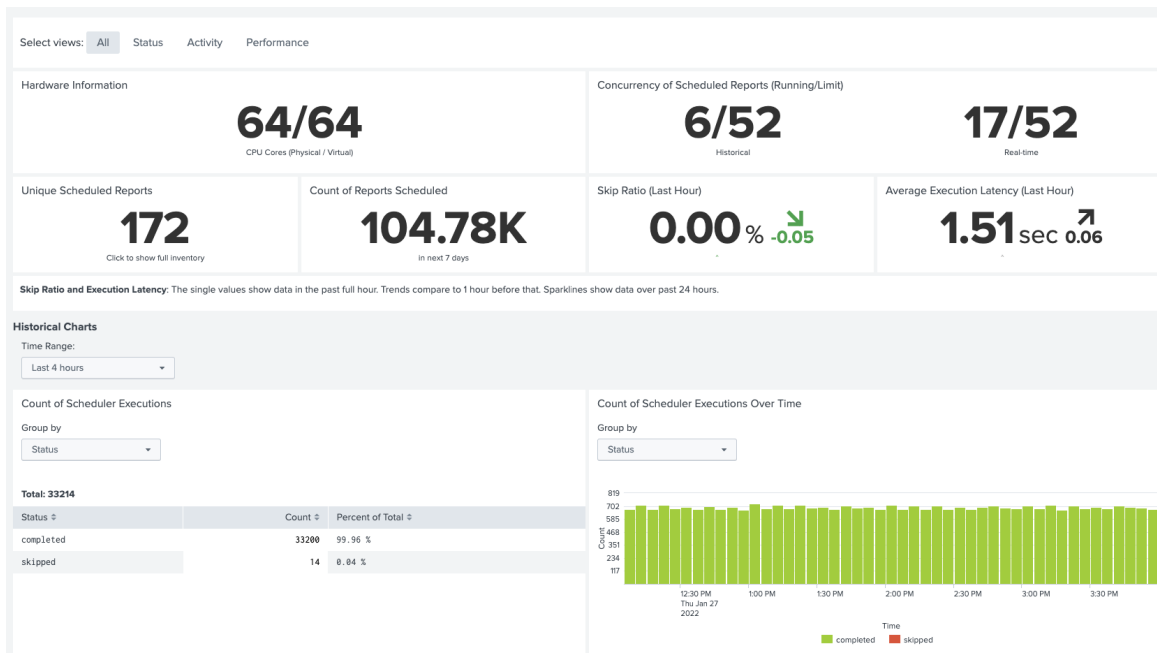


Figure 4: Scheduled search performance at 500 GB/day




Figure 5: SmartStore cache hits/misses at 500 GB/day

2 TB / Day Daily Ingest Rate

Similar to the 500 GB/day results shown above, figures 6 through 8 below show Splunk monitoring console output for cluster indexing performance, scheduled search performance, and SmartStore cache hits/misses at an ingestion rate of 2 TB/day. These indicate that 20 indexers were ingesting at or above the stated amount and that skipped searches remained at or near zero even with a high SmartStore cache miss ratio since the SmartStore cache settings were kept low to force more frequent downloads from SmartStore remote storage for testing.

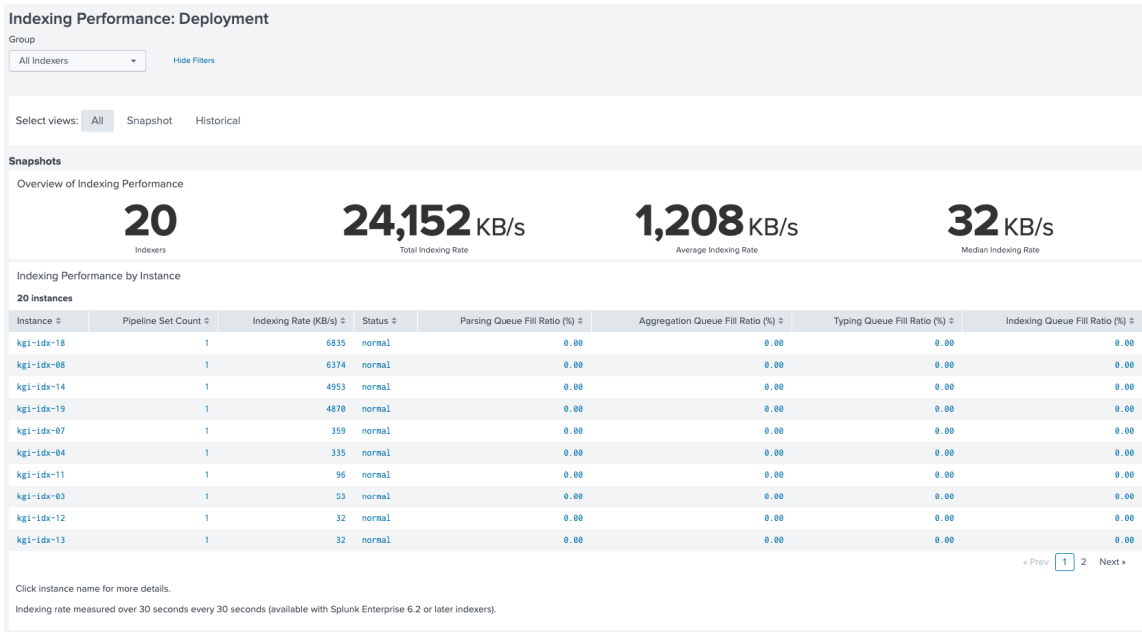


Figure 6: Indexing performance at 2 TB/day

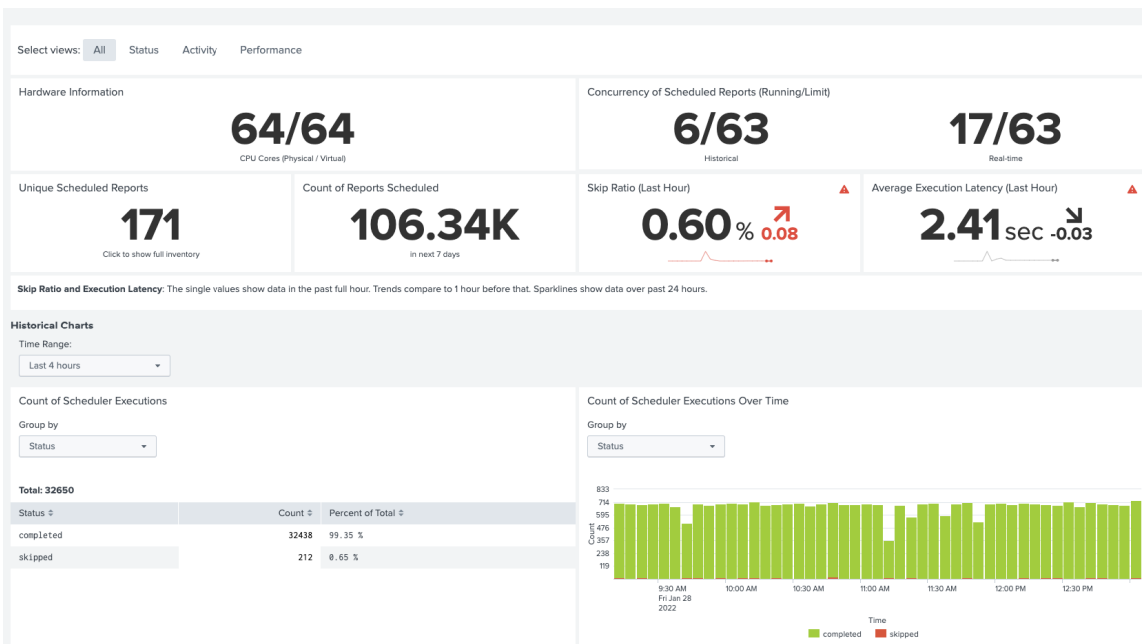


Figure 7: Scheduled search performance at 2 TB/day




Figure 8: SmartStore cache hits/misses at 2 TB/day

Additional Testing Observations

During testing, average CPU and memory usage for the virtual machines in the Splunk Enterprise cluster for each scenario was at or below 50%.

In all tested scenarios, disk IOPS were measured concurrently on all virtual machines to test vSAN storage I/O performance. Average disk I/O met or exceeded Splunk recommendations in all cases. To ensure the best performance for the SmartStore cache, NVMe should be used for vSAN.

For testing, the Dell EMC ECS EX300 was used for SmartStore storage and retention. Being a 7,200 RPM SATA solution, it performs well to meet retention and compliance requirements, assuming that frequently used Splunk searches are typically within the 30-day cache window outlined in this design. Splunk searches frequently made outside of this timeframe would result in additional SmartStore “cache misses”, resulting in more fetches from slower-tier storage on ECS and possibly longer search times. If this is expected, it would be beneficial to increase the SmartStore cache size by expanding the overall capacity of vSAN, or by moving to a more performant SmartStore solution using an all-flash capable ECS cluster, for example the Dell EMC ECS EXF900.

Conclusion

With this reference design, customers should be able to match current and future needs for a Splunk Enterprise deployment using Enterprise Security that meet or exceed Splunk recommendations. Dell EMC VxRail combined with Dell EMC ECS enterprise object storage provides a solution that can be scaled to handle future needs without the need for extensive upgrades. Utilizing Splunk SmartStore, customers also gain the flexibility to scale compute and storage of indexed data independently without sacrificing performance as data ingestion and compliance requirements change.



Copyright © 2022 Kinney Group, Inc., and Dell Inc. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA March 2022.

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Dell VxRail Splunk RD 2-20328