

Splunk Reference Architecture

Pure Storage FlashStack plus VMware Virtualization



Executive Summary

Kinney Group is a cloud solutions integrator specializing in analytics, automation, and hybrid cloud solutions. We design, build, and integrate IT infrastructure solutions for some of the most demanding government agencies and commercial organizations. By leveraging next-generation technologies, adopting proven engineering practices, and agile development principles, we create custom solutions and world-class environments for data.

Kinney Group, Inc. (KGI) is an award-winning, certified Splunk Elite Partner. Our team has experience working with Splunk deployments of all sizes, at various stages of execution, and across a variety of use cases. We've helped Commercial and Public Sector organizations design, develop, and implement Splunk at scale. This work includes guiding organizations with the design of their on-premise infrastructure for supporting the Splunk Enterprise platform.

KGI is leading the way in designing a virtualized reference architecture that can be utilized by organizations as guidelines for building their own resilient Splunk Enterprise environments. KGI has utilized their experience with the Pure Storage FlashBlade platform to create a virtualized design that operates on the Cisco-validated FlashStack infrastructure design.

This paper is intended to provide a framework for designing and sizing a high-performance, scalable, and resilient Splunk platform for core Splunk Enterprise and Splunk Enterprise Security (ES). The reference design in this paper utilizes the VMware platform for server virtualization and also uses the Splunk SmartStore technology for enabling streamlined movement of data between the on-premise FlashStack and AWS S3 storage resources. Using combination of VMware, Pure Storage FlashBlade, and Splunk SmartStore, customers can reduce storage complexity and datacenter footprint while maintaining platform performance, resiliency, and efficiency.

Splunk hardware specifications recommends 20 indexers for 2TB daily ingest. We observed acceptable user experience with only 12 indexers running on a FlashStack design using all FlashBlade storage. This resulted in a 40% improvement over Splunk's legacy bare-metal infrastructure design recommendations.

The below table shows our key findings:

Daily Ingest	Number of Indexers	Disk Latency	Search Latency	Skipped Searches
2TB	12	<10 ms	<3 seconds	0

GOALS AND OBJECTIVES

The goal of this reference architecture is to showcase the scalability, performance, manageability, and simplicity of the virtualized FlashStack-based solution for deploying a large scale Splunk Enterprise or Splunk ES environment.

The key objectives for this reference architecture:

- Design repeatable architecture that can be implemented quickly in production sites
- Utilize VMware to reduce datacenter footprint and scale environment quickly
- Utilize SmartStore to take advantage of volume reduction technologies and reduce overall disk requirements for the environment

AUDIENCE

The target audience for this document includes, but is not limited to, system administrators, storage administrators, IT managers, system architects, sales engineers, field consultants, professional services, and partners who are looking to design and deploy Splunk Enterprise on a virtualized FlashStack platform. A working knowledge of Splunk, VMware, Linux, server, storage, and networks is assumed but is not a prerequisite to read this document.


REFERENCE ARCHITECTURE DESIGN PRINCIPLES

The guiding principles for implementing this reference architecture are:

Simple	Using pre-built images and apps, we minimize the amount of manual configuration required.
Secure	Combination of solid security architecture and compliance concepts built in the design. Supports encrypted data-at-rest functionality and is FIPS 140-2 compliant.
Available	By using a combination of indexer clustering, Splunk SmartStore, and FlashBlade, we can create an environment that needs nearly zero downtime for upgrades and updates and is fault tolerant to unexpected failures.
Efficient	By utilizing VMware, SmartStore, and FlashBlade, we reduce the overall required datacenter footprint, saving power and cooling costs, and also reducing overall sustainment and operations costs.
Cost Effective	This design enables customers to easily move data between on-premise and cloud-based (i.e. S3) storage.
Elasticity	Through the use of virtualization combined with automation, this reference design can be scaled based on customer's daily data volume.

Solution Design

DESIGN TOPOLOGY



VIRTUAL SERVER CONFIGURATION

For ES, Splunk recommends sizing based on 80 to 100 GB ingest per indexer per day. This means an ES deployment with 2 TB daily ingest will require up to 20 indexers. Based on our test using this reference design and tuning, we were able to achieve similar, if not better, performance using 60% of the hardware. This reference design will require no more than 12 virtualized indexers to support an ES deployment with 2 TB daily ingest with up to 24 active users. The below table details the count and configuration of virtual machines for each component of Splunk. Long-term storage will be handled by SmartStore, with some local storage on indexers for data caching.


Component	Description	Count
Indexer	16 vCPU 64 GB vRAM 200 GB Local Storage Pure Storage SmartStore	12
Search Head (Enterprise Security)	16 vCPU 64 GB vRAM 200 GB Local Storage	1 ES 1 non-ES
Cluster Master	12 vCPU 32 GB vRAM 200 GB Local Storage	1
Deployer/Deployment Server	12 vCPU 12 GB vRAM 200 GB Local Storage	1
Heavy Forwarder	12 vCPU 32 GB vRAM 200 GB Local Storage	As needed

All servers will be run on a VMware stack hosted on the validated Pure FlashStack solution. For information about that solution, please refer to Pure documentation.

PHYSICAL TOPOLOGY

Transcending the conventional model of bare metal installs for Splunk, the FlashStack solution for Splunk involves all virtual machines. Apart from the benefits of server consolidation, rapid deployment & provisioning, and ease of management that VMs provide, the primary reason for choosing virtual machines is to allow for flexible workload positioning and scale out. By leveraging virtualization, it is possible to rapidly scale the compute layer, either at a resources per machine level or number of machines to match your required workloads. Leveraging the FlashStack reference design allows you to leverage an industry proven, fully documented hardware configuration to support your Splunk environment. By using Pure as the shared storage backbone, you get the benefits of Highly Performant storage with the business benefits such as Evergreen storage and non-disruptive upgrades.


Pure Storage FlashStack consists of a combined stack of hardware (storage, network, and compute) and software (Cisco UCS Manager, Splunk Enterprise & Splunk ES, Pure Storage GUI, Purity, Red Hat Enterprise Linux). The following diagram shows the architecture of Pure FlashStack with Splunk SmartStore.



SPLUNK ARCHITECTURE

The architecture chosen for this solution 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 Security, and multiple indexers in a cluster for data resiliency.



VMWARE BEST PRACTICES

The VMware configuration documented in their technology white paper published here and Splunk technical brief for deploying on VMware, found here, was used as a starting point to configure.

Based on our knowledge and experience with the Splunk platform, some of these configurations were further refined to achieve optimal data ingest and search performance

PERFORMANCE REPORTING APP

The Search Head will include a pre-build app to monitor the overall health of the Pure Storage array, along with measuring and reporting on performance metrics from the storage array. This app will be used in conjunction with pre-built searches from the Splunk Monitoring Console to determine performance.

TEST CONFIGURATION

TEST OVERVIEW

We are able to confirm the validity of a Splunk architecture utilizing virtualized Splunk indexers, Pure Storage, and Splunk SmartStore. This reference architecture is capable of handling a standard Splunk Enterprise Security load with a daily ingest of up to 2TB.

We were able to test this theory by loading a set amount of machine data into the indexing layer. The replication factor and search factor for the index cluster were both set to 3. Splunk SmartStore was enabled to archive historical data. To generate search load, datamodel acceleration was enabled on all data models with correlation searches scheduled to run through the duration of the test. In addition to these automated search lead, we had users run adhoc searches to measure response time and overall user experience.

This reference design includes Puppet Enterprise (PE) and automation modules to automate to automate deployment and enforce integrity of configurations and security settings.

HARDWARE USED

As shown in the physical topology, the test lab used pre-validated Pure FlashStack design using Pure FlashBlade and Cisco UCS. For more information on this design, please see Pure documentation.

SOFTWARE USED :

The following software packages are used in this design:

- Splunk 7.3
- Splunk Enterprise Security (latest version)
- SA-EventGen (KGI Custom Version)
- Splunk_TA_Stream (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)
- Splunk_TA_mcafee (latest version)
- Splunk_TA_nessus (latest version)
- TA-crowdstrike (latest version)
- TA-ps_flashblade (latest version)

VIRTUALIZATION SETTINGS :

Splunk has provided recommendations for virtualization in Deploying Splunk Enterprise Inside Virtual Environments. All these recommendations along with performance best practices guide for vSphere were followed while provisioning VM and allocating storage. Splunk VM uses RHEL 7.3.

The following guide was followed for all VMware configuration: <https://storagehub.vmware.com/t/vmware-vsan/splunk-on-vmware-vsan/splunk-virtual-machine-configuration/>

INDEX VOLUME AND SMARTSTORE CONFIGURATION

A single primary index volume was configured on each server. This volume was set to 100 GB. In addition, the SmartStore Cache was also set to 100 GB. While this worked in for early ingestion loads,

larger ingestion loads caused this space to fill up quickly, causing many bucket evictions (“cache thrash”) from the primary volume as seen below. To eliminate cache thrash, SmartStore cache had to be adjusted.

Index S	Repeat Download Percent S	All Downloads S	Repeated Downloads S	Excessively Repeated Downloads S
audit_summary	58.88 %	2	1	0
cim_reductions	58.88 %	2	1	0
risk	38.38 %	11	4	0
notable	35.71 %	14	5	0
_internal	26.58 %	79	21	0
main	22.81 %	268	59	0
_audit	16.67 %	48	8	0
endpoint_summary	14.29 %	7	1	0
_introspection	9.89 %	19	0	0
threat_activity	9.89 %	1	0	0

TESTING PROCEDURE:

Testing was designed to mimic a real-world customer environment with a single index cluster and Enterprise Security running on a single search head.

- Build Splunk Environment consisting of 1 SH, 4 IDX, 4 HF, 1 Cluster Master (see Design Topology for virtualized server sizing)
- Configure Splunk SmartStore to utilize Pure Storage
- Enable Event Gen on HFs to generate a level of data from the below chart
- Enable Splunk ES datamodel acceleration on all data models.
- Enable up to 5 Correlation searches that run on at least a 1-hour period
- Utilize Monitoring Console and PureStorage-TA to measure health of data ingestion queues, search scheduler skip ratio, search latency, storage IOPs and latency.

DATA INPUT

To mimic real-world data, we used EventGen to load sample data into Splunk. This allowed generating load in a nearly identical way as a production environment and ensure the data matches requirements for Enterprise Security CIM compliance.

The following data ingestion rates were tested using this reference architecture:


Volume Levels	Daily Ingest	Description
1	500 GB	Matches Splunk Reference Sizing for ES Indexers, and is a requirement to be met for any reference architecture.
2	960 GB	Almost double the Splunk Reference Sizing for ES Indexers. This represents a 50% reduction in required indexers.
3	1,920 GB	Almost 4x the Splunk Reference Sizing for ES Indexers. This represents a 75% reduction in required indexers.

ENTERPRISE SECURITY TUNING


Out of the box, Splunk ES contains many inefficiencies in the search configuration. Using KGI expertise,

Enterprise Security was tuned to avoid skipped searches while maintaining the level of searches in the environment. This included updating the timing of searches as well as providing additional search slots in the software.

Below shows how, out of the box, ES can already skip some scheduled searches due to inefficiencies:



Below we see the results after applying KGI expertise in tuning ES:



PERFORMANCE TESTS

DATA INGESTION

Data ingestion was set up at the following levels using eventgen: 500 GB per day, ~1 TB per day, and ~2 TB per day.

500 GB/day Results

When data ingestion is turned up to 500 GB/day – the Splunk recommended limit – we see that the indexers remain stable and ingestion queues are not maintaining high values (throughout testing, we noticed some momentary volume in ingestion queues, but these values are not maintained unless otherwise noted).

Data Volume - 1 Hour (Informational)				Indexing Rate (Informational)				
21.11 GB [↑] _{0.19}				Instance ↕	Pipeline Set Count ↕	Pipeline Set Selection Policy ↕	Indexing Rate (KB/s) ↕	Status ↕
				vsplunk-ix04	1		673	normal
				vsplunk-ix01	1		1360	normal
				vsplunk-ix03	1		509	normal
				vsplunk-ix02	1		1436	normal
				vsplunk-ix05	1		1777	normal
Current Queue Depths (Should = 0)								
Instance ↕	Pipeline Set Count ↕	Status ↕	Parsing Queue Fill Ratio (%) ↕	Aggregation Queue Fill Ratio (%) ↕	Typing Queue Fill Ratio (%) ↕	Indexing Queue Fill Ratio (%) ↕		
vsplunk-ix05	1	normal	0.00	0.00	0.00	0.43		
vsplunk-ix04	1	normal	0.00	0.00	0.00	0.00		
vsplunk-ix03	1	normal	0.00	0.00	0.00	0.00		
vsplunk-ix02	1	normal	0.00	0.00	0.00	0.00		
vsplunk-ix01	1	normal	0.00	0.00	0.00	2.65		

1 TB /day Results

When data ingestion is increased to 1 TB /day – double the Splunk recommended limit – we continue to see healthy indexers. Here we can see the queues are still not maintaining any high values.

Data Volume - 1 Hour (Informational)				Indexing Rate (Informational)				
41.76 GB [↓] _{-0.16}				Instance ↕	Pipeline Set Count ↕	Pipeline Set Selection Policy ↕	Indexing Rate (KB/s) ↕	Status ↕
				vsplunk-ix05	1		2185	normal
				vsplunk-ix04	1		3179	normal
				vsplunk-ix01	1		2154	normal
				vsplunk-ix02	1		3136	normal
				vsplunk-ix03	1		1499	normal
Current Queue Depths (Should = 0)								
Instance ↕	Pipeline Set Count ↕	Status ↕	Parsing Queue Fill Ratio (%) ↕	Aggregation Queue Fill Ratio (%) ↕	Typing Queue Fill Ratio (%) ↕	Indexing Queue Fill Ratio (%) ↕		
vsplunk-ix05	1	normal	0.00	0.00	0.00	0.00		
vsplunk-ix04	1	normal	0.00	0.00	0.00	0.00		
vsplunk-ix03	1	normal	0.00	0.00	0.00	0.00		
vsplunk-ix02	1	normal	0.00	0.00	0.00	0.00		
vsplunk-ix01	1	normal	0.00	0.00	0.00	0.00		

2 TB/day Results

When data ingestion is increased to 2 TB /day – four times the Splunk recommended limit – we begin to see indexer health reduced. Data ingestion queues begin to maintain high values, and the potential for data loss is introduced to the environment.



SEARCH PERFORMANCE

In addition to scheduling correlation searches. The following Splunk data models were accelerated, to represent the common customer use-cases for Enterprise Security:

- Authentication
- Email
- Web
- Malware
- Network Traffic
- Network Resolution
- Intrusion Detection
- Vulnerabilities
- Web

From the KGI Monitoring App, we looked for the following search metrics:

- Search Scheduler Skip Ratio of 0% throughout testing period
- Average search latency less than 3 seconds throughout testing period


Below are the results seen throughout testing.



As the charts above demonstrate, search performance was steady throughout the test duration. Increasing the number of correlation searches caused the first spike. At this time, we tuned to scheduled searches that caused the performance to return to normal acceptable levels. The second spike was observed when indexing rate was increased to 2TB per day. To mitigate this, we increase the number of indexers to 12 to distribute the overall load and responsiveness of the application.

STORAGE PERFORMANCE

Measuring overall IOPs on the storage array, we saw IOPs in excess of 1200 per indexer and a total storage read/write latency of less than 100 milliseconds. The Pure Storage benchmark is shown in the report, but only the Splunk benchmark is required for passing.



As we can see from the results, only at extremely high data volumes did the disk latency exceed the Pure benchmark of 1ms. In addition, we were able to achieve high IOPS values at the 2 TB/day data ingest.

Conclusions

REFERENCE ARCHITECTURE INDEXER SCALING

By utilizing the KGI and Pure design with FlashStack and VMware, a sizing capability of 180 GB per day per indexer of daily ingest is achievable without impacting indexing or search performance. As such, the following scaling table should be achievable with this architecture. This represents a 40% reduction in the number of indexers required for the workload.

Expected Total Daily Volume	Number of Indexers Recommended
2TB	12
6TB	36
12 TB	72

Appendix 1: Configuration Items

SMARTSTORE CONFIGURATION

The following details the configuration used for Splunk SmartStore on the index cluster. Some items may need to be adjusted for customer environments.

Volumes

The following volume configuration must be deployed to the indexers via the cluster master. The primary volume will be used for SmartStore cache on the indexers. The path should be set to the local path used for hot data (in this example /hot/splunk). The remote volume is used to point the indexers to the remote S3-compatible storage provider. The remote endpoint address, access key, secret key for the remote volume should be supplied by the Pure Storage Admin.

```
[volume:primary]
path = /hot/splunk

[volume:remote]
storageType = remote
path = s3://pk-smartstore
remote.s3.access_key = <access_key>
remote.s3.secret_key = <secret_key>
remote.s3.endpoint = http://<remote_ip>
```

Indexes

All indexes should be configured following the below template. This should be deployed to the indexers via the cluster master. The home path should be set to the primary volume where the SmartStore cache will be located. The remote path should be set to the remote volume. The cold path and thawed path must be set to avoid errors in Splunk, although they will not be used with SmartStore. The Splunk variable `$_index_name` is used for easy replication of individual index configuration, but cannot be used for the thawed path definition.

```
[customIndex]
homePath = volume:primary/$_index_name/db
remotePath = volume:remote/$_index_name
coldPath = $SPLUNK_DB/$_index_name/colddb
thawedPath = $SPLUNK_DB/customIndex/thaweddb
```