HITACHI
Inspire the Next

Hitachi Analytics Infrastructure for Splunk

Reference Architecture

MK-SL-204-01
November 2021

Legal Notices

Feedback

Hitachi Vantara welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@HitachiVantara.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

Revision history

Changes	Date
Support for Whitley processors on Hitachi Advanced Server DS120 G2 and Hitachi Advanced Server DS220 G2 servers.	November 1, 2021
Initial release	July 23, 2020

Reference Architecture

Hitachi Analytics Infrastructure for Splunk provides guidelines for deploying Splunk 8. Use this guide to implement an architecture that maximizes the return on your investment.

Splunk Enterprise is a software technology that is used for monitoring, searching, analyzing, and visualizing machine generated data in real time. It monitors and reads different types of log files and stores data as events in indexers. This tool allows you to visualize data in dashboards.

Splunk can consist of many different components such as the following:

- Indexer — This component stores the data.
- Indexer Cluster Master — This component coordinates the activities of an indexer cluster and distributes application configurations to the indexers.
- Forwarder — This component gathers the data and sends it to the search head.
  - Universal Forwarder — This lightweight component processes the runs on existing machines and forwards the data to the indexers.
  - Heavy Forwarder — This is a lightweight version of Splunk Enterprise that gathers data and forwards it to the indexers. It can store and manipulate the data before forwarding it.
- Search Head — This component is an instance of Splunk that distributes searches to the indexers.
- Search Head Captain — When configuring the Splunk deployment with a search cluster, this component coordinates job and replication activities among the search heads.
- Search Head Deployer — This component distributes applications and configurations to the search head cluster members.
- Deployment Server — This component distributes applications and configurations to other components, primarily forwarders.
- License Master — This component handles Splunk Enterprise licensing.
- HTTP Event Collector (HEC) — This component provides a means for Splunk to review data over HTTP or HTTPS.

Splunk can be deployed in multiple formats: standalone node, clustered node, and distributed with multiple independent nodes. It can consist of single-site or multi-site configurations. Some common deployments are as follows:

- Single Server Deployment — One server gathers the data, stores the data, and is used to search the data. Separate forwarders can be deployed.
- Distributed Indexer Non-Clustered Deployment — This deployment has a single search head, with multiple indexers working independent of each other. The forwarder forwards data directly to the individual indexers.
- Distributed Indexer Clustered Deployment — This deployment uses an indexer cluster master that forwards data directly to the indexer cluster master which then passes the data to the indexers.
- Distributed Indexer Clustered Deployment with a Search Head Cluster — Multiple search heads are added to the indexer cluster.

The following figure shows an example layout for a distributed non-cluster indexer with a single search head deployment.

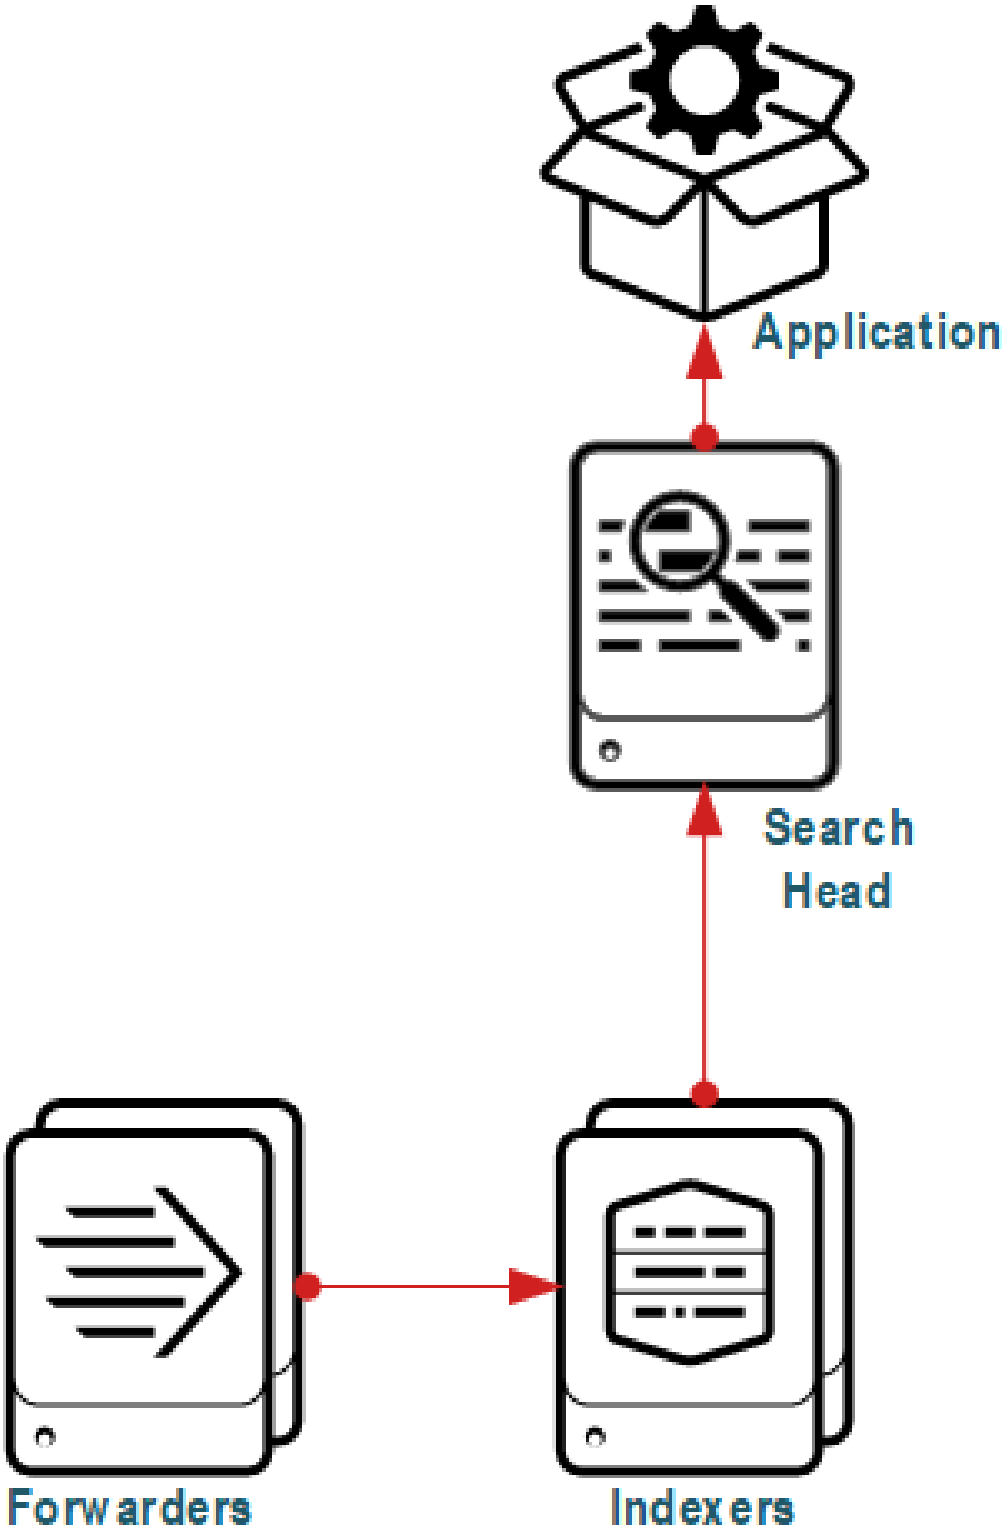

Splunk indexed data is stored in buckets, which are directories containing the data and its indexes. Splunk storage is built around the concepts of data temperature and has the following five data bucket tiers:

- Hot data — Data that is actively being written.
- Warm data — Data that is active but not being written to. Warm data uses the same storage as hot data. When a trigger is reached, either index size or indexer restart, and the hot data is rolled in place to a warm bucket.
- Cold data — After a condition is reached, data is moved from a warm bucket to a cold bucket. This move can be across different storage devices.
- Frozen Data — After a condition is reached, the data is moved to a frozen bucket. Frozen data is not searchable. Then default behavior is to delete frozen data. You can archive the frozen data for future retrieval.
- Thawed Data — Data retrieved from the frozen bucket. Depending on how the data is archived, retrieving this data can take a long time.

The following figure shows the data flow through the system.

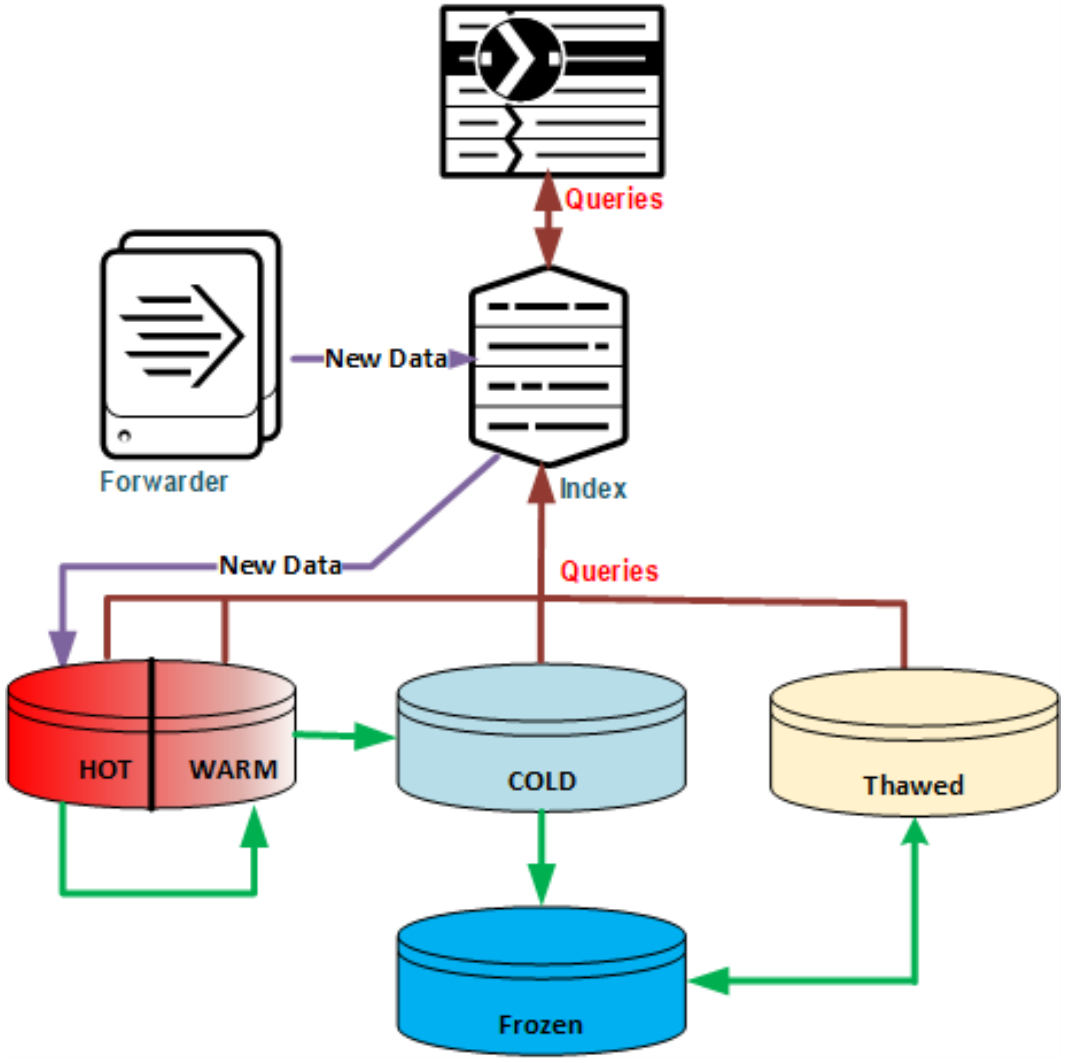

Data flow sequence

The sequence of the data flow through the system is as follows:

1. Forwarders send data to Splunk.
2. The data goes to an index on the indexer nodes. This stores the data in a bucket on the hot tier.
3. When certain conditions are met, Splunk performs an automatic logical move of data from the hot tier to the warm tier. This move keeps the data on the same storage device.
4. When another set of conditions are met, Splunk performs an automatic move of data from warm to cold. This move can be a physical move to different storage devices.
5. When the next set of conditions are met, the data moves to frozen storage. The default behavior is to delete this data. If it is being moved, its usually to different storage devices and could be remote. Frozen data cannot be queried.
6. To retrieve the data, Splunk uses a manual process to move the data from frozen to thawed and, when necessary, back again to frozen.

Other Splunk actions include the following:

- A query comes in from the search head asking the indexers to retrieve and process data.
- An indexer pulls data from buckets in the hot, warm, cold, and thawed tiers to resolve the query.

Designing and sizing a cluster is complex and depends on the following factors:

- Data digestion rate
- Data retention rate
- Number of concurrent queries
- High availability and disaster recover requirements
- Query response requirements

These and many other factors must be considered for any cluster design and architecture.

Data is moved from one tier to the next. To guarantee that you can still write data to a tier, the write performance of each tier must match the overall write performance of the previous tier.

If the hot data is written at 500 MBps, the system must be able to write cold data at 500 MBps or the hot/warm tier will fill up. In this example, the frozen tier must also be written at 500 MBps or the cold tier will fill up.

The hot/warm tier will have more reads than the cold tier and the frozen tier is not read in normal processing. When taking this into account, the combined read/write performance of a tier can be lower than the previous tiers.

See Splunk Validated Architectures, Splunk Enterprise Capacity Planning Manual, and Splunk Distributed Deployment Manual for details on designing a cluster and different cluster layouts.

Note: This reference architecture does not cover Splunk SmartStore, running Splunk in Containers, or any external storage.

Key solution components

These are the key components of this solution.

Hardware components

The following table describes standard configurations for indexers. See your Hitachi Vantara sales representative for a complete list.

Hardware	Details
DS120 G2 tiered	▪ 2 × Intel 6342 (24C, 2.7G, 220W) ▪ 2 × PSU 1600w AC Platinum x2 ▪ 8 × 32 GB 3200 RDIMM ▪ 2 × 256 GB m2 drives ▪ 1 × QS-3916 RAID 16i ▪ 1 × Intel VROC Standard License ▪ 1 × Mellanox CX-6 Lx EN Dual Port 25 GbE LP QSFP2 ▪ 2 × SSD 3DWPD 1.92 TB ▪ 10 × SFF SAS HDD 10K RPM 2.4 TB
DS120 G2 Flash	▪ 2 × Intel 6342 (24C, 2.7G, 220W) ▪ 2 × PSU 1600w AC Platinum x2 ▪ 8 × 32 GB 3200 RDIMM ▪ 2 × 256 GB m2 drives ▪ 1 × Intel VROC Premium license ▪ 1 × Mellanox CX-6 Lx EN Dual Port 25 GbE LP QSFP2 ▪ 12 × NVMe 3DWPD 2.0 TB
DS220 Tiered	▪ 2 × Intel 6342 (24C, 2.7G ,220W) ▪ 2 × PSU 1600w AC Platinum x2 ▪ 8 × 32 GB 3200 RDIMM ▪ 2 × 256 GB m2 drives ▪ 1 × SAS3916 4G RAID Mezzanine ▪ 1 × Intel VROC Premium license ▪ 1 × Mellanox CX-6 Lx EN Dual Port 25 GbE LP QSFP2

Hardware	Details
	▪ 2 × SSD 3DWPD 1.92 TB ▪ 20 × SFF SAS HDD 10K RPM 2.4 TB
DS220 Flash	▪ 2 × Intel 6342 (24C, 2.7G, 220W) ▪ 2 × PSU 1600w AC Platinum x2 ▪ 8 × 32 GB 3200 RDIMM ▪ 2 × 256 GB m2 drives ▪ 1 × Mellanox CX-6 Lx EN Dual Port 25 GbE LP QSFP2 ▪ 1 × Intel VROC Premium License ▪ 24 × NVMe 3DWPD 2.0 TB
Cisco Nexus 93180YC-FX switch	▪ 2 × ToR switches for data network per rack
Cisco Nexus 3332 switch	▪ 2 × Aggregate data networks as needed
Cisco Nexus 92348 switch	▪ 1 × Management network switch per rack
Power Supply Units	▪ 6 units per rack determined by region
Rack	▪ 1 × rack

Software components

The following table lists the solution software components.

Software	Version
Red Hat Enterprise Linux (RHEL)	8.2
SUSE Linux Enterprise Server (SLES)	15 SP2
Splunk Enterprise	8.2.1

Solution design

This section describes the design used for this solution.

Storage architecture

Splunk storage configuration depends on the individual deployment. Some of the considerations are as follows:

- Data size of each temperature tier
- Performance requirements for each temperature tier
- Retention for each temperature tier
- Storage devices used for each temperature tier
- The Splunk component deployed on the node

The following table lists the recommended storage types for different components.

Usage	Storage Type	Description
Search Head	SSD or HDD	SSDs are recommended. The storage should support at least 800 sustained IOPS and at least 300 GB of dedicated storage.
Indexer Hot/ Warm Tier	SSD or NVMe	This is the primary storage area in standard deployments. For availability purposes this storage should be RAID. Hot and warm tiers share the same storage area as RAID devices. SSD or NVMe drives are recommended.
Indexer: Cold Tiers	SSD/HDD, SAN, NAS, network file systems	This data is not used as often and can have lower performance requirements. This allows more storage options.
Indexer: Frozen storage	SAN, NAS, network file systems, HDD, archival devices	Frozen data is archived from the system. The default action for frozen data is to be deleted.
Indexer: Thawed Storage	SAN, NAS, network file systems, HDD	Thawed data storage requirements are similar to the cold data requirements, except that thawed data storage is usually short-lived and then deleted.
Forwarders	Any	Depending on your deployment forwarders can be deployed on existing devices with no extra storage requirements. Forwarders can also have their own devices and store data.

Application architecture

Sizing a Splunk solution is deployment specific. See Splunk's capacity planning documentation for information on sizing. The following table shows a starting point for machine sizing.

Component	Configuration Type	Sizing
Indexer	Minimum configuration	▪ 12 physical CPU cores, or 24 vCPUs at 2GHz ▪ 12 GB RAM ▪ 10/25 Gb NIC
	Mid-range configuration	▪ 24 physical CPU cores, or 48 vCPUs at 2GHz ▪ 64 GB RAM ▪ 10/25 Gb NIC
	High performance	▪ 48 physical CPU cores, or 96 vCPU at 2GHz ▪ 128 GB RAM ▪ 10/25 Gb NIC
Search Head	Any	▪ 16 physical CPU cores, or 32 vCPUs at 2GHz or greater speed per core ▪ 12 GB RAM ▪ 10/25 Gb NIC
Forwarders	Any	Depending on your deployment forwarders can be deployed on existing devices with no extra storage requirements. Forwarders can also have their own devices and store data.

The actual number of machines depends on the following factors:

- Data indexing volume
- Number of searches
- Total data size
- Replication factor
- Deployment architecture
- Number of sites

The following table provides a reference starting point based on the number of users and the volume of new data.

Number of Users	Indexing Volume < 1 TB per day	Indexing Volume 1 TB to 2 TB per day	Indexing Volume 2 TB to 3 TB per day
16	▪ 2 search heads ▪ 4 indexers	▪ 2 search heads ▪ 10 indexers	▪ 2 search heads ▪ 15 indexers
24	▪ 2 search heads ▪ 6 indexers	▪ 2 search heads ▪ 12 indexers	▪ 3 search heads ▪ 18 indexers
48	▪ 2 search heads ▪ 7 indexers	▪ 3 search heads ▪ 14 indexers	▪ 3 search heads ▪ 21 indexers

Network architecture

This solution uses two networks, a data network, and an out-of-band management network. The data network is uplinked to the rest of the client network. Because forwarders are deployed near the site where the data is generated, they use the client network to connect back to the indexers. The following figure illustrates this configuration.

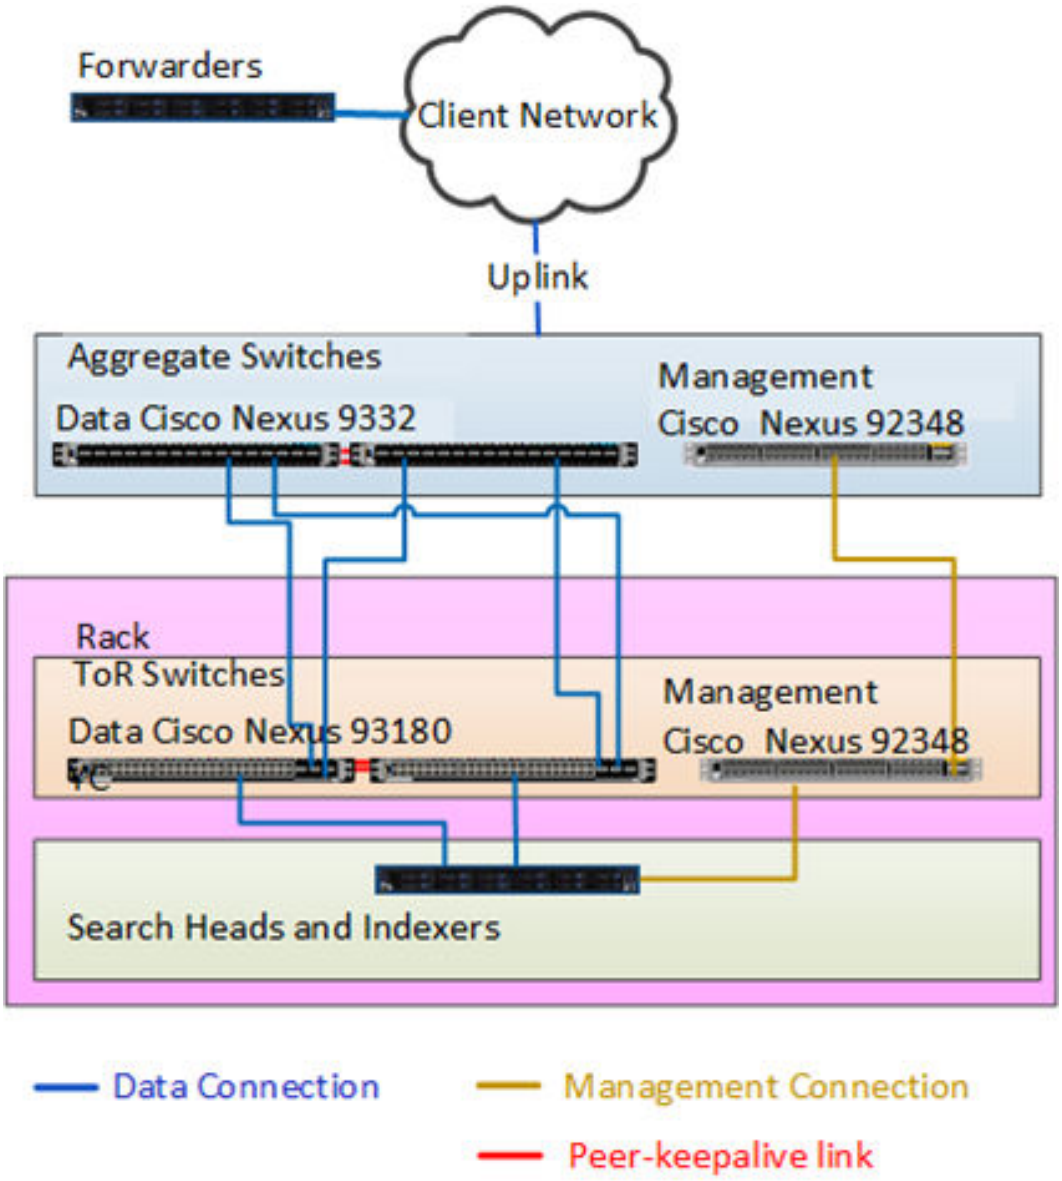

Engineering validation

Test hardware

This test was performed using one search head, three indexers, and multiple machines for the test driver. Its purpose is to validate the hardware, not the performance of the system. The configuration is listed in the following table.

Use	Machine Type	Configuration
Search head	1 × DS220 G2	▪ 2 × m2 NVMe drives RAID 1 for boot ▪ 6 × SSD for storage ▪ 521 GB memory ▪ 1 × Dual Port 25 GB NIC
Indexer	3 × DS220 G2	▪ 2 × m2 NVMe RAID 1 drives for boot ▪ 512 GB memory ▪ 1 × Dual Port 25 GB NIC Storage option 1 ▪ 4 × NVMe drives, RAID 0 Storage option 2 ▪ 6 × SSD, RAID 0 Storage option 3 ▪ 6 × SAS HDD, RAID 0
Test drivers	▪ 1 × test driver node ▪ 4 physical forward nodes ▪ 3 virtual machines	Multiple configurations using both physical machines and virtual machines

Test methodology

The purpose of this test was to generate a load for a small Splunk deployment. Different storage architectures were used to validate the design feasibility. Because large computers were used for the test, they were often idle during the test with many of their resources under-utilized.

The test harness used the following configuration:

- 3 × nodes used as indexers.
- 1 × node used as a search head.
- 1 × node used as a test controller node.
- 7 × nodes used as data generators/forwarders.
  - Each node had 50 forwarders running on each node.
  - 350 total forwarders were used.
- Data generation.
  - Each test generated data for 20 minutes.
  - No data met the requirements to move from hot/warm tier to cold storage.
- A single type of storage was used for Splunk hot, warm, and cold storage.
  - Run 1 used NVMe storage.
  - Run 2 used SSD storage.
  - Run 3 used HDD storage.
  - For all runs, the storage performance was well beyond the Splunk recommended minimum of 1200 IOPS.
- Searches.
  - These tests ran with no searches being performed.
  - These tests were repeated with searches being performed for every 10,000 records generated.

The following figure illustrates the test configuration.

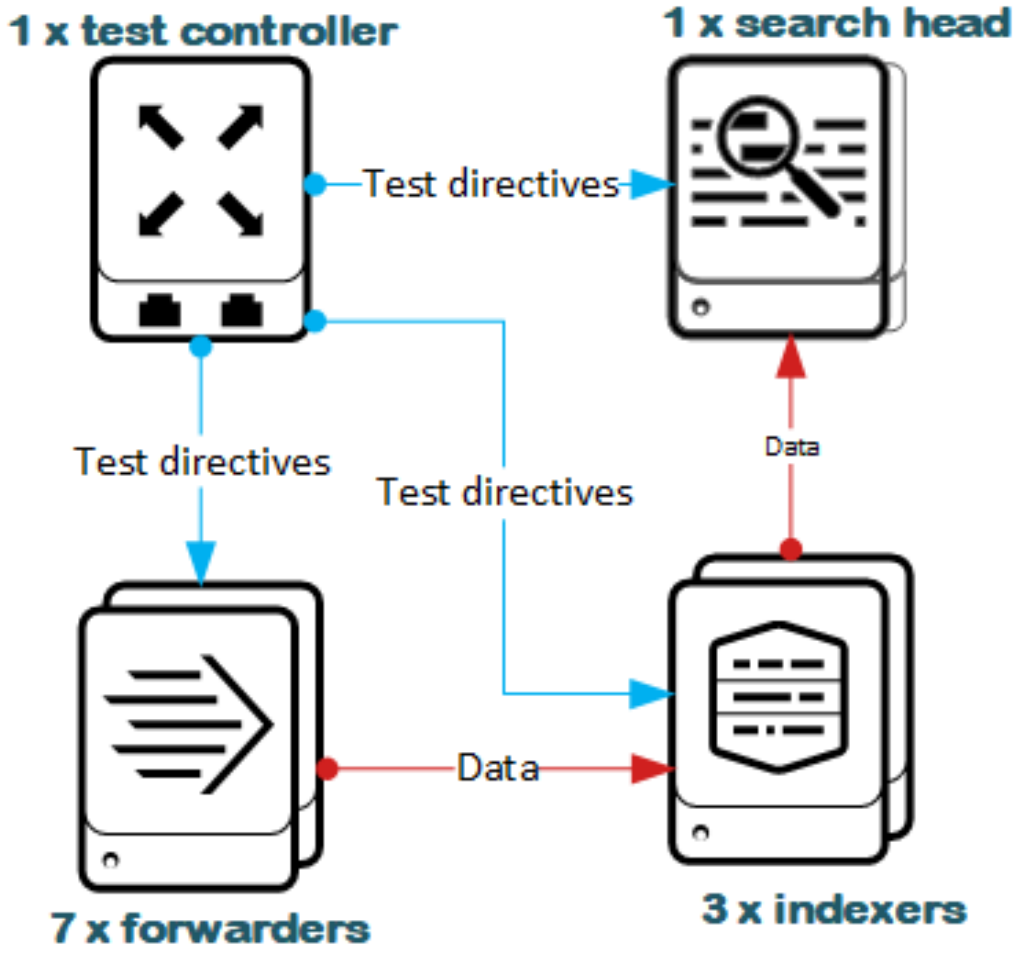

Test results

When results display stored data, an estimated value is displayed. This estimate is based upon extending the rate to show the sustained data stored and calculated based on uncompressed data.

Splunk is typically used for log or machine generated data that is in text format. The forwarders send the data in a compressed format and the indexer stores it in a compressed format. It is very common to see a 95% reduction in size.

The test results displayed are calculated by the test tool kit.

Note: These results are useful to show that the system is working; however, they do not provide performance information to compare the different configurations. Because of the size of the system being used, simplicity of the test, and the data set size, these results are skewed.

Results without performing searches

The following figures show the results of the test for all three storage types when searches were not performed.

The following figure shows the data indigested per day.

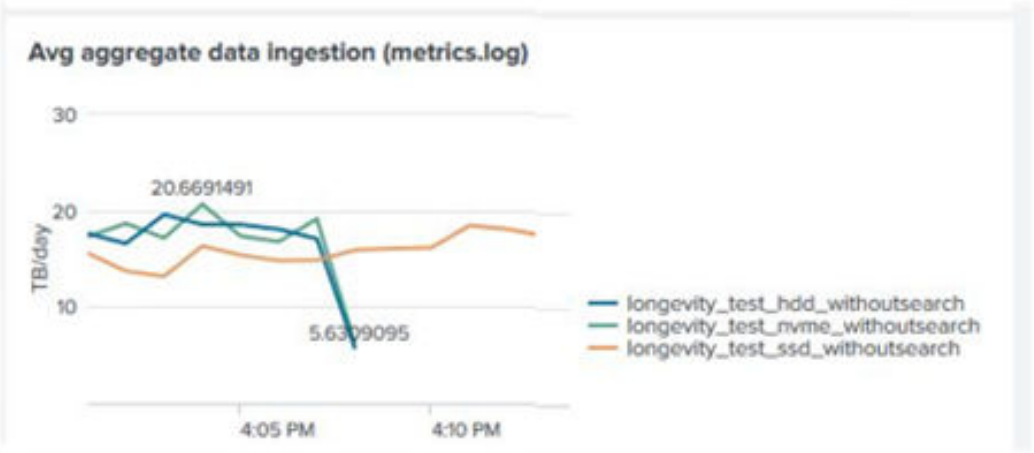

The following figure shows the average number of disk writes.

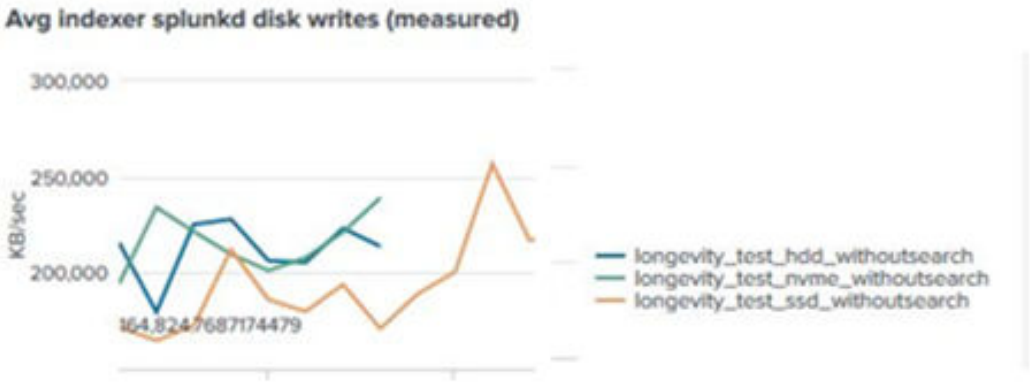

The following figure shows the average CPU usage. As you can see the system CPU usage is low.

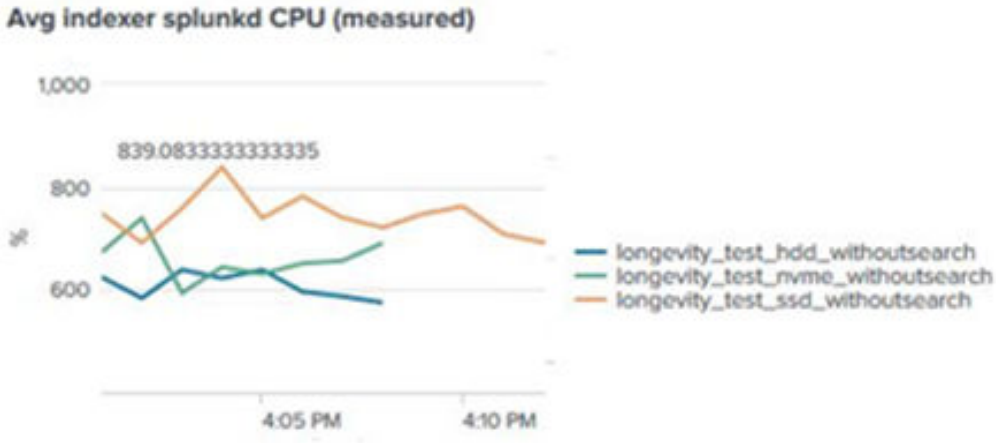

The following figure shows the average inbound network traffic.

Results when performing searches

The following figures show the results of the tests for all three storage types when searches were performed over 10,000 events. Again, these are not performance results. Instead they are to verify that the system works.

The following figure shows the data ingested per day.

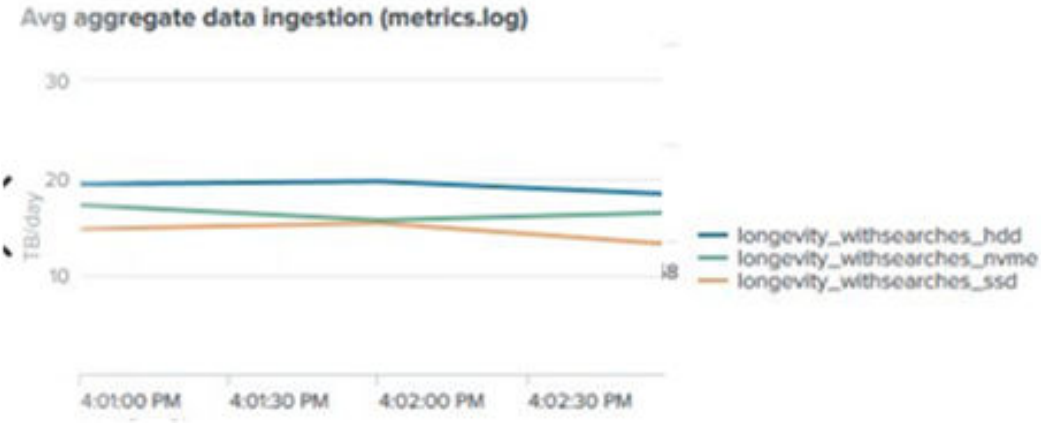

The following shows the average number of disk writes.

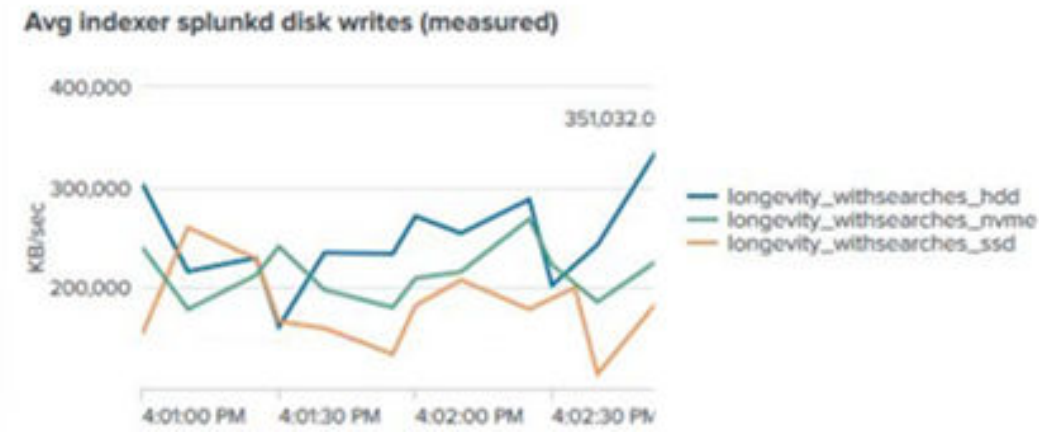

The following figure shows the average CPU usage. While the CPU usage is higher when performing searches, most of the cores are still idle.

The following figure shows the average inbound network traffic.

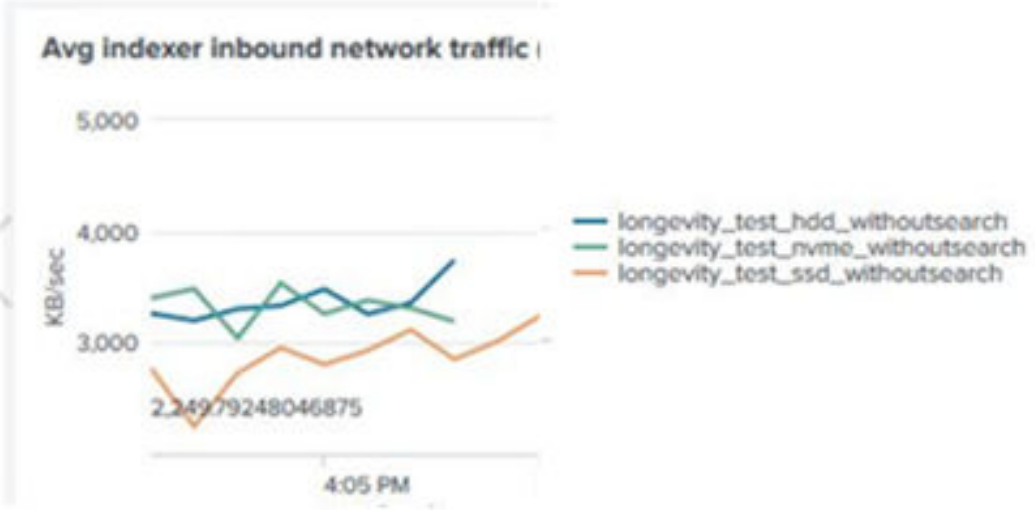

Product descriptions

This is information about the hardware and software components used in this solution for Hitachi Analytics Infrastructure for Splunk.

Splunk Enterprise

Splunk Enterprise is a data platform that allows you to investigate, monitor, analyze and act on your data with ease for enhanced security and operational efficiency.

Hitachi Advanced Server DS120 G2

With support for two Intel Xeon Scalable processors in just 1U of rack space, the Hitachi Advanced Server DS120 G2 delivers exceptional compute density. It provides flexible memory and storage options to meet the needs of converged and hyperconverged infrastructure solutions, as well as for dedicated application platforms such as internet of things (IoT) and data appliances.

The Intel Xeon Scalable processor family is optimized to address the growing demands on today's IT infrastructure. The server provides 32 slots for high-speed DDR4 memory, allowing up to 4 TB memory capacity with RDIMM population (128 GB × 32) or 8 TB (512 GB × 16) of Intel Optane Persistent Memory. DS120 G2 supports up to 12 hot-pluggable, front-side-accessible 2.5-inch non-volatile memory express (NVMe), serial-attached SCSI (SAS), serial-ATA (SATA) hard disk drive (HDD), or solid-state drives (SSD). The system also offers 2 onboard M.2 slots.

With these options, DS120 G2 can be flexibly configured to address both I/O performance and capacity requirements for a wide range of applications and solutions.

Hitachi Advanced Server DS220 G2

With a combination of two Intel Xeon Scalable processors and high storage capacity in a 2U rack-space package, Hitachi Advanced Server DS220 G2 delivers the storage and I/O to meet the needs of converged solutions and high-performance applications in the data center.

The Intel Xeon Scalable processor family is optimized to address the growing demands on today's IT infrastructure. The server provides 32 slots for high-speed DDR4 memory, allowing up to 4 TB memory capacity with RDIMM population (128 GB × 32) or 8TB (512 GB × 16) with Intel Optane Persistent Memory population.

DS220 G2 comes in three storage configurations to allow for end user flexibility. The first configuration supports 24 2.5-inch non-volatile memory express (NVMe) drives, the second supports 24 2.5-inch serial-attached SCSI (SAS), serial-ATA (SATA) and up to 8 NVMe drives, and the third supports 12 3.5-inch SAS or SATA and up to 8 NVMe drives. All the configurations support hot-pluggable, front-side-accessible drives as well as 2 optional 2.5-inch rear mounted drives. The DS220 G2 delivers high I/O performance and high capacity for demanding applications and solutions.

Cisco Nexus switches

The Cisco Nexus switch product line provides a series of solutions that make it easier to connect and manage disparate data center resources with software-defined networking (SDN). Leveraging the Cisco Unified Fabric, which unifies storage, data and networking (Ethernet/IP) services, the Nexus switches create an open, programmable network foundation built to support a virtualized data center environment.

Hitachi Unified Compute Platform Advisor

Hitachi Unified Compute Platform Advisor (UCP Advisor) is a comprehensive cloud infrastructure management and automation software that enables IT agility and simplifies day 0-N operations for edge, core, and cloud environments. The fourth-generation UCP Advisor accelerates application deployment and drastically simplifies converged and hyperconverged infrastructure deployment, configuration, life cycle management, and ongoing operations with advanced policy-based automation and orchestration for private and hybrid cloud environments.

The centralized management plane enables remote, federated management for the entire portfolio of converged, hyperconverged, and storage data center infrastructure solutions to improve operational efficiency and reduce management complexity. Its intelligent automation services accelerate infrastructure deployment and configuration, significantly minimizing deployment risk and reducing provisioning time and complexity, automating hundreds of mandatory tasks.

Red Hat Enterprise Linux

Using the stability and flexibility of Red Hat Enterprise Linux, reallocate your resources towards meeting the next challenges instead of maintaining the status quo. Deliver meaningful business results by providing exceptional reliability on military-grade security. Use Enterprise Linux to tailor your infrastructure as markets shift and technologies evolve.

SUSE Linux Enterprise High Availability Extension

Compete more effectively though improved uptime, better efficiency, and accelerated innovation using SUSE Linux Enterprise Server. This is a versatile server operating system for efficiently deploying highly available enterprise-class IT services in mixed IT environments with performance and reduced risk.

SUSE Linux Enterprise Server was the first Linux operating system certified for use with SAP HANA. It remains the operating system of choice for the vast majority of SAP HANA customers.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA
HitachiVantara.com | community.HitachiVantara.com

Contact Information
USA: 1-800-446-0744
Global: 1-858-547-4526
HitachiVantara.com/contact