HITACHI
Inspire the Next

WHITE PAPER

Optimize Splunk With Intelligent Data Tiering

Solution White Paper: Splunk SmartStore Single-Site With Hitachi Content Platform (HCP)

By Hitachi Vantara

August 2020

Contents

Executive Summary	3
Introduction	4
Solution Description	4
Test Methodology	5
Test Environment	5
Configuring Hitachi Content Platform for Splunk SmartStore	7
Test Cases	11
Test Results	11
Analysis	15
Performance Tuning and HCP Configuration Guidelines	15
Splunk Best Practices	16
Additional Resources	17

Executive Summary

Organizations are increasingly realizing the untapped potential in their business data. Machine data has especially become a key strategic asset, as it holds valuable information about customers, transactions, fraudulent activity and system behaviors. Larger datasets provide the foundation for effective analytics. However, as these datasets continue to grow, it becomes challenging for organizations to cost-effectively manage them. Splunk's traditional co-located compute and storage model was designed to help organizations quickly gain the insights that they seek from these massive datasets. However, the explosive data growth from the internet of things (IoT) has shifted the balance between compute and storage requirements. Traditional colocation models lack the flexibility to meet these new storage demands as compute must also be scaled to keep up with the ever-growing demands for storing machine data. Organizations need a more efficient and cost-effective solution to address these storage and analytics challenges for their machine and unstructured data.

Decouple Compute and Storage for Greater Scale and Efficiency

In today's era of digital transformation, organizations now recognize that decoupling compute and storage resources in their Splunk environment is the key. This approach allows them to reduce costs while gaining the ability to scale those resources independently for greater efficiency. Delivering performance at scale in this decoupled architecture requires on-demand tiering of data to the most appropriate storage resources to ensure that data continues to be utilized seamlessly and efficiently.

Cost-Effective, Scalable, Long-Term Data Retention for Machine Learning

Hitachi Content Platform (HCP), an on-premises cloud object storage solution, is a cost-effective, secure and robust storage platform that is purpose-built for scalability, compliance and long-term data retention.

With **Splunk SmartStore and HCP**, you gain greater flexibility with your compute and storage resources with incredible cost-efficiency and long-term data retention at scale. SmartStore enables you to reserve indexer local storage resources for hot and actively read data while tiering less frequently read datasets via S3 to HCP. This approach keeps data seamlessly searchable and retained longer at a much lower cost. As Splunk ingests more data, SmartStore storage can be scaled independently. As Splunk activity increases, performance can be scaled by adding more compute resources. This flexibility greatly simplifies Splunk management and greatly reduces operational costs.

Introduction

This document provides information on the interoperability between Splunk Enterprise's SmartStore feature and Hitachi Content Platform, in a single site configuration.

With Hitachi Content Platform (HCP) and Splunk SmartStore, customers can decouple compute and storage for greater scale and efficiency. HCP provides complete flexibility allowing you to start small and scale to more than 400PB+. With HCP in a hybrid configuration, organizations can get to unlimited cloud scale while maintaining full control, encryption protection, and data integrity.

With any cloud deployment, whether it's public, private or hybrid, security of data is paramount. HCP is an extremely secure and an efficient object storage solution with a comprehensive suite of data protection and security features. Information in a Hitachi Content Platform system can be written in a WORM format (Write Once, Read Many). HCP maintains the integrity and availability of stored object data by "fingerprinting" upon ingest and then self-repairing any defects. With WORM mode activated, data written to the repository cannot be modified, and thus is immune to ransomware attacks. HCP also provides flexible retention capabilities to prevent the accidental or malicious deletion of objects in the repository.

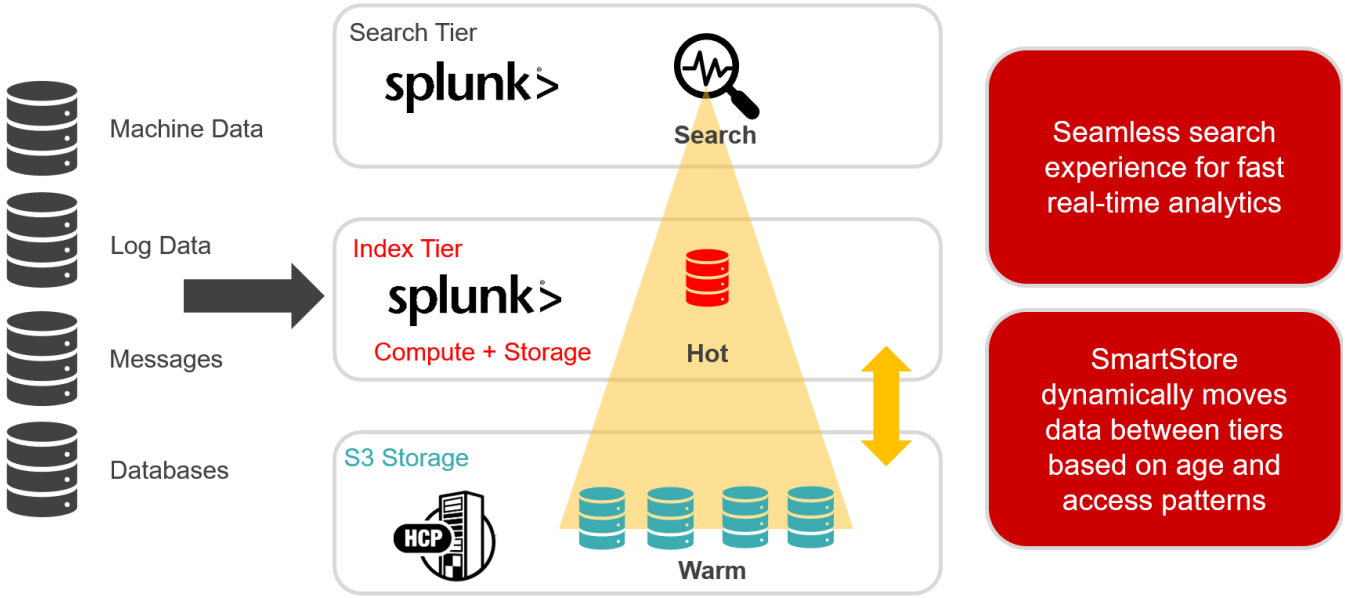

Solution Description

Splunk Enterprise is the leading enterprise data management, search, analysis, and visualization software. The traditional Splunk Enterprise indexer cluster deployment model (also known as the distributed scale-out model) provides high data availability and fidelity but presents significant cost challenges as data volumes grow. Splunk SmartStore provides efficient compute and storage elasticity along with simpler and more flexible data management infrastructure, which can help lower costs. The new model relies on an object data store located on remote storage. Because data availability and fidelity for less frequently read data sets are trusted to remote storage, the choice of platform for the remote object store is critical. Hitachi Content Platform (HCP), an on-premises cloud object storage solution, is a cost-effective, secure and robust storage platform that is purpose-built for scalability, compliance and long-term data retention. Splunk SmartStore on HCP accelerates access to critical, real-time data while reducing overhead costs, increasing availability, and improving operational efficiencies.

Key benefits of this solution are:

- **Performance:** SmartStore brings data closer to compute on demand, for fast performance and a seamless search experience. The distributed design of HCP is beneficial when retrieving data from remote store since all nodes participate, thus *maximizing throughput and overall search performance*.
- **Data Integrity:** The HCP design distributes cryptographic hashing computations made by HCP to *guarantee integrity of stored data*. HCP is constantly rechecking these hashes and self-repairs any defects that might develop over time. Data moved to HCP does not require any other form of data protection.
- **Reduce Storage Cost:** Reduced costs with relatively smaller indexer local storage footprint for hot and actively read data and cost-efficient HCP storage for less frequently read warm/cold data. HCP delivers over 60% lower total cost of ownership (TCO) versus public cloud alone.
- **Scalable:** HCP provides scalable, economical and highly available storage for Splunk. Scale SmartStore storage independently of compute while maintaining fast performance.
- **Availability:** When setting up Splunk SmartStore with HCP, you'll have several choices for backend data storage that include local disks (RAID6), Fibre Channel arrays (RAID6), or trays of commodity disks protected with EC (erasure codes). Through erasure coding, HCP protects data with storage durability mathematically computed to *exceed 15 nines of availability*. This is accomplished utilizing large-spindle, extremely cost-effective drives – which means lower $/GB for you. In all cases, HCP increases warm buckets availability by leveraging HCP built-in replication to one or more remote sites. In the event of a disaster, your DR HCP can serve up data in milliseconds.
- **Simplified Management:** As much of the data no longer resides on the Splunk nodes, you can eliminate much of the Splunk cluster management overhead related to data rebalancing and index maintenance. A node failure requires only metadata replication versus full data replication.
- **Security:** HCP also offers compliance options that make it impossible to accidentally overwrite or delete data.

Test Methodology

The goal of this SmartStore Single-Site testing was to validate functional and performance characteristics of the Splunk Enterprise v7.3 with Hitachi Content Platform V8.2 using Hitachi API for Amazon S3. The test suite was provided by Splunk and all results were submitted to Splunk for review and approval.

Test Environment

Hardware Components

This describes the configuration to be tested, following Splunk Enterprise 7.3 system requirements and HCP V8.2 system requirements.

- (1) Splunk Enterprise Cluster Master Server
  - Two Intel Xeon E5 2650 v3 2.3 GHz, 10 Core, 128GB RAM per processor (256GB total), 2 × 10 GbE NIC, 17.9 TB SSD
- (7) Splunk Enterprise Indexer Peer Servers
  - Two Intel Xeon E5 2698 v4 2.2 GHz, 20 Core, 128GB RAM per processor (256GB total), 2 × 10 GbE NIC, 17.9 TB SSD
- (1) Hitachi Content Platform Cluster
  - (8) G10-nodes with (2) S11 S-nodes (100TB minimum application usable storage)
  - SSD Acceleration with 10GbE end-end.
  - HCP Software V8.2

- (2) Brocade VDX 6740 Switches
  - 24-port, 10 GbE front-end network
  - 10GbE to be configured from Splunk Indexers to G-nodes, backend network between G-nodes and to S-nodes
- (1) Client
  - Intel Xeon E5-2670 @ 2.60 GHz, 192 GB RAM, 1 TB hard drive, 1 × 10 GbE NIC

Software Components

These software products were required to test the Splunk Enterprise v7.3 solution with HCP.

- Hitachi Content Platform V8.2+
- Splunk Enterprise Server 7.3
- Splunk Enterprise Console 7.3
- CentOS Linux Operating System 7.4.1708
- Microsoft Windows Server Standard 2012 R2

The figure below illustrates the test environment configuration.

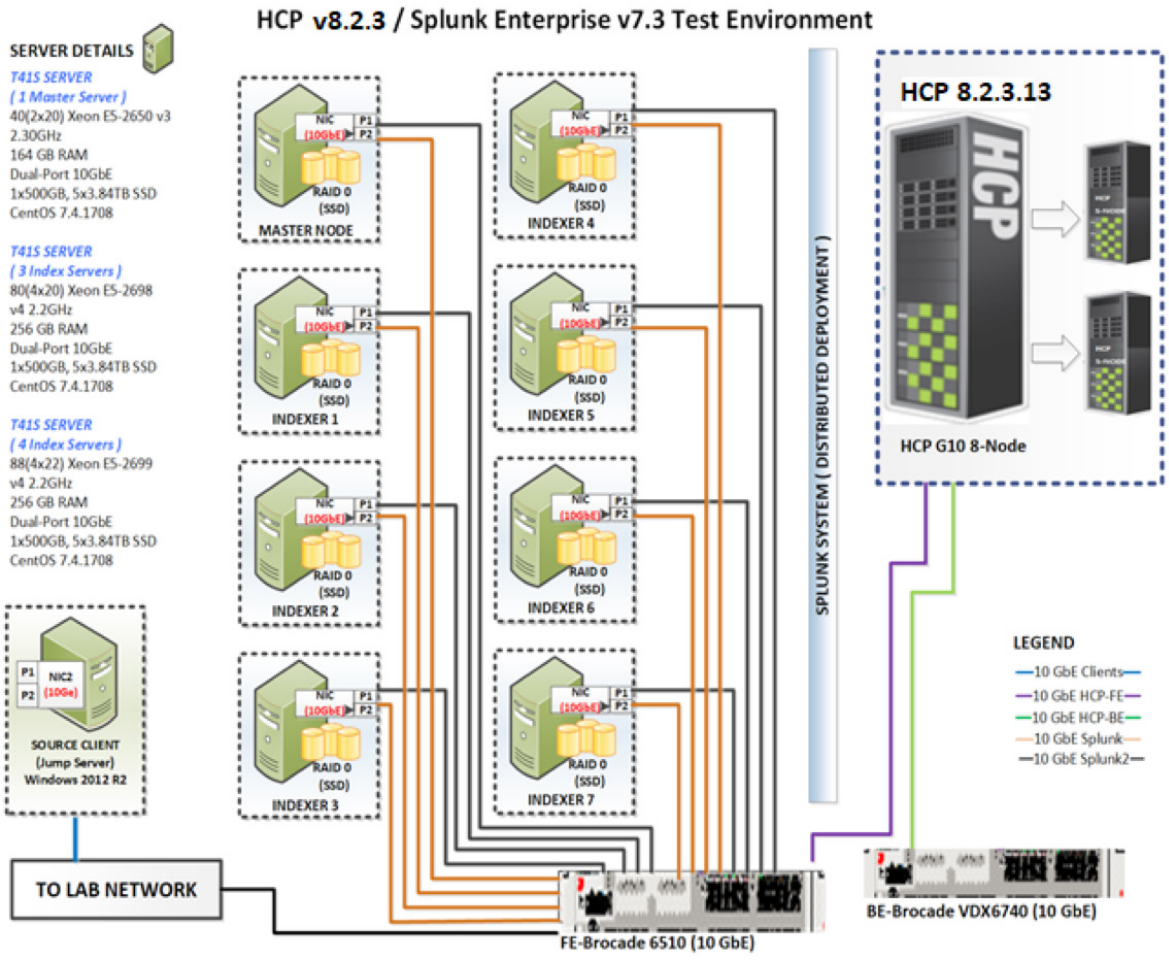

Configuring Hitachi Content Platform for Splunk SmartStore

HCP Tenant Setup

For this performance test, the tenant created for the Splunk SmartStore environment contained 1 namespace (Splunk). The namespace had versioning enabled, indexing enabled, and metadata XML verification enabled. In Protocols sections, HTTP, HTTPS and Hitachi API for Amazon S3 protocols should be enabled. Optimization for cloud only protocols should be set. We created a user for the Splunk to access the HCP tenant. The user was "Splunk" and this user had all permissions on all the namespaces.

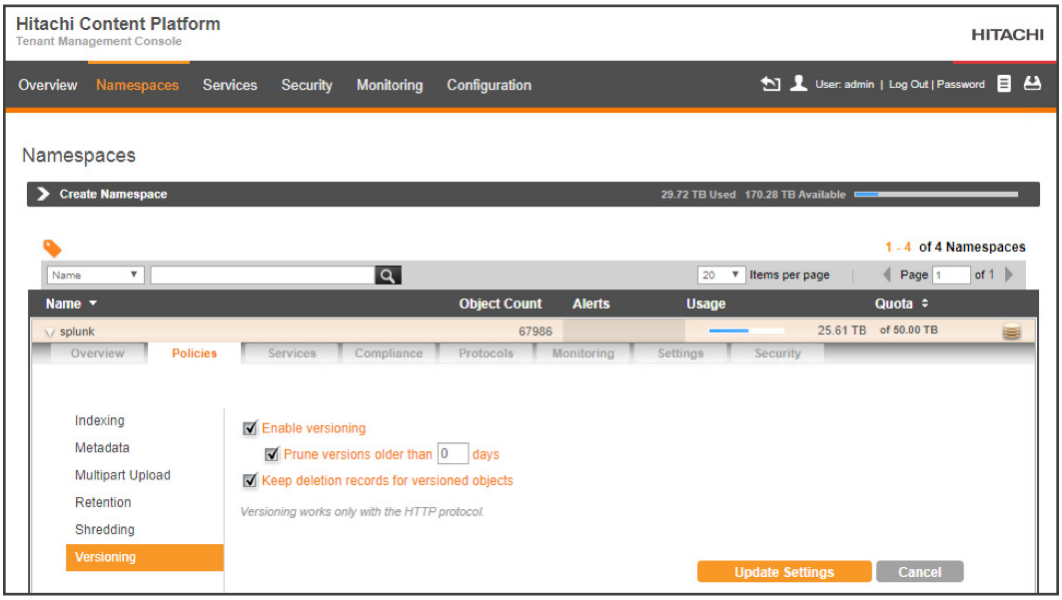

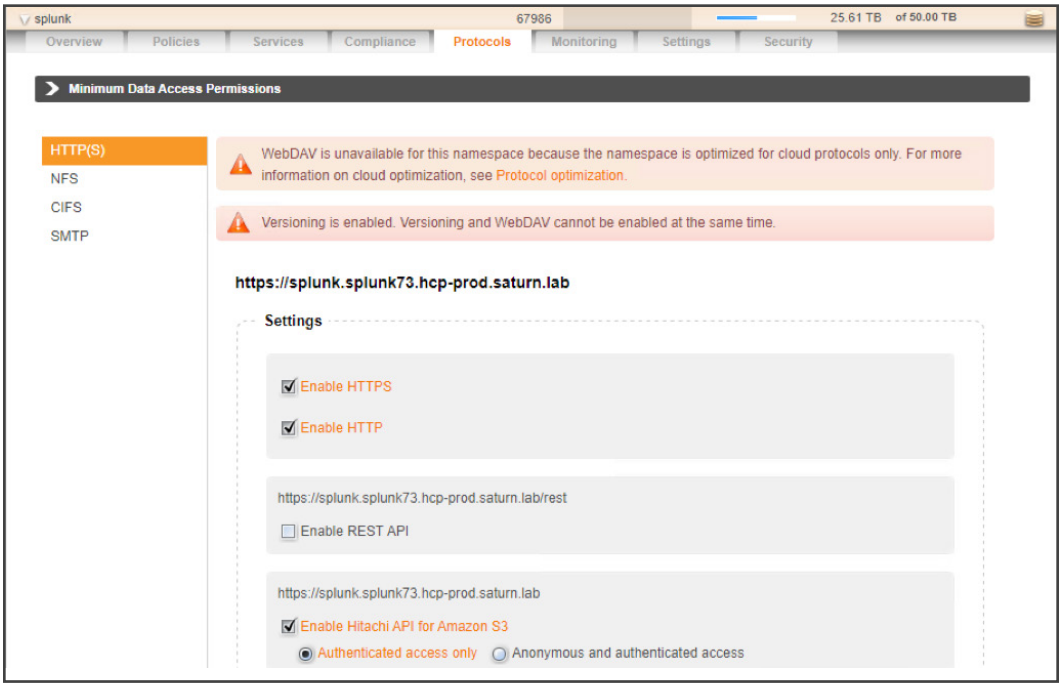

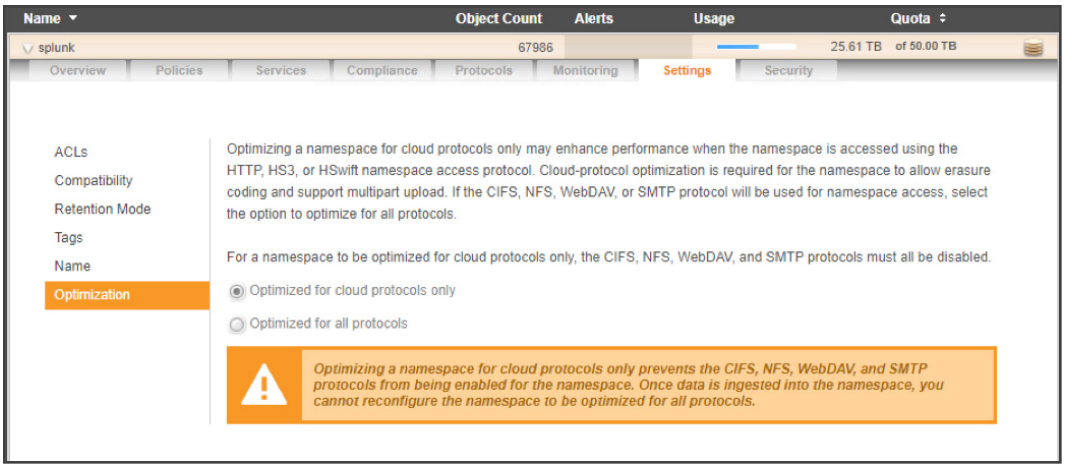

Splunk Configuration

7 Indexers, 1 Search Head, 1 Cluster Master

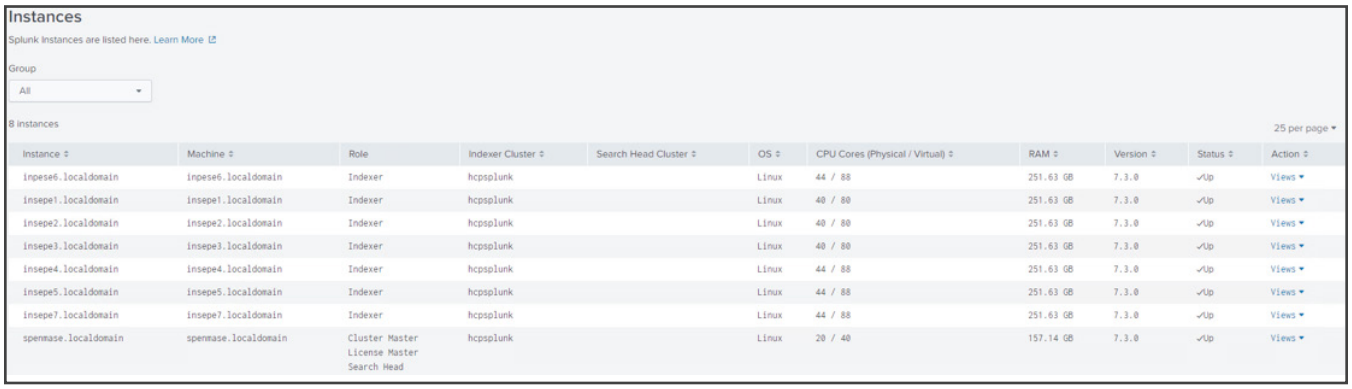

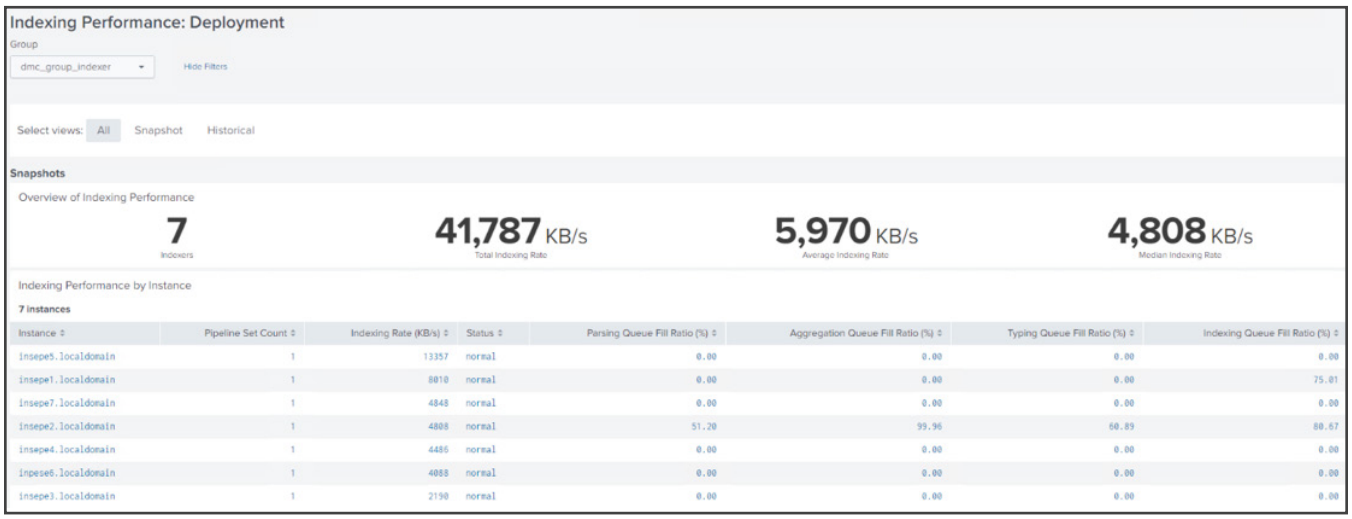

Configuring SmartStore Indexes With HCP

The SmartStore settings in indexes.conf enable and control SmartStore indexes. You can enable SmartStore for all of indexes, or you can enable it on an index-by-index basis, allowing a mix of SmartStore and non-SmartStore indexes on the same indexer.

Note: When you configure these settings on an indexer cluster's peer nodes, you must deploy the settings through the configuration bundle method. As with all settings in indexes.conf, SmartStore settings must be the same across all peer nodes.

This example configures SmartStore for an indexer cluster.

On the master node, cd to /$SPLUNK_HOME/etc/master/_cluster/local and create a file named indexes.conf.

The below SmartStore index example will store the Splunk index warm buckets (hcp_test, hcp_test2, etc.) on HCP.

```
# Configure the remote volume
# On the next line, the volume's path setting points to the remote storage location
# where indexes reside. Each SmartStore index resides directly below the location
# specified by the path setting. The <scheme> identifies a supported remote
# storage system type, such as S3. The <remote-location-specifier> is a
# string specific to the remote storage system that specifies the location
# of the indexes inside the remote system.
# This is an S3 example: "path = s3://mybucket/some/path".

[volume:splunk73]
storageType = remote
path = s3://splunk
remote.s3.access_key = YWRtaW4=
remote.s3.secret_key = 47b7bfb65fa83ac9a71dcb0f6296bb6e
remote.s3.endpoint = http://splunk.hcp-prod.saturn.lab

[hcp_test]
remotePath = volume:splunk73/$_index_name
homePath = $SPLUNK_DB/hcp_test/db
# SmartStore-enabled indexes do not use thawedPath or coldPath, but you must still specify them here.
coldPath = $SPLUNK_DB/hcp_test/colddb
thawedPath = $SPLUNK_DB/hcp_test/thaweddb
maxHotBuckets = 1000
maxDataSize = auto_high_volume
repFactor = auto
```

[hcp_test2]

remotePath = volume:splunk73/$_index_name

homePath = $SPLUNK_DB/hcp_test2/db

SmartStore-enabled indexes do not use thawedPath or coldPath, but you must still specify them here.

coldPath = $SPLUNK_DB/hcp_test2/colddb

thawedPath = $SPLUNK_DB/hcp_test2/thaweddb

maxHotBuckets = 1000

maxDataSize = auto_high_volume

repFactor = auto

[hcp_test3]

remotePath = volume:splunk73/$_index_name

homePath = $SPLUNK_DB/hcp_test3/db

SmartStore-enabled indexes do not use thawedPath or coldPath, but you must still specify them here.

coldPath = $SPLUNK_DB/hcp_test3/colddb

thawedPath = $SPLUNK_DB/hcp_test3/thaweddb

maxHotBuckets = 1000

maxDataSize = auto_high_volume

repFactor = auto

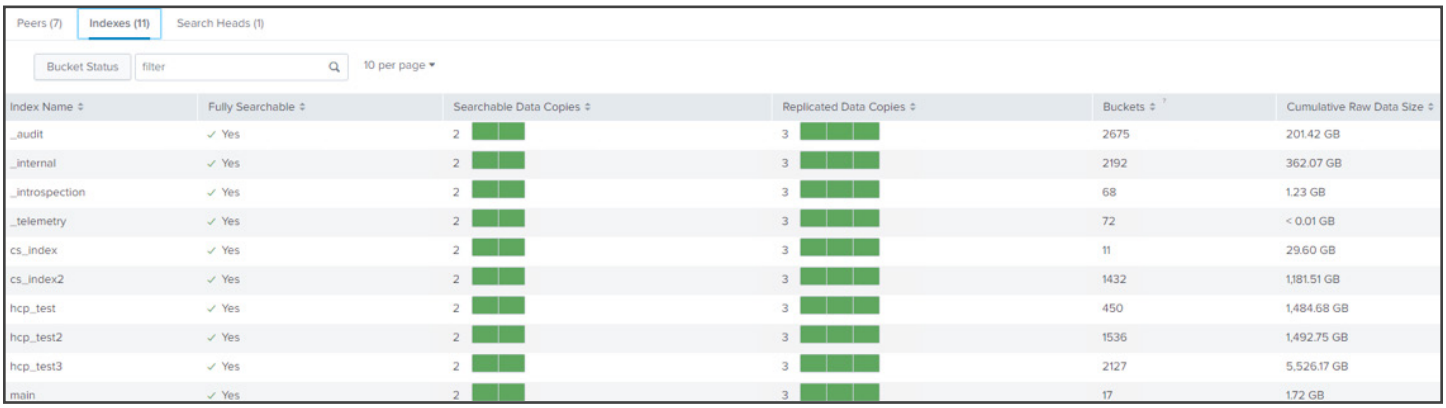

Index Name	Fully Searchable	Searchable Data Copies	Replicated Data Copies	Buckets	Cumulative Raw Data Size
_audit	Yes	2	3	2675	201.42 GB
_internal	Yes	2	3	2192	362.07 GB
_introspection	Yes	2	3	68	1.23 GB
_telemetry	Yes	2	3	72	< 0.01 GB
cs_index	Yes	2	3	11	29.60 GB
cs_index2	Yes	2	3	1432	1,181.51 GB
hcp_test	Yes	2	3	450	1,484.68 GB
hcp_test2	Yes	2	3	1536	1,492.75 GB
hcp_test3	Yes	2	3	2127	5,526.17 GB
main	Yes	2	3	17	1.72 GB

Verify Connectivity From SmartStore to HCP

SmartStore contains a CLI utility which can be used to verify connectivity to the remote store. The syntax for running the command is:

./splunk cmd splunkd rfs -- ls --starts-with volume:splunk73

The output of this command will list the contents of the remote store.

Test Cases

There were (4) specific test cases defined in the test plan as follow:

1. Upload Performance (10TB)
2. Download Performance (10TB)
3. Search Performance (10TB)
4. Scale Testing (Repeat 1,2 & 3 at 20-30TB)

Test Results

1. Upload Performance (10TB)

a. Configure the SmartStore cache to include at least 10 days' worth of data (may vary based on ingest volume per day)

b. Upload Performance

- Ingest Data (1TB/day or higher)
- Rollover all hot buckets to warm at the same time
- Measure upload performance through the SmartStore monitoring console dashboards

Upload Performance (MB/s)		
Indexer	10TB	20TB
1	127.7	238
2	127.6	275.8
3	127.8	258.1
4	131.3	219.9
5	117.1	215.33
6	133.8	194.9
7	131.7	182

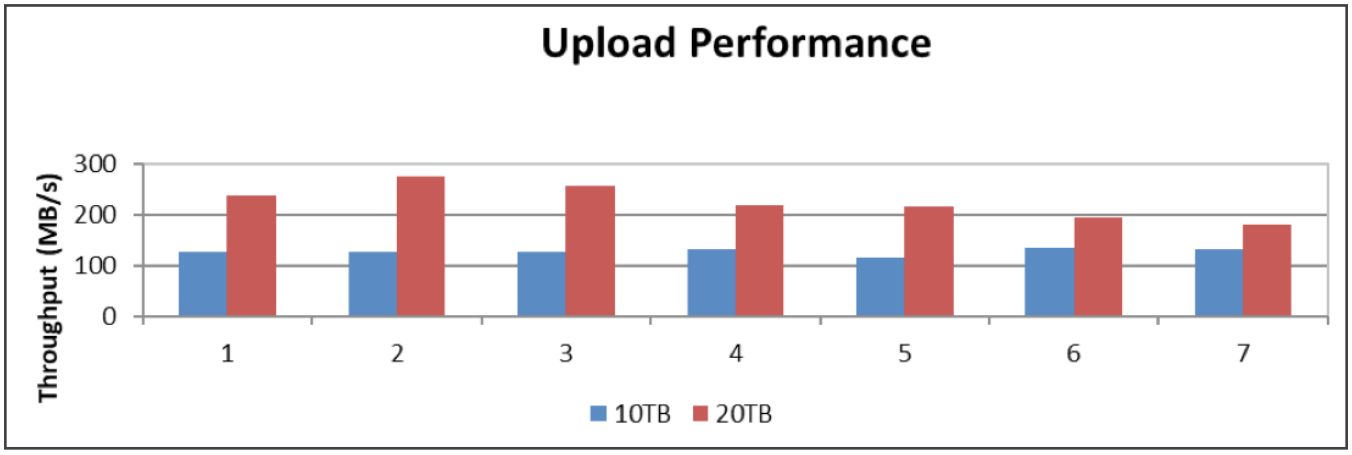

2. Download Performance (10TB)

a. Evict all data from the cache. Invoke the Cache Manager _evict call with a very high amount of free space, e.g., 1000000000 splunk _internal call /services/admin/cacheman/_evict -post:mb 1000000000 -post:path $SPLUNK_DB -method POST

b. Execute search that will download buckets from remote store (HCP)

c. Measure download performance through SmartStore monitoring console dashboards

Download Performance (MB/s)		
Indexer	10TB	20TB
1	192	278.1
2	275.1	261.4
3	383.1	273.1
4	764	260.4
5	601.2	239.4
6	269.4	245.5
7	200.5	204.6

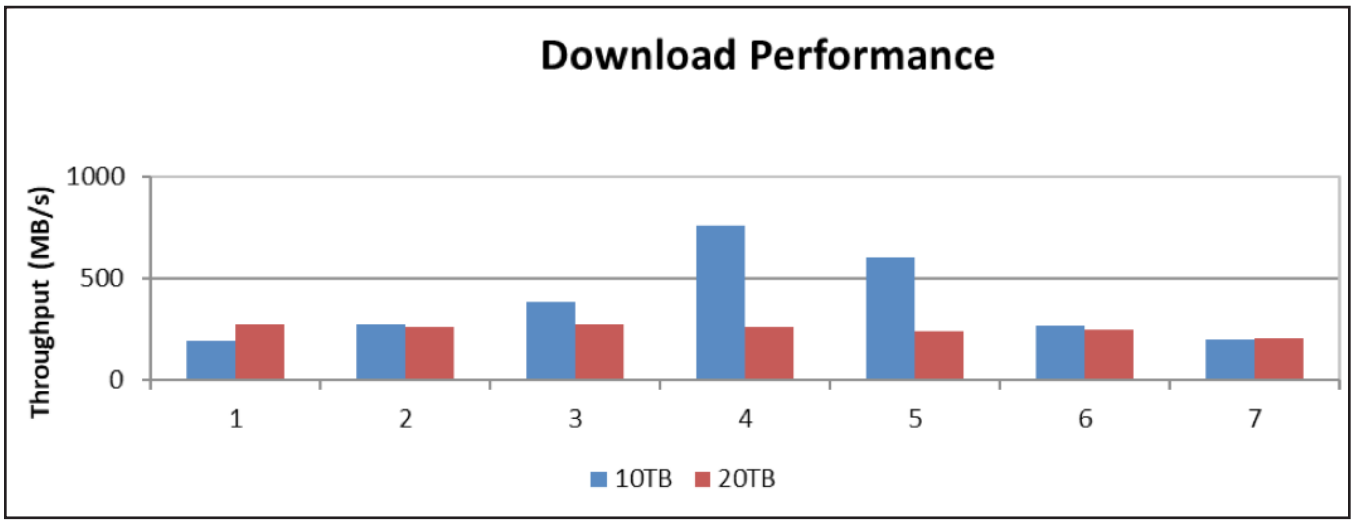

3. Search Performance (10TB)

a. Compare search performance when data is in the cache relative to data not in the cache

- Measure search latency using MC (MC->Search->Scheduler Activity->Deployment)
- Measure “Cache Hits Misses” using MC (MC->Indexing->SmartStore->SmartStore Activity Deployment)
- Compare “Remote Storage Search Overhead” using MC (MC->Indexing->SmartStore->SmartStore Activity Deployment)

b. Compare search performance when 25%, 50%, 75% in the cache

- Evict all data from the cache (see REST call below)
- Prepare a search that will look back in time-period with varying level of data in cache
- Measure search latency using MC (MC->Search->Scheduler Activity->Deployment).
- Measure “Cache Hits Misses” using MC (MC->Indexing->SmartStore->SmartStore Activity Deployment)

Search Performance (Run Time)		
Cache %	10TB	30TB
25	468.5	232.4
50	56.3	44.3
75	6.3	46.3

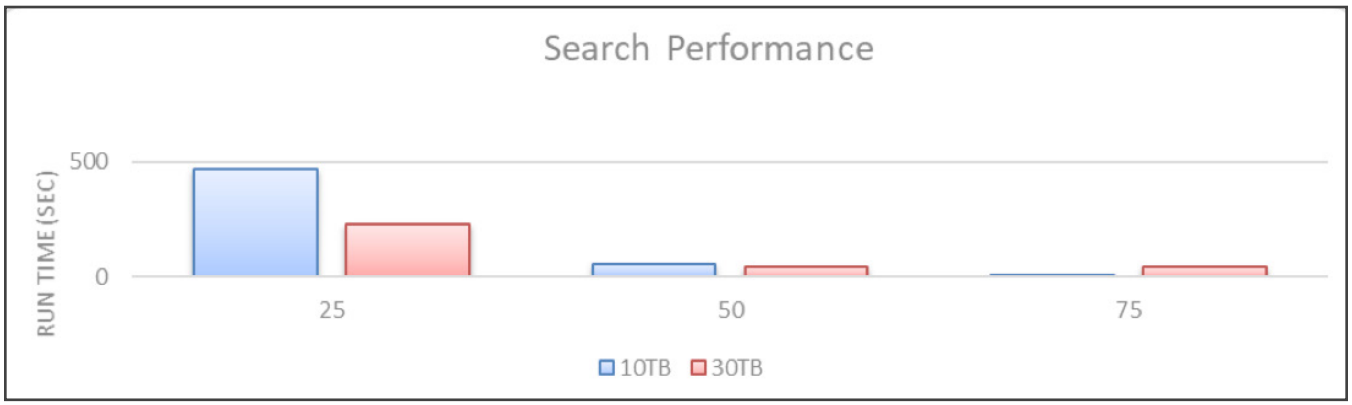

4. Scale Testing

a. Configure the SmartStore cache to include at least 10 days of data (20+TB)

b. Scale ingest volume to 2TB/day and track. Capture the below metrics from MC

- Upload and Download throughput
- Cache hits and misses before and after scaling ingest volume

Upload Performance (20TB)

Upload Performance (MB/s)		
Indexer	10TB	20TB
1	127.7	238
2	127.6	275.8
3	127.8	258.1
4	131.3	219.9
5	117.1	215.33
6	133.8	194.9
7	131.7	182

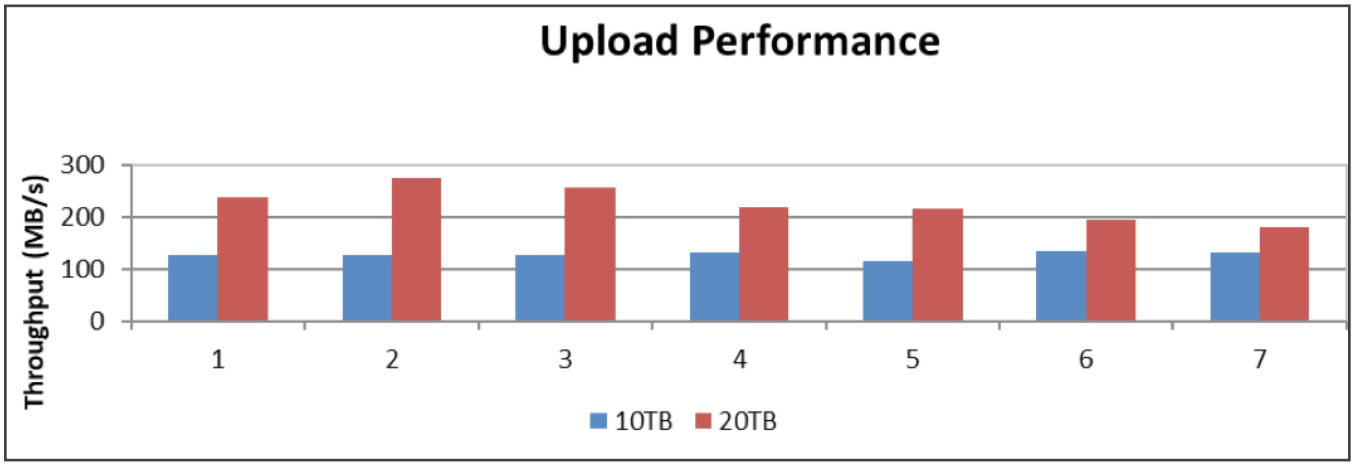

Download Performance (20TB)

Upload Performance (MB/s)		
Indexer	10TB	20TB
1	192	278.1
2	275.1	261.4
3	383.1	273.1
4	764	260.4
5	601.2	239.4
6	269.4	245.5
7	200.5	204.6

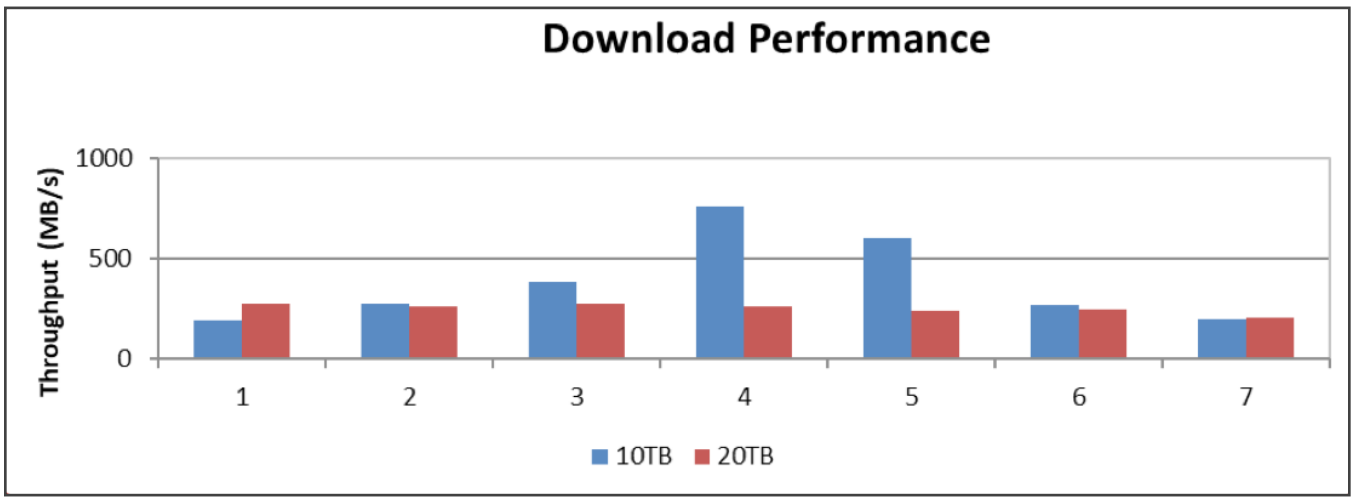

Search Performance (30TB)

Search Performance (Run Time)		
Cache %	10TB	30TB
25	468.5	232.4
50	56.3	44.3
75	6.3	46.3

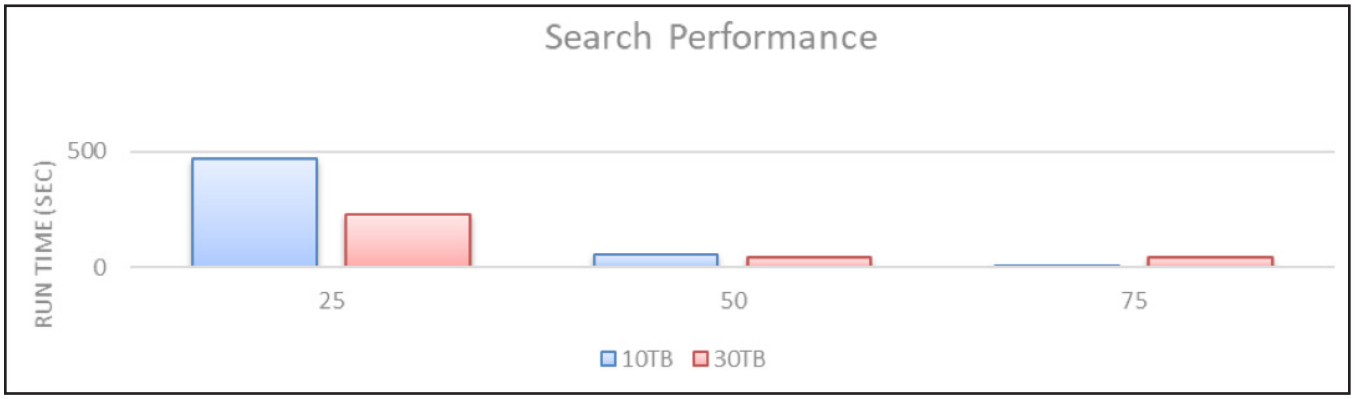

Analysis

This section includes the observations made during testing.

Observations

All the functional tests were successfully executed and met or exceeded our target throughput objectives. Upload and download throughput is better with larger datasets because Splunk has more data to be uploaded to HCP; provided proper network bandwidth is configured, Splunk SmartStore will maximize use of system resources to accommodate larger bucket rolls.

Performance Tuning and HCP Configuration Guidelines

- Multipart Upload
  - Splunk writes larger files on HCP via S3 multipart upload, which is very efficient. The default part size is 128MB. This is specified in the remote.s3.multipart_upload.part_size variable and can be changed via a configuration file. We used the default value in our testing.
- Multipart Download
  - SmartStore also utilized S3 ranged reads and the default value is 128MB. This is specified in the remote.s3.multipart_download.part_size variable and can also be changed via a configuration file. We used the default value in our testing.
- Simultaneous Connections
  - SmartStore by default utilizes eight simultaneous connections to HCP when performing Multipart Uploads to HCP as well as when performing ranged reads. This is specified in the remote.s3.multipart_max_connections variable and can be changed via a configuration file. We used the default value in our testing.
- Increase the Ulimit Values
  - Splunk Enterprise allocates system-wide resources like file descriptors and user processes on *nix systems for monitoring, forwarding, deploying, and searching. The ulimit command controls access to these resources which must be tuned to acceptable levels for Splunk Enterprise to perform adequately on *nix systems.

  - The more tasks your Splunk Enterprise instance performs, the more resources it needs. You should increase the ulimit values if you start to see your instance run into problems with low resource limits. For purposes of removing any potential performance bottlenecks in our testing, we increased the ulimit to 128,000 (Hard Limit).

- Network Topology and HCP Connectivity
  - Splunk SmartStore will consume as much network bandwidth as available from each indexer to the remote store HCP. As such, the number of 10GbE connections is a critical design point for each customer deployment. We configured multiple 10GbE paths to each HCP G-node to maximize performance. Assessing total bucket roll capacity and performance to meet specific service level requirements is necessary. Please utilize HCP performance configuration guidelines when configuring your SmartStore deployment and engage Hitachi Vantara technical experts to ensure success
  - As a basic design principal, the remote storage should accommodate rolling of warm buckets at a rate that is faster than peak ingestion rates into Hot bucket tier.

The following are key question to be answered, leading up to a proper system design:

1. How many Splunk clusters and individual applications will be supported. (one or more)
2. Describe major issues that you are trying to resolve. For example, search entire Splunk data set, without manually intervention for frozen data un-thaw while meeting retention requirements.
3. What is the current and future envisioned architecture?
   - Clustered or Non-Clustered; Single site or multi-site; Replication Factors (Local / Remote); Indexer Data Distribution (geographic, balanced, etc.); Bucket Roll Criteria (defaults or custom (time, size, global, by index)) and Archive / Freeze / Thaw Criteria.
4. Identify current and future Indexer and Search Head Configuration. Number of Indexers & Search Heads, Server Sizing (Physical or Virtual), network topology & connectivity between indexers and remote store to meet performance SLAs.
5. Workloads & Service levels. Please assess the time period for ingest and search activity, whether 7x24, 5x9 or otherwise. Daily Data Ingestion Rate and Ingest Throughput Averages and Peaks should be understood.
6. Search Profile. A minimum of 30 days of search activity should be assessed for proper sizing of Hot tier, cache residency time and corresponding warm bucket roll criteria. Please utilize Splunk search analysis and sizing tools as needed.
7. Retention Requirements. This will impact overall system design, in particular whether deletion, freezing, or archiving of specific data sets will be necessary. SmartStore in conjunction with HCP should be designed to meet company compliance and data retention policies while optimizing access searchable datasets.

Splunk Best Practices

1. For SmartStore use with Splunk Enterprise Security, confirm that you have enough local storage available to accommodate 90 days of indexed data, instead of the 30 days otherwise recommended. Review the local storage requirements section in the Splunk documentation for guidance.
2. Configure the data retention settings to ensure that the indexer follows your desired freezing behavior. This step is extremely important, to avoid unwanted bucket freezing and possible data loss. SmartStore bucket-freezing behavior and settings are different from the non-SmartStore behavior and settings.
3. The home path and cold path of each index must point to the same partition. Certain other indexes.conf settings are restricted with SmartStore. Review Settings in indexes.conf that are incompatible with SmartStore or restricted.
4. A SmartStore-enabled index cannot be converted to non-SmartStore. The coldPath setting for each SmartStore index requires a value, even though the setting is ignored except in the case of migrated indexes. The

thawedPath setting for each SmartStore index requires a value, even though the setting has no practical purpose because you cannot thaw data to a SmartStore index.

5. All index-related settings, including SmartStore-settings, must be configured identically across the peer nodes. You must enable SmartStore for the same set of indexes on all peer nodes. When you add a new index stanza, you must set repFactor to “auto”.
6. Each remote volume is limited to a single indexer cluster or standalone indexer. That is, each remote store holds buckets for only a single cluster or standalone indexer. From the configuration standpoint, the path setting within each remote volume stanza in indexes.conf must be unique to the cluster or indexer.
7. Each SmartStore index is limited to a single remote volume. All warm buckets for that index must reside in the same remote store.
8. All peer nodes on an indexer cluster must use the same SmartStore settings.

Additional Resources

Check out the resources below to learn how Hitachi Vantara and Splunk work together to enable organizations to optimize their Splunk environments for reduced storage costs and greater efficiency.

Learn how Hitachi Content Platform makes your data securely available anywhere, anytime.

Read about the power of the Hitachi Content Platform portfolio.

Learn how to use Splunk to gain insights from your machine data.

We Are Hitachi Vantara

We guide our customers from what’s now to what’s next by solving their digital challenges. Working alongside each customer, we apply our unmatched industrial and digital capabilities to their data and applications to benefit both business and society.

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

HITACHI is a trademark or registered trademark of Hitachi, Ltd. Microsoft and Windows Server are trademarks or registered trademarks of Microsoft Corporation. All other trademarks, service marks, and company names are properties of their respective owners.

WP-600-A BTD August 2020