FUNCTION1

Red Hat Storage Server as a Hybrid Storage Solution for Splunk Enterprise.

November 10, 2014

A publication by

Anshu Rastogi

Kevin Chu

Sandeep Khaneja

www.function1.com

Executive Summary.

Large-scale data analysis has seen tremendous growth in the last few years. As the data volumes and use cases expand, the tools and techniques to conduct these analyses also grow and improve. The identification of key trends to support these analyses is no longer measured through weeks and months, but instead over years and even decades. Terabytes and petabytes are now the industry standard for quantifying data volumes. Organizations retaining such high-levels of data now need an affordable, scalable, and flexible enterprise storage solution to manage this explosion of big, unorganized, and unstructured data growth.

Red Hat Storage Server is an open, software-defined storage platform that provides cost-effective storage for extremely large, historical data sets for enterprise class Splunk deployments. Red Hat maintains storage performance, capacity, and availability to meet the most demanding enterprise storage requirements. Red Hat Storage introduces a new Hybrid Storage model into the ecosystem, suggesting that customers leverage Red Hat Storage for their "cold" Splunk data, while leaving "hot/warm" data on Direct Attached Storage. This paper discusses the testing and validation of Red Hat Storage Server with different Splunk enterprise architectures.

Red Hat Storage Server is an open, software-defined storage platform that provides cost-effective storage for extremely large, historical data sets for enterprise class splunk> *deployments.*

Background.

Before discussing the testing and validation of the Red Hat Storage Server solution, it's important to understand the architectures for both a distributed Splunk deployment and a Red Hat Storage Server deployment.

Splunk Technical Architecture Overview.

An enterprise-level Splunk deployment typically consists of many nodes. A basic distributed deployment would consist of one search head and multiple indexers. A search head, a Splunk instance that a user interacts with to search and report on data, issues search queries and receives data back from the indexers in the deployment. The indexers are logically grouped together and have data sets stored across them.

splunk> Scales Across the Datacenter.

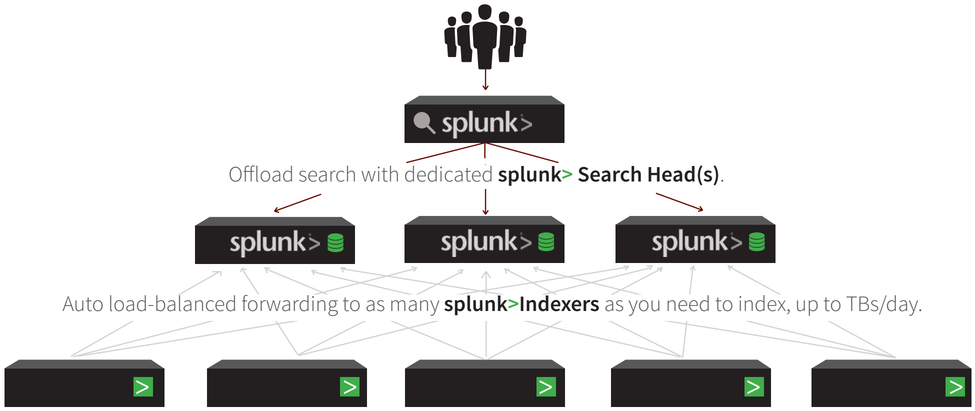

Offload search with dedicated **splunk> Search Head(s)**.

Auto load-balanced forwarding to as many **splunk>Indexers** as you need to index, up to TBs/day.

Send data from servers using combination of **splunk> Forwarders,** syslog, WMI, message queues, or other remote protocols.

Splunk indexes are analogous to databases and contain raw event data and supporting files. An index lives across multiple Splunk indexers and store data in units referred to as "buckets." New data received by Splunk indexers is stored in "hot" buckets. These are the only buckets that are actively writeable in a Splunk index. After the size of this hot bucket reaches a certain threshold, or another condition is met, the hot bucket "rolls" to "warm." After warm buckets are created, or other conditions are met, a warm bucket rolls to a "cold" bucket. Hot/warm and cold buckets are stored different file paths on separate storage tiers. The design is advantageous as it allows a search and reporting application the ability to store large amounts of legacy data in a more cost-effective manner, without sacrificing retention.

Red Hat Storage Server Architecture Overview.

A Red Hat Storage Server deployment consists of multiple nodes that are grouped together in "clusters." Members of the cluster are referred to as "peers." Each peer node, typically, provides at least one local file system that will be used to provide storage capacity to clients in the form of a 'brick'. Multiple bricks are then combined across the peers to form a 'volume'. Volumes are presented to clients over multiple protocols (as shown in the diagram below) enabling solution flexibility.

Other scale out file systems use meta data to track data placement across disparate systems, but this can lead to performance bottlenecks. Red Hat Storage avoids this issue by using an algorithm for data placement and retrieval across the cluster. This algorithmic approach to scale out storage improves performance, introduces linear scalability and improves reliability.

Red Hat Storage Server offers many advanced features such as local and remote data replication, automated self-heal and volume level snapshots.

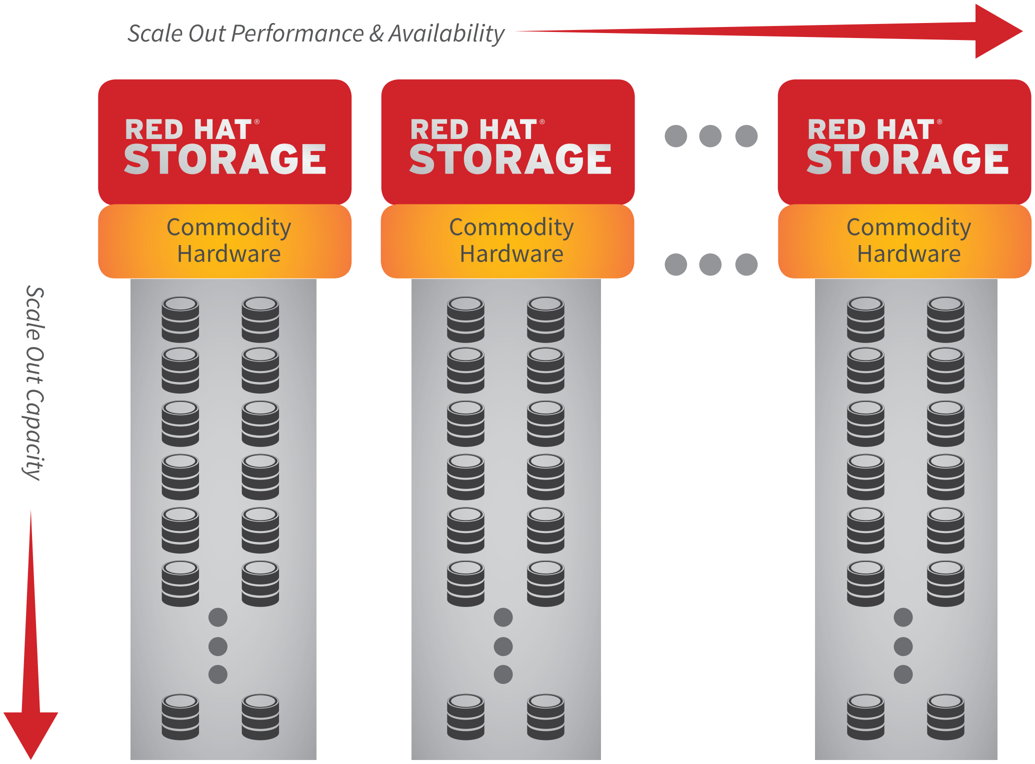

The Problem.

Organizations that store extremely large amounts of data in Splunk have to consider the cost and manageability of storage, which is not trivial when considering extremely large data sets.

The Solution.

Red Hat Storage Server aims at meeting this storage demand by providing a storage solution for "cold" data in Splunk, and unlike other shared storage solutions, leaving hot/warm data on Splunk indexers for optimal performance. Cold data in Splunk tends to be less actively searched and much larger in volume compared to hot and warm data. Due to these factors and the ability for Splunk indexers to store this data on separate storage tiers, Red Hat Storage Server is a cost-effective storage solution.

Because of the distributed nature of both Splunk Enterprise and Red Hat Storage Server deployments, it was important to test different scenarios to ensure that Red Hat Storage Server adequately met the needs of storing Splunk data. Below is a description of that testing and validation.

Testing.

Overview.

Testing, conducted in a lab environment, simulated various Splunk Enterprise architectures and operational scenarios, such as a Splunk indexer or a Red Hat Storage Server node becoming unavailable. During the trials, load was placed on the system using a Synthetic Benchmark Kit (SBK), a performance-benchmarking tool made by Splunk to generate both high data volume for indexing and concurrent searching.

The diagram below depicts, at a high level, the integration tested between Splunk and Red Hat Storage.

splunk> Single Site with Index Replication and RHS Cold Storage

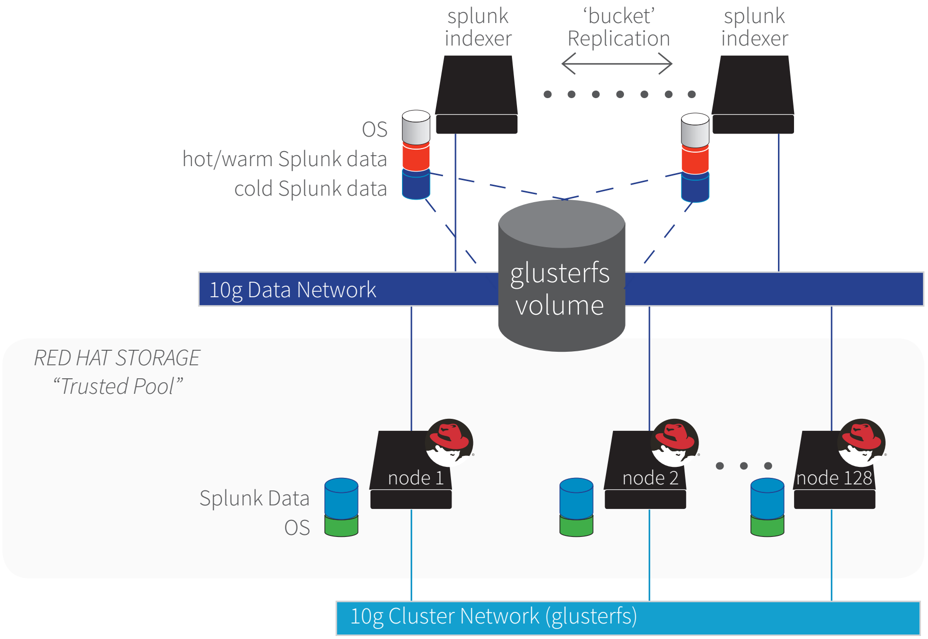

Disk Types

⬜ RAID-1 OS
🟥 RAID-10 15K
🟦 RHS Volume (NFS or glusterfs)
🟦 RAID-6 7k
🟩 Integrated SD Card

High Level Design Principles

>Data protection provided by Splunk and RHS

> Dedicated hot/warm and cold storage tiers enabling elastic cold storage and easier management

> Data movement to cold Storage managed by Red Hat, reducing the impact of Splunk's replication

> NFS or glusterfs connectivity provide transparent failover

> Server configuration flexibility Indexers: 2.5" 10k or 15k HDDs RHS: 3.5" 4TB or 6TB HDD

Scenario 1: Searching Availability during System Failure.

The test environment for this scenario consisted of a distributed Splunk deployment and a Red Storage Server cluster. The Splunk deployment consisted of multiple Splunk indexers searched by a Splunk search head. A replicated Red Hat Storage volume was presented to each of the Splunk indexer nodes. The cold paths of the test Splunk indexes pointed to the Red Hat Storage volume.

In order to test this design, four Red Hat Storage Server nodes were configured as storage peers. The replicated nodes were each configured with four bricks in a 2 x 2 fashion, to establish a distributed-replicated volume. The advantage of this storage configuration allows for any single node of Red Hat Storage Server to be taken down for maintenance, while still allowing Splunk to copy data buckets to its cold storage volume. The cold paths of two test Splunk indexes per indexer, test_static, and test_streaming were mounted to the Red Hat Storage Server nodes using NFSv3 hard mounts. The test_static index functioned as a search benchmark testbed, whilst the test_streaming index focused on data ingest, warm to cold migration and search load – effectively simulating as far as possible a typical live workload.

This test proves the ability to continue uninterrupted searching and indexing from Splunk during the shutdown or failure of any Splunk indexer or Red Hat Storage Server brick. In addition, the Red Hat Storage Server replicated volume proved the capability of online volume expansion, demonstrating that the Splunk indexers, without interruption, automatically utilized the increase of Red Hat Storage Server storage capacity.

.Scenario 2: Uninterrupted Scaling of Splunk Indexers.

The nature of Splunk's distributed architecture allows for the addition of any number of Splunk indexers. This unique trait ensures that as data volume grows, horizontal scaling can meet data volume demand with minimal disruption to a Splunk ecosystem. As achieved in a test environment, two Splunk indexers can be used to ensure that removing an indexer for maintenance purposes will not affect the overall system.

During the trials, a pair of indexers was configured to receive a streaming indexing load from the SBK. The indexers were taken down to simulate a maintenance or unexpected restart of the Splunk service. Throughout the shutdowns, all data streams automatically shifted to the other indexer, guaranteeing a zero loss of data. Once all indexers returned to operational status, testing through the search head confirmed that no data had been lost during the test window.

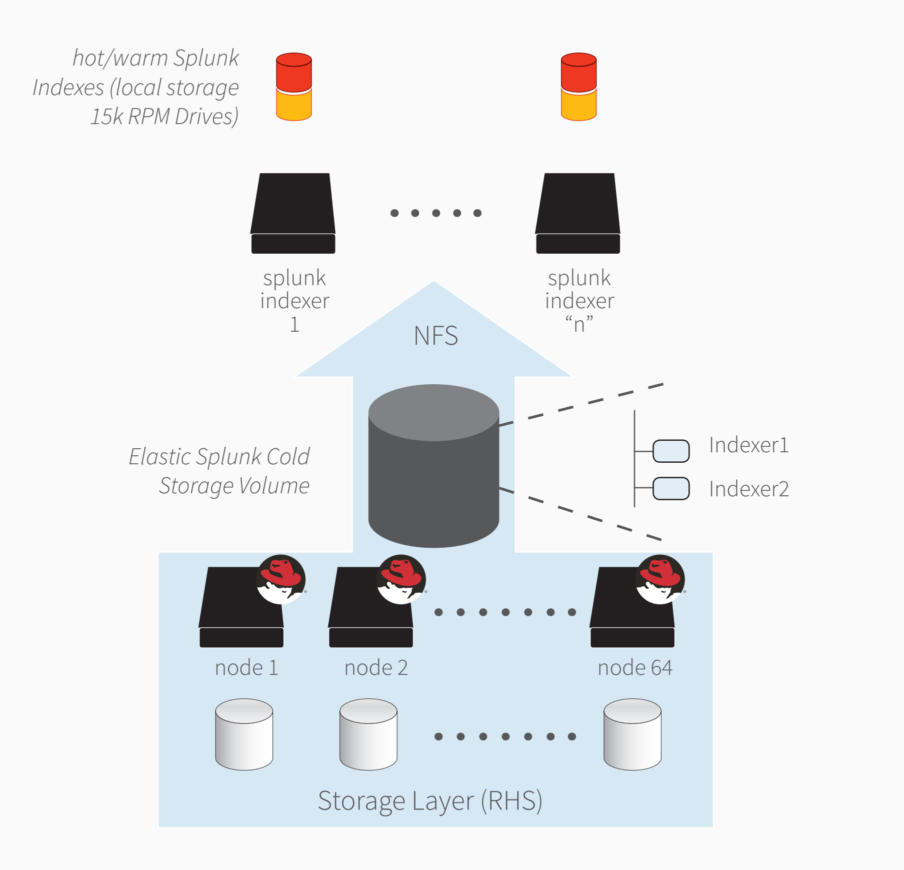

Scenario 3: Configuring Storage to be Non-Disruptive to Application Workload.

During testing, one of the four Red Hat Storage Server nodes was shut down while the Splunk indexers received a live streaming workload. Tracking the creation of ten buckets per indexer, and their copy over by Splunk's bucket copying process, proved that all data successfully migrated to the cold storage volume. When all nodes were restored, Red Hat Storage Server identified over 400 files created during the maintenance, and automatically replicated each file through Red Hat Storage Server's self heal process. The estimated time to complete the self-heal process was roughly half of the time taken to index the original test-streaming data. This indicates that Red Hat Storage Server is able to fully replicate following prolonged outages. During both a simulated outage and the subsequent self-heal process, Splunk logs reported no errors or warnings, and produced full search results without disruption.

Scenario 4: Splunk Bucket Management.

Splunk indexes are contained in units that it identifies as a bucket. Buckets of data are created and filled in a "hot/warm" database path within each index directory. As a preconfigured age is reached, they have the capability of being automatically migrated to a cold path for long-term storage.

The advantage of this design allows a search and reporting application the ability to store large amounts of older data in a more cost-effective manner, without sacrificing retention. In the event of a full failure of the entire cold storage mount, Splunk indexers are able to preserve multiple warm buckets until cold storage is restored, with zero data loss.

Splunk's bucket-mover process was verified, and resulted in the reset of the process with an automatic reattempt on a per minute basis. This was verified by causing an in-flight interruption. In our test environment, indexers were able to hold and manage upwards of dozens of buckets scheduled for cold migration without performance impact. Once access to the cold storage mount was restored all buckets were migrated to Red Hat Storage Server seamlessly.

Scenario 5: Achieving Non-Disruptive Storage Expansion.

By leveraging Red Hat Storage Server's software defined designs and replicated volumes, the addition of new bricks can be accomplished without impact to the application or its utilization of storage. During testing, the SBK was executed to simulate load, and was able to remain fully functional and unaware of any Red Hat Storage Server administrative changes. Additional bricks were added to the volume, doubling the available capacity to Splunk, and a rebalance performed to redistribute existing data across the new capacity. As all bricks were returned to fully operational status, the testing process verified that any Splunk buckets and files created during the configuration were automatically replicated to the resized bricks. In addition, testing was able to verify that Splunk encountered zero data loss, both during and after, maintenance, and post-expansion file replication.

Scenario 6: Warm and Cold Data Searching Availability throughout Different Stages of Red Hat Storage Server.

During testing, different stages of the Red Hat Storage Server were manipulated to reflect possible changes in the environment, and subsequently to record the availability of warm and cold data throughout the testing. These stages include, but are not limited to: a healthy environment, various nodes and indexers being down, moving buckets from warm to cold storage, manipulated cold storage space, and manipulated indexer peer nodes.

The following tests were performed:

- > Warm data is searchable when in a healthy environment;
- > Cold data is searchable when in a healthy environment;
- > Warm data is searchable when a clustered Indexer down;
- > Cold data is searchable when a clustered Indexer is down;
- > Warm data is searchable when a Red Hat Storage Server node is down;
- > Cold data is searchable when a Red Hat Storage Server node is down;
- > Warm data is searchable when both a clustered indexer and a Red Hat Storage Server node are down;
- > Cold data is searchable when both a clustered indexer and a Red Hat Storage Server node are down;
- > Impact to searching while cold storage space is expanded online; and
- > Data is searchable when an additional indexer peer node is added to the cluster.

The data throughout these test scenarios were tested and validated. The data continued to be available for search, as expected, without any service disruptions and without interruption.

Conclusion.

A distributed Splunk deployment using Red Hat Storage Server for cold storage was tested and validated, which proves that this solution is capable of performing the role as a hybrid storage solution for enterprise class Splunk deployments, without compromising on scale or performance. Red Hat Storage Server continues to help manage and maintain the explosion of large-scale data growth, meeting the most demanding enterprise storage requirements.

Just for Clicks.

Read our blog for further insight into our exciting world of Enterprise Technology: www.function1.com/blog

Download our Red Hat Storage App: https://apps.splunk.com/app/1830/

For More Information:

Sandeep Khaneja
sandeep@function1.com
202.486.4320