

TECHNOLOGY DETAIL

DELIVERING SPLUNK ENTERPRISE WITH RED HAT AND CISCO

Architecting elastic and resilient cold storage for large-scale cyber-security analytics



Quickly recognize security risk patterns by expanding analysis windows over longer periods of time.

Increase analysis windows by storing Splunk Enterprise cold data on software-defined storage from Red Hat and Cisco.

Unify operational management with the unique combination of Cisco UCS Manager and Red Hat Gluster Storage.

Improve operational elasticity of Splunk Enterprise deployments at a fraction of the cost of proprietary storage appliance solutions.

ABSTRACT

Splunk Enterprise deployed with Red Hat® Gluster Storage on Cisco Unified Computing System (Cisco UCS) servers represents a compelling cyber-analytics solution. Through this approach, organizations can cost-effectively collect, monitor, index, and analyze machine-generated data to detect and respond to threats—expanding analysis windows and increasing retention times that can greatly reduce risk for organizations under constant attack. Cisco UCS C240 M3 servers perform well under typical Splunk workloads, while software-defined storage from Red Hat and Cisco provides nondisruptive elastic cold storage expansion for Splunk Enterprise.

TABLE OF CONTENTS

1 INTRODUCTION	2
2 ANATOMY OF A SPLUNK DEPLOYMENT	2
3 DELIVERING SPLUNK ENTERPRISE WITH RED HAT SOFTWARE ON CISCO UCS SERVERS.....	8
4 ABOUT RED HAT GLUSTER STORAGE	13
5 ABOUT CISCO UCS C240 M3 SERVERS AND CISCO UCS MANAGER	14
6 SCALING THE SOLUTION ARCHITECTURE	16
7 SOLUTION ARCHITECTURE VALIDATION RESULTS	19
8 IMPLEMENTATION GUIDANCE	28
9 CONCLUSION	32



facebook.com/redhatinc
@redhatnews
linkedin.com/company/red-hat

INTRODUCTION

The very real danger from malicious cyber attacks has never been greater, and increasingly bold attacks are taking place on a constant basis. These attacks have grown exponentially in terms of frequency and sophistication, demanding an automated and near real time response to threats. In fact, the global cost of cybercrime has already grown to \$100 billion annually¹. These figures don't take into account the intangible damage to enterprise and government security and the immeasurable and sometimes irrevocable damage to brand and corporate credibility.

Every router, switch, firewall, web proxy, and intrusion protection system has a story to tell about the confidentiality, integrity, and availability of the IT environment. Relevant data from across these systems is critical to investigations as well as for continuous monitoring for situational awareness. The ability to intelligently analyze larger and larger amounts of this machine data is key to predicting and thwarting threats. Unfortunately, analysis windows are often artificially limited by the amount of data that can be cost-effectively retained for analysis. This constraint is particularly evident for larger Splunk Enterprise deployments, where significant amounts of data must be retained to develop a comprehensive level of threat awareness. Restricting analysis ultimately limits the ability of the organization to detect and respond to active threats.

Coupling Splunk Enterprise operational intelligence software with software-defined storage from Red Hat and Cisco represents a compelling approach. With its ability to collect, index, and report on terabytes of machine-generated data, Splunk Enterprise is ideally suited for security analytics, and its utility increases with more searchable data. Red Hat Gluster Storage complements Splunk Enterprise capabilities by providing elastic cold storage expansion, reducing the costs and complexity of retaining more data to expand Splunk Enterprise searches back in time. Moreover, by making an architectural decision to migrate cold data to an elastic storage pool, machine-generated data can be cost-effectively scaled out at will, without reconfiguring Splunk Enterprise indexers or disrupting the analytics process.

With high levels of CPU and I/O performance, versatile Cisco UCS C240 M3 servers represent an ideal platform for both Splunk Enterprise indexers as well as cold storage servers pooled with Red Hat Gluster Storage. This unique architecture also provides highly available enterprise or multi-site configurations, with a comprehensive management solution that includes Cisco UCS Manager to manage the infrastructure tier, while Red Hat Gluster Storage manages scale-out storage.

ANATOMY OF A SPLUNK DEPLOYMENT

Splunk Enterprise provides the ability to search, monitor, and analyze machine data with a powerful and flexible architecture. Understanding that architecture and the ways that Splunk Enterprise manages data is important background for deploying the solution with Red Hat Gluster Storage on Cisco UCS servers.

SPLUNK ENTERPRISE FOR SECURITY AND LOG ANALYTICS

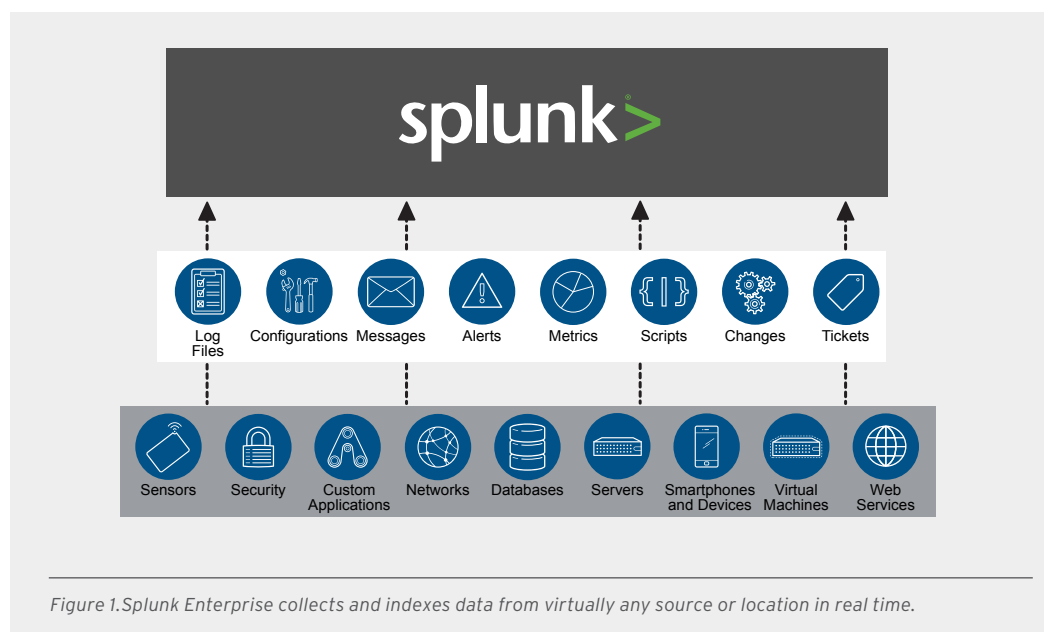
As an industry-leading platform for operational intelligence², Splunk Enterprise is a self-contained software package that supports storage on a variety of platforms, including Red Hat Gluster Storage. The software features an easy-to-use web interface and a powerful enterprise platform for analyzing machine data. Integration is also available for traditional relational databases and the Apache Hadoop open source software framework.

¹ "The Economic Impact of Cybercrime and Cyber-espionage," July 2013, Center for Strategic and International Studies.

² Gartner Magic Quadrant for Security Information and Event Management (<https://www.gartner.com/doc/2780017>).

As depicted in Figure 1, Splunk Enterprise allows organizations to do the following:

- **Collect and index any data from any source.** Splunk Enterprise enables organizations to collect and index machine-generated data from virtually any source or location in real time. This includes data streaming from packaged and custom applications, application servers, web servers, data-bases, networks, virtual machines, telecommunications equipment, operating systems, and much more.



- **Collect data from remote sources.** Splunk Enterprise forwarders deliver reliable, secure, real-time data collection from up to tens of thousands of sources, at no additional cost. Forwarding can monitor local data sources – such as applications, sensors, and endpoint devices – or collect the output of status commands on a schedule. Forwarders can also collect performance metrics from virtual or non-virtual sources, or watch the file system for configuration, permission, or attribute changes.
- **Correlate complex events.** With Splunk Enterprise, organizations can correlate complex events spanning many diverse data sources across their environments. Types of correlations include time-based and transaction-based correlations, sub-searches that take the result of one search and use them in another, as well as lookups and joins.
- **Achieve enterprise-class scale, resiliency, and interoperability.** Splunk Enterprise scales to collect and index tens of terabytes of data per day. In addition, because insights collected from that data are critical, Splunk Enterprise clustering technology provides essential availability, even as organizations scale out their low-cost distributed computing environments.

SPLUNK ENTERPRISE DATA AND STORAGE OPTIONS

Splunk Enterprise supports considerable ingenuity and flexibility in the way that data is handled as well as the supported options for storing and managing that data.

Splunk Enterprise data model

Splunk Enterprise does not manage its data through complex database schemas. Instead, the architecture is built on a concept of “buckets” and “indexes”. A bucket is a simple directory containing search indexes (metadata) and processed event data (termed “raw data”). Each bucket covers a finite time period. An index provides the higher-level construct that contains the physical buckets.

To start the analytics process, machine data is forwarded to Splunk Enterprise indexers. Indexed data then forms the basis for Splunk Enterprise to provide rich search and data-mining capabilities to the user. The Splunk Enterprise data life cycle follows a temperature-based paradigm, as described in the following and shown in Figure 2:

- **Hot data.** In the Splunk Enterprise indexer, newly indexed data first goes into a hot bucket that is searchable and being actively written. After the hot bucket reaches a certain size or time, it typically rolls to being a warm bucket, and a new hot bucket is created.
- **Warm data.** Warm buckets are searchable but are not being actively written. There are typically many warm buckets in the system. Once the indexer has created a configurable maximum number of warm buckets, it begins to roll the individual warm buckets to cold buckets based on their age.
- **Cold data.** Cold buckets are also searchable. For larger installations, cold buckets can optionally be placed on lower-cost media, such as that provided by Red Hat Gluster Storage. After a set period of time, cold buckets roll to frozen, at which point they are no longer searchable and are either archived or deleted.

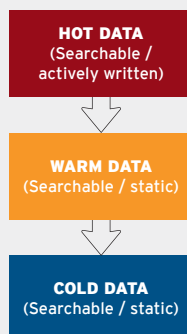


Figure 2. The Splunk Enterprise data hierarchy moves searchable data from hot to warm to cold buckets based on configurable settings.

Options for managing Splunk Enterprise data

A Splunk Enterprise solution is designed around an n-tier architecture that provides the ability to independently scale the various layers of the solution to meet data ingest and search requirements. This flexible approach gives architects a number of options for how to structure Splunk Enterprise data according to their needs. Figure 3 and the following list describe the basic options for managing Splunk Enterprise data:

- 1. All storage direct attached.** For smaller deployments, hot, warm, and cold data can all be housed on storage that is directly attached to the physical servers running Splunk Enterprise. While this approach provides fast access to storage, the configuration can be less straightforward to scale, particularly as the need for cold storage grows.
- 2. All storage remotely accessed.** As an alternative, hot, warm, and cold data can all be housed on external shared storage. Splunk strongly recommends block-level storage for indexing data and warns against file-level storage for indexing, due to network latencies that can be introduced for hot and warm data.³ However, like proprietary network-attached storage (NAS) solutions, Red Hat Gluster Storage can be used in this manner, and it performs in line with other NAS solutions.
- 3. Hybrid storage.** For many, hybrid storage represents a compelling alternative because it places hot and warm data on fast direct-attached media, while cold data is made cost effective and elastic by placing it on remote shared storage. This option is specifically supported and recommended by Splunk⁴, and it can be ideal for many larger Splunk Enterprise deployments where elastic cold storage can make a critical difference. Shown with a Red Hat Gluster Storage volume in Figure 3, hybrid storage yields fast ingest to direct attached storage for hot and warm data while allowing elastic growth and longer retention of data in low-cost less-complex cold storage infrastructure.

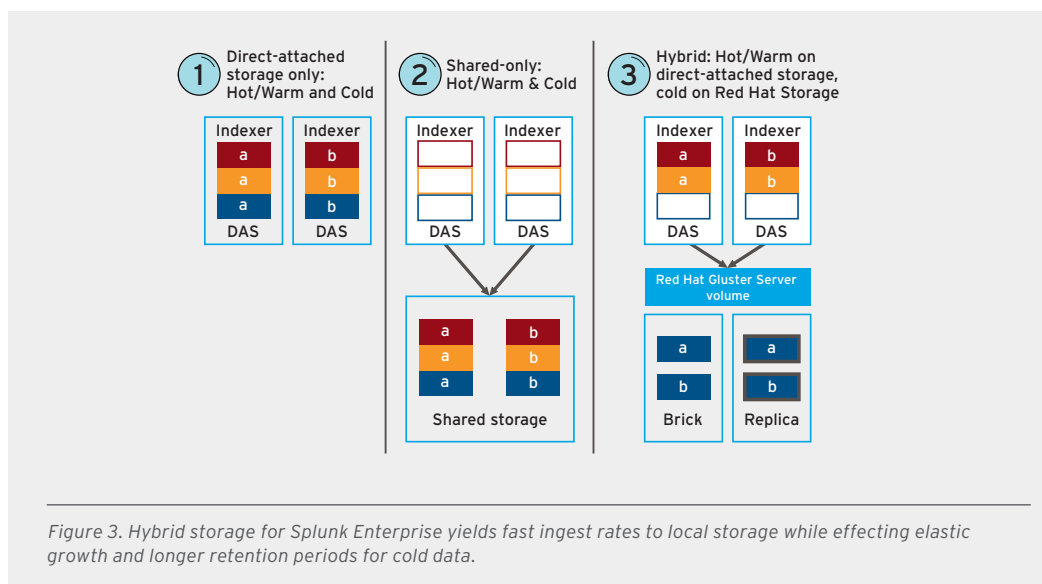


Figure 3. Hybrid storage for Splunk Enterprise yields fast ingest rates to local storage while effecting elastic growth and longer retention periods for cold data.

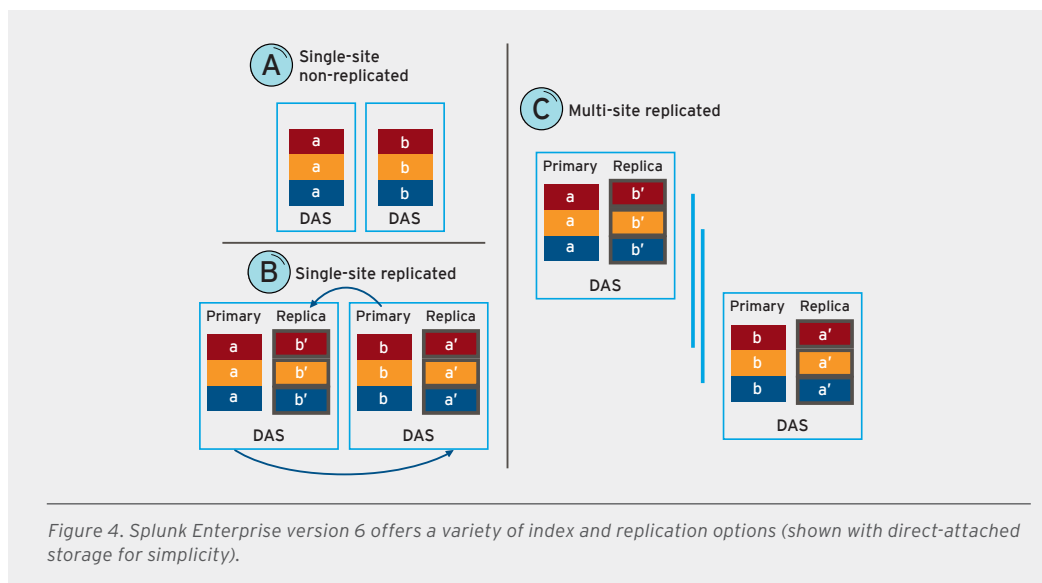
³ docs.splunk.com/Documentation/Splunk/6.1.3/installation/Systemrequirements

⁴ docs.splunk.com/Documentation/Splunk/6.1.3/indexer/Usemultiplepartitionsforindexdata

SOLUTION DEPLOYMENT ARCHITECTURES AND IMPLICATIONS

Organizations deploying Splunk Enterprise have diverse architectural objectives, and the platform is engineered to be flexible to meet a range of deployment scenarios. Splunk Enterprise architecture has been evolving over the past several years. New features recently introduced in Splunk Enterprise version 6 deliver application-level data replication and multisite awareness. In the past, these capabilities needed to be provided by the infrastructure layer. Figure 4 illustrates several Splunk Enterprise 6.1 deployment architecture options (shown with direct-attached storage for simplicity) as described in the following:

- A. Single site, no (index) replication.** This scenario implies a single Splunk Enterprise index and a single set of Splunk Enterprise hot/warm/cold buckets.
- B. Single site, with Splunk index replication.** Index replication enabled in Splunk Enterprise version 6 enables indexer server failover, offering data protection and high availability. In this scenario, the indexer starts the ingest process and then forwards data along to another indexer (or indexers) for parallel ingestion. Data is thus replicated at the point of ingest, rather than traditional approaches that replicate data at the back end.
- C. Multi-site, with Splunk index replication.** Splunk Enterprise version 6.1 can provide multisite support with index replication. This scenario allows for disaster recovery in addition to data protection and high availability. Indexers and data are replicated between sites, allowing continuous indexing and searching even in the face of a single site failure.



Red Hat Gluster Storage can provide elastic cold storage for all of these scenarios, helping to enable a choice of deployment architectures. In summary, Table 1 illustrates all of the storage approaches and Splunk Enterprise architectures where Red Hat Gluster Storage is recommended for use.

TABLE 1. WHEN TO USE SOFTWARE-DEFINED STORAGE BASED ON RED HAT GLUSTER STORAGE

	DIRECT-ATTACHED STORAGE ONLY (HOT/WARM AND COLD)	SHARED ONLY (HOT/WARM AND COLD)	HYBRID (HOT/WARM ON DIRECT-ATTACHED STORAGE, COLD ON RED HAT GLUSTER STORAGE)
SINGLE-SITE NON-REPLICATED	–	Yes	Yes
SINGLE-SITE REPLICATED	–	–	Yes
MULTISITE REPLICATED	–	–	Yes

The hybrid storage scenarios are the primary focus for this document because they are both recommended by Splunk and offer a number of benefits:

- Providing high-performance ingest and low-latency search on the most recent data located on direct-attached storage on Splunk Enterprise indexer nodes
- Offering low-cost, good-performance search on recent and older data for pattern identification across longer time windows, with elastic cold storage hosted on Red Hat Gluster Storage
- Scaling compute costs (Splunk Enterprise indexer nodes) independently from storage capacity costs (Red Hat Gluster Storage pools)
- Reducing administrative complexity and downtime with a single elastic pool of cold storage

While not generally recommended by Splunk, shared-only architectures can be constructed using Red Hat Gluster Storage. Technology assessment firm Principled Technologies conducted testing of Red Hat Gluster Storage using the performance benchmark kit SplunkIt for the shared-only scenario described previously, demonstrating highly favorable comparisons with competing NAS solutions.⁵

OPTIONAL TOOLS FOR REPLICATED DEPLOYMENTS ON RED HAT GLUSTER STORAGE

Since Splunk Enterprise replicates data upon ingest, both single-site and multisite replicated deployments can be further optimized when using Red Hat Gluster Storage for elastic cold storage. Since the Red Hat Gluster Storage volume is already replicated, naive multisite replication would result in an extra pair of copies of the cold storage data. In large configurations, this extra storage consumption would be cost prohibitive, and it is typically unnecessary.

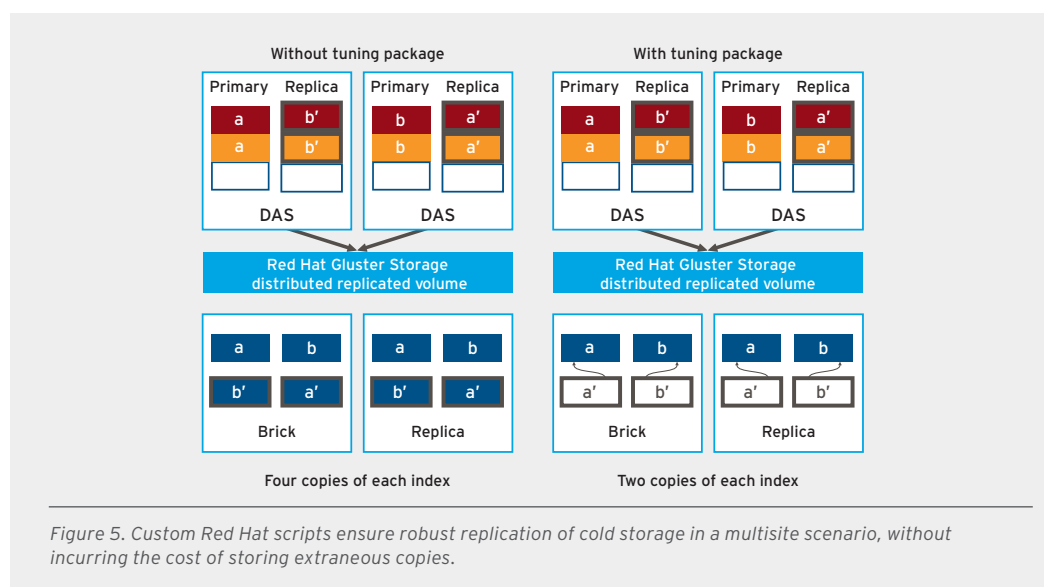
To retain the benefits of robust Splunk Enterprise index replication (such as high availability), while storing only two copies of the data, Red Hat offers an optional tuning package that helps to control costs while maintaining redundancy (Figure 5):

- **Optimized bucket movement.** Splunk Enterprise provides the capability to use a custom script to manage the movement of buckets from an index's warm path to cold storage. The Red Hat add-on utilizes this feature to determine whether a bucket should be moved to cold storage or whether links to an existing bucket can be used to save space. The distributed-replicated Red Hat Gluster Storage volume already writes two copies of the bucket across the storage pool, so fault tolerance is not impacted.

⁵ For more information on the shared-only testing, see: www.principledtechnologies.com/Red%20Hat/RedHatStorage_SplunkIt_0614.pdf.

- **Validating link integrity.** As described previously, the optimized bucket mover logic may use symbolic links to reduce the capacity usage incurred when Splunk Enterprise provides replicated buckets. This approach relies on the integrity of these links for both data availability and disaster recovery in a multisite scenario. To validate the integrity of any links within an index, Red Hat provides a tool that may be initiated by cron (8). The tool is designed to walk the index's cold path, checking each link by extracting the intended target and verifying that the target exists.

All the tools provided within the add-on package implement easily parsed log files. These log files enable Splunk administrators to monitor this aspect of the solution from within the standard Splunk web interface.



ROBUST DISASTER RECOVERY WITH MULTISITE OPERATION

Many larger operations may wish to spread their Splunk Enterprise infrastructure across multiple sites for performance or disaster recovery purposes. As discussed, Splunk Enterprise 6.1 supports this capability for Splunk Enterprise indexers. Likewise, Red Hat Gluster Storage can be used to stretch the GlusterFS volume being used for elastic cold storage between geographically separate sites. As with single-site deployments with Splunk Enterprise index replication, the add-on tuning package avoids replicating data that is already replicated, preventing unnecessary extra copies of cold data while preserving the high availability benefits of Splunk Enterprise index replication.

DELIVERING SPLUNK ENTERPRISE WITH RED HAT SOFTWARE ON CISCO UCS SERVERS

For any organization planning a large Splunk Enterprise deployment, it is important to adhere closely to recommended options and best practices. Deploying Splunk Enterprise on Cisco UCS servers with elastic cold storage provided by Red Hat Gluster Storage follows Splunk's recommended architectural practices and offers a wealth of advantages.

HYBRID SPLUNK STORAGE WITH RED HAT GLUSTER STORAGE

Splunk Enterprise stores its data in plain files, and it is specifically designed to support the configuration of an elastic cold storage tier. Splunk Enterprise allows different paths to be specified: one path for the hot and warm data, and another path for the cold data, allowing storage solutions to be selected that closely match the cybersecurity needs of a range of organizations.⁶

- A hot/warm storage tier deployed on fast local storage maximizes performance for Splunk Enterprise indexers, helping to ensure that they can keep pace with necessary ingest rates and support searches on hot/warm data.
- Deploying an elastic cold storage tier on Red Hat Gluster Storage improves the cost, capacity, scalability, and simplicity of the configuration through cost-effective scale-out storage.

Red Hat recommends configuring Red Hat Gluster Storage as a distributed-replicated volume, where the file system is distributed between multiple servers running Red Hat Gluster Storage and the data is replicated against loss. This strategy complements Splunk's optional replication of both indexers and hot/warm data.

THE ADVANTAGES OF DEPLOYING AN ELASTIC COLD STORAGE TIER

Smaller Splunk Enterprise deployments that are not expected to grow may well make due with direct-attached storage for hot, warm, and cold data. For larger deployments, however, elastic cold storage based on Red Hat Gluster Storage gives organizations the advantage of growing their searchable content at a commodity cost point – with low administrative overhead. The benefits of scale-out cold storage using Red Hat Gluster Storage on Cisco UCS C240 M3 servers include the following:

- **Performance.** Cisco UCS C240 M3 servers exhibit strong CPU and I/O performance required for running Splunk Enterprise. As shown later in this document, tested configurations using Red Hat Gluster Storage running on Cisco UCS C240 M3 servers surpass Splunk's chosen reference configurations⁷ to yield faster searches and quicker time to value.
- **Cost-effective scalability.** With a hybrid approach using Red Hat Gluster Storage, cold data can be scaled out independently from index servers without arbitrary limitations and expense imposed by proprietary storage systems.
- **Longer data retention, faster threat detection.** Depending on the use case, larger amounts of cold storage can offer the potential of improving the accuracy of analytics. The more data you can search, the more accurate you can be when it comes to determining attack patterns and vectors.
- **Simplified capacity planning and scalability.** Abstracting cold storage dramatically simplifies capacity planning because indexers only need to be sized to support desired ingest rates along with operational and tactical analysis requirements. Indexers also don't need to be taken down and reconfigured to add disk to scale cold storage.
- **Better availability and performance.** As cold storage data needs grow, reconfiguring Splunk indexers with all local data storage can result in downtime. In contrast, adding cold storage through Red Hat Gluster Storage requires no downtime and implies only doing a peer probe, adding in a storage brick, and launching a rebalance, with no disruption to indexers. Distributing cold storage across a great number of nodes and/or spindles also offers the potential of improved performance.

⁶ docs.splunk.com/Documentation/Splunk/6.1.3/Indexer/Usemultiplepartitionsforindexdata

⁷ docs.splunk.com/Documentation/Splunk/6.1.3/Installation/Referencehardware

- **Data availability and disaster recovery.** Combining Splunk Enterprise multisite support with that of Red Hat Gluster Storage provides redundancy at the indexer as well as cold storage. Red Hat Gluster Storage complements the multisite capabilities of Splunk Enterprise, allowing improved data availability, continuous operation, and recovery from otherwise disabling events.

SOFTWARE-DEFINED STORAGE ON CISCO C240M3 SERVERS

Cisco UCS C240 M3 servers combined with Red Hat Gluster Storage represent a flexible and compelling software-defined storage architecture. The servers feature strong CPU performance along with memory and I/O capacity that offers performance advantages for ingest and search. Recommended distributed-replicated volumes provided by Red Hat Gluster Storage deliver a cold storage layer that can be scaled up or down, without impacting Splunk Enterprise ingest or tactical search operations.

Mapping elastic cold storage to Splunk Enterprise deployment architectures.

The Red Hat and Cisco solution architecture is flexible and can easily map to the hybrid Splunk deployment architectures discussed earlier in this document. Common advantages of these deployment scenarios include the following:

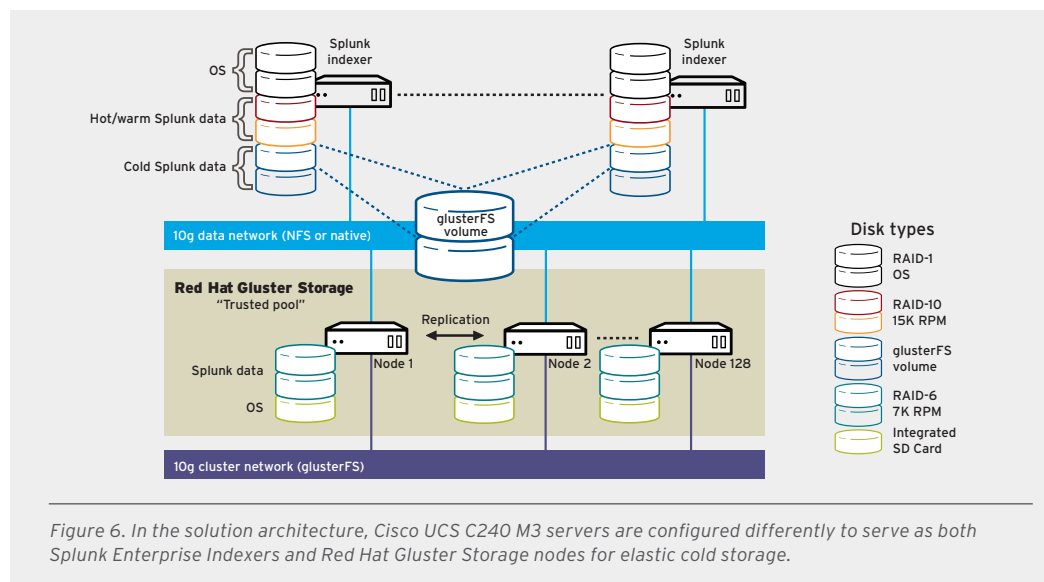
- Dedicated and separate hot/warm and cold storage tiers facilitate elastic cold storage and easier management.
- In Splunk Enterprise non-replicated scenarios, Red Hat Gluster Storage provides data protection for the cold volume through Red Hat Gluster Storage replicas.
- In replicated scenarios, Splunk Enterprise provides high-availability for Splunk indexer operations and data protection for Splunk data. For the cold data volume, Red Hat Gluster Storage assists by providing the data protection through Red Hat Gluster Storage replicas.
- Data movement to cold storage can be managed by a Red Hat Gluster Storage add-on, helping to ensure that Splunk Enterprise replication and Red Hat Gluster Storage replication complement each other (replicated configurations only).
- A choice of Network File System (NFS) or Gluster Native Client connectivity offers transparent failover.
- Server configuration flexibility allows faster and more expensive 10K or 15K RPM drives on Splunk indexers and more cost-effective 7.2K RPM drives on Red Hat Gluster Storage nodes.

Single site deployment without Splunk Enterprise index replication

Figure 6 illustrates a typical single-site scenario without Splunk-based index replication. In this scenario, each indexer is responsible for its own discrete ingest workload as well as hot/warm data storage. Cold storage is located on a shared GlusterFS volume provided by Red Hat Gluster Storage. Access to the GlusterFS volume is provided by either NFS or Gluster Native Client. Separate 10 Gigabit Ethernet networks isolate the GlusterFS traffic from the service-oriented traffic.

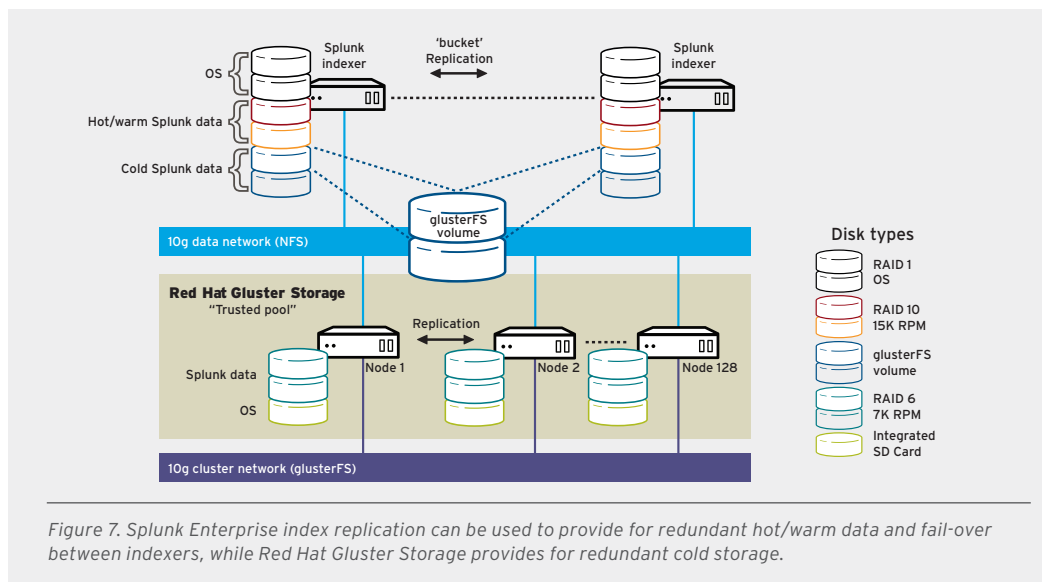
Though they are configured differently, indexers and Red Hat Gluster Storage nodes in the solution architecture are both served by Cisco UCS C240 M3 servers. To help ensure performance on the indexers, Splunk recommends faster 10K or 15K RPM direct-attached storage, configured as RAID 10. For cost-effective capacity, each Cisco UCS C240 M3 server running Red Hat Gluster Storage is configured with twelve 7K RPM disks in a RAID 6 configuration. The Red Hat Gluster Storage nodes boot from an integrated SD card.

Though this scenario does not include index replication, the cold data is replicated through Red Hat Gluster Storage. A distributed-replicated volume means that Red Hat Gluster Storage nodes are deployed in redundant replicated pairs. RAID 6 protects against any two concurrent disk failures on an individual Red Hat Gluster Storage node. As a distributed-replicated GlusterFS volume, cold data is also protected against the loss of one cold storage node in each pair of servers.



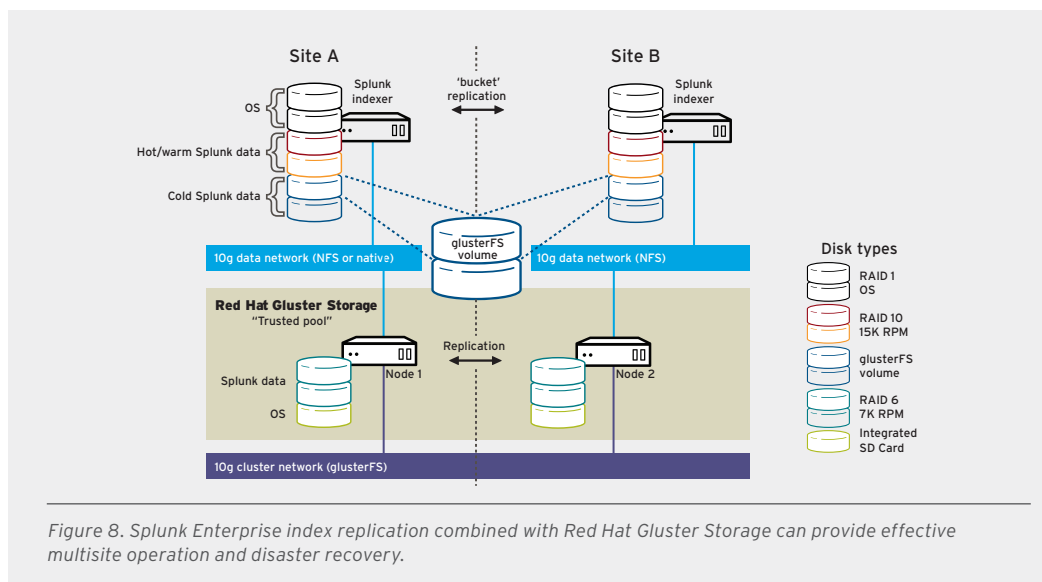
Single site deployment with Splunk Enterprise index replication

Figure 7 illustrates a single-site solution architecture that utilizes the index replication features provided in Splunk Enterprise 6 to allow for full indexer redundancy and failover. In this scenario, index replication is provided on the front end by Splunk Enterprise, with hot/warm bucket replication provided by at least a pair of indexers. As with the previous scenario, cold storage replication is provided by Red Hat Gluster Storage. As stated previously, the Red Hat add-on package minimizes extra data copies that would otherwise be created by Splunk Enterprise replication.



Multi-site deployment with Splunk Enterprise index replication

Many larger sites now require resiliency and disaster recovery for increasingly critical cybersecurity services. In this scenario, Splunk Enterprise provides multisite bucket replication while Red Hat Gluster Storage provides replication between sites for cold storage (Figure 8).



UNIFYING SERVER, STORAGE, AND ANALYTICS MANAGEMENT

Given the potential data volumes involved, analytics data can grow quickly and require extensive infrastructure. To be effective, organizations deploying large cybersecurity efforts need streamlined management that lets them focus on the task at hand. Beyond the technical merits of the solution architecture, the combination of Splunk Enterprise, Red Hat Gluster Storage, and Cisco UCS presents a compelling unified management solution:

- Cisco UCS Manager is embedded in the Cisco UCS fabric interconnects and provides true hardware state abstraction for servers, networking, and local storage. Cisco UCS Manager offers a model-based management approach and a highly available clustered configuration.
- Red Hat Gluster Storage complements Cisco UCS Manager by effectively managing all of the cold storage attached to the Cisco UCS C240 M3 servers running Red Hat Gluster Storage.
- The Splunk web interface provides a unified view of the Splunk Enterprise environment. The Red Hat Gluster Storage Application for Splunk Enterprise (apps.splunk.com) keeps the health of the cold storage layer visible to Splunk administrators.

ABOUT RED HAT GLUSTER STORAGE

Red Hat Gluster Storage is a software-defined open source solution that is designed to meet unstructured, semistructured, and big data storage requirements. At the heart of Red Hat Gluster Storage is a secure, open source, massively scalable distributed file system that allows organizations to combine large numbers of storage and compute resources into a high-performance, virtualized, and centrally managed storage pool.

Red Hat Gluster Storage is designed to achieve several major goals:

- **Elasticity.** Storage volumes are abstracted from the hardware and managed independently. Volumes can grow or shrink by adding or removing systems from the storage pool. Even as volumes change, data remains available without interrupting the application.
- **Petabyte scalability.** Today's organizations demand scalability from terabytes to multiple petabytes. Red Hat Gluster Storage lets organizations start small and grow to support multi-petabyte repositories as needed. Organizations that need substantial amounts of storage can deploy massive scale-out storage from the outset.
- **High performance.** Red Hat Gluster Storage provides fast file access by eliminating the typical centralized metadata server. Files are spread evenly throughout the system, eliminating hot spots, I/O bottlenecks, and high latency. Organizations can use commodity disk drives and 10 Gigabit Ethernet to maximize performance.
- **Reliability and high availability.** Red Hat Gluster Storage provides automatic replication that helps ensure high levels of data protection and resiliency. In addition to protecting from hardware failures, self-healing capabilities restore data to the correct state following recovery.
- **Industry-standard compatibility.** For any storage system to be useful, it must support a broad range of file formats. Red Hat Gluster Storage provides native Portable Operating System Interface (POSIX) compatibility as well as support for common protocols including Common Internet File System (CIFS), NFS, and HTTP. The software is readily supported by off-the-shelf storage management software.

- **Unified global namespace.** Red Hat Gluster Storage aggregates disk and memory resources into a single common pool. This flexible approach can simplify management of the storage environment and eliminate data silos. Global namespaces may be grown and shrunk dynamically, without interrupting client access.

Red Hat Gluster Storage provides distinct technical advantages over other technologies:

- **Software-defined storage.** Locking organizations into a particular storage hardware vendor or hardware configuration is a problem. Instead, Red Hat Gluster Storage works with a wide variety of industry-standard storage, networking, and compute server solutions.
- **Open source.** Red Hat Gluster Storage is based on the GlusterFS project. A worldwide community of developers, customers, and partners test and update GlusterFS in a wide range of environments and workloads, providing continuous and unbiased feedback to other users. This project is certified, secured, and made enterprise-ready in the Red Hat Gluster Storage distribution.
- **Complete storage operating system stack.** The storage product delivers more than just a distributed file system, adding distributed memory management, I/O scheduling, software RAID, self-healing, local n-way synchronous replication, and asynchronous long-distance replication.
- **User space.** Unlike traditional file systems, Red Hat Gluster Storage operates in user space, rather than kernel space. This makes installing and upgrading Red Hat Gluster Storage significantly easier and greatly simplifies development efforts because specialized kernel experience is not required.
- **Modular, stackable architecture.** Red Hat Gluster Storage is designed using a modular and stackable architecture approach. Configuring it for highly specialized environments is a simple matter of including or excluding particular modules.
- **Data stored in native formats.** Data is stored on disk using native formats (XFS) with various self-healing processes established for data. As a result, the system is extremely resilient and files stay naturally readable, even without the Red Hat Gluster Storage software. There is no proprietary or closed format used for storing file data.
- **No metadata with the elastic hash algorithm.** Unlike other storage systems with a distributed file system, Red Hat Gluster Storage does not create, store, or use a separate index of metadata in any way. Instead, it places and locates files algorithmically. The performance, availability, and stability advantages of not using metadata are significant and in some cases produce dramatic improvements.

ABOUT CISCO UCS C240M3 SERVERS AND CISCO UCS MANAGER

Cisco UCS is a next-generation datacenter platform that unites compute, network, storage access, and virtualization into a cohesive system designed to reduce total cost of ownership (TCO) and increase business agility. The system integrates a low-latency, lossless 10 Gigabit Ethernet unified network fabric with enterprise-class x86-architecture servers. The system is an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain. The Cisco UCS configuration deployed in this solution architecture includes the following:

- Cisco UCS 6200 Series Fabric Interconnects, for line-rate, low-latency, lossless 10Gbps Ethernet interconnect switches that consolidate I/O within the system
- Cisco UCS 2200 Series Fabric Extenders, which bring the unified fabric to rack-mount servers

- Cisco UCS C240 M3 Rack-Mount Servers, extending unified computing innovations to both Splunk Enterprise indexers and storage servers running Red Hat Gluster Storage
- Cisco UCS Virtual Interface Cards (VICs), providing a dual-port Enhanced Small Form Factor Pluggable (SFP+) 10 Gigabit Ethernet and Fibre Channel over Ethernet (FCoE)-capable PCI Express (PCIe) card designed exclusively for Cisco UCS C-Series Rack Servers
- Cisco UCS Manager, offering centralized infrastructure management capabilities through a unified management domain and serving as the central nervous system of Cisco UCS

CISCO UCS C240 M3 RACK SERVER

Cisco UCS C240 M3 Rack Server is designed for both performance and expandability over a wide range of storage-intensive infrastructure workloads. Occupying a 2 RU form factor, the server supports Intel Xeon E5 family processors, large memory, and a choice of either large form factor (LFF) or small form factor (SFF) disk drives. The choice of disk drive form factors makes the platform extremely flexible and ideal for the Splunk and Red Hat Gluster Storage solution architecture. Each Cisco UCS C240 M3 server includes the following:

- Up to two Intel Xeon E5-2600 or E5-2600 v2 processors
- Up to 768GB of RAM with 24 DIMM slots
- Choice of up to 12 LFF or 24 SFF front-accessible, hot-swappable SAS, SATA, or SSD drives for local storage, providing redundancy options and ease of serviceability
- Choice of RAID controllers to provide data protection for up to 12 or 24 SAS, SATA, or SSD drives
- Five PCI Express Gen 3 slots and four Gigabit Ethernet LAN interfaces on the motherboard
- Trusted Platform Module (TPM) for authentication and toolless access

The choice of LFF or SFF disk drives in the same physical server form factor makes the Cisco UCS C240 M3 server particularly effective for this solution architecture. The server can be configured for a range of performance as a Splunk indexer node using more and faster SFF hard disk drives (HDDs). The server can also be configured for capacity as a Red Hat Gluster Storage node utilizing a smaller number of cost-effective, large-capacity LFF HDDs. Sample configurations are provided in the section “Scaling the Solution Architecture.”

CISCO UCS SERVICE PROFILES

To understand how the Cisco Unified Computing System delivers these benefits, it is necessary to understand the concept of a service profile, the fundamental mechanism by which the Cisco Unified Computing System models the necessary abstractions of server, storage, and networking. Conceptually, a service profile is an extension of the virtual machine abstraction applied to physical servers. The definition has been expanded to include elements of the environment that span the entire datacenter, encapsulating the server identity (LAN and SAN addressing, I/O configurations, firmware versions, boot order, network VLAN, physical port, and quality-of-service [QoS] policies) in logical service profiles that can be dynamically created and associated with any physical server in the system within minutes rather than hours or days. The association of service profiles with physical servers is performed as a simple, single operation. It enables migration of identities between servers in the environment without requiring any physical configuration changes and facilitates rapid bare-metal provisioning of replacements for failed servers.

CISCO UCS MANAGER

Cisco UCS Manager provides infrastructure management for the solution architecture, offering unified embedded management of all software and hardware components in Cisco UCS. Embedded in the fabric interconnects and utilizing Cisco SingleConnect technology, Cisco UCS Manager controls multiple racks of blade and rack-mount servers, managing resources for both virtual and bare-metal deployments. Through its unified, embedded, and policy-based approach, Cisco UCS Manager helps reduce management and administration expenses. Cisco UCS Manager can manage up to 160 servers and thousands of Cisco UCS components. With Cisco UCS Central, management can be extended globally to thousands of servers in multiple domains.

SCALING THE SOLUTION ARCHITECTURE

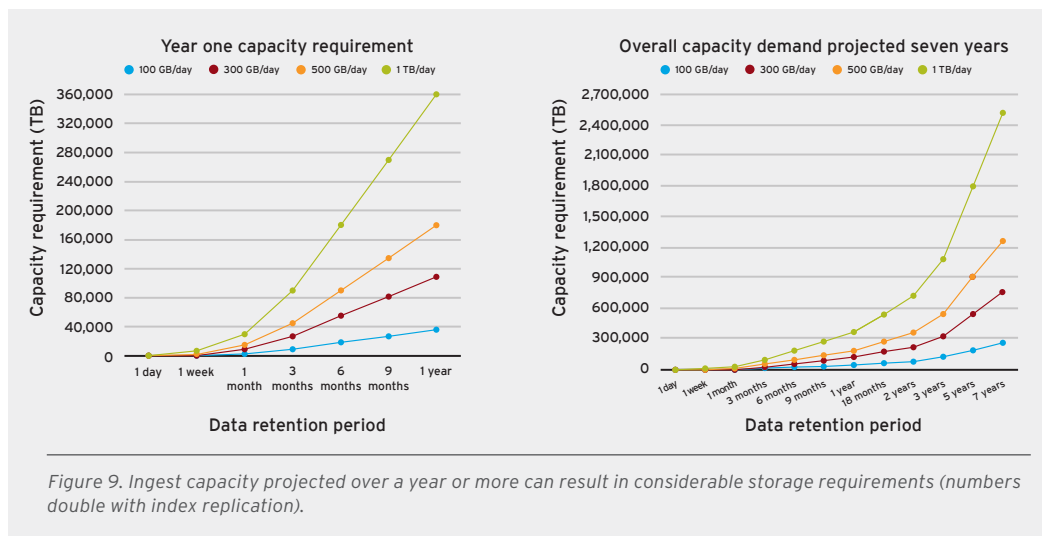
While providing specific scaling and sizing recommendations is beyond the scope of this paper, the Red Hat and Cisco solution architecture is highly scalable and can be adapted to a wide variety of situations. The distributed software-defined nature of both Splunk Enterprise and scale-out storage provided by Red Hat Gluster Storage means that configurations can grow without arbitrary limitations. In addition, the considerable power flexibility of the Cisco UCS C240 M3 server means that it can be appropriately configured as either an indexer or a Red Hat Gluster Storage node.

INGEST VOLUME AND RETENTION PERIOD

No two organizations will have the same needs in terms of sizing their Splunk Enterprise deployments. Designing and sizing a Splunk Enterprise configuration is highly dependent on the answers to several important questions. For a full treatment of these questions, refer to the Splunk Installation Manual⁸ and the Splunk Distributed Deployment Manual⁹. Primary among these considerations are the anticipated daily ingest volume and the retention period for the data. Figure 9 illustrates capacity requirements based on various levels of ingest volume, projected across retention periods over a full year and for a seven-year period. Together, these charts demonstrate the considerable cold storage volumes that can result. Importantly, these numbers double if the Splunk Enterprise implementation is configured for index replication.

⁸ docs.splunk.com/Documentation/Splunk/6.1.3/Installation/CapacityplanningforalargerSplunkdeployment

⁹ docs.splunk.com/Documentation/Splunk/latest/Deploy/Summaryofperformancerecommendations



SPLUNK ENTERPRISE INDEX SERVER CONFIGURATIONS

In hybrid configurations, Splunk Enterprise indexers must be sized to handle the CPU-intensive indexing and search workload as well as provide sufficient storage for hot and warm data. The Cisco UCS C240 M3 server can be configured to provide indexers appropriate for varying levels of ingest volume. For a basic medium configuration, a Cisco UCS C240 M3 server could be configured with the following:

- Two 2.8GHz Intel Xeon E5-2680 v2 processors, each with 10 cores and 25MB of cache
- 64GB of memory
- Twenty-four 300GB 6Gb SAS 15K RPM SFF disk drives (7.2TB raw capacity)
- Cisco VIC dual port SFP+ converged network adapter

Configurations expecting higher levels of ingest or search would benefit from more processor cache, more system memory, and larger internal storage capacity for hot and warm data, as follows:

- Two 2.7GHz Intel Xeon E5-2697 v2 processors, each with 12 cores and 30MB of cache
- 128GB of memory
- Twenty-four 1.2TB 6Gb SAS 10K RPM SFF disk drives (28.8TB raw capacity)
- Cisco VIC dual port SFP+ converged network adapter

Medium configuration indexer bill of materials

Table 2 lists the bill of materials for a Cisco UCS C240 M3 indexer for medium-sized configurations.

TABLE 2. MEDIUM INDEXER CONFIGURATION

CISCO PART NUMBER	QUANTITY	DESCRIPTION
UCSC-C240-M3S	1	UCS C240 M3 SFF w/o CPU, mem, HD, PCIe, w/ rail kit, expdr
CON-SNT-C240M3SF	1	SMARTNET 8X5XNBD UCS C240 M3 SFF w/o
UCSC-RAIL-2U	1	2U Rail Kit for UCS C-Series servers
UCS-CPU-E52680B	2	2.80 GHz E5-2680 v2/115W 10C/25MB Cache/DDR3 1866MHz
UCS-MR-1X162RY-A	4	16GB DDR3-1600-MHz RDIMM/PC3-12800/dual rank/1.35v
UCS-HDD300GI2F105	24	300GB 6Gb SAS 15K RPM SFF HDD/hot plug/drive sled mounted
UCSC-PCIE-CSC-02	1	Cisco VIC 1225 Dual Port 10Gb SFP+ CNA
UCSC-PSU-650W	2	650W power supply for C-series rack servers
CAB-N5K6A-NA	2	Power Cord, 200/240V 6A North America
UCSC-HS-C240M3	2	Heat Sink for UCS C240 M3 Rack Server
UCSC-PCIF-01F	3	Full height PCIe filler for C-Series
UCS-M3-V2-LBL	1	Cisco M3 - v2 CPU asset tab ID label (Auto-Expand)
UCS-RAID9271CV-8I	1	MegaRAID 9271CV with 8 internal SAS/SATA ports with Supercap

Large configuration indexer bill of materials

Table 3 lists the bill of materials for a Cisco UCS C240 M3 indexer for a larger sized configuration.

TABLE 3. LARGE INDEXER CONFIGURATION

CISCO PART NUMBER	QUANTITY	DESCRIPTION
UCSC-C240-M3S	1	UCS C240 M3 SFF w/o CPU, mem, HD, PCIe, w/ rail kit,
CON-SNT-C240M3SF	1	SMARTNET 8X5XNBD UCS C240 M3 SFF w/o
UCSC-RAIL-2U	1	2U Rail Kit for UCS C-Series servers
UCS-CPU-E52697B	2	2.70 GHz E5-2697 v2/130W 12C/30MB Cache/DDR3 1866MHZ
UCS-MR-1X162RY-A	8	16GB DDR3-1600-MHz RDIMM/PC3-12800/dual rank/1.35v
UCS-HD12T10KS2-E	24	1.2 TB 6G SAS 10K rpm SFF HDD
UCSC-PCIE-CSC-02	1	Cisco VIC 1225 Dual Port 10Gb SFP+ CNA
UCSC-PSU-650W	2	650W power supply for C-series rack servers
CAB-N5K6A-NA	2	Power Cord, 200/240V 6A North America
UCSC-HS-C240M3	2	Heat Sink for UCS C240 M3 Rack Server
UCSC-PCIF-01F	3	Full height PCIe filler for C-Series
UCS-M3-V2-LBL	1	Cisco M3 - v2 CPU asset tab ID label (Auto-Expand)
UCS-RAID9271CV-8I	1	MegaRAID 9271CV with 8 internal SAS/SATA ports with Supercap

RED HAT GLUSTER STORAGE NODES

In contrast to indexers that are optimized primarily for performance, Red Hat Gluster Storage nodes deployed for elastic cold storage need to be optimized for capacity. With support for either SFF or LFF disk drives, the versatile Cisco UCS C240 M3 server platform can be configured to optimally support both Splunk Enterprise indexers and Red Hat Gluster Storage nodes. Table 4 provides a bill of materials for configuring a Red Hat Gluster Storage node with the following properties:

- Two 2.9GHz Intel Xeon E5-2690 v2 processors, each with eight cores and 20MB of cache
- 128GB of memory
- Twelve 4TB disk drives (48TB raw capacity)
- Cisco VIC dual port SFP+ converged network adapter

TABLE 4. RED HAT GLUSTER STORAGE NODE CONFIGURATION

CISCO PART NUMBER	QUANTITY	DESCRIPTION
UCSC-C240-M3L	1	UCS C240 M3 LFF w/oCPU,mem,HD,PCIe,PSU w/ rail kit, expdr
CON-SNT-C240M3LFF	1	SMARTNET 8X5XNBD UCS C240 M3 Server - LFF
UCSC-RAIL-2U	1	2U Rail Kit for UCS C-Series servers
UCS-CPU-E5-2690	2	2.90 GHz E5-2690/135W 8C/20MB Cache/DDR3 1600 MHz
UCS-MR-1X162RY-A	8	16GB DDR3-1600-MHz RDIMM/PC3-12800/dual rank/1.35v
UCS-HD4T7KS3-E	12	4TB SAS 7.2K RPM 3.5 inch HDD/hot plug/drive sled mounted
UCSC-PCIE-C10T-02	1	Cisco VIC 1225T Dual Port 10GBaseT CNA
UCSC-PSU-650W	2	650W power supply for C-series rack servers
CAB-N5K6A-NA	2	Power Cord, 200/240V 6A North America
UCSC-HS-C240M3	2	Heat Sink for UCS C240 M3 Rack Server
UCSC-PCIF-01F	3	Full height PCIe filler for C-Series
UCS-RAID9271CV-8I	1	MegaRAID 9271CV with 8 internal SAS/SATA ports with Supercap

SOLUTION ARCHITECTURE VALIDATION RESULTS

Red Hat understands the importance of careful testing and validation before deploying any storage solution for production and mission-critical cybersecurity operations. For this reason, Red Hat partnered with Splunk, Cisco, and enterprise technology company Function1 (function1.com) to perform extensive testing that evaluated both the performance and reliability of providing elastic cold storage on Red Hat Gluster Storage using Cisco UCS C240 M3 servers. Testing of the solution architecture was performed using Expanded Benchmark Framework (EBF) on Cisco UCS C240 M3 servers running Red Hat Gluster Storage 2.1. Subsequent testing was performed with Red Hat Gluster Storage 3.0 for verification.

SPLUNKIT AND THE EXPANDED BENCHMARK FRAMEWORK (EBF)

Most Splunk Enterprise users are aware of the SplunkIt benchmark, which has been publicly available for some time. SplunkIt was designed to be quick and simple and give a rough comparative performance measurement. Unfortunately, SplunkIt is less useful for emulating real-world situations, particularly for larger or distributed multi-server Splunk Enterprise deployments. For example, the default SplunkIt single-server test only indexes and searches 50GB of syslog data, and the test runs for only 80 minutes.

In contrast, EBF was designed by a team of Splunk certified architects and tested by Function1 (a certified Splunk partner) to be substantially more robust to reflect challenging real-world conditions. The purpose of the tool is to benchmark the performance of a Splunk Enterprise environment with specific hardware and hardware combinations, which can then be compared against a Splunk reference configuration.¹⁰

The typical EBF workflow is as follows:

- 1. Deploy a single-server or multi-server test bed.
- 2. Install a full Splunk Enterprise deployment.
- 3. Generate a 1TB custom, high-cardinality, syslog-like dataset.
- 4. Index the generated dataset..

Note that the 1TB dataset is large enough to effectively exercise cold storage, which is the goal of the testing. A number of types of search tests are performed as a part of EBF. Both static (no indexing load) and streaming (searching while indexing) tests are performed as a part of EBF. The following types of searches are evaluated:

- **Rare searches.** These searches are typically I/O bound and are performed in EBF with term density of one in 1 million, one in 10 million, and one in 100 million.
- **Dense searches.** These searches are typically CPU bound, and EBF performs dense searches with a term density of one in 100 and one in 1,000. Random search terms are carefully selected to fit within the correct densities. Various search concurrency levels are run as a part of the EBF, simulating various numbers of concurrent search users. Search concurrency starts at four, and is increased in increments of four as a part of the testing. Multiple runs are also performed.

A number of system metrics is collected as a part of EBF testing to evaluate performance compared to the reference configuration:

- Indexing throughput is measured for all indexers.
- Metrics are evaluated over both concurrency and term density.
- For streaming tests, indexing throughput is evaluated over both concurrency and term density.

As a part of testing conducted by Red Hat on Cisco UCS servers, a full set of EBF test results was audited by Splunk. In addition, engineers used EBF to generate load on the Cisco UCS based Splunk Enterprise configuration to validate and evaluate the behavior and performance of Red Hat Gluster Storage across a number of configuration variants.

¹⁰ docs.splunk.com/Documentation/Splunk/6.1.3/Installation/Referencehardware

TESTED HARDWARE CONFIGURATION

To run EBF, Red Hat and Cisco engineers utilized a laboratory configuration consisting of a small Cisco UCS cluster comprised of the following:

- Two Cisco UCS C240 M3 servers acting as Splunk Enterprise indexers
- Four Cisco UCS C240 M3 servers acting as Red Hat Gluster Storage nodes
- A Cisco UCS C240 M3 server acting as search head
- Forwarders running as virtual machines under the KVM hypervisor on a Cisco C240 M3 server

Splunk Enterprise indexer nodes

The two Splunk Enterprise indexer nodes were configured as follows:

- **CPU.** Two 2.40GHz Intel Xeon E5-2665 processors
- **RAM.** 256GB
- **RAID.** LSI MegaRAID SAS 9266-8i card
- **Network.** Two 10 Gb interfaces (VIC)
- **HDDs.** An initial audited configuration was used to evaluate Red Hat Gluster Storage against direct-attached storage, which used eighteen 1.2TB 7K RPM HDDs configured as a 16+2 RAID group. The index servers were subsequently reconfigured as per Splunk's best practice recommendations with twelve 300GB 15K RPM HDDs configured as RAID 10.

Red Hat Gluster Storage nodes

The two Red Hat Gluster Storage nodes were configured as follows:

- **CPU.** Two 2.40GHz Intel Xeon E5-2665 processors
- **RAM.** 256GB installed, with different configuration options tested (64GB-256GB)
- **RAID.** Nytro MegaRAID8110-4i - Nytro cache (SSD) disabled, FBWC active
- **Network.** Two 10Gb interfaces (Cisco)
- **HDDs.** Twenty-four 1.2TB 7K RPM HDDs configured as two RAID 6 groups (10+2)

Search head

The search head was configured as follows:

- **CPU.** Two 2.4GHz Intel Xeon E5-2665 processors
- **RAM.** 256GB
- **RAID.** LSI MegaRAID SAS 9266-8i
- **Network.** Two 10Gb interfaces (vNIC)
- **HDDs:** Eighteen 1.2TB 7K RPM HDDs configured as a 16+2 RAID group

Forwarders

Forwarders were configured as virtual servers running on a KVM hypervisor on a Cisco UCS C240 M3 server as follows:

- **CPU.** Four virtual CPUs
- **RAM.** 8GB vRAM
- **Network.** One vNIC
- **Disk.** 20GB vDisk for the OS, 2TB vDisk for data used to load the EBF environment

Network configuration

The physical network was comprised of 10Gbps optical connections to a Cisco UCS 6296UP fabric interconnect. Two separate vLANS were configured using Cisco UCS Manager to provide the following:

- Front-end vLAN for communication between Splunk Enterprise indexers and Red Hat Gluster Storage nodes
- Back-end vLAN for communication between Red Hat Gluster Storage nodes

Jumbo frames were enabled on all interfaces, and MTU was set to 9,000 on all associated interfaces. Network bonding (mode=6) was used on both of the indexers to aggregate multiple 10Gb links. The eNIC driver was used. Recommended DNS entries were configured for Red Hat Gluster Storage nodes:

- A private network was configured to be the same as the host name.
- Virtual IPs were configured for NFS.

EVALUATING PERFORMANCE BETWEEN LOCAL AND REMOTE STORAGE

One testing goal was to evaluate any performance impact from deploying remote and distributed cold storage as compared with cold storage directly attached to the indexers. To perform this evaluation, Red Hat engineers conducted initial performance testing to establish a baseline and identify any performance differences between the following configurations:

- Hot, warm, and cold storage configured on direct-attached storage on the Cisco UCS C240 M3 servers that were acting as indexers
- Hybrid configuration with hot and warm storage on the indexers and elastic cold storage accessed on a GlusterFS volume implemented using Cisco UCS C240 M3 servers and accessed via NFS
- Splunk Enterprise reference configuration¹¹

¹¹ The Splunk reference configuration utilized a RAID-0 configuration for locally-attached storage. This is different than the stated reference configuration of RAID 10 (docs.splunk.com/Documentation/Splunk/6.1.3/Installation/Referencehardware). Though a RAID-0 configuration (blocks striped, no mirror, no parity) would provide good performance, it would not be useful in a real-world scenario since there is no redundancy and recovery path for the data. In contrast, Red Hat Gluster Storage employed RAID-6, allowing the volume to continue to execute read and write requests in the presence of any two concurrent disk failures.

To make the testing even-handed, indexers and Red Hat Gluster Storage nodes were all equipped with 7K RPM drives. This was done to observe any differences from having local vs. distributed elastic cold storage. As cited previously, the Splunk Enterprise reference configuration utilized direct-attached 15K RPM drives, which is a Splunk best practice for hot/warm data.

Indexing throughput test

To evaluate indexing performance across different types of storage, a 2x 1TB dataset was generated and forwarded to the two index servers. The results in Table 5 compare the Splunk reference configuration with locally attached 15K RPM drives against two configurations of Cisco UCS C240 M3 servers—a direct-attached configuration and one using Red Hat Gluster Storage for elastic cold storage.

TABLE 5. EBF THROUGHPUT COMPARISON AGAINST THE SPLUNK REFERENCE CONFIGURATION

TEST RUN	KBPS95	KBPSAVG	KBPSMED	EPS95	EPSAVG	EPSMED
SPLUNK REFERENCE PLATFORM	30600	23200	24800	81800	62059	66300
DIRECT-ATTACHED STORAGE	35119 ↑	30720 ↑	32800 ↑	77244 ↓	63182 ↑	67657 ↑
RED HAT GLUSTER STORAGE	32753 ↑	30108 ↑	31459 ↑	72705 ↓	55176 ↓	56533 ↓

A number of observations can be drawn from these results.

- The Cisco UCS C240 M3 servers in the local storage and using Red Hat Gluster Storage largely exceeded the performance of the Splunk Enterprise reference configuration (despite the RAID 0 configuration of the reference configuration). This is largely due to the superior CPU resources of the Cisco UCS C240 M3 servers.
- The cases where throughput is lower for the Cisco UCS servers (events per second) is likely due to the use of slower 7.2K RPM disks, compared to the faster 15K RPM disks used for the reference configuration.
- As expected, distributing cold storage via Red Hat Gluster Storage impacted performance, but only slightly. Throughput numbers for direct-attached storage and storage distributed via Red Hat Gluster Storage are very close, yielding a small performance cost for the benefits of having elastic cold storage.

Rare term streaming search tests

Searches were run against the entire dataset for terms that are as rare as one in 1 million (term length 6), one in 10 million (term length 7), and one in 100 million (term length 8), and for concurrent searches ranging from 4-36 in steps of four. Streaming tests were conducted under a non-throttled indexing load. As shown in Figure 10, all of the rare term streaming tests conducted using Red Hat Gluster Storage for cold storage exceeded the Splunk Enterprise reference configuration in terms of searches per minute, despite using slower 7.2K RPM disk drives.

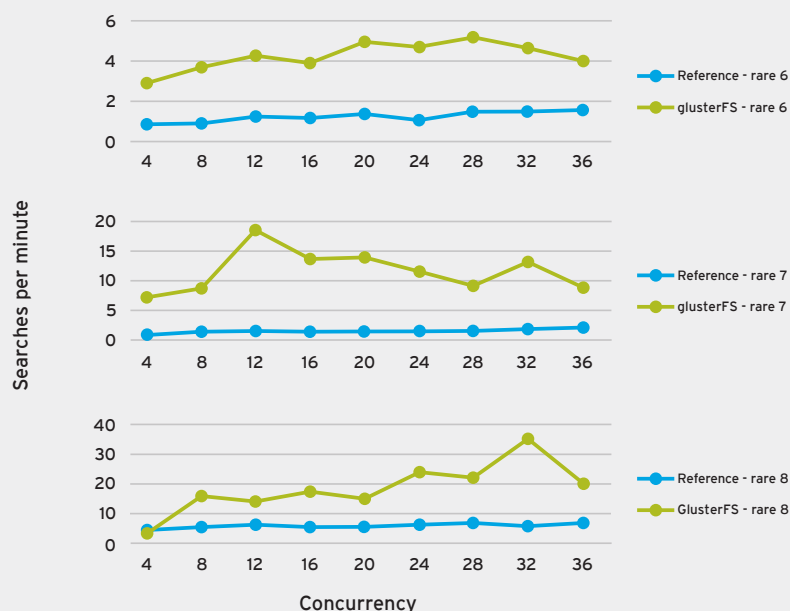


Figure 10. In EBF testing, Cisco UCS C240 M3 servers with elastic cold storage on Red Hat Gluster Storage uniformly exceeded reference configuration on rare search terms under streaming load.

Dense term streaming search tests

To complement the rare term streaming tests, EBF also provides streaming testing for dense term searches. Searches were run against a preselected time range for terms that are as dense as one in 100 (term length 2) and one in 1,000 (term length 3) with concurrent searches ranging from 4-36 in steps of four. As with the rare term streaming tests, dense term streaming tests were conducted under a non-throttled indexing load.

As shown in Figure 11, the dense term streaming tests run with Red Hat Gluster Storage elastic cold storage offer similar performance to the reference configuration that used direct-attached storage for cold storage.

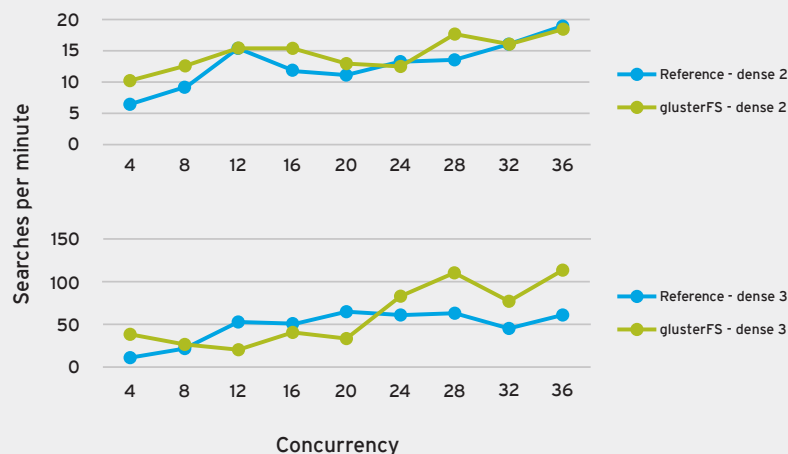


Figure 11. In EBF testing, Cisco UCS C240 M3 servers with elastic cold storage on Red Hat Gluster Storage offered similar performance to reference configurations under streaming load for dense term searches.

CPU percentage vs. concurrency

Figure 12 illustrates CPU percentage plotted against concurrency for the dense term tests using direct-attached storage as well as hybrid storage with elastic cold storage on Red Hat Gluster Storage. In all cases, CPU utilization ramps at about the same rate, and peak CPU utilization is reached at around 16 concurrent searches. These results were consistent with the reference configuration as well. The observation is that for the densest types of searches—such as when searching one in 100 events—each search consumes a full CPU core for the duration of that test. Importantly, using Red Hat Gluster Storage for elastic cold storage did not noticeably impact CPU utilization over that of direct-attached local storage.

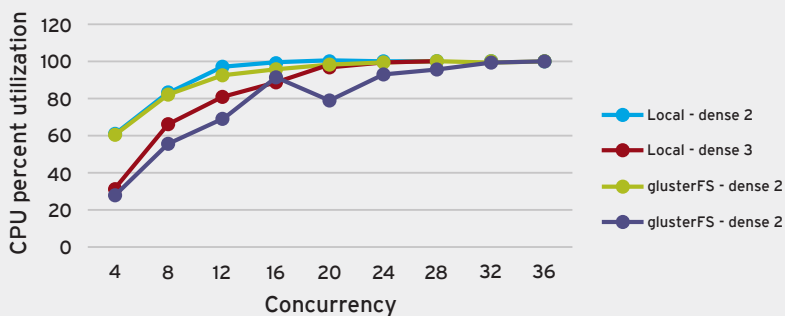


Figure 12. In EBF testing, Cisco UCS C240 M3 servers with elastic cold storage on Red Hat Gluster Storage showed similar CPU utilization profiles to servers with direct-attached storage only.

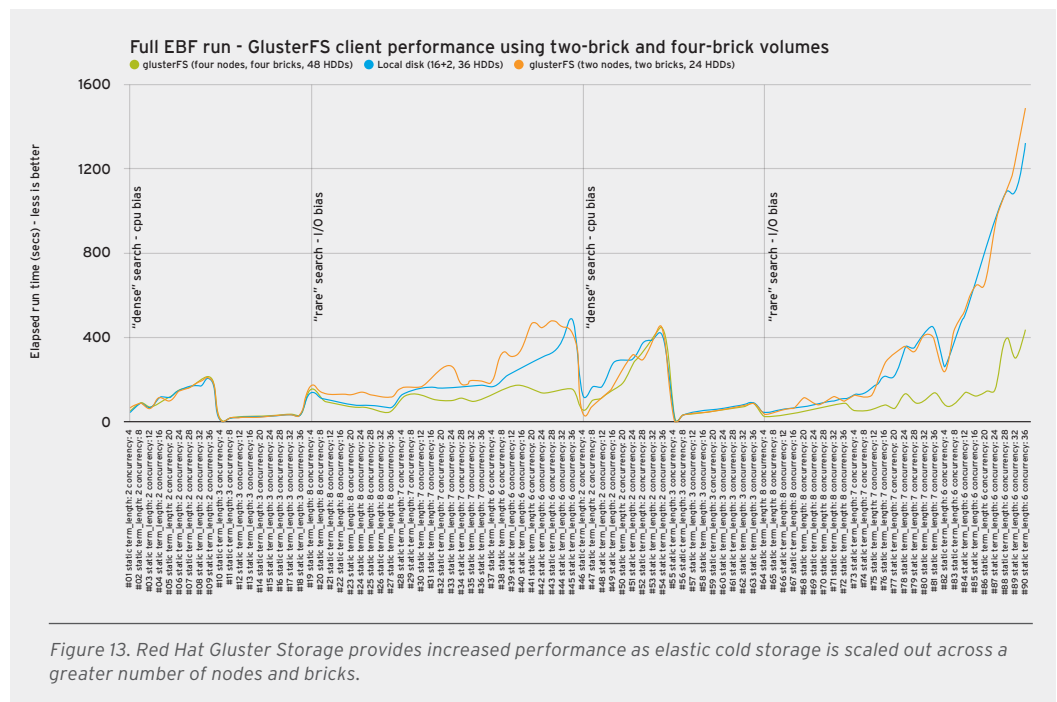
PERFORMANCE TESTING USING SPLUNK BEST PRACTICE RECOMMENDATIONS

Once performance parity was verified between the Splunk Enterprise reference and a hybrid approach using Red Hat Gluster Storage, Red Hat engineers wanted to evaluate scalability to see how cold storage performs as the number of GlusterFS nodes and bricks are scaled out. For this testing

- 15K RPM drives were configured for the indexers for hot and warm dat
- 7.2K RPM drives were configured for the systems running Red Hat Gluster Storage

When deploying Red Hat Gluster Storage for cold storage, organizations can choose either NFS or Gluster Native Client. Gluster Native Client is a POSIX-conformant, FUSE-based client running in user space. Gluster Native Client is the recommended method for accessing volumes when high concurrency and high write performance is required. When reading many small files, the native NFS client may provide better performance. Additionally, IP failover is not required for Gluster Native Client, simplifying deployment.

To examine scalability, testing was conducted using EBF to generate load on the system. Using Gluster Native Client, a full benchmark run was conducted using a GlusterFS volume built from two nodes and two bricks (total of 24 HDDs). A subsequent EBF run was then performed on a scaled-out configuration with a GlusterFS volume built from four nodes and four bricks (for a total of 48 HDDs). The results of this scalability testing clearly demonstrate the scale-out nature of Red Hat Gluster Storage (Figure 13). Adding more capacity through additional Red Hat Gluster Storage nodes increases the CPU, cache, and spindle count. As cold storage scales, data is distributed across more bricks, improving search responsiveness, particularly for rare searches.



OPERATIONAL ACCEPTANCE TESTING

As a part of operational acceptance testing performed by Red Hat, Function1 (a Splunk certified partner) was retained to provide independent verification of results. Function1 verified that Red Hat Gluster Storage outages don't disrupt Splunk Enterprise operations. It likewise confirmed that extending cold storage with Red Hat Gluster Storage is a transparent activity, with no disruption to Splunk Enterprise indexing and search activities.

Red Hat Gluster Storage failover

Under typical operation with hot, warm, and cold data on local file systems, Splunk Enterprise 6.1 manages index replication and failover. Within this solution architecture, Red Hat Gluster Storage provides redundancy and failover for the Splunk cold data through a distributed-replicated GlusterFS volume. This configuration essentially allows the system to lose both an indexer and a Red Hat Gluster Storage node without interruption, and without losing data.

To evaluate stability and resiliency, one of the Red Hat Gluster Storage nodes providing cold storage to the solution was intentionally powered off in an NFS-based configuration. The node was then brought back online and allowed to self-heal. As expected, the Splunk Enterprise configuration continued to operate uninterrupted. A number of conclusions were drawn from the testing:

- During the outage, the two forwarders maintained an ingest rate of approximately 30MB per second to each of the indexers.
- There was little or no impact to Splunk Enterprise during the IP failover process that takes place when a node within a Red Hat Gluster Storage cluster is taken offline during ingest or migration activity.
- Data written to the cluster during the outage period was tracked successfully by the remaining Red Hat Gluster Storage nodes, enabling replica copies to be created once the offline node returned to an operational state.
- The self-heal process restored the environment to full redundancy within minutes of the server being back online.

Red Hat Gluster Storage nondisruptive volume growth

Beyond adding nodes and failing nodes, engineers also wanted to understand the impact of the rebalance process that takes place as a GlusterFS volume is scaled out, either to accommodate more nodes or more bricks. To this end, the EBF test_streaming workload was the test harness, using NFS version 3 between the indexers and the Red Hat Gluster Storage cluster. Cold storage was expanded nondisruptively under typical Splunk Enterprise operating conditions, which include the following:

- Concurrent searches for up to 36 rare search sessions
- Data migration from warm to cold
- Each indexer ingesting data at approximately 30MB-35MB per second

After adding 143GB of new data to the cold storage volume, the volume was grown nondisruptively by adding four more bricks. In addition to the increase in capacity, search responsiveness also improved as data rebalanced across the available spindles (Figure 14).

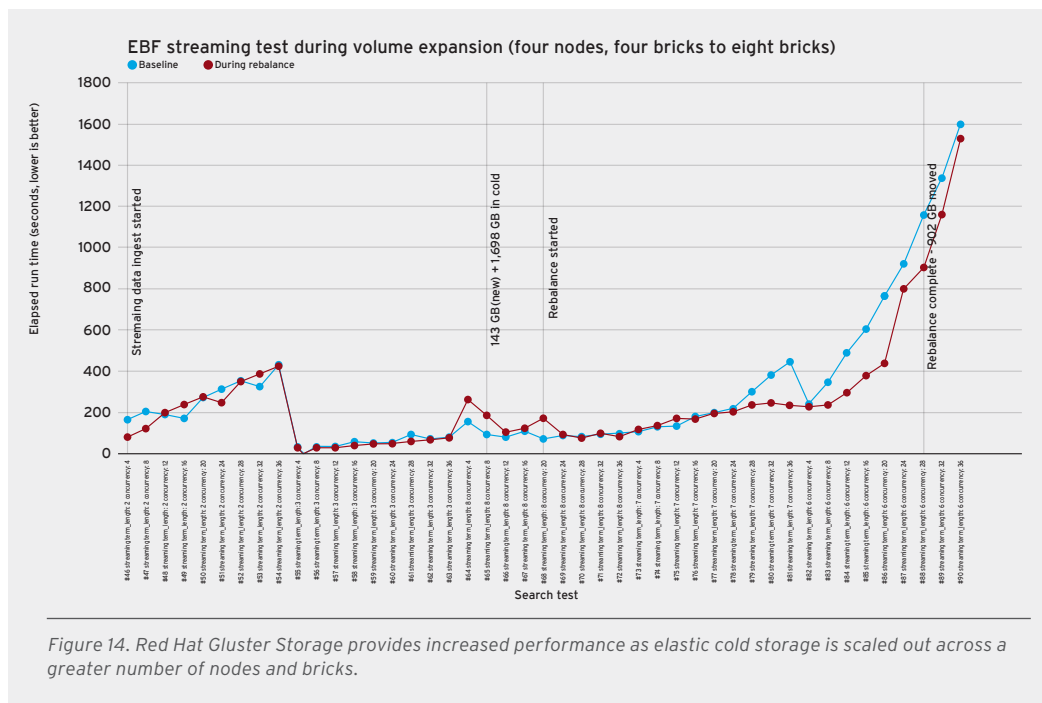


Figure 14. Red Hat Gluster Storage provides increased performance as elastic cold storage is scaled out across a greater number of nodes and bricks.

IMPLEMENTATION GUIDANCE

The sections that follow provide recommended deployment principles as well as general implementation guidance for the solution.

RECOMMENDED DEPLOYMENT PRINCIPLES

To derive the most benefit from a hybrid Splunk Enterprise deployment using Red Hat Gluster Storage for elastic cold storage, the following deployment principles are recommended:

- **Adopt the recommended file system layout.** Red Hat Gluster Storage provides a single volume to all indexers, which simplifies capacity planning and management. However, to exploit the capacity savings offered by the Red Hat add-on package, the layout of the cold storage paths must be done in a specific manner as described by the Red Hat add-on package documentation.
- **Use a Red Hat Gluster Storage distributed-replicated volume.** Both distributed and distributed-replicated volume types are supported for cold storage with Splunk Enterprise. By using a distributed-replicated volume, the solution supports highly available environments and provides a platform that enables rolling upgrades throughout the lifespan of the cold data.
- **Use CTDB to provide highly available NFS.** When choosing to use NFS to connect the Splunk Enterprise indexers to a Red Hat Gluster Storage volume, Clustered Trivial Database (CTDB) should be configured to deliver high availability failover to NFS traffic. Enabling Highly Available NFS is key to Splunk service availability. Each indexer should attach to a separate Red Hat Gluster Storage node to eliminate the potential of CPU-related NFS bottlenecks. To help ensure that virtual IP failover does not lead to multiple indexers hitting the same IP address, consider aligning the virtual IP failover scope to the Red Hat Gluster Storage replication pairs.

- **Reduce capacity usage with the Red Hat add-on package.** Splunk Enterprise supports custom scripts to manage the transition of data from warm storage to the cold storage layer. By using the scripts provided in the Red Hat add-on package, data movement decisions are taken to either copy the data to the GlusterFS volume or to link to existing data on the cold storage volume. Linking to existing primary data reduces capacity requirements, lowering CAPEX/OPEX; brings capacity utilization closer to a native share-nothing Splunk architecture; and helps assure data availability through Red Hat Gluster Storage synchronous replication.
- **Protect the Red Hat Gluster Storage volume with snapshots.** Since moving data to the Red Hat Gluster Storage volume through the add-on package can result in either a copy or a link, the availability of the primary copy of the data is paramount. Although Red Hat Gluster Storage distributed-replicated volumes maintain two copies of all data, snapshots (available with Red Hat Gluster Storage 3.0) provide an additional layer of protection to help insure against malicious or accidental deletion of primary data.
- **Schedule periodic checks of the links within the cold storage volume.** Developing a strategy that validates the links used between Splunk replicated data and primary data stored on Red Hat Gluster Storage is a key deliverable. Scheduling the optional symbolic link checking tool on each indexer enables all links to be validated on a periodic basis.

INSTALLING AND CONFIGURING THE CLUSTER

Detailed instructions for installing and configuring the cluster are beyond the scope of this document. Basic pointers to documentation are provided herein, along with exceptions or special tunings that were done as a part of testing performed by Red Hat engineers. Respective installation manuals and guides can be referenced for more in-depth information.

Installing and configuring Cisco UCS and Red Hat Gluster Storage

In testing performed by Red Hat, Cisco UCS service profile templates were used to define the storage and network configurations. Both Red Hat Gluster Storage and Splunk Enterprise were installed through ISO and Red Hat Package Manager respectively. ISO installations were performed through the Cisco UCS Manager KVM console using virtual media support. See the Red Hat Storage Installation Guide¹² for more information.

LSI RAID configuration

Red Hat Gluster Storage nodes were configured as a RAID 6 10+2 RAID group. Virtual drive settings were configured as follows:

- Write-back enabled
- No read ahead
- Direct I/O
- Stripe size of 256KB
- Disk cache policy disabled

The actual storage configuration was performed with the WebBIOS interface before boot. The StorCLI interface was then used to confirm system state prior to conducting performance runs.

¹² access.redhat.com/documentation/en-US/Red_Hat_Storage/2.0/html/Installation_Guide/chap-Installation_Guide-Install_RHS.html

Kernel tuning

Kernel tuning can help optimize Red Hat Gluster Storage. The `rhs-virtualization` setting was used on all Red Hat Gluster Storage nodes because the number of clients connecting was relatively low in Red Hat testing. A number of `sysctl.conf` updates were performed to optimize the network for high throughput over 10 Gigabit Ethernet. For more information on these settings, reference the SplunkIt Red Hat Gluster Storage testing conducted by Principled Technologies.¹³

The following settings were added to `sysctl.conf` as per the Principled Technology testing:

```
net.ipv4.tcp_timestamps=0
net.ipv4.tcp_sack=1
net.core.netdev_max_backlog=250000
net.core.rmem_max=4194304
net.core.wmem_max=4194304
net.core.rmem_default=4194304
net.core.wmem_default=4194304
net.core.optmem_max=4194304
net.ipv4.tcp_rmem=4096 87380 8388608
net.ipv4.tcp_wmem=4096 65536 4194304
net.ipv4.tcp_low_latency=1
```

Anticipating a heavy NFS network workload, a page fault workaround was applied as follows:

```
vm.min_free_kbytes=2097152
```

Installing Splunk Enterprise 6.1

As a part of testing conducted by Red Hat, the Splunk installation was performed by EBF. EBF pushes out a deployment package over ssh to each of the components that run Splunk Enterprise, with installs occurring locally. Similarly, the indexes used by the test workloads `test_static` and `test_streaming` were created by EBF. The whole environment is encapsulated in a YAML configuration file that the deployment tool uses for configuration. For more typical non-EBF installations, see the Splunk Enterprise Installation Manual.¹⁴

RED HAT GLUSTER STORAGE 3.0 CONFIGURATION

Both Red Hat Gluster Storage 2.1 and 3.0 were utilized in the testing performed by Red Hat. The sections that follow detail the Red Hat Gluster Storage 3.0 configuration.

¹³ principledtechnologies.com/Red%20Hat/RedHatStorage_SplunkIt_0614.pdf

¹⁴ docs.splunk.com/Documentation/Splunk/6.1.3/Installation/InstallonLinux

Brick configuration

Thinly provisioned logical volumes (using the LVM dm-thinp) and XFS file system were used on all bricks in the Red Hat Gluster Storage configuration. Physical volumes and XFS file systems were aligned with the RAID configuration according to established best practices. The `gluster-deploy` tool was used to automate the deployment (forge.gluster.org/gluster-deploy).

The following options were used to mount the brick, as defined through the `gluster-deploy` tool:

```
Allocsize=4096,inode64,logbsize=256K,logbufs=8,noatime
```

CTDB configuration

The CTDB configuration for this solution architecture is based on dm-thinp for the CTDB bricks. Within a `replica 3` volume, discrete partitions were configured for virtual IPs used by the NFS connections, aligning to the replica sets within the volume. This practice is followed to avoid indexers overloading an NFS process on a single Red Hat Gluster Storage node and works as follows:

- Each Splunk Enterprise indexer mounts the Red Hat Gluster Storage volume using a specified virtual IP address.
- Each virtual IP address is managed by the nodes within a replication set.
- IP failover within a replication set provides redundancy and eliminates the potential of two indexers accessing the Red Hat Gluster Storage volume through the same node.

For the full documentation relating to CTDB configuration, refer to the Red Hat Gluster Storage Administration guide.

The following was performed on three of the nodes in the trusted pool:

```
> lvcreate -V 1g -T <vg>/<thin_pool_name> -n ctdb
> mkfs.xfs /dev/<vg>/ctdb
> mkdir -p <brick Path>
```

After updating `/etc/fstab` as appropriate, the following commands were issued:

```
> mount -a
> mkdir -p <brick path>/ctdb
```

Volume configuration

Table 6 defines the volumes created, together with their specific settings. A minimum of four bricks is recommended within the volume. Four or more nodes and bricks are required to deliver optimum performance to search over two local RAID groups on Splunk Enterprise indexers.

TABLE 6. VOLUME CONFIGURATION SETTINGS.

NAME	TYPE	TUNING
SPLUNK	distributed-replicated (replica 2)	small-file-perf group selected nfs.disable off user.cifs: disable network.ping-timeout: 15 storage.owner-uid <splunk> storage.owner-gid <splunk>
CTDB	distributed-replicated (replica 3)	nfs.disable on user.cifs disable network.ping-timeout: 15

ADDITIONAL SUPPORT PACKAGES AND OPTIONAL PACKAGES

The StorCLI package can be obtained by searching for “StorCLI” on the LSI downloads page (lsi.com). If the environment will be monitored with the Red Hat Gluster Storage Splunk Application, the gstatus package will need to be installed on each node. An RPM package of the gstatus package is available from Gluster Community Forge (forge.gluster.org/gstatus).

DOWNLOAD CUSTOM SCRIPTS

To acquire the optional Red Hat add-on package containing the optimized bucket movement and link-checking scripts, contact a Red Hat technical field representative.

CONCLUSION

The combination of Splunk Enterprise and Red Hat Gluster Storage on Cisco UCS C240 M3 servers represents a compelling hybrid Splunk configuration. Testing of this solution architecture has demonstrated that a hybrid solution can provide performant and nondisruptive elastic cold storage scalability for Splunk Enterprise. The Cisco UCS C240 M3 server provides characteristics and flexibility that allow it to serve effectively as both a performance-oriented Splunk indexer and a capacity-oriented node running Red Hat Gluster Storage. Providing scale-out cold storage lets organizations cost-effectively retain their cybersecurity data for longer periods, allowing more effective security analytics and intrusion detection.

ABOUT RED HAT

Red Hat is the world's leading provider of open source solutions, using a community-powered approach to provide reliable and high-performing cloud, virtualization, storage, Linux, and middleware technologies. Red Hat also offers award-winning support, training, and consulting services. Red Hat is an S&P company with more than 80 offices spanning the globe, empowering its customers' businesses.



facebook.com/redhatinc
@redhatnews
linkedin.com/company/red-hat

redhat.com
#12329847_INC0210625_v2_0215

NORTH AMERICA
1 888 REDHAT1

EUROPE, MIDDLE EAST,
AFRICA
00800 7334 2835
europe@redhat.com

ASIA PACIFIC
+65 6490 4200
apac@redhat.com

LATIN AMERICA
+54 11 4329 7300
info-latam@redhat.com