Accelerate Your Splunk Data-to-Everything Platform with Cisco UCS and Scality RING

Contents

Cisco Unified Computing System™ (Cisco UCS®), Splunk, and Scality deliver a scalable unified infrastructure platform for operational intelligence.

Highlights

Cisco® Data Intelligence Platform for modernizing Splunk Enterprise deployments

- Modern disaggregated architecture dramatically reduces the total cost of ownership (TCO) and increases the return on investment (ROI) of your data center infrastructure.
- The proven Cisco Unified Computing System™ (Cisco UCS®) foundation offers complete integration of computing, networking, and storage resources with unified management.
- Splunk Enterprise deployments provide industry-leading performance, capacity, and scalability.
- Cloud-scale architecture offers seamless scalability and performance for the Splunk Data-to-Everything platform, with smart tiered storage powered by Cisco UCS and Scality.
- Independently scale computing, network, and storage resources on demand to manage and derive insights from multiple petabytes (PB) of any data.

Consistent, rapid deployment and out-of-the-box performance with Cisco UCS, Cisco UCS Manager, and Cisco Intersight™ platform

- Cisco UCS Manager simplifies infrastructure deployment with an automated, policy-based mechanism that helps reduce configuration errors and system downtime.
- Gain proven, high-performance, scalable architecture with single- and multiple-rack deployments.
- The Cisco Intersight platform provides intelligent cloud-powered infrastructure management for the Cisco UCS and Cisco HyperFlex™ platforms.

Real-time insights with Splunk Enterprise

- Monitor and analyze data from any source, including customer click streams and transactions, network activity, and call records, turning machine-generated data into business insight.
- Splunk Enterprise provides powerful search, analysis, and visualization capabilities.
- Gain an easy, fast, and secure way to analyze massive streams of data generated by IT systems, security devices, and technical infrastructure.

Splunk SmartStore and Scality RING

- Splunk SmartStore, the latest evolution of distributed scale-out architecture, provides a data management model that brings data closer to computing on demand.
- Splunk SmartStore provides a high degree of computing and storage elasticity and makes cost efficiency easy to achieve, with longer data retention at scale.
- Scality's market-leading software-defined file and object platform is designed for on-premises, hybrid, and multiple-cloud environments.
- Scality RING can be seamlessly deployed on various types of Cisco UCS rack or storage servers with different form factors. This flexibility makes it possible to construct both capacity-optimized and performance-optimized RING structures.

- Scality RING and Splunk SmartStore together render indexers stateless for warm and cold data, boosting operational flexibility and agility and simplifying deployment and management with greater elasticity.
- With Scality RING, warm and cold data no longer needs to be backed up, making recovery-time objective (RTO) and recovery-point objective (RPO) targets easier to meet.
- The extreme resiliency, high availability, and data durability of Scality RING (through erasure coding and geographical distribution) provide peace of mind, keeping data protected and available for search and analytics.

Today's data center needs new tools

Today's data center has evolved into a complex mix of layered and interconnected systems with blended boundaries to support modern applications. When problems arise, finding the root cause and gaining visibility across the infrastructure to proactively identify and prevent outages is a huge challenge. Meanwhile, virtualization and cloud infrastructures introduce additional complexity and create an environment that is more difficult to control and manage.

Traditional tools for managing and monitoring IT and security infrastructure are out of step with the environments they are meant to control because the environment is constantly changing. These tools are inflexible, costly, usually not scalable, and not explicitly designed for the complexity of today's environments and application demands. Designed for individual specific IT functions, traditional tools do not work across multiple data center technologies to help solve problems. When problems arise, these tools typically lack the capability to provide targeted, detailed analysis of IT and security data. Traditional monitoring tools built on relational databases cannot handle the complexity or massive scale of today's machine data.

The Splunk Enterprise advantage

Machine data is one of the fastest-growing and most complex varieties of big data. It is also one of the most valuable, containing a definitive record of user transactions, customer activity, sensor readings, machine behavior, security threats, and fraudulent activity. Splunk Enterprise is the industry-leading platform for big data analytics.

With Splunk Enterprise, you can troubleshoot problems and speed up investigations to just minutes, not hours or days. Splunk Enterprise scales linearly to collect and index petabytes of machine data generated across your entire data center, including cloud, on-premises, and hybrid environments. It enables you to search, monitor, and analyze your data from one place in real time. You can see across your entire infrastructure stack to avoid service degradation and outages. You can get answers from your data with proactive monitoring and real-time visibility into the most complex IT and security systems.

The Scality RING advantage

Splunk SmartStore mitigates the challenges that arise as incoming data sets get larger and retention time grows longer and reins in storage costs. SmartStore introduces a remote storage tier with a native cache manager, allowing data to reside either locally on indexer storage or remotely on Scality RING. Data movement between indexer storage and Scality RING is managed transparently by the cache manager, which resides on the indexer.

With SmartStore, storage is decoupled from computing, allowing these resources to be scaled independently and cost effectively. The resulting TCO reduction is further compounded by higher utilization. SmartStore also reduces storage costs by offloading older, less frequently used data (warm buckets) to Scality RING, without affecting search performance. Scality RING holds master copies of warm buckets, and the indexer local storage is used for hot and cache copies of warm buckets.

With most data residing on Scality RING, the indexer maintains a local cache that contains a minimal amount of data: hot buckets and copies of warm buckets participating in active or recent searches, plus bucket metadata. This approach renders indexers stateless for warm and cold data, boosting operational flexibility and agility.

Cisco Data Intelligence Platform for Splunk Enterprise with Splunk SmartStore

Cisco Data Intelligence Platform is a cloud-scale architecture that brings together big data, artificial intelligence (AI) and computing farms, and storage tiers to work together as a single entity while also allowing them to scale independently to address IT challenges in the modern data center.

Data Intelligence Platform supports today's evolving architecture for big data analytics. It combines a fully scalable infrastructure with centralized management and a fully supported software stack (in partnership with industry leaders) to each of three independently scalable components of the architecture: the data lake, AI and machine-learning (ML) technologies, and object stores.

This general-purpose architecture for data-intensive workloads has been customized for Splunk Enterprise with Splunk SmartStore based on a Splunk Validated Architectures design.

Splunk hot data and cache data are kept on the computing tier and work in conjunction with the Splunk SmartStore, with warm storage on Scality RING. These two tiers can be scaled independently based on demand. This solution is designed to meet a variety of needs for modern Splunk infrastructure, including support for high performance, high capacity, high availability, massive scalability, ease of management, and integration capabilities. With Scality RING, indexes containing large data sets can be retained for longer periods of time at lower costs, with reduced complexity and increased flexibility.

With hybrid-cloud data management capabilities, enterprises can use public cloud services for data stored on RING. This approach supports a new class of hybrid-cloud workflows using applications such as Splunk SmartStore.

Cisco Data Intelligence Platform features

Cisco Data Intelligence Platform offers the following main features:

- Intelligent multidomain management with the Cisco Intersight platform: The Cisco Intersight platform enables IT to operationalize heterogeneous infrastructure and the application platform at scale to seamlessly function as a single cohesive unit through single-pane management.
- Powered by the latest generation of Intel® CPUs: The latest generation of processors from Intel (Cascade Lake Refresh) provides the foundation for powerful data center platforms with an evolutionary leap in agility and scalability.
- Elimination of infrastructure silos with Data Intelligence Platform: Data Intelligence Platform is a highly modular platform that brings big data, AI computing farms, and object storage to work together as a single entity, with each component able to scale independently to address IT challenges in the modern data center.

- Disaggregated architecture: Data Intelligence Platform is a disaggregated architecture that brings together a more integrated and scalable solution for big data analytics and AI. It is specifically designed to improve resource utilization, elasticity, heterogeneity, and failure handling. It can also consume continuously evolving AI and ML frameworks and landscapes.
- Prevalidated and fully supported: Cisco Validated Designs facilitate faster, more reliable, and more predictable customer deployments by providing a blueprint for configuration and integration of all components into a fully working optimized design. Cisco Validated Designs also provide scalability and performance recommendations.

Reference architecture

The reference architectures for the solution include server configurations such as CPU, memory, and I/O subsystems settings configured appropriately to address the specific resource requirements of Splunk Enterprise and Splunk SmartStore. Cisco, Splunk, and Scality together have created this reference architecture to accelerate deployment and reduce risk. Figures 1 and 2 show the solution design.

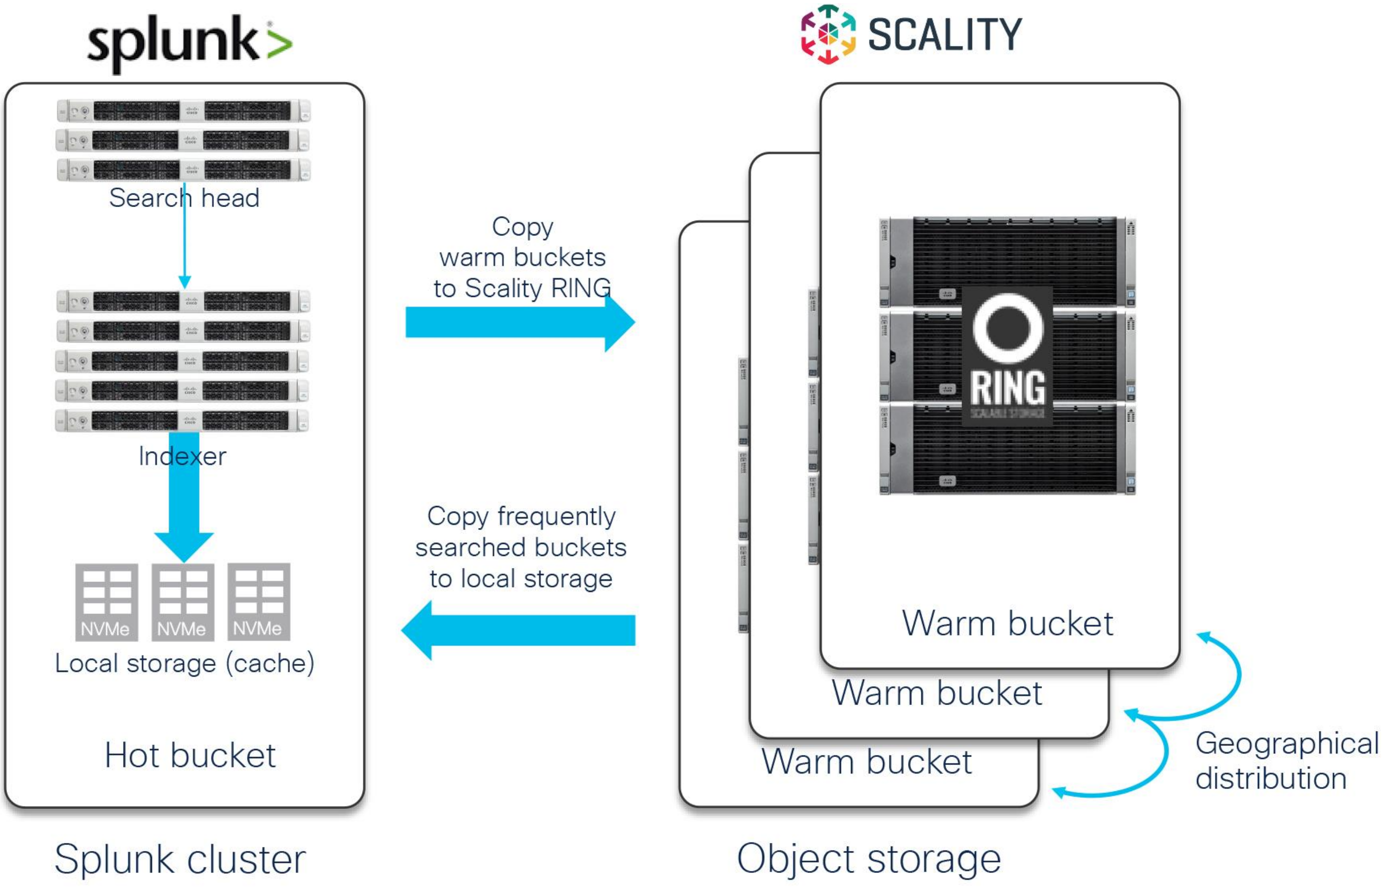

Figure 1.
Splunk Enterprise with Splunk SmartStore on Cisco UCS powered by Scality RING

Configuration tips

Note the following tips when configuring a solution:

- Three or more servers are required for management nodes
- Three or more servers are required for the search heads.
- Eight or more servers are required for indexing (hot tier along with SmartStore cache).
- Three or more servers are required for storage (warm tier).
- The Splunk Enterprise Security application may require a dedicated search head (or cluster).
- Storage capacity and retention are inversely related, and a smaller indexing volume enables a greater retention capacity.
- Splunk premium solutions (such as Splunk Enterprise Security and Splunk IT Services Intelligence) may require greater hardware resources than the reference configuration. Before designing for a Splunk premium solution, you should adjust the configuration accordingly.

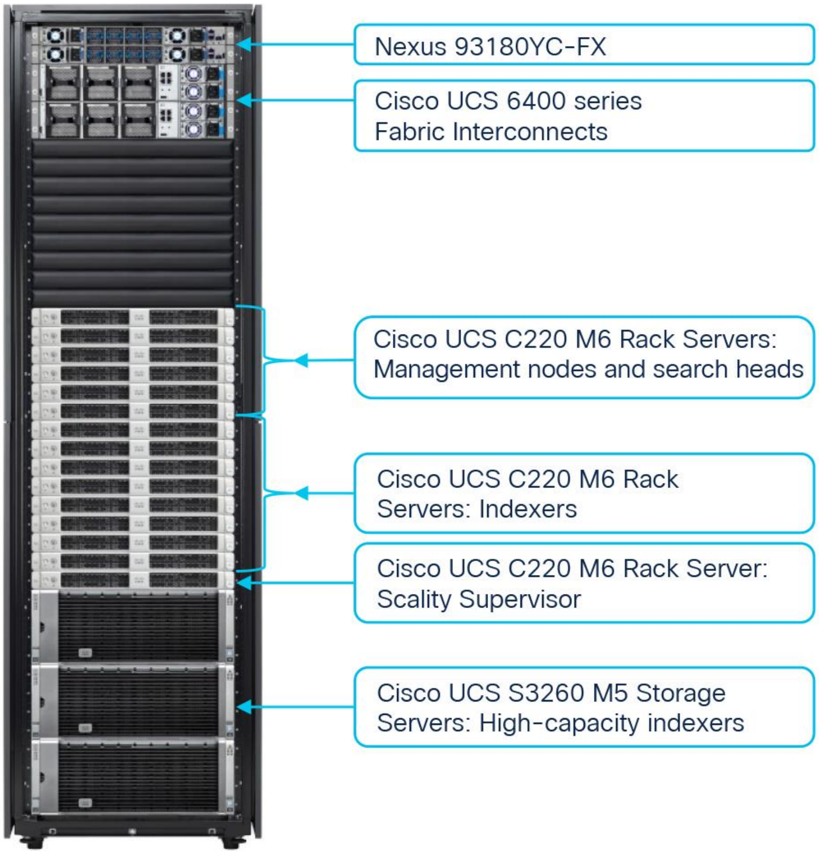
Nexus 93180YC-FX
Cisco UCS 6400 series Fabric Interconnects
Cisco UCS C220 M6 Rack Servers: Management nodes and search heads
Cisco UCS C220 M6 Rack Servers: Indexers
Cisco UCS C220 M6 Rack Server: Scality Supervisor
Cisco UCS S3260 M5 Storage Servers: High-capacity indexers

Figure 2.
Reference architecture for Splunk Enterprise with Splunk SmartStore on Cisco UCS powered by Scality RING

Cisco UCS 6400 Series Fabric Interconnects

Cisco UCS fabric interconnects establish a single point of connectivity and management for the entire system. They provide high-bandwidth, low-latency connectivity for Cisco UCS servers, with integrated, unified management for all connected devices provided by Cisco UCS Manager, which is embedded within each fabric interconnect. Deployed in redundant pairs, Cisco UCS fabric interconnects offer full active-active redundancy, high performance, and the exceptional scalability needed to support the large number of servers that are typical in clusters serving big data applications. Cisco UCS Manager enables rapid and consistent server configuration using Cisco UCS service profiles, advanced health monitoring, and automation of ongoing system maintenance activities across the entire cluster as a single operation.

Cisco UCS C-Series Rack Servers

Cisco UCS C220 M6 [one-rack-unit (1RU)] and Cisco UCS C240 M6 (2RU) Rack Servers support the latest Intel® Xeon® Scalable processor family, up to 3200 MHz of DDR4 memory, and Non-Volatile Memory Express (NVMe) PCI Express (PCIe) solid-state disks (SSDs) with significant I/O performance and efficiency, thereby improving application performance.

The Cisco UCS C240 M6 Rack Server is a 2-socket, 2RU rack server offering industry-leading performance and expandability. It supports a wide range of storage and I/O-intensive infrastructure workloads, from big data and analytics to collaboration.

The Cisco UCS C220 M6 Rack Server (1RU) is among the most versatile general-purpose enterprise infrastructure and application servers in the industry. It is a high-density 2-socket rack server that delivers industry-leading performance and efficiency for a wide range of workloads, including virtualization, collaboration, and bare-metal applications.

Cisco UCS C-Series Rack Servers can be deployed as standalone servers or as part of a Cisco UCS managed environment to take advantage of Cisco's standards-based unified computing innovations that help reduce customers' TCO and increase their business agility.

For more information, refer to the Cisco UCS C-Series Rack Server data sheet.

Cisco UCS S3260 Storage Servers

The Cisco UCS S3260 Storage Server is a high-density modular storage server designed to deliver efficient, industry-leading storage for data-intensive workloads. The Cisco UCS S3260 is a modular chassis with dual server nodes (two servers per chassis) and up to 60 large-form-factor (LFF) drives with up to 18 TB of capacity per HDD in a 4RU form factor. The server uses the latest Intel Xeon Scalable processor family, with up to 40 cores per socket, and supports up to 12 TB of main memory and a range of hard-disk-drive (HDD), SSD, and NVMe options.

The modular Cisco UCS S3260 chassis offers flexibility with more computing, storage, and PCIe expansion on the second slot in the chassis. This second slot can be used for:

- An additional server node
- Four additional LFF HDDs, with up to 18 TB of capacity per HDD
- PCIe expansion tray with up to two x8 half-height, half-width PCIe slots that can use any industry-standard PCIe card, including Fibre Channel and Ethernet cards

The Cisco UCS S3260 chassis includes a Cisco UCS Virtual Interface Card (VIC) 1400 platform chip onboard the system I/O controller, offering high-performance bandwidth with up to 100 Gigabit Ethernet and Fibre Channel over Ethernet (FCoE) interfaces in each system I/O controller.

For more information, refer to the Cisco UCS S-Series Storage Server data sheet.

Reference configuration

Table 1 summarizes the reference configuration for the validated design.

Table 1. Reference architecture

Server role	Server configuration
Splunk management nodes: • Monitoring console • Cluster manager • Deployer • Deployment servers • License manager	3 Cisco UCS C220 M6 Rack Servers, each with: • 2 Intel Xeon Scalable 4310 CPUs (24 cores) at 2.1 GHz • 4 x 32 GB at 2933 MHz (128 GB) • 2 x 240-GB M.2 SSDs for OS with Cisco boot-optimized M.2 RAID controller • 2 x 480-GB SSD (for data) • Cisco 12-Gbps RAID controller with 4-GB flash-based write cache (FBWC) • Cisco UCS VIC 1457
Search heads with search-head clustering	3 Cisco UCS C220 M6 Rack Servers, each with: • 2 Intel Xeon Scalable 5320 CPUs (52 cores) at 2.2 GHz • 8 x 32 GB at 2933 MHz (256 GB) • 2 x 240-GB M.2 SSDs for OS with Cisco boot-optimized M.2 RAID controller • 2 x 480-GB SSD (for data) • Cisco 12-Gbps RAID controller with 4-GB FBWC • Cisco UCS VIC 1457
Indexers[1,2]	8 Cisco UCS C220 M6 or C240 M6 Rack Servers, each with: • 2 Intel Xeon Scalable 5320 CPUs (56 cores) at 2.2 GHz • 8 x 32 GB at 2933 MHz (192 GB) • 2 x 240-GB M.2 SSDs for OS with Cisco boot-optimized M.2 RAID controller • Cisco UCS VIC 1457 • 1 x 1.6-TB NVMe SSD configured as JBOD
Syslog-ng servers and heavy forwarders (optional)	3 Cisco UCS C220 M6 Rack Servers, each with: • 2 Intel Xeon Scalable 5320 CPUs (52 cores) at 2.2 GHz • 4 x 32 GB at 2933 MHz (128 GB) • 2 x 240-GB M.2 SSDs for OS with Cisco boot-optimized M.2 RAID controller • 2 x 480-GB SSD (for data) • Cisco 12-Gbps RAID controller with 4-GB FBWC • Cisco UCS VIC 1457
Storage nodes	3 Cisco UCS S3260 M5 Storage Servers with 2 processing nodes, each with: • 2 Intel Xeon Scalable 4214R CPUs (24 cores) at 2.4 GHz • 6 x 32 GB at 2933 MHz (192 GB) • 2 x 1.6-TB SSDs for OS • Cisco 12-Gbps dual RAID controller with 4-GB FBWC • System I/O controller with Cisco VIC 1455 10- and 25-Gbps quad port • 1 x 2-TB NVMe SSD configured as JBOD • 28 x 12-TB 7200-rpm HDDs configured as individual RAID 0
Storage capacity per indexer[3]	1.6 TB
Total indexer storage	12 TB

Sample cache[1] (IT operations analytics [ITOA])	2.4 TB per day with up to 10 days of hot and SmartStore cache
Sample cache[1] (enterprise security)	800 GB per day with up to 30 days of hot and SmartStore cache
Total usable storage on Scality RING	1.4 PB
Administration node for Scality	1 Cisco UCS C220 M6 Rack Server, with: 2 Intel Xeon Scalable 4310 CPUs (24 cores) at 2.1 GHz 6 x 32 GB at 2933 MHz (192 GB) 2 x 240-GB M.2 SSDs for OS with Cisco boot-optimized M.2 RAID controller 2 x 480-GB SSD (for data) Cisco 12-Gbps RAID controller with 2-GB FBWC Cisco UCS VIC 1457
Connectivity	2 Cisco UCS 64108 Fabric Interconnects (10- and 25-Gbps ports)
Rack space	32 RU

Notes:

1. The suggested maximum indexing capacities per indexer node are up to 300 GB per day for IT operational analytics, up to 200 GB per day for IT services intelligence (ITSI), and up to 100 GB per day for enterprise security.
2. Sample cache durations were calculated with the assumption of 50% compression of original data without any data replication.
3. The total storage capacity per server is based on the unformatted storage. The actual available storage space varies depending on the file system used.

Conclusion: A solution for massive scalability

Splunk Enterprise makes machine data accessible, usable, and valuable for any organization. Cisco Data Intelligence Platform for Splunk Enterprise with Splunk SmartStore powered by Scality RING, with its computing, storage, connectivity, and unified management features, simplifies deployment and offers a dependable, massively scalable integrated infrastructure that delivers predictable performance and high availability for your Splunk Enterprise platform with reduced TCO. Cisco and Scality solve the problem of massive storage with a solution that manages data effectively. Cisco, with Cisco UCS, provides enterprise-class computing, network, and storage infrastructure, building the foundation for the Scality RING storage platform.

Our reference architectures are carefully designed, optimized, and tested with Splunk Enterprise and Scality RING in a clustered distributed search environment to reduce risk and accelerate deployment. These architectures allow you to achieve a high-performance Splunk Enterprise deployment to meet your current needs, and they scale as your needs grow. You can deploy these configurations as is or use them as templates for building custom configurations. The Cisco UCS reference architectures for Splunk Enterprise support the massive scalability that Splunk deployments demand.

For more information

For additional information, refer to the following links:

- For more information about Cisco Data Intelligence Platform, visit https://www.cisco.com/go/bigdata.
- For more information about Cisco UCS, visit https://www.cisco.com/go/ucs.
- For more information about Splunk, visit http://www.splunk.com.
- For more information about Scality, visit http://www.scality.com.
- For more information about the Cisco UCS S3260 Storage Server, visit https://www.cisco.com/go/storage.
- For more information about Cisco's big data validated designs, visit https://www.cisco.com/go/bigdata_design.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **https://www.cisco.com/go/offices**.

Printed in USA 221181.2 07/21