The bridge to possible

Get Answers from Your Data with Cisco UCS Integrated Infrastructure for Splunk Enterprise

Cisco and Splunk deliver a scalable unified infrastructure platform for operational intelligence.

Highlights

Proven platform for operational intelligence

- Based on the fifth-generation of the converged infrastructure for operational intelligence
- Deployed across major industry- specific environments

Cisco UCS Reference Architectures for Splunk Enterprise

- Provides industry-leading performance, capacity, and scalability for Splunk Enterprise deployments
- Designed to scale linearly to handle Multiple Petabytes (PB) of storage

Real-time insights with Splunk Enterprise

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

Cisco Unified Computing System foundation

- Provides unified fabric, unified management, and advanced monitoring capabilities
- Using service profiles, delivers consistent and rapid deployment for out-of-the-box performance

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 consciously 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 issues 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-prem, and hybrid environments. It enables you to search, monitor, and analyze your data from one place in real time. See across your entire infrastructure stack to avoid service degradation and outages. Get answers from your data with proactive monitoring and real-time visibility into the most complex IT and security systems.

Cisco UCS Integrated Infrastructure for Splunk Enterprise

Cisco UCS® Integrated Infrastructure for Splunk Enterprise is based on the fifth generation of industry-leading architectures known as Cisco UCS Integrated Infrastructure for Big Data and Analytics. We designed these solutions to meet a variety of scale-out application demands such as support for high performance, high capacity, high availability, massive scalability, ease of management, and integration capabilities.

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 used for data intensive use cases such as Splunk Enterprise. 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 Rack and Storage Servers

The Cisco UCS C240 M5 Rack Server is a dual-socket, 2-Rack-Unit (2RU) server offering industry-leading performance and expandability for a wide range of storage and I/O-intensive infrastructure workloads, from big data analytics to collaboration. This server uses the latest Intel® Xeon® Refresh Processor Scalable Family with up to 28 cores per socket. It supports up to 24 DDR4 DIMMs for improved performance and lower power consumption. The DIMM slots are also 3D XPoint ready, supporting next-generation nonvolatile memory technology.

The server offers a range of storage options, with up to 26 Small-Form-Factor. (SFF) 2.5-inch drives (with support for up to 10 Non- Volatile Memory Express [NVMe] PCIe Solid-State Disks [SSDs] on the NVMe-optimized chassis version) with a Cisco® 12-Gbps SAS Module RAID Controller. Additionally, it has two modular M.2 cards that can be used for boot. A modular LAN-on-motherboard (mLOM) slot supports the Cisco UCS Virtual Interface Card (VIC) 1457 with four 10/25-Gbps network connectivity.

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 in a 4RU form factor. The server uses the latest Intel® Xeon® Processor Scalable Family with up to 28 cores per socket, and supports up to 1.5 TB of main memory and a range of Hard-Disk-Drive (HDD) and Solid-State Disk (SSD) options.

The Cisco UCS S3260 chassis has 56 top-load LFF HDDs with a maximum capacity of 12 TB per HDD and can be mixed with up to 28 SSDs with maximum capacity of 3.2 TB per SSD. 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 12 TB capacity per HDD
- New 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 quad-port 10/25 Gigabit Ethernet and FCoE interfaces per system I/O controller.

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. Cisco and Splunk together have created reference architectures to accelerate deployment and reduce risk. Figure 1 shows the configurations available.

Figure 1. Two options give you flexibility in performance and capacity

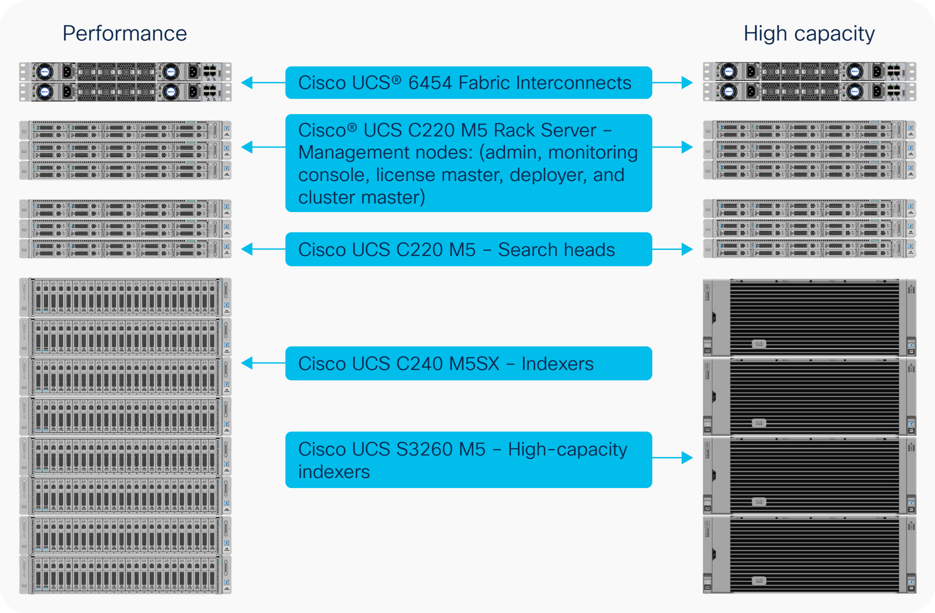
Performance
High capacity
Cisco UCS® 6454 Fabric Interconnects
Cisco® UCS C220 M5 Rack Server – Management nodes: (admin, monitoring console, license master, deployer, and cluster master)
Cisco UCS C220 M5 – Search heads
Cisco UCS C240 M5SX – Indexers
Cisco UCS S3260 M5 – High-capacity indexers

The two Cisco Unified Computing System™ (Cisco UCS) reference architectures for Splunk are based on Cisco UCS Integrated Infrastructure for Big Data and Analytics. The architectures vary in CPU, memory, disk capacity, and performance. Table 1 shows the reference architectures.

Configuration tips

Note the following tips when configuring a solution:

- Three or more servers are required for search-head clusters
- Splunk Enterprise Security applications 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 Service Intelligence) can require greater hardware resources than a reference configuration. Before designing a deployment for a Splunk premium solution, you should adjust the configuration accordingly

Table 1. Reference architectures

Configuration	Performance	High capacity
Search heads	1 to 3 Cisco UCS C220 M5 Rack Servers, each with: ▪ 2 Intel Xeon Processor Scalable Family 5220R CPUs (48 cores) at 2.2 GHz ▪ 6 x 32 GB 2933 MHz (192 GB) ▪ 2 x 240-GB M.2 SSDs for OS with Cisco Boot Optimized M.2 RAID Controller ▪ 4 x 600-GB 10K SAS HDD (OR) 2 x 480-GB or larger SSDs (for data) ▪ Cisco 12-Gbps RAID Controller with 2-GB FBWC ▪ Cisco UCS VIC 1457	
Syslog-ng servers and heavy forwarders (optional)	2 Cisco UCS C220 M5 Rack Servers, each with: ▪ 2 Intel Xeon Processor Scalable Family 5220R CPUs (48 cores) at 2.2 GHz ▪ 6 x 32 GB 2933 MHz (192 GB) ▪ 2 x 240-GB M.2 SSDs for OS with Cisco Boot Optimized M.2 RAID Controller ▪ 4 x 1.2 TB 10K SAS HDD (for data) ▪ Cisco 12-Gbps RAID controller with 2-GB Flash-Based Write Cache (FBWC) ▪ Cisco UCS VIC 1457	

Configuration	Performance	High capacity
Indexers[2,3]	8 Cisco UCS C240 M5 Rack Servers, each with: ▪ 2 Intel Xeon Processor Scalable Family 5220R CPUs (48 cores) at 2.2 GHz ▪ 6 x 32 GB 2933 MHz (192 GB) ▪ 2 x 240-GB M.2 SSDs for OS with Cisco Boot Optimized M.2 RAID Controller ▪ Cisco 12-Gbps RAID Controller with 4-GB FBWC ▪ Cisco UCS VIC 1457 ▪ [1]26 x 960-GB SSDs configured as RAID5	4 Cisco UCS S3260 Storage Servers with 2 processing nodes, each with: ▪ 2 Intel Xeon Processor Scalable Family 5220R CPUs (48 cores) at 2.2 GHz ▪ 6 x 32 GB 2666 MHz (192 GB) ▪ 2 x 480-GB SSDs for boot ▪ Cisco 12-Gbps RAID controller with 4-GB FBWC ▪ System I/O Controller with Cisco VIC 1455 10/25G Quad Port ▪ 8 x 1.6-TB SSDs configured as RAID5 ▪ 20 x 10-TB 7200-rpm HDDs configured as RAID10
Storage capacity per indexer [4]	24 TB	Hot Tier: 11 TB Cold Tier: 100 TB
Total Storage	192 TB	Hot Tier: 88 TB Cold Tier: 800 TB
Sample retention[5] **(IT operations analytics [ITOA]) per indexer**	2.4 TB per day for 5 months	2.4 TB per day for 70 days hot and 22 months cold
Sample retention[5] **(enterprise security) per indexer**	800 GB per day for 15 months	800 GB per day for 7 months of hot and over 5 years of cold
Administration nodes	1 to 3 Cisco UCS C220 M5 Rack Servers, each with: ▪ 2 Intel Xeon Processor Scalable Family 4210 CPUs (20 cores) at 2.1 GHz ▪ 6 x 32 GB 2933 MHz (192 GB) ▪ 2 x 240-GB M.2 SSDs for OS with Cisco Boot Optimized M.2 RAID Controller ▪ 4 x 600-GB 10K SAS HDD (OR) 2 x 480-GB or larger SSDs (for data) ▪ Cisco 12-Gbps RAID Controller with 2-GB FBWC ▪ Cisco UCS VIC 1457	
Connectivity	2 Cisco UCS 6454 Fabric Interconnects (10/25-Gbps ports)	
Rack space	35RU	21RU

Configuration	Performance	High capacity

Notes:

1. Other storage options:
   a. Larger SSDs may be used instead of the 960-GB SSDs.
   b. A combination of SSD for HOT/WARM data and HDDs for COLD data are supported. For example, 6 x 1.6-TB or larger SSDs configured as RAID5 for HOT/WARM data and 20 x 1.8-TB or 2.4-TB 10,000-rpm SAS HDD configured as RAID10 for COLD data (OR) 10 x 800GB SSD-EP configured as RAID5 for HOT/WARM data and 16 x 1.9 SSD-EV configured as RAID5 for COLD data. When HDDs are used in the COLD tier, it is important to configure them as RAID10 - 12 or more HDDs are recommended in this tier.
2. The indexers can be used in standalone or distributed searches. In the distributed architecture, both the indexers and search heads can be configured as clustered or non-clustered. You can scale by adding search heads and indexers to the cluster.
3. 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.
4. The total storage capacity per server is the unformatted available storage space based on the parity used for the RAID group. The actual available storage space varies depending on the file system used.
5. Sample retention durations were calculated with the assumption of 50% compression of original data without any data replication.

Conclusion: A solution for massive scalability

Splunk Enterprise makes machine data accessible, usable, and valuable for any organization. Cisco UCS Integrated Infrastructure for Splunk Enterprise, 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 Total Cost of Ownership (TCO).

Our reference architectures are carefully designed, optimized, and tested with Splunk Enterprise 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 reference architectures described in this document can easily scale to thousands of servers through the use of Cisco Nexus® 9000 Series Switches.

Reference

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 the Cisco UCS S3260 Storage Server, visit https://www.cisco.com/go/storage.

For more information about Cisco UCS Big Data Solutions, visit https://www.cisco.com/go/bigdata.

For more information on Cisco's big data validated designs, visit https://www.cisco.com/go/bigdata_design.