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

© 2022 Cisco and/or its affiliates. All rights reserved.

Contents

Highlights 3

The Splunk Enterprise advantage 4

Cisco UCS Integrated Infrastructures for Splunk Enterprise 4

Reference architecture 6

Conclusion: A solution for massive scalability 9

Reference 10

© 2022 Cisco and/or its affiliates. All rights reserved.

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

Highlights

Proven platform for operational intelligence

- Based on the sixth generation of the unified 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
- Offers 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

Simplify and unify with cloud delivery platform

- The Cisco Intersight™ solution is a cloud-based infrastructure management platform.
- Intersight Managed Mode (IMM) is a new architecture that manages the Cisco Unified Computing System™ (Cisco UCS®) fabric interconnected systems through a Redfish-based standard model.
- You can use all the capabilities of Cisco UCS Manager from the Cisco Intersight platform to manage the platform completely in the cloud.
- Cisco UCS Domain profiles and server profiles aid in rapid deployment by enabling resource management by streamlining policy alignment and server configuration.
- Cisco Intersight™ Cloud Orchestrator (ICO) is a powerful automation tool that enables IT operations teams not only to move at the speed of the business but also to provide a consistent cloud-like experience for users.

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 environments constantly change. These tools are

© 2022 Cisco and/or its affiliates. All rights reserved.

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 problems and reduce 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. 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 Infrastructures 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 that Cisco UCS Manager provides, 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 C-Series Rack Servers

Cisco UCS C220 M6 1-rack-unit (1RU) and Cisco UCS C240 M6 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 is a 2-socket, 2RU rack server that offers 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.

© 2022 Cisco and/or its affiliates. All rights reserved.

The Cisco UCS C220 M6 is a 1RU rack server that is among the most versatile general-purpose enterprise infrastructure and application servers in the industry. This high-density 2-socket rack server 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 are deployable as standalone servers or as part of a Cisco UCS managed environment to take advantage of Cisco standards-based unified computing innovations that help reduce customers' total cost of ownership (TCO) and increase their business agility.

For more information, refer to the Cisco UCS C220/C240 Series Rack Server data sheet.

Cisco UCS X-Series Modular System

The Cisco UCS X-Series Modular System simplifies your data center, adapting to the unpredictable needs of modern applications while also providing for traditional scale-out and enterprise workloads. Powered by the Cisco Intersight cloud-operations platform, it shifts your IT focus from administrative details to business outcomes—with hybrid cloud infrastructure that is assembled from the cloud, shaped to your workloads, and continuously optimized.

Cisco UCS X210c M6 Compute Node

The Cisco UCS X210c M6 Compute Node is the first computing device to integrate into the Cisco UCS X-Series Modular System. Up to eight compute nodes can reside in the 7RU Cisco UCS X9508 chassis, offering one of the highest densities of compute, I/O, and storage per rack unit in the industry. With six large-capacity drives, the Cisco UCS X210c M6 can be used for many workloads that used to require a rack server simply because of the storage requirements.

Cisco Intersight platform

The Cisco Intersight platform manages on-premises and cloud-based infrastructure as one cohesive environment. It is a Software-as-a-Service (SaaS) based hybrid operating model that uses high-level automation to make your hybrid cloud environment simpler, more efficient, and more secure. The platform supports Cisco UCS and Cisco HyperFlex™ hyperconverged infrastructure, other Cisco Intersight connected devices, third-party Intersight connected devices, cloud platforms and services, and other integration endpoints.

Cisco Intersight Cloud Orchestrator

Cisco Intersight Cloud Orchestrator (ICO) is a powerful automation tool that can save time and streamline automation by extending orchestration across any infrastructure and workload. It presents a user-friendly Graphical User Interface (GUI)-based designer that makes it easy to create and execute complex workflows without having coding expertise and allows you to standardize your deployment process with a selection of validated blueprints.

Cisco Intersight Management Mode

Cisco Intersight Managed Mode (IMM) combines the impressive capabilities of the Cisco UCS systems with the cloud-based flexibility and resiliency of the Cisco Intersight platform, streamlining the infrastructure management experience for both standalone and fabric interconnect-attached systems.

© 2022 Cisco and/or its affiliates. All rights reserved.

Reference architecture

The reference architecture for the Splunk solution includes server configurations such as CPU, memory, and I/O subsystems settings configured appropriately to address the specific resource requirements of Splunk Enterprise. This reference architecture is based on traditional Splunk Architecture, which is based on scale-out server configurations. This architecture is designed to scale horizontally to meet the growing needs of data ingest. It can scale to meet any type of performance needs and can be completely managed by the Intersight platform.

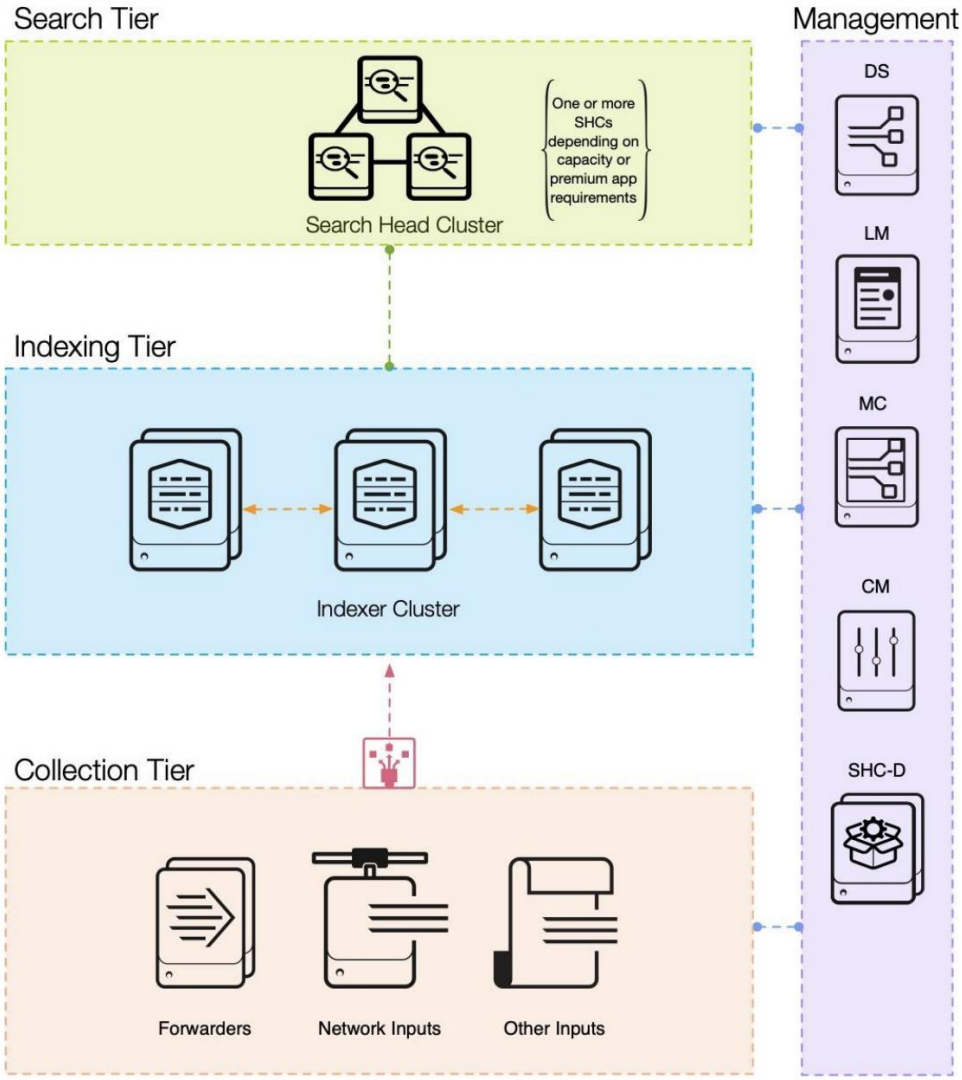

Figure 1.
Splunk reference architecture - Distributed clustered deployment with searchhead cluster

Distributed deployments are designed to separate the index and search functions into dedicated tiers that you can size and scale independently without disrupting the other tier.

© 2022 Cisco and/or its affiliates. All rights reserved.

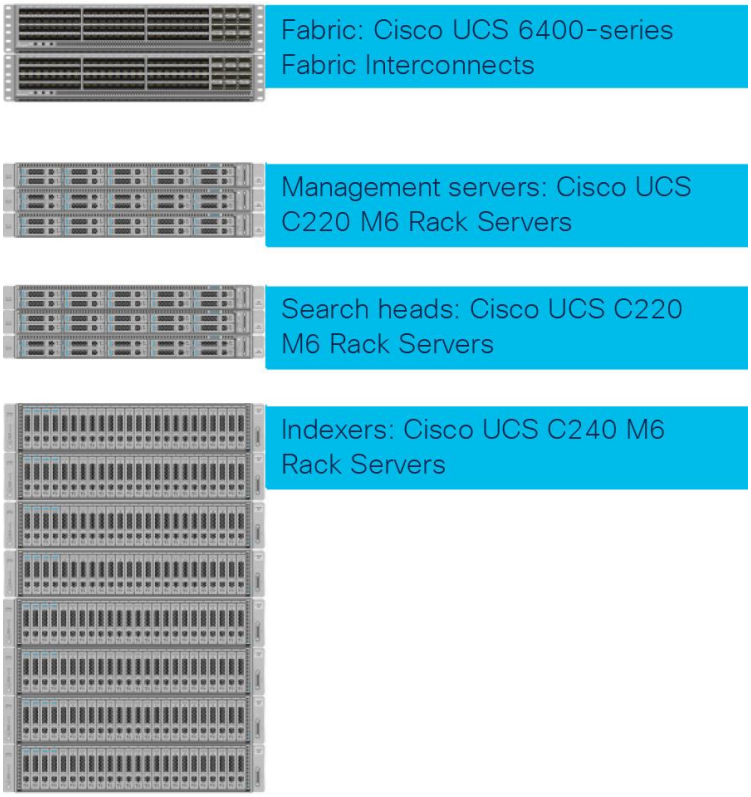

Figure 2.
High-performance scale-out design

Table 1 lists various Splunk roles with their resource allocation.

Table 1. Splunk tiers and resource allocation

	Server type	CPU options	Memory options	Storage	Notes
Indexers	Recommended: Cisco UCS C240 M6 SX Alternate: Cisco UCS C220 M6 SX Cisco UCS X210C	5318Y (24 cores)	256 GB	Boot: Two 240 GB M.2 (Cisco Boot-Optimized M.2 RAID Controller) Hot/warm: NVMe Cold: SATA SSD	RAID 1 – Boot RAID 1 – Hot/warm data RAID 5 – Cold data
Search heads (3 or more)	Recommended: Cisco UCS C220 M6 SX Cisco UCS X210C	5318Y (24 cores)	256 GB 128 GB	Boot: Two 240 GB M.2 Data: Two 480 GB SSDs	RAID 1 – Boot RAID 1 – Data For Splunk Indexers to searchhead ratio, please visit: https://docs.splunk.com/Documentation/Splunk/8.2.3/Capacity/Referencehardware#Ratio_of_indexers_to_search_heads

© 2022 Cisco and/or its affiliates. All rights reserved.

	Server type	CPU options	Memory options	Storage	Notes
Management servers (3 or more)	Recommended: Cisco UCS C220 M6 SX Cisco UCS X210C	4310 (12 cores)	256 GB 128 GB	Boot: Two 240 GB M.2 Data: Two 480GB SSDs	RAID 1 – Boot RAID 1 – Data

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 searchhead (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 2 provides details on the hardware components for each Splunk roles.

Table 2. Hardware components

Configuration	Performance
Search heads	Three Cisco UCS C220 M6 Rack Servers, each with: • Two Intel Xeon processor scalable family 5318Y CPUs (48 cores) at 2.1 GHz • Eight 32 GB 3200 MHz (256 GB) Memory • Two 240 GB M.2 SSDs for OS with Cisco Boot-Optimized M.2 RAID Controller • Two 480 GB SSDs configured as RAID1 • Cisco 12 Gbps RAID Controller with 4 GB Flash-Backed Write Cache (FBWC) • Cisco UCS VIC 1457
Management servers[6]	• Three Cisco UCS C220 M6 Rack Servers, each with: • Two Intel Xeon processor scalable family 4310 CPUs (12 cores) at 2.1 GHz • Eight 32 GB 3200 MHz (256 GB) Memory • Two 240 GB M.2 SSDs for OS with Cisco Boot-Optimized M.2 RAID Controller • Two 480 GB larger SSDs (for data) configured as RAID1 • Cisco 12 Gbps RAID Controller with 4 GB FBWC • Cisco UCS VIC 1457
Indexers[2,3]	Eight Cisco UCS C240 M6 Rack Servers, each with: • Two Intel Xeon processor scalable family 5318Y CPUs (48 cores) at 2.1 GHz • Eight 32 GB 3200 MHz (256 GB) Memory • Two 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 • One 3.2 TB NVMe for hot/warm data • Twentyfive 960 GB SSDs configured as RAID5 for cold data

© 2022 Cisco and/or its affiliates. All rights reserved.

Configuration	Performance
Storage capacity per indexer [4]	Hot/warm: 3.2 TB (RAID 1) Cold: 23 TB (RAID 5)
Total storage	Hot/warm: 25.6 TB Cold: 184 TB
Sample retention[5] (IT Operations Analytics [ITOA]) per indexer	Hot/warm: 20 days Cold: 5 months
Sample retention[5] (enterprise security) per indexer	Hot/warm: 60 days Cold: 15 months

Notes:

1. Other storage options:
   - Larger SSDs may be used instead of the 960 GB SSDs.
   - A combination of SSD for hot/warm data and HDDs for cold data are supported. For example, six 1.6 TB or larger SSDs configured as RAID5 for hot/warm data and twenty 1.8 TB or 2.4 TB 10,000 rpm SAS HDDs configured as RAID10 for cold data (or) ten 800 GB SSD EP configured as RAID5 for hot/warm data and sixteen 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 or more HDDs are recommended in this tier.
2. The indexers can be in standalone or distributed mode. In the distributed architecture, you can configure both the indexers and search heads 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. General recommendation is to use Replication Factor (RF) of 2 and Search Factor (SF) of 2 in both hot/warm and cold tiers.
6. Management servers include all the management roles of Splunk such as cluster master, search head deployer, deployment server, monitoring console, and license master as shown in Figure 1.

Cisco UCS Sizer

As per your average daily amount of data ingestion, you can size your resources using the Cisco UCS Sizer tool. Splunk sizing can get complicated with scale, this tool can help you choose the resources appropriately based on your input accuracy: https://www.cisco.com/go/ucsappsizer

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 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

© 2022 Cisco and/or its affiliates. All rights reserved.

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 with Cisco Nexus® 9000 Series Switches.

Reference

For more information on Cisco UCS, visit https://www.cisco.com/go/ucs

For more information on Splunk, visit http://www.splunk.com

For more information on Cisco UCS big data solutions, visit https://www.cisco.com/go/bigdata

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

For more information on Splunk validated architectures, refer https://www.splunk.com/pdfs/technical-briefs/splunk-validated-architectures.pdf

© 2022 Cisco and/or its affiliates. All rights reserved.