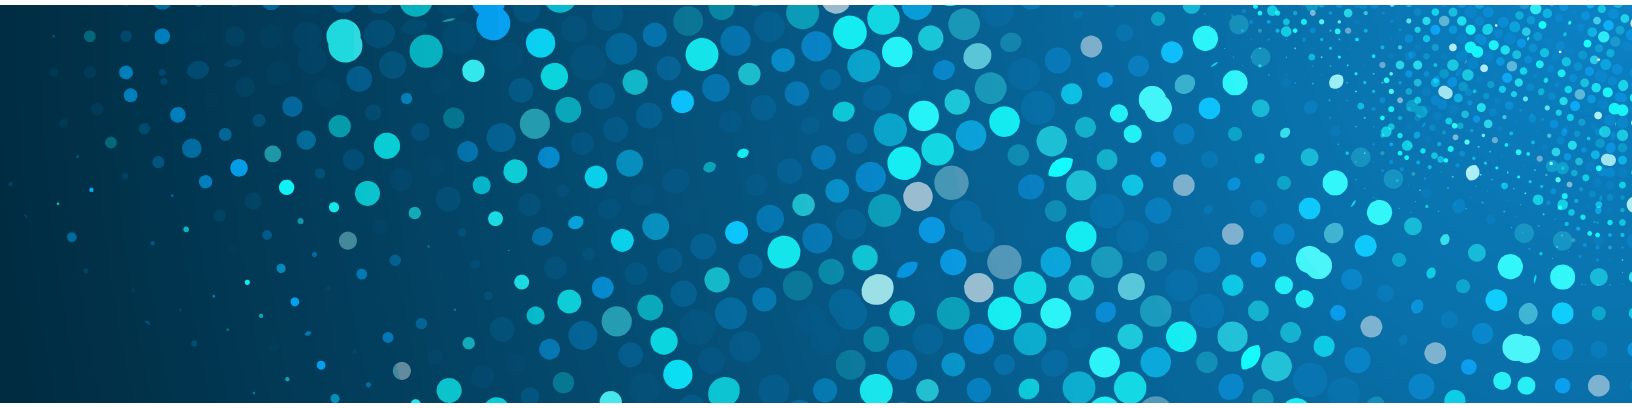

Cisco UCS Integrated Infrastructure for Big Data with the Elastic Stack

Cisco and Elastic deliver a powerful, scalable, and programmable IT operations and security analytics platform with near-real-time search, visualization, and monitoring capabilities.

Organizations today are overwhelmed with data. This data arrives in multiple formats and includes many associated documents, and its huge volume dictates a need for a high-performance analytics platform. Organizations must manage increasing data volume and variety and the need for ever quicker responsiveness to new data with the continuously increasing complexity of the IT infrastructure.

Organizations also are challenged by the need to deploy a fast, reliable, and scalable hardware and powerful and flexible software platform to store and gain actionable insight from large volumes of data. Systems must provide the facility to mine data for valuable information that could augment the business and help make important decisions. The technology used must be able to correlate data and documents from multiple sources to get a full understanding of the health of digital infrastructure and applications.

IT operations analytics is a set of methodologies and processes for collecting, organizing, and identifying patterns in network, computing, and application data across multiple tiers. It also includes the capability to automate routine analyses and health checks through interactive dashboards.

The Elastic Stack for comprehensive operational analytics

The Elastic Stack (previously known as the ELK Stack) is a combination of powerful, open-source software products for logging, storage, search, data visualization, and monitoring. Although each product can be used separately, together they make a powerful, comprehensive, flexible, and programmable IT operational analytics platform that addresses a variety of use cases, such as those listed in Table 1, and puts an organization's data into action.

Table 1 The Elastic Stack use cases

Use case	Industry or vertical market
Operational log analytics: Gain real-time operational insight, reduce Mean Time To Resolution (MTTR), improve customer satisfaction and overall experience, and protect sensitive personal information.	Telecommunications, retail, finance, gaming, e-commerce, research & development, technology, entertainment, and public sector
Security analytics: Protect against cyber threats, secure internal networks and applications, monitor IT activity, empower centralized Security Operations Centers (SOCs) and Security Information and Event Management (SIEM), and detect insider threats and cyber intruders including Advanced Persistent Threats (APTs).	Software, finance, telecommunications, and government
Business analytics: Improve marketing campaigns, detect customer purchasing patterns, analyze customer sentiment, provide recommendations, get credit scores, and analyze medical information.	Healthcare, education, telecommunications, technology, retail, online dating, finance, news, and gaming
Metrics analytics: Monitor sensors, improve manufacturing efficiency, and power large-scale data analysis.	Technology, research, and gaming
Application search: Perform e-commerce searches and media searches, improve research and development, perform real-time plagiarism detection, improve customer conversion and engagement, perform job searches and archive searches, get 360-degree customer views, and perform web searches.	E-commerce, news, research, technology, pharmaceutics, healthcare, retail, government, finance, telecommunications, and education
Enterprise search: Perform document searches, integrate content repositories, and improve internal knowledge portals.	Technology, software, logistics, healthcare, finance, education, and telecommunications

The Elastic Stack consists of the following components:

- Elasticsearch is a distributed search and analytics engine and data store. It can scale from a single node to hundreds of servers and store many terabytes to petabytes of data. It maintains near-real-time query performance by efficiently sharding data and distributing query processing. Its open, JavaScript Object Notation (JSON)–based Representational State Transfer (REST) API makes it easy to use in any environment and provides both full text search and analytics capabilities. Queries include simple term searches, wildcard queries, aggregations including cardinality calculations at scale, and predictive analytics.
- Kibana is an intuitive and visual web-based user interface built on top of the Elasticsearch APIs. As the window into the whole Elastic Stack, it allows organizations to perform impromptu data exploration, build and share visualizations and dashboards, configure and run machine learning jobs, and monitor and manage the full stack.
- Logstash is the data transformation pipeline. It collects data from a multitude of sources and then transforms and outputs the results to Elasticsearch or other sources. More than 200 plug-ins are available, including

input plug-ins for reading data from sources such as text files, message queues, syslog, and various databases; filter plug-ins for field extraction and data enrichment from SQL data sources; and codec plug-ins for formats such as Avro, Comma-Separated Values (CSV), and NetFlow.

- Beats is a collection of lightweight data shippers or agents. It can collect data from log files, general system metrics, and specific metrics from many open-source applications (for example, databases). It can capture network traffic in real time and dozens of other types of data. It is can be deployed across many hosts to capture and send data to Logstash and Elasticsearch.
- ES-Hadoop integrates many Apache Hadoop technologies with Elasticsearch, allowing big data developers to combine and process data from Elasticsearch in native Hadoop technologies such as MapReduce, Cascading, Spark, Pig, Hive, and Storm.
- X-Pack adds critical enterprise features to the Elastic Stack, such as enterprise-class security, unsupervised anomaly detection, alerting to indicate changes in data instantly, constant monitoring, graph-based exploration for behavioral analytics, and PDF reporting.

Cisco UCS Integrated Infrastructure for Elastic

Cisco UCS Integrated Infrastructure for Elastic is based on the fifth generation of the industry-leading converged infrastructure solution: Cisco UCS Integrated Infrastructure for Big Data. Two reference configurations are available:

- The high-performance configuration is suitable for deployments with short retention periods. It offers fast access to data. See Figure 1 and Table 2.
- The performance- and capacity-balanced configuration is suitable for deployments that require a tiered storage model with longer data retention. It provides faster access to frequently accessed data. See Figure 2 and Table 3.

These configurations can be deployed as is or used as templates for building custom configurations. The solution can be customized based on workload demands, including expansion to thousands of servers through the use of Cisco Nexus® 9000 Series Switches. With its extremely fast computing and memory and flexible storage options, this next-generation infrastructure can be used to power fast data access to the large storage resources required for modern applications.

Figure 1 High-performance configuration

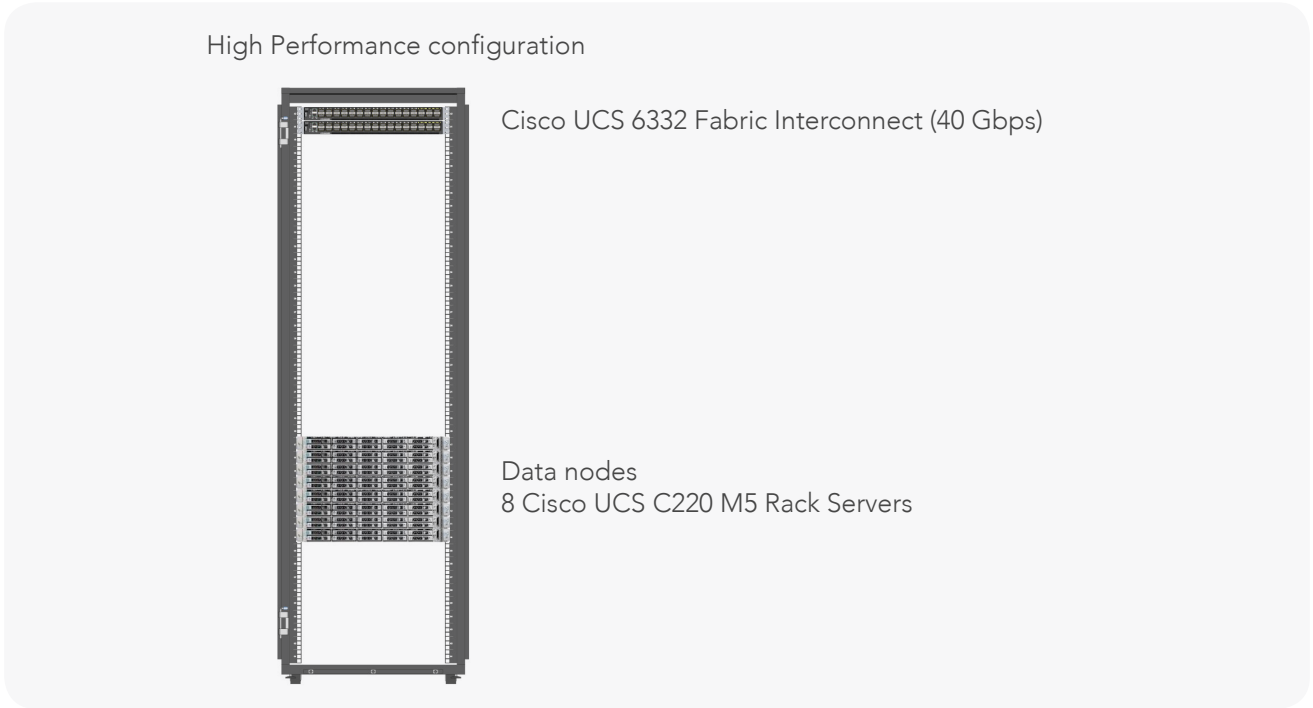

Table 1 Reference configuration 1: High performance

Item	Configuration
Servers	8 x Cisco UCS C220 M5 Rack Servers, each with: • 2 Intel® Xeon® Processor Scalable Family 6132 CPUs (2 x 14 cores at 2.6 GHz) • 192 GB of memory at 2666 MHz (DDR4) • 10 x 1.9-TB Enterprise Value SATA Solid-State Disks (SSDs)[1] configured as RAID 5 • M.2 with 2 x 480-GB SSDs • Cisco® 12-Gbps SAS Modular RAID Controller with 2-GB Flash-Based Write Cache (FBWC) • 40-Gbps card (Cisco UCS Virtual Interface Card [VIC] 1387)
Network connectivity	Cisco UCS 6332 Fabric Interconnect
Storage available[2]	17 TB per server (136 TB total)
Elasticsearch ingest capacity[3]	• More than 2 TB per day • More than 100,000 events per second

Figure 2 Performance- and capacity-balanced configuration

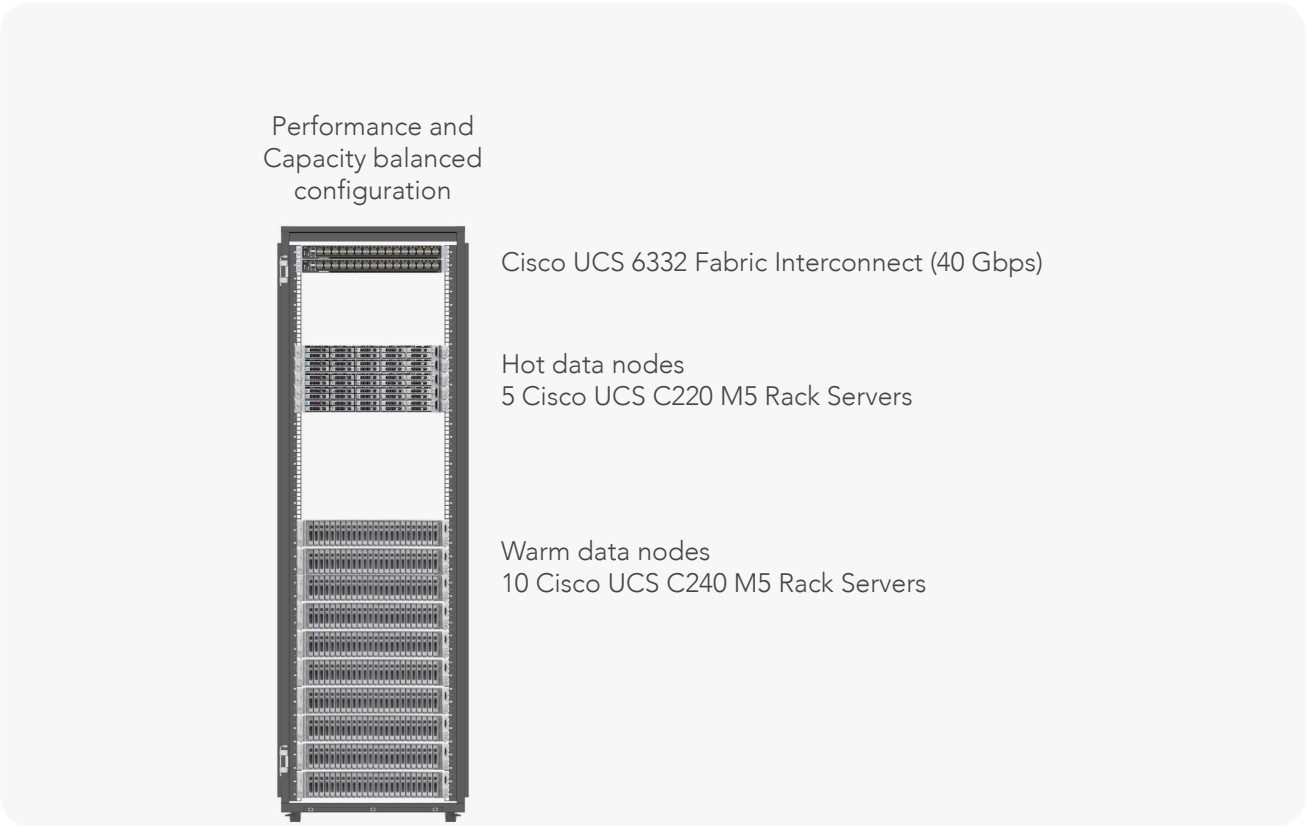

Table 3 Reference configuration 2: performance and capacity balanced

Item	Configuration
Servers (hot tier)	5 x Cisco UCS C220 M5 Rack Servers, each with: • 2 Intel® Xeon® Processor Scalable Family 6132 CPUs (2 x 14 cores at 2.6 GHz) • 192 GB of memory at 2666 MHz (DDR4) • 10 x 960-GB Enterprise Value SATA SSDs[1] configured as RAID 5 • M.2 with 2 x 480-GB SSDs • Cisco 12-Gbps SAS Modular RAID Controller with 2-GB FBWC • 40-Gbps card (Cisco UCS VIC 1387)
Servers (warm tier)	10 x Cisco UCS C240 M5 Rack Servers, each with: • 2 Intel® Xeon® Processor Scalable Family 6132 CPUs (2 x 14 cores at 2.6 GHz) • 192 GB of memory at 2666 MHz (DDR4) • 24 x 1.8-TB 10,000-rpm Hard-Disk Drives (HDDs)[1] configured as RAID 10 • M.2 with 2 x 480-GB SSDs • Cisco 12-Gbps SAS Modular RAID Controller with 4-GB FBWC • 40-Gbps card (Cisco UCS VIC 1387)
Network connectivity	Cisco UCS 6332 Fabric Interconnect
Storage available[2]	• Hot tier: 8 TB per server (40 TB total) • Warm tier: 21 TB per server (210 TB total)
Elasticsearch ingest capacity[3]	• More than 2 TB per day • More than 100,000 events per second

Notes:

1. Other SSD and HDD options are available: 1.6- and 3.8-TB SSDs, and 1.8-TB 10,000-rpm SAS HDD.
2. The total storage capacity per server is the unformatted available storage based on the parity used for the RAID group. The actual available storage space varies depending on the file system used.
3. Actual ingest capacity depends on the retention period of the data, the average event size, and other factors.
4. Each server can host several Elasticsearch data nodes with the X-Pack extension.
5. Elasticsearch master nodes can be deployed on the same servers or in a separate virtual environment.
6. Any Elasticsearch ingest nodes as well as Logstash and Kibana services can be deployed on additional Cisco UCS Integrated Infrastructure.
7. For cold storage, any of the Hadoop solutions based on Cisco UCS Integrated Infrastructure for Big Data can be used with the Elastic ES-Hadoop connector.

Cisco UCS 6300 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 C240 and C220 M5 Rack Servers

The Cisco UCS C240 and C220 M5 Rack Servers are dual-socket servers offering industry-leading performance and expandability for a wide range of storage and I/O-intensive infrastructure workloads, from big data analytics to collaboration. These servers use the new Intel Xeon Scalable processors with up to 28 cores per socket. The servers support 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 2-rack-unit (2RU) Cisco UCS C240 M5 supports up to 26 Small-Form-Factor (SFF) 2.5-inch drives (with support for up to 10 Non-Volatile Memory Express [NVMe] PCIe SSDs on the NVMe-optimized chassis version) with a Cisco 12-Gbps SAS Module RAID Controller.

The 1RU Cisco UCS C220 M5 supports up to 10 SFF 2.5-inch drives (with support for up to 10 NVMe PCIe SSDs on the NVMe-optimized chassis version). 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 VIC 1387 with dual 40-Gbps network connectivity.

Conclusion

Big data technology has become a valuable and compelling asset for organizations of all sizes. Harnessing the value of big data analytics in real time from a stable, manageable, and scalable infrastructure has been a longstanding challenge. The Elastic Stack on Cisco UCS solution meets this challenge by delivering a real-time data analytics and search platform that can be rapidly deployed and scaled on demand and is easy to use. It also reduces the Total Cost of Ownership (TCO) by requiring fewer infrastructure components and reducing operating expenses associated with staff time. Organizations can rely on this solution to provide real-time data intelligence that can help them make critical business decisions.

For more information

- For more information about Cisco UCS, visit https://www.cisco.com/go/ucs.
- For more information about Elastic, visit https://www.elastic.co/.
- For more information about Cisco UCS Big Data Solutions, visit https://www.cisco.com/go/bigdata.
- For more information about Cisco Validated Designs for big data, visit http://www.cisco.com/go/bigdata_design

C22-739529-00 09/17