SEC455: **SIEM Design & Implementation**

2 Day Course	14 CPEs	Laptop Required

You Will Be Able To

- Architect and design a SIEM solution
- Design a SIEM focused on speed and efficiency
- Deploy an open-source SIEM solution meant for enterprise workloads
- Size and use a SIEM based on any budget (from shoestring budgets to unlimited funding)
- Collect and parse logs of any type or source
- Scale log collection, ingest, and search capabilities
- Enrich logs to provide advanced detection as well as context to analysis
- Build a compliance and tactical SIEM, whether a single system or dual stack (multiple SIEMs)
- Know when, why, and how to deploy multiple SIEM solutions and how to integrate them
- Deploy an alert engine and setting up alert rules
- Implement tiered storage with aging policies to handle data retention and disk speeds
- Enhance logs to add context
- Implement searches that do not take coffee breaks to finish
- Know when and when not to augment logs
- Find meaningful log sources and how to automate data collection
- Identify common SIEM deployment pitfalls and hurdles

Course Preview
available at: **sans.org/demo**

Security Information and Event Management (SIEM) can be an extraordinary benefit to an organization's security posture, but understanding and maintaining it can be difficult. Many solutions require complex infrastructure and software that necessitate professional services for installation. The use of professional services can leave security teams feeling as if they do not truly own or understand how their SIEM operates. Combine this situation of complicated solutions with a shortage of available skills, a lack of simple documentation, and the high costs of software and labor, and it is not surprising that deployments often fail to meet expectations. A SIEM can be the most powerful tool a cyber defense team can wield, but only when it is used to its fullest potential. This course is designed to address this problem by demystifying SIEMs and simplifying the process of implementing a solution that is usable, scalable, and simple to maintain.

The goal of this course is to teach students how to build a SIEM from the ground up using the Elastic Stack. Throughout the course, students will learn about the required stages of log collection. We will cover endpoint agent selection, logging formats, parsing, enrichment, storage, and alerting, and we will combine these components to make a flexible, high-performance SIEM solution. Using this approach empowers SIEM engineers and analysts to understand the complete system, make the best use of technology purchases, and supplement current underperforming deployments. This process allows organizations to save money on professional services, increase the efficiency of internal labor, and develop a nimbler solution than many existing deployments. For example, many organizations pay thousands of dollars in consulting fees when a unique log source needs a custom parser. This course will train students how to easily parse any log source without requiring consulting services, saving their organizations both time and money, and facilitating faster collection and use of new log sources.

SEC455 serves as an important primer to those who are unfamiliar with the architecture of an Elastic-based SIEM. Students that have taken or plan to take additional cyber defense courses may find SEC455 to be a helpful supplement to the advanced concepts they will encounter in courses such as SEC555. In addition, the material discussed in this course will enable students to not only build a new SIEM, but improve and supplement their already existing implementations, producing a more efficient solution that provides the answers they need more quickly and at a lower cost. The overall goal is to educate students on what they need to know to design and modify a SIEM, improve upon their current solution, and enable them to reach their original defensive goal - catching adversary activity in their environment.

"SEC455 has made me rethink how I do event log monitoring and what I can do to improve."

— Roger Christopher, **Bureau of Land Management**

Available Training Formats

Live Training

Live Events
sans.org/information-security-training/by-location/all

Summit Events
sans.org/cyber-security-summit

Online Training

OnDemand
sans.org/ondemand

Simulcast
sans.org/simulcast

Section Descriptions

SECTION 1: Distributed Search and Visualization

Day one focuses on Elasticsearch and Kibana and will take students on a journey from their first steps in the Elastic stack, to having a secured and production-ready Elasticsearch and Kibana instance by the end of the section. Students will learn the skills required to install, configure, and use Elasticsearch, and will become comfortable with using Kibana to visualize imported data in multiple useful ways.

The section begins with an introduction to the components of a SIEM and how each relates to the pieces of the Elastic stack. After a quick, high-level view, Elasticsearch receives a deep dive with a focus on the core practical concepts of node types, indexes, shards, and data type mapping. Also, administrative activities such as cluster creation, management, data retention and optimization are covered and put into practice with hands-on labs. Through these activities, students will become comfortable creating, modifying, and managing their Elasticsearch cluster. The Elasticsearch lesson also includes recommendations and calculations to ensure the capacity of the cluster meets storage and event-per-second requirements.

The second part of the section features a similar deep dive on how to install, setup, and use Kibana. Students will become familiar with the search, visualization, and dashboard interfaces, and will learn how to use these tools to explore log data. Also, students will learn how to secure access to their Elastic stack and to lock down indexes and documents with role-based permission schemes.

TOPICS: What is ELK?; Elasticsearch; Kibana; Securing the Elastic Stack

SECTION 2: Enriching and Managing Logs

Building on the infrastructure prepared during day one, day two focuses on how to efficiently move logs from your edge devices, and then transport, parse, and enrich them. Any organization can create an enormous amount of log events in a short period, so the creation of an efficient and dependable pipeline is crucial to maintaining the integrity and stability of any logging solution. The multitude of log formats and transport protocols are discussed, as well as how to decide on the best configuration for any given situation. Traditionally, log parsing has been painful and full of potential error, but the techniques shown throughout this course day will reduce or eliminate this pain and teach students how to substitute legacy solutions with more modern and efficient solutions. By the end of day 2, students will be familiar with optimal logging formats and with new and effective ways to parse legacy or difficult-to-handle formats.

While having perfectly parsed logs is great on its own, we can go much further. The value of a parsed log can be improved hundreds of times over with proper enrichment and with nominal performance impacts on log ingestion rates. Log enrichment includes adding context to logs and various other techniques used to increase your detection capabilities. Additionally, conditional logic and strategies for log filtering are discussed to ensure that the system will not be slowed by processing unneeded information.

The final piece of SIEM architecture is collecting logs off edge devices. Many organizations are unwilling or unable to deploy agent-based log collection, so both agent and agentless methods of log collection are discussed so that students can identify their ideal deployment. Although many students may already have a SIEM system in their environment, the Elastic set of tools can also be used to further supplement and improve the performance of other commercial SIEMs. We'll explain new trends such as the dual-stack SIEM environment, and examine how to use Logstash to supplement pre-existing SIEM deployments that struggle with high volume issues and poor data enrichment features. Alerting based on logs is also covered, with a review of both Elastic and third-party solutions.

TOPICS: Log Aggregation; Traditional Parsing; Modern Parsing; Log Enrichment; Agents and Log Collection; Third-Party Integration and Dual-Stack SIEM; Alerting

Author Statement

Security professionals who want to "After seeing Elasticsearch continue to pop up in SANS courses across the curriculum, I have noticed students are consistently curious and excited by the search features the open-source Elastic Stack provides. Numerous security tools, projects, and even commercial SIEMs have moved to using the lightning-fast distributed search tool as the cornerstone of their functionality. The trend is clear - Elasticsearch is emerging as a great solution to the "needle in a haystack" problems we often face in information security, and its inclusion in professional products shows it is indeed ready for primetime. Elasticsearch has an enormous number of possible uses, however, many of which are considerably different than the security use case. Understanding which features are important for a specific use is not a simple task given the extensive documentation. Considering the rapid pace of development throughout the past few years, much of the existing information online has rapidly become outdated as the software has changed. SEC455 was written from day 1 with an eye toward the future using only the newest version of Elasticsearch in mind. We have reduced the documentation to the most important information and simplified learning the Elastic Stack to the items relevant for security use. Taking this class is guaranteed to save you numerous hours documentation reading, experimentation, and frustration, and will give you a shortcut to the front of the Elasticsearch trend. If you're wondering what the distributed search platform can do for you, and want to learn it with a focus on understanding, improving, and building a SIEM, this is the course for you!"

— John Hubbard