SEC545:™ GenAI and LLM Application Security™

1	7	Laptop
Day Course	CPEs	Required

You Will Be Able To

- Understand key concepts and terminologies—gain a deep understanding of GenAI, LLM architectures, and their application in real-world scenarios
- Explore various models and tools—examine the types of models and tools available for building and deploying GenAI applications
- Explore fine-tuning and customization—learn how to fine-tune and customize models for specific use cases
- Assess risks and mitigation strategies—identify security risks unique to GenAI applications and explore effective mitigation techniques
- Secure RAG, embeddings, and vector databases—understand retrieval-augmented generation (RAG), embeddings, and VectorDB, and how to securely configure different components
- Explore operations and security controls—explore the operational aspects of building and deploying GenAI applications and learn about the relevant security controls
- Compare hosting options—understand the various GenAI hosting options and their differences from a security perspective
- Leverage cloud security controls—learn about the security controls offered by cloud providers for LLM hosting services
- Explore GenAI adjacent technologies—examine technologies such as LangChain and agents, and understand the security risks they introduce
- Integrate GenAI into security frameworks: Learn how to build or integrate GenAI security practices into existing organizational security frameworks

Business Takeaways:

- Understand GenAI applications
- Identify potential security risks associated with GenAI applications
- Learn how to mitigate GenAI security risks effectively

SEC545 training provides an in-depth exploration of GenAI technologies, starting with core principles and underlying technologies. It will assess security risks by identifying and analyzing real-world threats impacting GenAI applications. As students progress, they will learn to establish security best practices by exploring different measures for securing GenAI applications effectively.

The course begins with a brief introduction to the fundamentals of GenAI, covering key concepts and terminologies such as Large Language Models (LLMs), embeddings, and Retrieval-Augmented Generation (RAG). It then examines the security risks associated with GenAI, including prompt injection attacks, malicious models, and third-party supply chain vulnerabilities. Following this, the course dives into the essential components needed to build a secure GenAI application, including coverage of vector databases, LangChain, and AI agents. The course concludes with a comprehensive overview of hosting GenAI applications, discussing options for local deployment, cloud solutions, and platforms like AWS Bedrock.

Hands-On GenAI and LLM Application Security Training

This course covers essential GenAI concepts, technologies, and security risks, featuring hands-on labs designed to illustrate how attackers can exploit specific vulnerabilities and the strategies to mitigate them. Throughout the course, we will reference the OWASP Top 10 for LLMs and Generative AI applications to highlight prevalent security issues and their solutions.

The labs are conducted using a chat application hosted on AWS EKS, with agents capable of assisting in evaluating resumes for potential candidates and interacting with various AWS infrastructure components. The app also integrates with a Weaviate Vector Database to demonstrate both attack scenarios and defense mechanisms. Participants will work with multiple LLM providers, including OpenAI, a locally hosted Llama 3.2, and AWS Bedrock, providing a comprehensive hands-on experience in securing GenAI applications.

Author Statement

Emerging technologies often bring substantial value, transforming industries and opening new possibilities. However, their rapid adoption also introduces complex risks that are frequently not fully understood at the outset. As these technologies evolve, the nature and scale of associated risks can shift in unexpected ways, making it challenging to anticipate their full impact. This pattern has been clear with technologies like cloud computing, where the pace of innovation often surpasses our understanding of its security implications. The greater the potential of a technology, the more complex its associated risks.

AI, particularly generative AI, represents the next major wave of transformation, with the potential to reshape nearly every application. This course aims to deepen students' understanding of GenAI and its security challenges, equipping them with the skills to proactively manage and mitigate these risks.

As the industry evolves, so will this course, ensuring that our approach to securing GenAI applications remains at the forefront.

—Ahmed Abugharbia

Section Descriptions

GenAI and LLMs Application Security

The course begins with an introduction to Generative AI (GenAI) concepts, including General AI, Large Language Models (LLMs), Vector Databases, and Embeddings. Students will explore prompt injection techniques and their implications.

Next, the focus shifts to security risks associated with GenAI, such as prompt and instruction poisoning, as well as malicious models. Students will learn how to compromise Vector Databases and corrupt data.

The course then delves into GenAI application architecture, covering LLM customization, integration with tools like LangChain and AI agents, and prompt engineering. Students will gain hands-on experience in compromising the LLM supply chain and pivoting to other components within the infrastructure.

The course concludes with a discussion on hosting GenAI applications, utilizing local models, OpenAI, AWS Bedrock, and Hugging Face. This practical experience equips students with essential skills for securing GenAI applications and deployment.

TOPICS: GenAI Introduction and Concepts; Augmenting GenAI Knowledge; Hosting GenAI applications; GenAI Application Architecture

Who Should Attend

- **Application security engineers:** Professionals seeking to understand how LLMs and GenAI components impact traditional applications or differ from them, and learn about the unique security challenges posed by GenAI and tools available to secure the entire GenAI application lifecycle
- **Cloud security engineers:** Cloud professionals who need to understand how hosting GenAI applications affects their security posture and how to identify new risks, learn mitigation techniques, and apply security controls to protect GenAI workloads in the cloud
- **SOC analysts, incident handlers, and threat intelligence professionals:** Experts responsible for monitoring, investigating, and hunting threats—they need to understand GenAI components, hosting environments, and related systems, and the ability to analyze GenAI logs and alerts to detect anomalies and conduct thorough investigations
- **Security professionals:** Experts responsible for securing an organization's network and infrastructure—they need to understand GenAI's impact within their environment, the systems it interacts with, and related risks and vulnerabilities
- **Security auditors, compliance, and risk managers:** Professionals focused on understanding the risks of GenAI adoption, assessing their impact, and developing risk management strategies, integrating GenAI risks into current auditing, compliance, and risk frameworks

Prerequisites

- Familiarity with Linux command shells and associated commands
- Familiarity Python and Bash scripting
- Basic understanding of common application attacks and vulnerabilities