SEC545: **GenAI and LLM Application Security™**

5 Day Program	30 CPEs	Laptop Required

You Will Be Able To

- Understand GenAI and LLMs, and examine LangChain agents and MCP
- Explore models, tools, fine-tuning, and customization options
- Identify GenAI-specific risks and mitigations
- Secure RAG pipelines, embeddings, and vector databases
- Apply security controls in GenAI operations
- Compare secure hosting, deployment, and cloud options
- Run Threat Modeling exercises for AI applications and infrastructure
- Incorporate security throughout data pipelines, model training, and deployment
- Integrate GenAI into existing security frameworks

Business Takeaways

- Understand GenAI applications
- Identify potential risks associated with GenAI applications and MLOps tools
- Learn how to mitigate GenAI applications and Infrastructure risks effectively
- Learn to implement end-to-end security for GenAI applications, infrastructure, and operations

SEC545 training explores GenAI security, from core concepts like LLMs and RAG to real-world risks like prompt injection and supply chain threats. Students learn to build, secure, and deploy GenAI apps using best practices for tools like LangChain, agents, and cloud platforms such as AWS Bedrock.

Hands-On GenAI and LLM Application Security Training

With the anticipated transformative impact of Generative AI (GenAI) on industries and technologies, the need for robust security practices to address its risks has never been more critical. SEC545 training equips students with the necessary knowledge to secure GenAI applications.

SEC545 training offers a comprehensive exploration of GenAI technologies, starting with foundational principles and underlying frameworks. Students rigorously evaluate security risks by identifying and analyzing real-world threats affecting GenAI applications, and will progressively learn to implement security best practices by exploring strategies to safeguard GenAI systems effectively.

By the end of this training, students will possess a holistic understanding of GenAI security, empowering them to design, deploy, and defend GenAI systems in a rapidly evolving technological landscape.

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. Organizations are increasingly focusing on the full lifecycle of AI applications, including data pipelines, model training, deployment, and overall MLOps. As a result, security leaders must expand their focus beyond traditional application and system security to also encompass these operational areas, ensuring that the entire AI workflow is protected.

"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, infrastructure and operations remains at the forefront."

—Ahmed Abugharbia

Section Descriptions

SECTION 1: GenAI, Large Language Models, and Security Risks

The course starts with GenAI fundamentals, covering key concepts like large language models (LLMs), embeddings, and retrieval-augmented generation (RAG). Students will explore security risks unique to GenAI, including prompt injection, malicious models, and third-party supply chain vulnerabilities.

TOPICS: GenAI Introduction and Concepts; Fine-Tuning Models; Augmenting GenAI Knowledge; Safe Use and Moderation

SECTION 2: Securing GenAI Applications

Building on Section 1, students explore core components for GenAI apps, like vector databases, LangChain, and AI agents. Section 2 also covers deployment strategies, comparing cloud and on-premises setups with a focus on the security risks unique to each. Section 2 concludes by introducing agents communication protocols such as MCP.

TOPICS: AI Agents; GenAI Applications Architecture; AI Development Frameworks Security; Agents Communication Protocols

SECTION 3: Agentic AI Security

In the third section, students continue exploring MCP security before diving into Transformers, the foundational technology behind LLMs. Explaining the Mathematical Foundations of Predictive Modeling. The section then examines hosting options for AI applications and their associated security considerations. It concludes with a discussion of data orchestration pipelines and related security risks, including tools such as Airflow.

TOPICS: MCP Attacks and OAuth Security; Transformer Architecture Fundamentals; Hosting GenAI Applications; Data Workflow Orchestration

SECTION 4: MLSecOps and Securing GenAI Applications Lifecycle

Section 4 focuses on MLOps operations and integrating security controls throughout the pipelines. It begins by introducing MLOps and its relationship to DevOps, then discusses model-specific attacks such as model serialization vulnerabilities and backdooring models with malicious code. The section covers how to secure pipelines by implementing controls such as model signatures and automated scanning. It concludes with a discussion of AI threat modeling, where students use the infrastructure they have built to perform a hands-on threat modeling exercise using the MAESTRO framework.

TOPICS: Machine Learning Ops (MLOps); Hosting Models; MLSecOps; AI Threat Modeling

SECTION 5: AI for Security

Section 5 focuses on two major components. The first covers using AI for security, where students leverage the infrastructure they have built to conduct investigations and perform threat hunting using AI. The second component is a Capture-the-Flag (CTF) exercise, where students apply everything they have learned over the previous four days to identify and remediate security issues within their AI infrastructure. This infrastructure includes technologies they are likely to encounter in real life, such as Kubernetes, cloud environments, containers using Docker Compose, MCP servers, Airflow, SageMaker, AWS Bedrock, and AWS as a cloud provider.

TOPICS: Incident handling and Investigation with AI; Threat hunting with AI

Who Should Attend

- **Application security engineers:** Professionals seeking to understand how LLMs and GenAI components impact traditional applications or differ from them. They are interested in learning about the unique security challenges posed by GenAI and exploring the tools available to secure the entire GenAI application lifecycle.
- **Cloud security engineers:** Cloud professionals who need to understand how hosting GenAI applications impacts their security posture. They want to identify the new risks introduced by these applications, learn how to mitigate them, and explore the security controls that can be applied to protect GenAI workloads in the cloud.
- **SOC analysts, incident handlers, and threat intelligence professionals:** Professionals responsible for monitoring, investigating security alerts, and hunting for threats. They need to understand the components of GenAI applications, how these applications are hosted, and the internal and external systems they interact with. Additionally, they must be able to analyze logs and alerts generated by various GenAI components, detect anomalies, and conduct thorough investigations.
- **Security professionals:** Experts responsible for securing an organization's network and infrastructure. They need to understand the impact of GenAI applications within their environment, the internal systems these applications may interact with, and the potential risks and vulnerabilities from a security perspective.
- **Security auditors, compliance, and risk managers:** Professionals focused on understanding the risks introduced by adopting GenAI, assessing their potential impact, and developing effective risk management strategies. They also need to learn how to integrate GenAI-related risks into their current auditing, compliance, and risk management frameworks.

Prerequisites

- Familiarity with Linux command shells and associated commands
- Familiarity with Python and Bash scripting
- Basic understanding of common application attacks and vulnerabilities
- Basic knowledge of public cloud services and cloud-native technologies, including Kubernetes