SEC545:™ GenAI and LLM Application Security™

3 Day Program	18 CPEs	Laptop 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 focuses on understanding the security risks associated with Generative AI (GenAI) applications and implementing security controls throughout their lifecycle—from development to hosting and deployment. The course begins with an introduction to core GenAI concepts, covering popular tools and vendors. It then explores specific topics such as large language models (LLMs), agents, retrieval-augmented generation (RAG), and best practices for hosting GenAI applications. Security controls and risk mitigation strategies are examined at each stage. The course concludes with guidance on establishing a GenAI security practice or integrating it into existing security frameworks.

The course begins with an 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 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 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. It rigorously evaluates security risks by identifying and analyzing real-world threats affecting GenAI applications. Students 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. 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

SECTION 1: GenAI, Large Language Models (LLMs), and Security Risks

This course begins with a thorough introduction to GenAI fundamentals, covering essential concepts such as Large Language Models (LLMs), embeddings, and Retrieval-Augmented Generation (RAG). Students will dive into the security risks unique to GenAI, including prompt injection attacks, malicious model manipulation, and vulnerabilities within third-party supply chains.

EXERCISES:

- Lab 1.1: LLMs and Prompt Injection
- Lab 1.2: Fine-tuning LLMs
- Lab 1.3: Compromising Vector Database
- Lab 1.4: Moderation

TOPICS:

- GenAI Introduction and Concepts
  - General AI and Generative AI
  - Large Language Models (LLMs)
  - Retrieval-Augmented Generation (RAG)
  - GenAI Application Components Security
  - Prompt Injection
- Fine-Tuning Models:
  - OpenAI fine-tuning
  - File-tuning risks and models' access
  - Augmenting GenAI Knowledge
- Vector Databases
  - Knowledge Sources
  - Poisoning Data Sources
  - Prompt and Instruction Poisoning

SECTION 2: Securing GenAI Applications

Building on the foundation of Section 1, students will examine the key components needed to develop GenAI applications, including vector databases, LangChain, and AI agents. The course extends to deployment strategies, offering a comparative analysis of cloud-based solutions and on-premises setups, with an emphasis on the specific security risks inherent to each option.

EXERCISES:

- Lab 2.1: AWS Bedrock
- Lab 2.2: Pivoting from LLMs
- Lab 2.3: Compromising LLM Supply Chain
- Lab 2.4: Using Langchain

TOPICS:

- GenAI Applications Architecture
  - Building and deploying GenAI applications
  - GenAI Architecture Security
- Hosting GenAI Applications
  - AWS Bedrock and it's security features
  - Running Local Models Securely
  - LLM customization
  - Models hosting and Supply Chain Attacks
- Agentic AI
  - Agents' design and capabilities
  - Agents' security risks

SECTION 3: MLSecOps and Securing GenAI Applications Lifecycle

In the third and final section, this course shifts its focus to MLSecOps—the integration of security operations into the machine learning lifecycle—and concludes with advanced threat modeling techniques aimed at identifying, assessing, and comprehensively mitigating risks.

EXERCISES:

- Lab 3.1: MLSecOps
- Lab 3.2: Data Security
- Lab 3.3: Threat Modeling

TOPICS:

- Machine Learning Security Operations (MLSecOps)
  - GenAI application Lifecycles
  - Data Protections
  - Security Operations
- Threat Modeling
  - Summering risk associated with GenAI applications
  - Developing GenAI Security Program

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