SEC411: **AI Security Principles and Practices: GenAI and LLM Defense**

3 Day Course	24 CPEs	Laptop Required

You Will Be Able To

- Master AI security fundamentals including tokenization, attack surface analysis, and OWASP Top 10 for LLMs
- Exploit and defend against prompt injection, jailbreaking, RAG manipulation, MCP security, and context hijacking across 125+ hands-on lab tasks
- Implement production-ready defenses for inference runtime, system prompts, and RAG applications
- Secure agentic workflows, autonomous systems, and reasoning models against manipulation and misuse
- Apply OWASP Top 10 for LLMs, MITRE ATLAS, and NIST AI RMF in real security operations scenarios
- Bridge traditional cybersecurity skills to AI-specific requirements and SOC integration
- Design defense-in-depth strategies that balance security and usability across the AI lifecycle
- Protect Model Context Protocol (MCP) environments by identifying attack surfaces and enforcing controls
- Master prompt-injection techniques from direct instruction attacks and encoding-based obfuscation to context-window manipulation and social-engineering
- Build a vendor-neutral detection and incident-response capability spanning prompt, retrieval, tool, and memory-based attacks, grounded in OWASP's Agentic Top 10 threat categories including Tool Misuse and Memory Poisoning
- Critically evaluate Anthropic's Fable 5 and Mythos 5
  safety disclosures and government-driven access changes, translating a fast-moving capability and access-governance story into concrete defensive action

SEC411 is a living, practitioner-focused AI security course for cybersecurity professionals entering GenAI and large language models (LLM) defense. No prior AI experience required.

How Does This Course Differ From Other OnDemand Courses?

The course includes 24+ hours of expert-led content and five hands-on labs with 125+ tasks, including expanded coverage of Model Context Protocol (MCP) security and emerging AI attack surfaces. Unlike other OnDemand courses, the curriculum will be dynamically updated during your four-month access window. Depending on when you enroll, you will receive new videos and you might also get new modules and lab challenges, so your training keeps pace with the rapidly evolving AI threat landscape. To facilitate this approach, this course does not have a print coursebook or a traditional PDF coursebook.

What Is AI Security And Why Is It Important?

AI Security Principles form the foundation of AI risk management and LLM security, helping cybersecurity professionals identify, assess, and defend against threats unique to Generative AI systems. These principles include AI threat modeling, input and output validation, data integrity protection, and AI governance controls that keep models secure, transparent, and compliant. As organizations integrate AI and LLM technologies into critical operations, applying strong AI security practices is essential to prevent model abuse, data leakage, and adversarial manipulation.

How Will This Course Benefit My Career?

SEC411 provides a direct path for cybersecurity professionals to develop in-demand expertise in AI security. You'll gain practical experience identifying and mitigating GenAI and LLM threats, bridging traditional cybersecurity operations with emerging AI defense strategies. Graduates will be prepared to lead AI risk assessments, advise on secure deployments, and contribute to building resilient, compliant AI ecosystems within their organizations.

Business Takeaways

This course will help your organization:

- Close critical skill gaps quickly with progressive learning paths building on existing expertise
- Maximize training ROI with a living curriculum that adds new modules and labs throughout four-month access
- Apply learning immediately with Docker-based hands-on labs deployable in real production environments
- Accelerate team readiness with gamified labs and an integrated Learning Assistant for faster competency
- Reduce deployment risk with comprehensive OWASP Top 10 for LLMs and production security controls coverage
- Support compliance with NIST AI RMF, EU AI Act, and other standards connected to existing security ops
- Future-proof security capabilities spanning current and emerging AI threats with continuous curriculum updates
- Reduce response gaps with a unified monitoring and incident-response framework for the major ways untrusted content reaches AI systems
- Stay ahead of the frontier: SEC411's living curriculum keeps your team current on capability and access shifts like Anthropic's Fable 5/Mythos 5, not months after the fact

Section Descriptions

SECTION 1: KNOW — Understanding the AI Threat Landscape

Build essential AI literacy for security professionals. Learn how LLMs operate, identify AI-specific attack surfaces, and develop intuition about GenAI and LLM applications. This foundation bridges traditional security experience to AI threats with hands-on exploration of tokenization security.

TOPICS:

- Security implications of LLM architecture, training processes, and inference mechanisms
- AI system attack surface mapping and threat modeling using MITRE ATLAS
- Tokenization vulnerabilities and exploitation techniques across different model families
- AI supply chain security, training data poisoning, and model manipulation tactics

SECTION 2: DEFEND — Securing the AI Lifecycle

Secure AI systems from training to runtime. Implement practical defenses for training pipelines, inference environments, and RAG systems. Master input and output filtering, guardrail implementation, and RAG-specific security controls through progressive attack and defense challenges.

TOPICS:

- Training pipeline security with data validation, poisoning detection, and secure infrastructure
- Inference runtime defenses using input sanitization, output validation, and semantic analysis
- OWASP Top 10 for LLM Applications, from prompt injection to model theft
- Advanced prompt attack techniques such as jailbreaking, system prompt extraction, and context manipulation
- Advanced prompt-injection techniques, including instruction manipulation, encoding and obfuscation, conversational-window exploitation, and social-engineering tactics
- RAG system security including vector database protection, retrieval poisoning, and access control
- Production security controls that balance protection and usability

SECTION 3: DEPLOY — Integration, Autonomy, and Advanced AI

Integrate AI security into enterprise architecture and protect autonomous systems. Deploy secure LLM applications, implement robust API security, connect AI monitoring with SOC operations, and address new threats in agentic systems and reasoning models using production-ready strategies.

TOPICS:

- Secure AI deployment architectures and API protections for LLM services
- Authentication, authorization, and monitoring integrated with existing security operations
- Security for agentic systems, including permission models, tool validation, and agent control
- Reasoning model security, focusing on chain-of-thought protection and defending against reasoning manipulation
- Incident response for AI security events and development of enterprise AI security programs
- MCP protocol security, including attack surface and defensive controls
- Unified detection, monitoring, and incident-response practices for context-injection attacks, covering direct prompts, RAG/retrieval poisoning, tool and protocol abuse, and persistent agent-memory poisoning
- Current-events case study examining Anthropic's Fable 5 and Mythos 5 safety disclosures alongside a real US government access suspension, translating a fast-moving capability and governance story into concrete organizational defenses

Who Should Attend?

Cybersecurity professionals entering GenAI and LLM defense. No prior AI background needed. Intermediate skill level — geared for practitioners with hands-on security experience.

What Are The Laptop Requirements?

Docker-based hands-on labs are used in this course. Students will need a laptop capable of running Docker. See the SANS website for full laptop requirements.

How Does This Course Differ From Other OnDemand Courses?

The course features 24+ hours of expert-led videos and five progressive hands-on labs. Unlike other OnDemand courses, the curriculum will be dynamically updated during your four-month access window. Depending on when you enroll, you will receive new videos and you also might get new modules and lab challenges, so your training keeps pace with the rapidly evolving AI threat landscape. To facilitate this approach, this course does not have a print coursebook or a traditional PDF coursebook.

What Will You Get?

- 24+ hours of expert-led videos
- Five progressive hands-on labs
- Continuous updates during four-month access window
- Access to a student Slack community for support and collaboration

What Are The Prerequisites For This Course?

- There are no prerequisites for this course.

What Learning Path Is This Course Part Of?

- The SEC411 course is part of the "Artificial Intelligence" learning path, designed to train cybersecurity professionals in AI security essentials.
- Depending on your current or desired future role, one of these courses is a great next step in your cybersecurity journey:
- SEC595: Applied Data Science and AI/Machine Learning for Cybersecurity Professionals | GMLE
- SEC545: GenAI and LLM Application Security | GAIPS