COURSE SYLLABUS

AI Security Foundations Certificate Program (Self-Paced)

Secure and govern AI in a federal environment — from the threat landscape to a working agency AI security and governance program.

> Group classes (Live Online) and onsite training are available for this course. For more information, email onsite@graduateschool.edu or visit: https://www.graduateschool.edu/certificates/ai-security-foundations-certificate-program-self-paced



support@graduateschool.edu •
(888) 744-4723

Course Outline

This package includes these courses

- AI Security Fundamentals for the Cyber Workforce Course (Self-Paced) (6 Hours)
- AI Security & Governance for Government Course (Self-Paced) (12 Hours)

AI Security Fundamentals for the Cyber Workforce Course (Self-Paced)

Most organizations now use AI rather than build it, so every security professional needs to treat AI as part of the attack surface. This foundational course—no AI/ML background needed—covers the AI threat landscape (MITRE ATLAS), the risks of LLM apps (OWASP Top 10 for LLMs), and the dangers of AI agents, tying it to the NIST AI Risk Management Framework and the everyday defenses that keep AI systems secure.

- Explain what AI/ML systems are as a security attack surface
- Distinguish traditional ML, generative AI, LLMs, and AI agents
- Navigate the AI threat landscape using MITRE ATLAS
- Threat-model LLM applications with the OWASP Top 10 for LLM Applications
- Assess the risks of autonomous AI agents and excessive agency
- Apply the NIST AI Risk Management Framework (Govern, Map, Measure, Manage)
- Apply baseline AI defenses: inventory, input and output filtering, least-privilege tool access, monitoring, and human review gates

AI Security & Governance for Government Course (Self-Paced)

Federal agencies are deploying AI faster than the practices meant to govern and secure it. This intermediate course helps federal staff and their partners build an agency AI governance and security program grounded in federal AI policy and the NIST AI Risk Management Framework—covering AI inventories and risk classification, data governance and privacy, the AI development and acquisition lifecycle, machine learning and generative AI threats, and the oversight and monitoring that keep

AI trustworthy, ending with a capstone plan you build yourself.

- Navigate federal AI policy, roles, and governance structures
- Apply the NIST AI Risk Management Framework (Govern, Map, Measure, Manage)
- Build AI system inventories and classify AI use cases by risk
- Apply data governance, privacy, and information-protection practices to AI
- Secure the AI development and acquisition lifecycle
- Recognize threats to machine learning and generative AI systems
- Establish testing, evaluation, validation, and human oversight
- Manage third-party AI, supply-chain risk, and federal procurement
- Run continuous monitoring, incident response, and accountability for AI