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 in Live Online and onsite training is 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) (5 hours)
- AI Security for Federal Systems 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 for Federal Systems Course (Self-Paced)

Learn to govern, secure, and audit federal AI systems using the NIST AI Risk Management Framework, EO 14110, and OMB M-24-10. This self-paced course covers the full lifecycle: GOVERN, MAP, MEASURE, and MANAGE functions, 800-53 Rev 5 controls, adversarial threat identification, vendor evaluation, and continuous monitoring.

- Explain EO 14110, OMB M-24-10, and the NIST AI RMF 1.0 structure — and how they create specific, enforceable federal AI

governance requirements

- Apply the NIST AI RMF GOVERN, MAP, MEASURE, and MANAGE functions to a federal AI deployment scenario
- Identify adversarial AI threats — including data poisoning, adversarial examples, and LLM prompt injection — and recommend appropriate technical countermeasures
- Apply NIST 800-53 Rev 5 AI-relevant controls to AI system security architectures
- Evaluate AI vendor security documentation and draft contract language to fill identified gaps
- Conduct a structured AI red team exercise against an LLM-powered federal application
- Develop an AI audit plan applying GAGAS standards to AI accuracy, fairness, security, and transparency objectives
- Design a federal AI risk governance program with a use case inventory, risk tiering framework, and continuous monitoring plan