COURSE SYLLABUS

AI Security Foundations Certificate Program

Secure and govern AI inside a federal environment — from the threat landscape through 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

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

Course Outline

This package includes these courses

- AI Security Fundamentals for the Cyber Workforce Course (6 Hours)
- AI Security & Governance for Government Course (14 Hours)

AI Security Fundamentals for the Cyber Workforce Course

Most organizations now use AI rather than build it, so every security professional has 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 LLM Applications), 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 makes AI/ML systems a security attack surface
- Tell traditional ML, generative AI, LLMs, and AI agents apart
- Find your way around the AI threat landscape using MITRE ATLAS
- Threat-model LLM applications against the OWASP Top 10 for LLM Applications
- Weigh the risks that autonomous AI agents and excessive agency introduce
- Put the NIST AI Risk Management Framework to work (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

Federal agencies are putting AI into service 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 rooted 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.

- Find your way through federal AI policy, roles, and governance structures
- Put the NIST AI Risk Management Framework to work (Govern, Map, Measure, Manage)
- Build AI system inventories and sort AI use cases by risk
- Bring data governance, privacy, and information-protection practices to bear on AI
- Secure the AI development and acquisition lifecycle
- Recognize the threats facing 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