Cloud Contraptions
TRAINING & CONSULTING

To discuss this course and customizations:
Call: 434-509-5680 or Email: sales@cloudcontraptions.com

Securing AI-Assisted Development Pipelines

Class Duration

7 hours of live training delivered over 1-2 days to accommodate your scheduling needs.

Student Prerequisites

- Professional software development experience
- Basic familiarity with security concepts (OWASP, authentication, secrets management)

Target Audience

Software engineers, security engineers, and DevSecOps practitioners responsible for keeping AI-assisted development workflows secure. Relevant for any team using AI coding assistants, agentic CI pipelines, or LLM-powered developer tooling in a context where code quality, intellectual property, or regulated data is at stake.

Description

AI-assisted development introduces a new class of security risks that traditional AppSec tooling doesn't cover. This course provides a practitioner's guide to securing the AI layer: prompt injection attacks (direct and indirect), secret and credential leakage through AI context, supply-chain risks from AI-generated dependencies, data exfiltration through model APIs, and policy guardrails for autonomous agents. We ground each topic in realistic attack scenarios and pair them with concrete mitigations applicable at the developer, team, and platform level.

Learning Outcomes

- Identify and demonstrate direct and indirect prompt injection attacks in realistic scenarios.
- Implement mitigations for prompt injection in LLM-powered applications and agent pipelines.
- Audit AI assistant configurations for secret and credential leakage risk.
- Apply input/output validation to reduce data exfiltration risk through model APIs.

To discuss this course and customizations:
Call: 434-509-5680 or Email: sales@cloudcontraptions.com

- Assess AI-generated dependency and package suggestions for supply-chain risk.
- Design organizational policy guardrails for autonomous agent workflows.
- Map AI-assisted development risks to OWASP LLM Top 10, OWASP Agentic AI Threats and Mitigations (2025), and NIST AI RMF controls.

Training Materials

Comprehensive courseware is distributed online at the start of class. All students receive a downloadable MP4 recording of the training.

Software Requirements

A modern IDE with an AI coding assistant, Python 3.12+ or Node.js 20+ for lab exercises, and Git.

Training Topics

AI-Assisted Development Attack Surface

- What changes about your attack surface when AI is in the loop
- OWASP LLM Top 10 for developers
- OWASP Agentic AI Threats and Mitigations (2025): goal hijack, tool misuse, delegated trust, persistent memory, inter-agent communication
- Threat model for coding assistant and agentic CI workflows

Prompt Injection

- Direct prompt injection: attacking user-facing LLM features
- Indirect prompt injection: malicious content in retrieved context
- Agentic prompt injection: hijacking autonomous agent tasks
- Mitigation patterns: input sanitization, output validation, privilege separation

Secret and Credential Leakage

- How secrets end up in AI context windows
- IDE assistant context scope and exclusion configuration
- Detecting secrets in prompts before transmission
- `.copilotignore`, `.cursorignore`, and equivalent controls

Cloud Contraptions
TRAINING & CONSULTING

To discuss this course and customizations:
Call: 434-509-5680 or Email: sales@cloudcontraptions.com

AI-Generated Code Security

- Common vulnerability patterns in AI-generated code
- Hallucinated dependency attacks (package confusion)
- Static analysis integration for AI-generated changes
- Code review checklist for AI output security

Data Exfiltration and Privacy

- What data leaves your environment through model APIs
- API telemetry and training data use policies
- Avoiding PII and proprietary data in prompts
- Air-gapped and private model deployment options

Supply-Chain Risk

- AI-suggested package and library risks
- Evaluating new dependencies before accepting AI suggestions
- SBOM implications of AI-assisted development

Policy Guardrails for Agents

- Capability scoping for autonomous agents
- Approval gates for file system, network, and command execution
- Audit logging requirements for agentic actions
- Kill switch and circuit breaker design

Compliance and Governance Frameworks

- OWASP LLM Top 10 mapping
- OWASP Agentic AI Threats and Mitigations (2025) mapping
- NIST AI RMF application to developer tooling
- SOC 2 and ISO 27001 considerations for AI tooling

Workshop

- Prompt injection attack and mitigation lab
- Secret leakage audit exercise
- Agent policy design exercise
- Q&A session