



AI Agent Security Engineering and Assurance Training Course

8 – 12 August 2027
Johannesburg

AI Agent Security Engineering and Assurance Training Course

Ref: 199_118059 **Date:** 8 – 12 August 2027 **Location:** Johannesburg **Fees:** 4500 Euro

Overview

AI Agent Security Engineering and Assurance Training Course is a five-day advanced course for AI security engineers, application security specialists, agent developers, architects, red teams, and technology risk leads, who leave with an AI Agent Security Assurance Plan. Participants examine agent threat modeling, prompt injection, tool misuse, identity, least privilege, memory and retrieval poisoning, inter-agent trust, data leakage, sandboxing, monitoring, incident response, adversarial testing, and security evaluation. Agile Leaders Training Center delivers training in AI agent security assurance.

Who Should Attend

- AI security personnel responsible for agent threat models, controls, and assurance evidence
- Application security personnel responsible for design review, testing, and release gates
- Agent engineering personnel responsible for tools, memory, retrieval, orchestration, and deployment
- Security testing personnel responsible for adversarial exercises and failure discovery
- Technology risk personnel responsible for control ownership, monitoring, and incident readiness

The course assumes participants can assess application architecture and security controls, and it leaves out introductory AI literacy, general agent construction, and vendor-platform administration.

Departments and Industries

The course supports AI security across software, finance, health care, retail, logistics, public services, and professional services.

- Application security and product security
- AI engineering and solution architecture
- Identity, access, and data protection
- Security operations and incident response
- Technology risk across regulated and digital-service organizations

Learning Objectives

By the end of this course, participants will be able to:



- Analyze agent assets, trust boundaries, and attack paths
- Evaluate prompt injection, tool misuse, and poisoning risks
- Apply identity, least-privilege, isolation, and approval controls
- Diagnose monitoring signals and agent security incidents
- Build adversarial tests and security evaluation gates
- Build an AI Agent Security Assurance Plan

Course Agenda

Day 1: Agent Threat Modeling

- Agent Asset and Action Inventory
- Trust Boundary and Data-Flow Diagram
- Agent Attack-Path Tree
- Autonomy and Impact Classification Matrix
- Security Requirement Traceability Register

Day 2: Injection, Tools, and Data Boundaries

- Direct and Indirect Prompt-Injection Test Map
- Untrusted Content Isolation Pattern
- Tool Misuse and Chained-Action Scenario
- Data Leakage and Exfiltration Control Grid
- Human Confirmation and Transaction Boundary

Day 3: Identity, Memory, and Inter-Agent Trust

- Dedicated Agent Identity and Ownership Card
- Least-Privilege Tool and Resource Matrix
- Time-Bound Access and Revocation Workflow
- Retrieval and Memory Poisoning Test Set
- Inter-Agent Handoff Trust Protocol

Day 4: Monitoring, Testing, and Response

- Agent Action and Authorization Audit Trail
- Security Signal and Anomaly Detection Board
- Adversarial Agent Test Case Library
- Containment and Kill-Switch Runbook
- Agent Incident Triage and Recovery Playbook



Day 5: Security Assurance Practice and Capstone

- Suggested Exercise: Build an Agent Threat Model
- Suggested Exercise: Test Injection and Tool Misuse
- Suggested Exercise: Challenge Identity and Memory Controls
- Suggested Exercise: Rehearse Detection and Containment
- Capstone Exercise: AI Agent Security Assurance Plan

Practical Exercises

The course uses suggested activities that turn an agent architecture into testable security controls and assurance evidence.

- Suggested activity: map assets, data flows, trust boundaries, actions, identities, tools, memory stores, and impact levels
- Suggested activity: run injection, tool abuse, data leakage, retrieval poisoning, and memory poisoning scenarios against defined controls
- Suggested activity: test least privilege, approval gates, sandbox boundaries, audit records, revocation, and kill-switch paths
- Suggested activity: assemble findings, owners, remediation priorities, regression tests, monitoring signals, and response actions

FAQs

Who suits the AI Agent Security Engineering and Assurance Training Course, and what does it assume?

The course suits security engineers, application security specialists, agent developers, architects, red teams, and risk leads. It assumes experience assessing application architecture and controls.

How does AI agent security training differ from general AI agent development training?

AI agent security training concentrates on attack paths, identities, permissions, untrusted content, poisoning, leakage, adversarial tests, monitoring, and response rather than building general agent features.

What threats should an AI agent security assessment examine?

An assessment should examine prompt injection, unsafe tool calls, excessive privilege, identity ambiguity, data leakage, retrieval and memory poisoning, insecure handoffs, weak audit trails, uncontrolled autonomy, and delayed containment.



How should teams test AI agent security before deployment?

Teams should test representative and adversarial tasks, direct and indirect injection, tool policies, privilege boundaries, approval gates, data handling, memory writes, inter-agent messages, logging, revocation, containment, and recovery.

What evidence supports AI agent security assurance?

Assurance evidence includes a threat model, control traceability, identity and permission records, adversarial test results, regression cases, action logs, monitoring rules, remediation ownership, incident runbooks, and release decisions.

Conclusion

Participants take back an AI Agent Security Assurance Plan linking threats, trust boundaries, controls, tests, monitoring, and response actions. It changes how teams assess and release agent systems. The plan supports traceable security decisions, control ownership, adversarial validation, rapid containment, and continuing regression testing after material changes.