JUPITER PRIME

Consulting • Training • Advisory

جوبيتر برايم للاستشارات والتدريب

COURSE BROCHURE

AI Threat Detection and SOC Response Automation Training Course

jupiterprime.uk

info@jupiterprime.uk

Registered in the United Kingdom and Saudi Arabia

AI Threat Detection and SOC Response Automation Training Course

AI Threat Detection and SOC Response Automation is a 1-week training course for security operations leads and IT risk managers who decide how machine learning detectors and automated playbooks enter a security operations centre SOC, ending with an AI defence architecture plan. Attackers now script reconnaissance and phishing at machine speed, so alert queues outgrow analyst capacity. Jupiter Prime Training approaches AI Threat Detection and SOC Response Automation from the shift floor.

Why AI Threat Detection and SOC Response Automation Matters Now

Three pressures meet in the SOC. Alert volumes from endpoint, identity and cloud telemetry exceed what tier-one analysts can triage; attackers probe the detection models themselves with evasion and poisoning; and boards ask how much response can safely run without a human click. A centre that tunes its detectors and automates containment with clear limits shortens dwell time. The course draws on the NIST AI Risk Management Framework NIST AI RMF and ISO/IEC 42001, the AI management system standard, for model oversight and accountability.

Who Gains Most from AI Threat Detection and SOC Response Automation

- Leads who own detection content, alert thresholds and triage queues
- Managers who approve which containment steps run without analyst sign-off
- Risk owners who report residual cyber exposure to executives
- Engineers who connect log pipelines, security information and event management SIEM and orchestration platforms
- Specialists who test detection models against evasion and data poisoning

A SOC leadership group gains most when its detection lead, response lead and risk owner attend together and agree one automation boundary.

AI Threat Detection and SOC Response Automation Capabilities

Participants judge detectors by precision on their own telemetry and draw the line between machine action and analyst judgement.

- Assess anomaly detection and supervised classifiers against labelled incident history — detection leads
- Calculate precision, recall and alert cost for each detector — SOC managers
- Design user and entity behaviour baselines with drift monitoring — identity owners
- Specify playbook steps that run automatically and those held for approval — response leads
- Test models against evasion, poisoning and prompt manipulation — red team leads
- Communicate residual exposure and automation limits — chief information security officers

Key AI Threat Detection and SOC Response Automation Terms

- User and entity behaviour analytics UEBA — models that learn normal activity for accounts and devices and flag deviations.
- Security orchestration, automation and response SOAR — a platform that runs response playbooks across security tools.
- Alert fatigue — declining analyst attention caused by high volumes of low-value alerts.
- Evasion attack — input crafted so that a detection model misclassifies malicious activity as benign.
- Data poisoning — tampering with training data so that a model learns a blind spot.
- Mean time to contain — the average period between detecting an incident and isolating its source.

AI Threat Detection and SOC Response Automation Day by Day

Day 1: The Machine-Speed Threat Picture

- Attacker Use of Automation and Generative Tools — Changing Alert Volumes
- SOC Maturity Assessment — Readiness for Model-Driven Detection
- Telemetry Inventory Across Endpoint, Identity and Cloud — Coverage Gaps
- Alert Fatigue Diagnosis — Where Analyst Hours Are Lost
- Model Risk Framing Under NIST AI RMF — Accountable Detection Choices

Day 2: Detection Models and Behavioural Baselines

- Anomaly Detection Versus Supervised Classifiers — Fit to Threat Type
- UEBA Baselines for Accounts and Devices — Insider and Takeover Signals
- Feature Selection From Log Sources — Signal Quality Over Volume
- Precision, Recall and Alert Cost Scoring — Detector Acceptance Thresholds
- Model Drift Monitoring — Retraining Triggers After Environment Change

Day 3: Automating Triage and Response

- Alert Enrichment and Risk Scoring — Queue Order for Analysts
- SOAR Playbook Design — Containment Steps and Approval Gates
- Generative Assistants for Investigation Notes — Analyst Review Duties
- Automation Boundary Matrix — Machine Action Versus Human Decision
- Rollback and Kill-Switch Design — Limiting Damage From Wrong Actions

Day 4: Defending the Detection Models

- Evasion Testing of Classifiers — Hardening Against Crafted Inputs
- Data Poisoning Controls — Protecting Training Pipelines and Labels
- Prompt Manipulation of Security Assistants — Guardrails and Logging
- Model Inventory Under ISO/IEC 42001 — Ownership and Change Records
- Purple Team Exercises on AI Detectors — Measured Coverage Gains

Day 5: Practising the AI Defence Architecture

- Detector Scoring Applied to a Retail Bank SOC Case — Keep or Retire
- Behaviour Baselines Applied to a Utility Operator Case — Priority Signals
- Playbook Boundaries for Ransomware Containment — Approval Points Tested
- Exposure Briefing for Executives — Residual Cyber Risk Statement
- AI Defence Architecture Plan Assembly — Detectors, Playbooks, Controls and Owners

Applying AI Threat Detection and SOC Response Automation After the Course

- Score the three noisiest detectors in the SIEM for precision and alert cost and retire or retune the weakest.
- Draft an automation boundary matrix for one containment playbook and agree its approval gates with the risk owner.

AI Threat Detection and SOC Response Automation Questions FAQ

What is AI threat detection and SOC response automation, and why does it matter now?

It is the use of machine learning to spot attacks and of playbooks to contain them inside a SOC. It matters now because automated attacks produce more alerts than analysts can triage by hand.

Which roles gain most from AI threat detection and SOC response automation?

SOC managers, detection and response leads and IT risk owners gain most. Security engineers who run SIEM and SOAR platforms also benefit, because they turn detector and playbook choices into working pipelines.

How does AI threat detection and SOC response automation differ from an AI agent security course?

It defends the organisation with AI rather than securing AI agents themselves. An agent security course hardens autonomous assistants, while this one tunes detectors, automates containment and protects those detection models from evasion.

What do you bring back from AI threat detection and SOC response automation training?

You bring back an architecture plan for your own centre. It lists detectors with acceptance thresholds, behaviour baselines, playbooks with approval gates, model defences and named owners.

AI Threat Detection and SOC Response Automation: The Way Forward

The AI defence architecture plan first informs which detectors stay in production and which containment steps may run unattended. SOC managers use it to brief executives. After the first major incident handled under it, revisit the automation boundaries: actions held for approval may prove safe to automate, and others may need a human gate.