Training: EC-Council
COASP - Certified Offensive AI Security Professional

EC-Council
Building A Culture Of Security

TRAINING GOALS:

C|OASP
Certified Offensive AI Security Professional

The Certified Offensive AI Security Professional (C|OASP) credential is designed to transform professionals into enterprise-ready offensive AI security specialists. This program focuses on real-world AI attack techniques and adversary methodologies, equipping professionals to identify, exploit, and remediate vulnerabilities across LLMs, agentic AI systems, and AI pipelines using industry frameworks like OWASP LLM Top 10 and MITRE ATLAS.

The C|OASP credential validates your ability to:

- Execute prompt injection, jailbreaking, and prompt chaining attacks
- Red-team AI agents, including memory corruption, tool misdirection, and checkpoint manipulation
- Apply OWASP LLM Top 10 and MITRE ATLAS frameworks
- Conduct adversarial ML attacks, including data poisoning and model extraction
- Build detection rules and hardening strategies for AI systems

C|OASP trains you to perform end-to-end adversarial testing and deliver defensive validation evidence, including:

- Simulate adversarial AI kill chains
  - reconnaissance > mapping > exploitation > manipulation > exfiltration
- Harden AI architectures

Compendium Education Center Ltd.
ul. Tatarska 5, 30-103 Kraków, tel.: (12) 298 47 77
ul. Postępu 18B, 02-676 Warszawa, tel.: (22) 417 41 70
BNP Paribas Bank Polska S.A.
EUR: PL 79 1600 1462 1853 6451 3000 0006
USD: PL 52 1600 1462 1853 6451 3000 0007

    - secure system prompts, context windows, tool integrations, RAG pipelines, and agent memory
- Conduct AI security assessments
    - aligned to MITRE ATLAS, OWASP LLM/ML Top 10, NIST AI RMF, and DoD Test & Evaluation practices
- Build SOC-ready capabilities
    - AI-focused detection logic, incident playbooks, and forensic procedures
- Execute prompt injection, adversarial prompting
    - data poisoning against LLMs/ML systems to identify training-time and inference-time weaknesses
- Assess AI supply-chain risk
    - across models, datasets, dependencies, and third-party integrations using SBOM/MBOM approaches
- Implement defensive engineering controls
    - filtering, sandboxing, rate limiting, anomaly detection, and drift monitoring
- Produce assurance and compliance artifacts
    - mapped to NIST AI RMF, ISO/IEC 42001, and emerging AI regulatory expectations

Enterprise Impact of Verifiable Skills from Certified Offensive AI Security Professional:

- Helps organizations identify and neutralize AI-specific threats before attackers do.
- Bridges security, engineering, and data science so controls exist across the full AI life cycle.
- Addresses the lack of standardized AI red-teaming methodology by applying OWASP LLM Top 10 and MITRE ATLAS frameworks.
- Strengthens resilience across plugins, APIs, and vendor ecosystems by exposing third-party risk.
- Improves monitoring and response because defenders understand attacker tactics at the model, application, and system level.
- Reinforces secure and ethical AI deployment, supporting innovation without sacrificing trust.

Who COASP is for:

- OFFENSIVE SECURITY
    - Penetration Tester/Ethical Hacker
    - Red Team Operator/Red Team Lead
    - Offensive Security Engineer
    - Adversary Emulation/Purple Team Specialist

Compendium Education Center Ltd.
ul. Tatarska 5, 30-103 Kraków, tel.: (12) 298 47 77
ul. Postępu 18B, 02-676 Warszawa, tel.: (22) 417 41 70
BNP Paribas Bank Polska S.A.
EUR: PL 79 1600 1462 1853 6451 3000 0006
USD: PL 52 1600 1462 1853 6451 3000 0007

- DEFENSIVE SECURITY
  - SOC Analyst (Tier 2/3)/Detection Engineer
  - Blue Team Engineer/Threat Detection Engineer
  - Incident Responder (IR)/DFIR Analyst
  - Security Operations Manager (SOC Lead)
- THREAT INTELLIGENCE
  - Malware Analyst/Threat Researcher
  - Cyber Threat Intelligence (CTI) Analyst – AI Focus
  - Fraud/Abuse Detection Analyst (AI-enabled threats)
- AI/ML ENGINEERING
  - ML Engineer/Applied AI Engineer
  - GenAI Engineer (RAG/Agents)
  - AI/LLM Application Developer
  - MLOps/AI Platform Engineer
- SECURITY ENGINEERING
  - DevSecOps/Secure DevOps Specialist
  - Application Security Engineer (LLM Apps/APIs)
  - Product Security Engineer/AI Product Security
- AI SECURITY ARCHITECTURE
  - Secure AI Engineer/AI Security Architect
  - LLM Systems Engineer

Each participant in an authorized training COASP - Certified Offensive AI Security Professional held in Compendium CE will receive a free COASP certification exam voucher.

CONSPECT:

- Module 1 - Offensive AI and AI System Hacking Methodology
  - AI & ML Fundamentals
  - AI A1ack Surface and Threat Landscape (ATLAS-Aligned)
  - AI A1ack Taxonomy and Classification
  - OWASP LLM and ML Top 10 (2025) - Overview & Mapping
  - AI System Hacking Methodology
  - Securing AI Systems - Foundations (Defensive Anchor)

Compendium Education Center Ltd.
ul. Tatarska 5, 30-103 Kraków, tel.: (12) 298 47 77
ul. Postępu 18B, 02-676 Warszawa, tel.: (22) 417 41 70

BNP Paribas Bank Polska S.A.
EUR: PL 79 1600 1462 1853 6451 3000 0006
USD: PL 52 1600 1462 1853 6451 3000 0007

- AI Security Governance and Compliance
- Module 2 - AI Reconnaissance and Attack Surface Mapping
  - OSINT for AI Assets
  - Tools and Techniques for AI OSINT
  - Data & Training Pipeline Intel Gathering
  - Mapping AI A1ack Surfaces from OSINT
  - Discovering AI Endpoints & Services
  - AI API & Parameter Enumeration
  - Model & Vector Store Enumeration
  - Defensive - Reducing AI OSINT Exposure
  - Defensive - Hardening Enumerated Surfaces
  - AI Threat Intelligence & Continuous Monitoring
- Module 3 - AI-Specific Vulnerability Scanning and Fuzzing
  - Fundamentals of AI Vulnerability Assessment
  - Tools and Techniques for Vulnerability Scanning
  - Fuzzing Techniques for AI Systems
  - Defensive - Integrating Scanning & Fuzzing
- Module 4 - Prompt-Based and LLM Application Attacks
  - LLM Architecture & Trust Boundaries
  - Prompt Injection & Jailbreaking
  - Sensitive Information Disclosure and System Prompt Leakage
  - Improper Output Handling and Misinformation
  - Advanced Prompt A1ack Techniques
  - Defensive - Secure LLM Application Design
- Module 5 - Adversarial Machine Learning and Model Privacy Attacks
  - Adversarial ML A1acks
  - Practical Adversarial Input A1acks
  - Privacy & Model Extraction A1acks
  - Evaluating Robustness & Trustworthiness
  - Emerging Model A1ack Techniques
  - Defensive - Privacy & Robustness Mitigations
- Module 6 - Data and Training Pipeline Attacks
  - Understanding AI Data & Training Pipelines
  - Data Poisoning A1acks
  - Backdoor / Trojan A1acks in Training Pipelines
  - AI Supply Chain A1ack Vectors

Compendium Education Center Ltd.
ul. Tatarska 5, 30-103 Kraków, tel.: (12) 298 47 77
ul. Postępu 18B, 02-676 Warszawa, tel.: (22) 417 41 70

BNP Paribas Bank Polska S.A.
EUR: PL 79 1600 1462 1853 6451 3000 0006
USD: PL 52 1600 1462 1853 6451 3000 0007

- Defensive - Securing Data & Training Pipelines
- Module 7 - Agentic AI and Model-to-Model Attacks
  - Agentic AI Architecture & A1ack Surface
  - Excessive Agency & Autonomy
  - Model-to-Model and Cross-LLM A1acks
  - Unbounded Consumption and Denial of Wallet
  - AI Workflow and Orchestration A1acks
  - Defensive - Securing Agentic Applications
- Module 8 - AI Infrastructure and Supply Chain Attacks
  - AI Infrastructure & Integration Landscape
  - System and Framework Exploits
  - Tool and API Abuse in AI Apps
  - Supply Chain Threats (Deep Dive)
  - Defensive - Hardening AI Infra & Supply Chain
- Module 9 - AI Security Testing, Evaluation, and Hardening
  - AI Security Test & Evaluation Fundamentals
  - Designing AI Security Test Plans
  - Executing AI Security Tests
  - Reporting, Assurance & Risk Management
  - Defensive - Embedding T&E into MLOps/DevSecOps
- Module 10 - AI Incident Response, Forensics, and Capstone Red Team
  - Detecting & Responding to AI-Specific Incidents
  - Logging, Telemetry & Evidence Collection
  - AI Forensics & Post-Incident Analysis
  - Capstone: Full-Scope AI Red Team Engagement
  - Course Wrap-Up & Professional Practice

REQUIREMENTS:

Participants are highly recommended to have:

- Minimum of 2 years of professional experience in the cybersecurity/Information Security (InfoSec) domain.
- A strong foundational understanding of security operations, computer networks, and application security.
- Familiarity with ethical hacking or penetration testing concepts. If you are completely new to offensive security, it is advised to complete foundational certifications like CEH (Certified Ethical Hacker)

Compendium Education Center Ltd.
ul. Tatarska 5, 30-103 Kraków, tel.: (12) 298 47 77
ul. Postępu 18B, 02-676 Warszawa, tel.: (22) 417 41 70

BNP Paribas Bank Polska S.A.
EUR: PL 79 1600 1462 1853 6451 3000 0006
USD: PL 52 1600 1462 1853 6451 3000 0007

Difficulty level

CERTIFICATE:

The participants will obtain certificates signed by EC-Council (course completion). This course will help prepare you also for the COASP certification exam.

Certified Offensive AI Security Professional (COASP) exam details:

- Exam Code : 312-52
- Number of Questions : 70 (65 MCQ + 5 PBQ)
- Duration : 6 hours
- Availability: ECC Exam Portal
- Passing Score: 70-80%
- Test Format : Multiple Choice Question (MCQs) and Performance-Based Questions (PBQs)

Each participant in an authorized training COASP - Certified Offensive AI Security Professional held in Compendium CE will receive a free COASP certification exam voucher.

TRAINER:

Certified EC-Council Instructor (CEI)

ADDITIONAL INFORMATION:

The training materials include official EC-Council electronic courseware, 180-day access to iLabs, and an exam voucher.

Compendium Education Center Ltd.
ul. Tatarska 5, 30-103 Kraków, tel.: (12) 298 47 77
ul. Postępu 18B, 02-676 Warszawa, tel.: (22) 417 41 70

BNP Paribas Bank Polska S.A.
EUR: PL 79 1600 1462 1853 6451 3000 0006
USD: PL 52 1600 1462 1853 6451 3000 0007