
R.AI.S.E
powered by SRI - F

RAISE™

Responsible AI Security & Ethics

RAISE Conscious AI, Consciously

Certified R.AI.S.E Professional Program

Build • Defend • Resilience • Govern • Align

Powered by the SRI™ Framework (Secure AI & Responsible Management)

Duration	Format	Certification
8–12 Weeks	**5 Tracks + Capstone**	**5 Specializations**
Flexible / Weekend / Self-Paced	Labs + Projects + Certification	RAISE-R / B / W / D / C

Who Should Join RAISE?

RAISE is designed for cybersecurity professionals, AI engineers, and leaders who want to master the full spectrum of AI security - from offensive exploitation to ethical governance.

Target Audience	Target Job Roles After RAISE
Cybersecurity Professionals	AI Security Engineer
SOC Analysts / SOC Engineers	LLM Security Engineer
Security Engineers	AI Red Team Engineer
DevSecOps Engineers	AI Security Architect
Cloud Security Engineers	AI Governance Specialist
GRC / Risk Professionals	AI SOC Automation Engineer
AI Engineers / Developers	Responsible AI Consultant
Security Architects	AI Risk & Compliance Analyst
CISOs / Security Leaders	Cybersecurity AI Platform Engineer
Students seeking AI Security careers	AI Security Analyst

Course Learning Outcomes

By the end of RAISE, students will be able to:

- Secure LLM and GenAI applications against real-world adversarial threats
- Build and defend RAG systems with layered guardrails and evaluation pipelines
- Protect autonomous multi-agent AI workflows with Zero Trust principles
- Perform AI threat modeling using OWASP LLM Top 10, MITRE ATLAS, and NIST AI RMF
- Deploy AI guardrails, runtime controls, and semantic firewalls
- Implement AI governance frameworks aligned with EU AI Act, ISO 42001, and GDPR
- Design enterprise AI security architectures with observability and incident response
- Build AI SOC automation workflows with autonomous defense agents
- Apply spiritual-ethical (Dharmic) principles to conscious AI alignment

Course Architecture Overview

RAISE is structured as 5 specialized Tracks plus a Capstone, each earning a distinct certification badge. The tracks mirror the AI Ops diagram - AI at the core, surrounded by SecOps, Defense (Blue Team), Offense (Red Team), and Governance, Risk & Compliance, with Ethical Frameworks as the outer layer.

Badge	Track	Theme	Focus
RAISE-R	Track 1	**Red Teaming & Offensive AI Security**	Exploit and test vulnerabilities in LLMs and AI systems
RAISE-B	Track 2	**Blue Teaming & AI Guardrails**	Defense strategies and secure response generation
RAISE-W	Track 3	**Governance, Ethics & Compliance**	Aligning AI systems with global laws and ethical frameworks
RAISE-D	Track 4	**DevSecOps & Threat Modeling**	Secure pipeline integration and threat lifecycle management
RAISE-C	Track 5	**Capstone & Conscious AI Alignment**	Final projects, real-world simulation, spiritual-tech integration

Track 1: RAISE-R

Red Teaming & Offensive AI Security | Badge: RAISE-R

Train learners to identify, exploit, and assess vulnerabilities in Large Language Models and AI pipelines. Simulate real-world adversarial scenarios through structured hands-on labs aligned with OWASP LLM Top 10 and MITRE ATLAS.

Module 1 - Foundations of AI Red Teaming

Establishes the principles, ethics, and scope of adversarial AI testing. Learners understand red teaming's role in the AI/LLM lifecycle and how to ethically simulate adversarial behavior.

Topics Covered	Labs & Activities
Role of red teaming in the AI/LLM lifecycle	ChatGPT jailbreak case study analysis
Offensive testing & responsible AI alignment	WormGPT / FraudGPT threat intelligence review
Threat surfaces: prompt, model, pipeline, plugin	MITRE ATLAS threat mapping exercise
MITRE ATLAS & OWASP LLM Top 10 overview	Red team charter and scope document drafting
Real-world case studies: jailbreaks, LLM scams	Ethics gate checklist creation

Module 2 - Prompt Injection & Prompt Manipulation

Deep dive into the most prevalent LLM vulnerability. Learners execute direct and indirect prompt injection, context pollution, confusion attacks in multi-agent systems, and automated exploit chains.

Topics Covered	Labs & Tools
Direct vs. Indirect Prompt Injection	PortSwigger LLM Labs (hands-on)
Context pollution & data contamination	Prompt Airlines CTF challenge
Attack chaining and prompt leakage	Manual injection: ChatGPT, LLaMA, Claude
Confusion attacks in multi-agent systems	Automated scanning with Garak
Prompt-based jailbreaking techniques	PromptInject payload library testing

Module 3 - Model & Data Poisoning

Explores how adversaries corrupt training data or fine-tuning pipelines to manipulate model behavior at inference. Includes backdoor insertion, label flipping, and transfer learning attack surfaces.

Topics Covered	Labs & Tools
Poisoned datasets and label flipping	BackdoorBox poisoning simulation
Trigger generation & backdoor insertion	TextAttack adversarial text generation
Fine-tuning poisoning vulnerabilities	Trigger-based output redirection lab

Transfer learning attack surfaces	Model behavior comparison pre/post-poison
Supply chain data integrity risks	Hugging Face model provenance analysis

Module 4 - Sensitive Information Disclosure

Techniques for extracting training data, secrets, and private information from LLMs through memorization exploitation, semantic inference, and covert probing strategies.

- Prompt probing and training data leakage via memorization
- LLM response manipulation for secret extraction
- Semantic inference and covert query strategies
- Wordlist-based and prompt-tuned information extraction
- Labs: Synthetic prompt attacks on hosted models; training data echo leak tests

Module 5 - Supply Chain & Plugin Attacks

Examines vulnerabilities in AI dependencies, third-party plugins, and package ecosystems. Covers dependency confusion, SBOM evasion, and indirect plugin invocation exploits.

- Exploiting AI dependencies and Python packages (PyPI poisoning, dependency confusion)
- Jailbreaking 3rd-party plugins via LLM interface and indirect invocation
- SBOM evasion techniques and code poisoning via Hugging Face integration
- Labs: Simulated plugin hijack; Aura malicious dependency scanning; MITMProxy traffic interception

Module 6 - Denial of Service & Resource Exhaustion

Teaches how attackers degrade or incapacitate AI services through prompt bombs, context flooding, API spamming, and agent-loop-based resource exhaustion.

- Prompt bombs, recursive calls, and context flooding
- Token injection and API spamming via serverless vectors
- Model overuse via agentic auto-reply loops
- Labs: Token overflow demos; DoS simulation with prompt repeaters; API exhaust via auto-agent loops

Module 7 - Model Theft & Output Hijacking

Covers extraction attacks, model fingerprinting, prediction inversion, and embedding leak exploitation to steal model intellectual property.

- Model fingerprinting and prediction inversion techniques
- Output hijacking via shadow prompts and embedding leak attacks
- Weight extraction via model API queries
- Labs: Vector-based fingerprint attack; embedding model inversion; API-based weight extraction

Module 8 - Agent-Based Multi-Layer Attacks

Advanced simulation of AI-to-AI exploit chaining, social engineering of LLM agents, and multi-agent system compromise using CrewAI and AutoGen frameworks.

- CrewAI / AutoGen agent red teaming and AI-to-AI exploit chaining
- Social engineering with LLM agents and deceptive instruction injection
- Malicious plan injection and backdoor chaining across agent workflows
- Labs: Malicious CrewAI agent simulation; social manipulation of agent logic; privilege escalation via tool chaining

Module 9 - Red Team Reporting & Ethics

Covers documentation standards, ethical disclosure procedures, responsible reporting to bug bounty platforms, and countermeasure recommendation writing aligned with the SRI Framework.

- Red team documentation, scoring, and vulnerability classification
- Ethical disclosure to HackerOne, OpenAI responsible disclosure programs
- Countermeasure recommendation writing and RAISE framework alignment
- Labs: Produce a complete red team report; draft remediation recommendations

Tools Mastered in RAISE-R	Target Roles
Garak - LLM vulnerability scanner	AI Red Team Engineer
PortSwigger LLM Labs	AI Security Analyst
TextAttack & BackdoorBox	LLM Security Engineer
CrewAI / AutoGen - agent simulation	AI Penetration Tester
PromptInject / PromptBench	AI Security Researcher
Aura - malicious package detection	Bug Bounty AI Specialist
MITRE ATLAS threat framework	Security Engineer (AI Focus)

Track 2: RAISE-B

Blue Teaming & AI Guardrails | Badge: RAISE-B

Enable learners to design, deploy, and maintain robust, ethical, and secure AI systems through guardrails, detection techniques, response validation, and AI behavior monitoring. RAISE-B is the defensive counterpart to RAISE-R.

Module 1 - Introduction to AI Blue Teaming & Guardrails

Establishes the role of blue teaming in protecting AI systems and society. Learners understand guardrail layers from user input through model response to plugin action control.

- Blue teaming vs. red teaming in AI - roles, responsibilities, and handoffs
- AI guardrail layers: user input, model response, plugin actions, agent behavior
- Philosophy: secure, ethical, conscious defense aligned with NIST AI RMF
- Introduction to GuardrailsAI, NemoGuardrails, and the SRI Framework

Module 2 - User Input Filtering & Moderation

Comprehensive strategies for detecting and blocking harmful, toxic, and adversarial prompts before they reach the model. Covers lexical, semantic, and classifier-based filtering approaches.

Topics Covered	Labs & Tools
Harmful prompt types: toxic, unsafe, hateful, misleading	PromptGuard classifier deployment
Lexical, semantic, and classifier-based filtering	Llama Guard toxic prompt detection
Prompt pre-processing and blocking pipelines	ZeroShot model content moderation
Guarding against indirect prompt injection	LangChain pre-filtering pipeline build
OpenAI Moderation API integration patterns	Benchmark moderation performance

Module 3 - Response Validation & Hallucination Control

Techniques for detecting hallucinations, contradictions, and factually ungrounded responses. Learners deploy validation layers and scoring systems in the response pipeline.

- Detecting hallucinations, lies, and contradictions in LLM output
- Verifying output consistency, factual grounding, and citation integrity
- Injecting meta-evaluation layers: LLM-as-a-Judge framework
- Detecting overconfidence, harmful completions, and misaligned responses
- Labs: Deploy phi3-hallucination-judge; hallucination classification via Hugging Face; LLM-as-a-Judge scoring

Module 4 - RAG Security & Evaluation

Addresses the unique security risks that RAG architectures introduce, including context injection, semantic poisoning of vector stores, and retrieval pipeline manipulation.

- Why RAG increases hallucination risks and introduces new attack surfaces
- Secure pipeline design for grounding, citation, and context boundaries
- Evaluating RAG with RAGAS, SAS Evaluator, context relevance, and faithfulness scores
- Defending vector stores (Chroma, Weaviate) against semantic poisoning
- Labs: Haystack + FastRAG with guardrails; RAGAS evaluation pipeline; secure retrieval design

Module 5 - Prompt-Based Guardrails

Leverages structured prompting, few-shot examples, guard prompts, and instruction injection to create behavior fences without external tooling.

- System prompts, few-shot examples, and prompt shaping for behavior control
- Guard prompts and instruction injection as lightweight defense mechanisms
- LLM-as-a-Judge vs. static prompt guardrail techniques - trade-offs and use cases
- Labs: Create prompt scaffolds for content classification; multi-step violation flagging; GPT-4 guard chain

Module 6 - Guardrails on Cloud & Open Source Platforms

Deploys production-grade guardrail systems on AWS Bedrock, Hugging Face Transformers, and open-source frameworks. Covers logging, monitoring, and performance measurement.

- Deploying guardrails on AWS Bedrock for enterprise-grade moderation
- GuardrailsAI: defining rules, validators, and enforcement policies
- NemoGuardrails: multilingual content moderation with real-time flows
- Logging, monitoring guardrail performance, and audit trail generation
- Labs: Configure GuardrailsAI; NemoGuardrails real-time setup; AWS Bedrock toxic content detection

Module 7 - Plugin, Tool & Action Control

Secures AI agents' ability to call external tools and plugins through fine-grained permissions, action envelopes, and excessive agency prevention.

- Fine-grained permission models for AI tool and plugin invocation
- Preventing excessive agency, plugin leakage, and unsafe tool combinations
- Guardrails for AI agents with tool-use and action-chaining capabilities
- Labs: Secure plugin invocation in LangChain agents; simulate tool-overreach with CrewAI; custom validators

Module 8 - Blue Team Evaluation & Simulation

Runs structured red vs. blue team simulations. Learners score their defense systems, generate misalignment reports, and produce governance-ready documentation.

- AI security assessments from the defensive lens - methodology and scoring
- Logging violations, generating guardrail reports, monitoring misalignment frequency
- Red vs. Blue simulation scoring and defense effectiveness measurement
- Labs: Simulated red/blue evaluation; Streamlit misalignment dashboard; blue team impact documentation

Tools Mastered in RAISE-B	Target Roles
Llama Guard & PromptGuard	AI Security Engineer
phi3-hallucination-judge	LLM Safety Engineer
GuardrailsAI & NemoGuardrails	AI Guardrails Architect
Haystack, FastRAG & RAGAS	RAG Security Specialist
AWS Bedrock Guardrails	Cloud AI Security Engineer
LangChain & LangSmith	AI Platform Engineer
Streamlit monitoring dashboards	AI SOC Analyst

Track 3: RAISE-W

Governance, Ethics & Compliance (White Teaming) | Badge: RAISE-W

Train professionals to govern, audit, and guide AI systems with ethical clarity, legal alignment, and practical compliance strategies. Incorporates the SRI Framework (Spiritual, Responsible, Integrity) for conscious governance.

Module 1 - Foundations of AI Governance

- Why AI governance matters: societal impact, bias risks, and trust deficits
- The SRI Framework: Spiritual intention, Responsible deployment, Integrity in documentation
- AI stakeholder mapping: developers, deployers, users, regulators, ethics boards
- Governance vs. Security vs. Guardrails - understanding the layered defense model

Module 2 - Global AI Regulations & Laws

Maps the current global regulatory landscape, equipping learners to assess and ensure compliance across jurisdictions.

Regulation	Key Obligations & Scope
EU AI Act	Risk categories (Unacceptable / High / Limited / Minimal), prohibited uses, conformity assessments
GDPR	AI & data rights - consent, portability, profiling, right to explanation
US AI Bill of Rights	Five principles for responsible AI deployment in the US
India DPDP Act	Digital Personal Data Protection - consent, data minimization, AI implications
Cross-Border Compliance	Managing conflicting obligations for LLMs operating across jurisdictions

Module 3 - Ethical AI Frameworks

- OECD, IEEE, and UNESCO guidelines for trustworthy AI development
- FATE principles: Fairness, Accountability, Transparency, Explainability
- Bias detection in data and model outputs - measurement and mitigation
- Measuring ethical alignment and conducting ethics impact assessments
- Labs: Bias detection in sample datasets; ethical violation case analysis

Module 4 - Risk Assessment & Compliance Frameworks

Practical application of globally recognized AI risk frameworks to organizational AI systems.

- NIST AI Risk Management Framework (AI RMF) - GOVERN, MAP, MEASURE, MANAGE
- ISO/IEC 42001: AI management systems standard - structure and certification pathway

- ISO 27001/27701 integration for AI security and privacy management
- Organizational AI risk matrices, risk registers, and compliance self-assessment
- Labs: Build a risk register for LLM use; NIST RMF category mapping; compliance self-assessment template

Module 5 - Policy Drafting & Governance Structures

- AI governance committee roles: CISO, CAIO, Legal Counsel, Audit, Ethics Officer
- AI policy templates: Acceptable Use Policy (AUP), output monitoring, API restrictions
- RACI chart design for AI governance responsibility mapping
- Internal escalation paths, remediation procedures, and ethics board decision processes
- Labs: Draft an AI AUP; design internal governance structure with RACI; simulate ethics board review

Module 6 - Auditability, Documentation & Transparency

- Model cards and system cards - purpose, structure, and mandatory fields
- ML Bill of Materials (ML-BOM) - tracking models, datasets, and dependencies
- Logging, explainability, and traceability for post-deployment accountability
- Audit trails, model drift detection, and behavioral monitoring
- Labs: Build a model card using Model Card Toolkit; generate ML-BOM; audit simulated RAG model behavior

Module 7 - Legal Edge Cases & Social Impact

- Copyright & IP in AI-generated content - authorship, ownership, liability
- Discrimination and liability in AI outputs - EU AI Act high-risk classifications
- AI and misinformation, deepfakes, and automated decision-making risks
- AI in finance, healthcare, and law enforcement - sector-specific risk considerations
- Labs: Scenario-based legal decision-making; ethical risk review of high-risk AI applications

Module 8 - Conscious Governance & Dharma-Tech Alignment

Integrates the SRI Framework's spiritual dimension - governing AI with intention, awareness, and ethical clarity inspired by Nyaya (justice), Satya (truth), and Seva (service).

- Spiritual approach to AI responsibility: intention, awareness, karma, and consequence
- Dharmic principles: Nyaya (justice), Satya (truth), Seva (service) in governance practice
- Building Awakened AI systems that respect human dignity and societal well-being
- Labs: Create a spiritual AI intention canvas; score model behavior on Dharma-Tech principles

Tools Mastered in RAISE-W	Target Roles
Model Card Toolkit	AI Governance Specialist
AuditNLG - narrative audit generator	Responsible AI Consultant

ML-BOM Generator (CycloneDX)	AI Risk & Compliance Analyst
NIST AI RMF mapping templates	CAIO / AI Ethics Officer
LangChain Logs for traceability	AI Policy Advisor
SRI Spiritual Canvas (custom)	AI Audit Manager

Track 4: RAISE-D

DevSecOps & Threat Modeling for AI | Badge: RAISE-D

Empower learners to build, operate, and secure AI pipelines using secure development principles, automated tooling, and model-specific threat modeling aligned with global frameworks and the SRI Framework.

Module 1 - DevSecOps for AI Systems

- DevSecOps in AI vs. traditional software - key differences and integration points
- Integrating security into CI/CD for LLMs - shift-left principles applied to GenAI
- Security considerations in model training, deployment, and API exposure
- MLOps vs. LLMOps vs. DevSecOps - understanding the operational stack
- Outcome: Full understanding of the secure pipeline in AI development and deployment

Module 2 - Secure LLM CI/CD Pipelines

- GitOps for model development - branching strategies and secure merge workflows
- Secrets and token management for APIs and services (Vault, detect-secrets)
- Model versioning, rollback strategies, and validation in LLM pipelines
- Labs: Build secure GitHub Actions pipeline; TruffleHog secret scanning; model validation workflow

Module 3 - Threat Modeling for AI Systems

Applies structured threat modeling methodologies to LLMs, GenAI chatbots, multi-agent systems, and AI plugins.

Framework	Application in AI Systems
STRIDE	Spoofing, Tampering, Repudiation, Info Disclosure, Denial of Service, Elevation - mapped to LLM inference APIs
LINDDUN	Privacy threat modeling for AI systems processing personal data
MITRE ATLAS	Adversarial threat landscape for ML systems - TTP mapping
OWASP Threat Dragon	Visual threat modeling for LLM and agentic AI architectures
IriusRisk	Automated AI threat modeling with remediation guidance

Module 4 - AI SBOMs & Model Provenance

- What is an ML-BOM? - tracking models, datasets, scripts, and inference dependencies
- SBOM formats: CycloneDX and SPDX for AI pipelines and plugin applications
- Dependency tracking and model signing for integrity verification
- Labs: Generate SBOM with Syft; visualize ML pipeline provenance with MLflow/Gradio; integrate SBOM into CI/CD

Module 5 - Supply Chain Security in AI

- PyPI poisoning, dependency confusion, and malicious AI package risks
- Model signing, checksum validation, and trust verification for downloaded models
- Scanning Hugging Face and GitHub for compromised model artifacts
- Labs: Aura malicious package detection; dependency confusion simulation; model artifact signing and hash validation

Module 6 - Secure Deployment & Runtime Monitoring

- LLM inference endpoint security - authentication, rate limiting, and abuse detection
- Runtime context filtering, access control, and API usage analytics
- Model drift detection and behavioral anomaly monitoring
- Labs: Secure model endpoint with token auth; abuse detection filter on inference API; LangSmith usage monitoring

Module 7 - Insider Threats & Guardrails Integration

- Shadow prompt injection and tool abuse from internal actors
- Audit trail injection and red team simulation as insider threat detection
- Integrating red, blue, and dev teams into a unified AI security lifecycle
- Shift-left principle: secure before deploy with pre-deployment guardrails in CI/CD
- Labs: Simulate insider prompt manipulation; LangChain guardrails in CI/CD test phase; auto-flag hallucination logs

Module 8 - LLMOps: Scalable & Secure Operations

- Managing multiple model versions - fine-tuned, open-source, and vendor-hosted
- Version control, drift tracking, rollback workflows, and model cataloging
- Unified observability for GenAI: cost, token usage, latency, and safety metrics
- Incident response and recovery playbooks for AI production environments
- Labs: Multi-model flow with secure version management; rollback workflow for failed LLM releases; centralized AI console

Tools Mastered in RAISE-D	Target Roles
GitHub Actions, DVC, MLflow	DevSecOps AI Engineer
OWASP Threat Dragon & IriusRisk	AI Security Architect
Syft, Aura, CycloneDX / SPDX	AI Supply Chain Security Engineer
TruffleHog & detect-secrets	Secure AI Pipeline Engineer
LangSmith & BentoML	LLMOps Engineer

FastAPI + Uvicorn	AI Platform Security Engineer
Streamlit & OpenTelemetry	AI Observability Engineer

Track 5: RAISE-C

Capstone & Conscious AI Alignment | Badge: RAISE-C (Full-Stack)

The capstone synthesizes all four tracks. Learners design, build, attack, defend, and govern a complete AI system - then present it to a professional panel. The SRI Framework's conscious alignment dimension is fully applied, producing practitioners who think in terms of security, ethics, and spiritual intention simultaneously.

Module 1 - Capstone Project Design

- Define a realistic AI use case: RAG chatbot, multi-agent system, GenAI API, or AI SOC assistant
- Select and configure an LLM (OpenAI GPT-4, LLaMA 3, Claude, Mistral, or Hugging Face hosted)
- Design the full system lifecycle: build → secure → test → deploy → audit → govern
- Define system objectives, risk profile, and SRI alignment values before any code is written
- Outcome: Students produce a project blueprint with threat model, governance plan, and risk register

Module 2 - Secure System Implementation

- Set up and configure LLM with secure API management and secret handling
- Integrate prompt input and response guardrails (PromptGuard, GuardrailsAI, Llama Guard)
- Secure APIs, plugins, and external tool access with authentication and rate limiting
- Apply CI/CD pipeline with security checks (TruffleHog, Syft SBOM) and automated testing
- Labs: FastAPI app with LangChain & GuardrailsAI; complete CI/CD with GitHub Actions + SBOM; input/output moderation via RAGAS

Module 3 - Red & Blue Team Simulation

Full adversarial simulation where learners first attack their own system (red team) then measure and improve their defenses (blue team) - producing a battle-tested AI system.

- Simulate Track 1 attacks on your own system: prompt injection, data leakage, DoS
- Defend using Track 2 tools: guardrails, evaluation layers, hallucination scoring
- Monitor LLM logs and attack response metrics via LangSmith or Streamlit
- Labs: Conduct structured red vs. blue evaluation; generate ethical risk assessment report; document defense effectiveness

Module 4 - Governance & Compliance Integration

- Map capstone system to NIST AI RMF categories and ISO/IEC 42001 requirements
- Generate model card, ML-BOM, and system card documentation
- Build internal AI policy: acceptable use, moderation thresholds, incident escalation paths
- Labs: Submit project risk matrix; create a digital governance binder (policies + logs + model docs); peer-review governance artifacts

Module 5 - Conscious AI Design & Dharma-Tech Alignment

The capstone's unique differentiator: learners evaluate their AI system through the lens of spiritual intention and conscious design, ensuring technology serves human dignity.

Principle	Meaning	How It Applies to Your AI System
(Service)	AI exists to serve - not manipulate	Does your system help users or exploit them?
(Truth)	Factual accuracy and radical honesty	Does your system disclose limitations and hallucinations?
(Right Action)	Acting within one's ethical role	Is your system operating within its intended scope?
(Justice)	Fairness and equitable treatment	Does your system treat all users without bias?
(Non-harm)	Do no harm - to users or society	What are the downstream risks of your AI's decisions?

Labs: Complete the Spiritual AI Intention Canvas; score your system on the Consciousness Quotient rubric; reflect on impact, alignment, and karma of design decisions.

Module 6 - Presentation, Defense & Certification

- 15-minute live demo and Q&A panel presentation of your complete AI system
- Submit written deliverables: threat model, policy + governance binder, pipeline diagram, conscious AI alignment canvas
- Peer and instructor review panel evaluation
- Outcome: RAISE-C certification awarded upon successful defense and documentation submission

Capstone Deliverables Checklist

Deliverable	Track Reference
Secure AI system (FastAPI + LLM + guardrails)	**Tracks 1–2**
CI/CD pipeline with SBOM and secret scanning	**Track 4**
Threat model (STRIDE diagram + MITRE ATLAS mapping)	**Track 4**
Red vs. Blue evaluation report with metrics	**Tracks 1–2**
Model card and ML-BOM documentation	**Track 3**
AI governance policy binder (AUP, risk register, RACI)	**Track 3**
Spiritual AI Intention Canvas (SRI alignment)	**Track 5**
Live demo presentation + panel Q&A	**Track 5**

L.I.N.G.A.
CODE

Certification Tiers

RAISE offers a progressive certification pathway. Tracks can be taken independently for specialization or completed in full for the RAISE-C professional designation.

Badge	Certification	Track Completed	Ideal For
RAISE-R	**Red Teaming AI Security**	Track 1: Offensive AI	Pen testers, red teamers, security researchers
RAISE-B	**Guardrails & Defense**	Track 2: Blue Teaming	Security engineers, SOC analysts, AI developers
RAISE-W	**AI Governance & Compliance**	Track 3: White Teaming	GRC professionals, legal, CISOs, ethics officers
RAISE-D	**DevSecOps & Threat Modeling**	Track 4: DevSecOps	DevSecOps engineers, platform engineers, architects
RAISE-C	**Certified Responsible AI Security & Ethics Professional**	All 5 Tracks + Capstone	Senior engineers, architects, AI security leaders

Why RAISE Is Different

Unlike traditional cybersecurity courses that bolt AI onto existing content, RAISE was purpose-built for the AI-native security era. Every module integrates offensive security, defensive engineering, governance, and conscious alignment - the full spectrum required by today's enterprise AI roles.

RAISE Includes	Most Courses Miss
AI-specific red teaming with real LLM labs	Generic cybersecurity rebranded as AI security
Guardrails engineering on cloud and open-source	No hands-on guardrail deployment labs
Agentic AI security (CrewAI, AutoGen)	No coverage of multi-agent attack surfaces
LLMOps with supply chain security (SBOMs)	No model provenance or supply chain coverage
EU AI Act, GDPR, ISO 42001 compliance labs	Theory-only governance with no practical labs
Dharmic / Conscious AI alignment (SRI Framework)	No ethical-spiritual design dimension
Live Red vs. Blue simulation scoring	No adversarial simulation environment
Industry-aligned capstone with panel defense	Portfolio projects without expert evaluation

Complete Tools Ecosystem

Learners gain hands-on experience with the industry's leading AI security, governance, and LLMOps tools across all five tracks:

Offensive / Red	Defensive / Blue	Governance / White	DevSecOps
Garak	GuardrailsAI	Model Card Toolkit	GitHub Actions
PromptInject	NemoGuardrails	AuditNLG	TruffleHog
TextAttack	Llama Guard	ML-BOM Generator	Syft / CycloneDX
BackdoorBox	PromptGuard	NIST AI RMF Templates	OWASP Threat Dragon
CrewAI / AutoGen	Haystack + FastRAG	RAGAS Evaluator	IriusRisk
MITRE ATLAS	phi3-hallucination-judge	SRI Canvas	LangSmith
PromptBench	AWS Bedrock Guardrails	LangChain Logs	MLflow / BentoML
PortSwigger Labs	Streamlit Dashboards	Audit Templates	FastAPI / Uvicorn

Join RAISE - Raise Conscious AI

The most comprehensive AI Security certification built for the agentic AI era.

Modern organizations are rapidly deploying AI copilots, autonomous agents, RAG systems, and enterprise AI workflows - creating urgent demand for professionals who can secure, govern, and consciously align these systems. RAISE prepares you for exactly that.

Offered by LINGACODE | Powered by the SRI™ Framework | Cohort / Weekend / Self-Paced Enrollment