AI SECURITY SERVICES

Secure the AI you use, build, and deploy.

Wherever you are on your AI journey, whether assessing exposure, building new capabilities, or scaling production workloads, Arctiq meets you there. We assess where you need to be, design the controls, build them, and run them. Platform-agnostic by design, we fit in your environment instead of you adhering to our model.

Talk to an AI security Advisor →

THE PROBLEM

AI security isn't a product. It's a matrix.

AI doesn't fit cleanly into one aspect of your security program. It touches governance, data, identity, application security, security operations, and incident response, all at once. Most organizations arrive at one of three crossroads:

> *We're already using AI and we don't know our exposure.*

> *We're building AI products and need to secure them.*

> *We want to leverage AI, but we don't know where to start.*

Regardless of the path you're on, we meet you there and walk you through the full maturity continuum. **Once stable, we can hand you the keys when you're ready, or we can run it for you.**

THE EIGHT DOMAINS

The eight domains of AI security

Mature AI security programs eventually address the following eight domains. We have services that map to each, and we'll advise how to prioritize the engagements detailed below.

01
AI Governance & Risk Management
Policy, acceptable-use, model risk management, and alignment to NIST AI RMF, ISO 42001, and the EU AI Act. Third-party AI risk and ethics committee enablement included.

02
AI Discovery & Posture (AI-SPM)
Find the shadow AI. Build an AI bill of materials, harden your model registry, and catch misconfigurations across Bedrock, Vertex, Azure OpenAI, and SageMaker.

03
Data Security for AI
Classify training data. Protect RAG corpora. Apply DLP to prompts and responses. Track data lineage so sensitive information doesn't leak through your models.

04
Identity & Access for AI
Govern human-to-AI access and the non-human identities behind autonomous agents. OAuth and MCP scope design, least-privilege boundaries, and secrets management for model APIs.

05
Model & Application Security
OWASP LLM Top 10 coverage, secure SDLC for AI-enabled apps, model supply chain scanning, and AI red teaming against adversarial inputs.

06
Runtime Protection
LLM firewalls and guardrails. Prompt injection defense, output filtering, jailbreak detection, and validation of every action your agents take.

07
Detection & Response
SIEM use cases tuned to AI-specific TTPs. UEBA for AI service abuse. Incident response playbooks for model theft, prompt injection chains, data exfiltration via LLMs, and agent compromise.

08
Trust, Safety & Assurance
Bias monitoring, hallucination detection, content moderation, human-in-the-loop controls, and continuous model evaluation.

HOW WE ENGAGE

The five-step service lifecycle

One framework that scales from a single advisory engagement to a multi-year managed program.

1

Assess: Know where you are.

Current-state maturity, AI inventory and discovery sprint, threat modeling against your actual use cases, gap analysis against a chosen framework (NIST AI RMF is our defensible default).

OUTPUT maturity score, prioritized risk register, target state.

2

Strategize: Decide where you're going.

Governance framework selection, target operating model (who owns what across security, data, ML, legal), risk appetite, multi-year roadmap with budget envelopes, platform selection criteria.

OUTPUT strategy document, roadmap, business case.

3

Architect: Design the controls.

Reference architectures for every AI use case pattern: RAG, fine-tuning, agentic, SaaS AI consumption. Control mappings, platform evaluation, integration design with your existing security stack.

OUTPUT architecture artifacts, control matrix, platform shortlist.

4

Implement: Build the controls.

Hands-on engineering. Governance program build-out, AI-SPM deployment, guardrail integration, detection engineering, IR playbook authoring and tabletop exercises, identity controls for agents.

OUTPUT working controls, validated detections, exercised playbooks.

5

Operate: Run the controls.

Co-managed or fully managed. We run the platform, tune the detections, handle AI-specific incidents, lead ongoing red teaming, and operate the governance program: model intake reviews, risk assessments, the works.

OUTPUT continuous coverage and reporting.

PLATFORM-AGNOSTIC BY DESIGN

Your stack, your call.

Every AI security domain has an evolving and maturing market category supporting it, and we are unbiased when solutioning. Arctiq's methodology, integration expertise, and operational muscle stay constant. The tooling we develop integrates based on your existing tech stack, budget, and cloud footprint. If you already have preferences, we work with them. If you don't, we'll recommend what fits.

> **" We bring the methodology, the integration expertise, and the operational muscle. You bring the environment, and we fit into it.**

DOMAIN	PLATFORM CATEGORY
Governance	AI GRC platforms
Discovery & Posture	AI Security Posture Management (AI-SPM)
Data Security	DSPM with AI module
Identity (Agents & NHI)	Non-human identity platforms
Model & AppSec	Model scanning, AI SAST
Runtime	LLM firewall / guardrails
Detection & Response	SIEM with AI content packs
Red Teaming	AI red team platforms
Observability	LLM observability

ENGAGEMENT PACKAGES

Pick your starting point.

Most customers start with a narrow focus and expand.

01

AI Security Quick Start

Discovery sprint plus maturity assessment, delivered with a 90-day roadmap. Low commitment, high signal. The fastest way to know where you are on the AI continuum.

2–4 weeks

02

AI Governance Foundation

Policy, acceptable-use, model intake process, risk framework, and executive committee enablement. Often the package the board or audit committee asks for by name.

4–8 weeks

03

AI-SPM Deployment & Tuning

Platform stand-up, integration with your environment, and an initial findings remediation playbook so the first wave of alerts are actionable, not noise.

6–10 weeks

04

AI Detection & Response Build

SIEM content development, IR playbook authoring, tabletop exercise, and integration with your existing SOC runbooks. Built to trigger only when it should.

8–12 weeks

05

Managed AI Security

We operate the platform, run the governance cadence, tune detections, lead IR, and run periodic red teaming. Continuous coverage, single escalation.

Ongoing

Wherever you are in your AI security journey, we'll meet you there.

Assess where you need to be. Design the controls. Build them. Run them. Platform-agnostic, end to end.

Book a Quick Start Assessment →

arctiq.com/ai-security-quick-start-assessment

ARCTIQ