AI in the Enterprise: A Framework-Aligned Roadmap for AI & Security Leaders

For AI-native SaaS & product teams · enterprises deploying AI internally · anyone whose AI must clear enterprise security & procurement reviews

> AI systems do things they are not supposed to. Regulators and enterprise buyers now expect you to **prove they don't - against the EU AI Act, NIST AI RMF and ISO/IEC 42001**.

THE FRAMEWORKS THAT NOW DEFINE AI ACCOUNTABILITY

EU AI Act

BINDING · EXTRATERRITORIAL

GPAI rules in force (Aug 2025); high-risk by Dec 2027. Fines up to **€35M / 7%** of turnover. Applies if your AI touches the EU.

NIST AI RMF + GenAI Profile

US BASELINE · VOLUNTARY

Govern · Map · Measure · Manage. GenAI Profile (AI 600-1) adds 12 risks incl. prompt injection & data poisoning.

ISO/IEC 42001

AUDITABLE · PROCUREMENT BENCHMARK

The certifiable AI management system. Crosswalks to NIST & EU AI Act - and now a common **vendor-review requirement**.

WHAT EVERY FRAMEWORK CONVERGES ON

- Accountability sits with **you**, not the model
- Risk-based controls across the lifecycle
- Human oversight of AI decisions
- Documented evidence, not assertions

THE AI IMPLEMENTATION ROADMAP - 5 PHASES BUYERS & AUDITORS WILL INSPECT

PHASE 0 — Govern

Own the risk before AI ships

AI policy & roles (ISO 42001) · NIST **GOVERN** · EU AI Act risk tiering · AI inventory

PHASE 1 — Build & Test

Prove it's safe before it ships

Red-teaming - OWASP LLM Top 10 & MITRE ATLAS · NIST **MAP/MEASURE** · EU AI Act technical docs

PHASE 2 — Deploy & Enforce

Control what the AI can access & do

Human oversight (EU AI Act Art. 14) · access control & agent guardrails · shadow-AI controls · NIST **MANAGE**

PHASE 3 — Monitor & Respond

See it, contain it, report it

Logging & post-market monitoring (EU AI Act Art. 12/73) · serious-incident reporting · drift & output checks

CONTINUOUS — Assure & Disclose

Sustain trust; answer buyers & auditors

ISO 42001 audit & continual improvement · transparency (Art. 50) · disclosures · questionnaire evidence

THE READINESS GAP - WHY THIS IS URGENT, NOT THEORETICAL

88% had a confirmed or suspected AI-agent security incident this year

80% report moderate-to-pervasive shadow AI

18% have a formal AI security policy

21% have mature governance for AI agents

Source: Gravitee (survey of 919 organizations), ISACA & industry surveys, 2026. Adoption has outpaced control - that gap is the exposure.

THE GAP YOUR EXISTING STACK CANNOT CLOSE

Your IAM/PAM, WAF and API gateways, DLP/CASB, SIEM and EDR are built for **deterministic** software. None can catch a **legitimately-authorised AI agent, driven by untrusted input, doing something inside its permissions but outside your policy** - the dominant 2026 failure mode (prompt injection, the "lethal trifecta", agentic data exfiltration).

IAM / PAM	WAF / API gateway	DLP / CASB	SIEM / EDR
identity authenticated, in scope	request well-formed	sanctioned tool & channel	no known-bad signature
WAVED THROUGH	WAVED THROUGH	WAVED THROUGH	WAVED THROUGH

Every layer says "yes." None can see the AI's **intent vs your policy**.

Trampolyne makes sure your AI remains within intended bounds - and proves it - from pre-deployment testing, through runtime, to audit-ready evidence.

COVERAGE MAP - TRAMPOLYNE AI PRODUCTS × ROADMAP PHASE

PRODUCT	0 · GOVERN	1 · BUILD & TEST	2 · DEPLOY & ENFORCE	3 · MONITOR & RESPOND	CONTINUOUS · ASSURE
AI Red-Teaming		●			●
Shadow AI Detection & Runtime Control			●	●	○
Enterprise AI Security & Runtime Control			●	●	○
AI Compliance Assistant				○	●

● core control & evidence Trampolyne AI provides · ○ supporting input. **Each phase also demands policies, documentation and processes you own** - Trampolyne AI supplies the technical controls and audit evidence that make them credible, not the entire framework. Phase 0 (Govern) stays with you; we feed its inventory and risk classification.

WHAT EACH PRODUCT SOLVES - MAPPED TO THE FRAMEWORKS

AI Red-Teaming ↗ AWS MARKETPLACE

Continuous adversarial testing - scheduled cadence + on every material change. 27+ attack classes (prompt injection, jailbreak, tenant-isolation bypass, MCP/tool-call hijack, exfiltration). Working exploits + fixes, mapped to OWASP LLM Top 10 & MITRE ATLAS - the report your enterprise buyers ask for. No code changes.

Supports: OWASP LLM Top 10 · MITRE ATLAS · NIST MAP + MEASURE · EU AI Act tech docs · Security reviews

Enterprise AI Security & Runtime Control ↗

Sits inline, evaluates every request **before the model or agent acts** - RBAC/ABAC/PBAC/NGAC enforcement in milliseconds. Allow / block / modify / redact. A deterministic guardrail over a non-deterministic model. No model rewrites.

Supports: EU AI Act Art. 14 Human Oversight · NIST MANAGE · ISO 42001 controls · Agent access control

Shadow AI Detection & Runtime Control ↗

Sits between employees and every public AI tool (ChatGPT, Claude, Gemini, Copilot, APIs, MCP). Classifies data by type / provenance / role, blocks sensitive data before it leaves, logs every interaction. Exception workflow preserves productivity.

Supports: Data leakage to public AI · GDPR / DPDP · NIST MANAGE · ISO 42001 data controls

AI Compliance Assistant ↗

Turns enforcement + red-team + runtime evidence into review-ready answers. Continuous per-decision audit trail; checks alignment; packages evidence for auditors, your board and enterprise-buyer security questionnaires.

Supports: ISO 42001 audit evidence · EU AI Act logging & transparency · NIST GOVERN evidence · Security questionnaires

Outcomes, not threat counts.
We stay to help you turn outputs into evidence.

We're your team.
Partners take us into their enterprise-client CISO meetings.

Enforcement + proof.
Human oversight, accountability with you, evidence on demand.

See if we're a fit - a 20-minute call is enough to map your AI estate to the roadmap.

Book a fit call →

AI Red-Teaming · Enterprise Runtime Control · Shadow AI Control · AI Compliance Assistant · How it works

Sources: EU AI Act (Reg. 2024/1689) - GPAI in force Aug 2025, high-risk obligations Dec 2027, fines to €35M / 7% · NIST AI RMF 1.0 + Generative AI Profile (NIST AI 600-1) · ISO/IEC 42001:2023 AI Management System · US state laws - Texas TRAIGA (Jan 2026), Colorado & California ADMT (Jan 2027) · India DPDP Act rules (Nov 2025) + India AI Governance Guidelines (Nov 2025) · OWASP LLM Top 10 (2025) & Agentic Top 10 (2026) · MITRE ATLAS. Readiness stats: Gravitee / ISACA / industry surveys, 2026.

Positioning collateral, not legal advice. Trampolyne AI helps you align with regulatory expectations; it does not issue certifications.