OneCompliant

Runtime AI Governance for Regulated Enterprises

EXECUTIVE BRIEFING

AI Governance for Regulated Enterprises

Governing AI adoption under the EU AI Act, NIS2, GDPR, and the NIST AI RMF — without surrendering control.

Kevin Wade Stout
AI Security Expert · Founder, OneCompliant

Version 1.0

EXECUTIVE SUMMARY

AI is already inside your enterprise.

Artificial intelligence has moved from pilot projects to daily use inside regulated enterprises — often faster than governance can adapt. Employees use public AI tools, business units embed AI in workflows, and development teams ship AI-enabled features. Meanwhile, traditional security controls remain blind to prompts, models, autonomous agents, and AI-driven data movement.

For boards and executives in regulated sectors, the question is no longer **whether** to adopt AI, but whether the organisation can **demonstrate that its AI is governed, controlled, and defensible** — under the EU AI Act, NIS2, GDPR, and the NIST AI RMF.

This briefing sets out the governance gap, the regulatory drivers, and a practical, sequenced operating model — **Assess, Govern, Enforce** — that enables responsible AI adoption while maintaining accountability and operational resilience.

The shift boards are feeling

"The question has moved from *Can we use AI?* to *Who owns the AI risk — and can we prove it is controlled?*"

THE GOVERNANCE GAP

Your existing controls were not designed for AI.

Decades of investment secured email, endpoints, servers, and networks. None of it was built to see — let alone stop — what AI introduces. As a result, traditional controls quietly pass AI risk straight through.

Built to protect	Blind to
• Email	• Prompt injection
• Endpoints	• Model misuse
• Servers	• Autonomous AI agents
• Networks	• AI data leakage
	• Runtime AI governance

Three exposures every regulated enterprise now carries

- **Shadow AI** — staff use unsanctioned AI tools for sensitive work, with no visibility and no audit trail.
- **No runtime enforcement** — policy documents do not stop source code or customer data being pasted into a public model. Only runtime controls do.
- **Zero audit visibility** — when a regulator asks what data went to which model and when, most enterprises cannot answer.

Most organisations already have AI policies, awareness, and risk assessments — but no runtime control. That is the gap.

THE REGULATORY LANDSCAPE

Four frameworks now converge on AI accountability.

EU AI Act	Risk-based obligations for AI systems. **AI literacy duties (Article 4) are already in force** for organisations deploying AI.
NIS2	Essential and important entities must manage and report cyber risk — including risk introduced by AI in operational systems.
GDPR	Personal data flowing into AI engages records of processing (Art. 30), security of processing (Art. 32), and DPIAs (Art. 35) for high-risk use — including employee monitoring.
NIST AI RMF	A widely adopted reference for AI risk: Govern, Map, Measure, Manage — useful as a common language with boards and partners.

Regulatory obligations apply regardless of who — or what — wrote the software, or which business unit deployed it. Compliance without operational controls is not compliance.

A PRACTICAL OPERATING MODEL

Assess. Govern. Enforce.

A defined, repeatable sequence takes an organisation from AI risk exposure to governed, audit-ready AI operations. Each stage builds on the one before — you cannot enforce at runtime what you have not first governed.

01 · ASSESS

Understand exposure and gaps

Identify where AI is used, what data is at risk, and where you are exposed under the EU AI Act and NIST AI RMF.

02 · GOVERN

Establish policies, controls, accountability

Define policy domains, authorisation boundaries, and the rules that are enforced at runtime.

03 · ENFORCE

Apply controls at runtime, generate evidence

A governed AI gateway inspects, controls, and logs every prompt, model interaction, and data flow in real time.

A fourth step: Insure

Once AI is governed and controlled, governance maturity and AI attack surface can be translated into an underwriting-grade risk score — so residual AI risk can be transferred, not just managed. This is an unusual and valuable position: most AI governance stops at compliance.

WHAT GOOD LOOKS LIKE

A defensible AI posture, in practice.

- **Visibility** of AI usage across business units, including shadow AI.
- **A sanctioned AI gateway** — the approved entry point for enterprise AI use.
- **Runtime inspection and redaction** of sensitive data before it reaches a model.
- **An immutable audit trail** that can answer the regulator's questions.
- **Controls mapped** to the EU AI Act, NIS2, GDPR, and the NIST AI RMF.
- **AI literacy** embedded across the workforce (EU AI Act, Article 4).

A pragmatic first 90 days

- **Weeks 1–4:** AI risk assessment — map usage, data exposure, and regulatory gaps.
- **Weeks 4–8:** quick-win controls — a governed gateway and high-risk data safeguards.
- **Weeks 8–12:** governance framework and awareness — policy, accountability, and AI literacy.

HOW TO ENGAGE

Start with an assessment.

OneCompliant provides runtime AI governance and governed AI access for regulated enterprises. We work directly with CISOs, CIOs, and security leaders — no sales intermediary — and begin with a structured assessment of your AI risk posture, regulatory gaps, and governance readiness.

About the author

Kevin Wade Stout is the founder of OneCompliant and an AI security practitioner with 20+ years in regulated enterprise and critical-infrastructure security. His background includes leading a 60+ person global information security organisation at Merck KGaA, embedded AI security leadership at a Tier-1 European telecommunications operator, and senior roles at the U.S. Department of Defense and Department of Homeland Security with Top Secret/SCI clearance. CISM-certified.

Contact kevin@onecompliant.net

Web onecompliant.ai

Location Bratislava, Slovak Republic — EU