CYBERSECURITY RISK & COMPLIANCE

AI Will Move Fast. **Governance Must Move First.**

Governance-First AI Security for Every Organization Deploying AI

AI Security Governance – a structured, executive-ready program to establish clear ownership, policy, risk management, infrastructure controls, and technical assurance for AI systems.

$4.9M	89%	9 seconds	70%
average cost of a data breach	of organizations lack formal AI governance policies	for an AI agent to destroy 3 months of company data	of breached orgs reported major operational disruption

THE CHALLENGE

Most organizations are deploying AI faster than they are governing it.

> AI incidents are not technical failures alone. They are leadership failures when ownership, controls & decision rights are unclear.

Real Incident, April 2026: An AI coding agent destroyed an automotive SaaS company's entire production database and 90 days of backups in 9 seconds. No cyberattack, no malicious actor—just an AI agent with unconstrained permissions and no guardrails.

Artificial intelligence is rapidly embedded in core clinical, financial, academic, and operational processes. Unlike traditional IT risk, AI risk combines cybersecurity, privacy, legal, and regulatory exposure—and requires governance before scale.

Regulators, plaintiffs' attorneys, auditors, and insurers are already asking hard questions. Organizations that answer them confidently are those that have built governance structures before they were forced to.

THE SERVICE

AI Security Governance Program

Executives are asking:

- Who owns AI risk and who is accountable?
- What AI use is acceptable for employees?
- How do we assess third-party AI vendors?
- How do we evaluate model risk and data leakage?
- Is our infrastructure configured to prevent autonomous AI actions?

What makes this different:

- Executive-level governance, not experimental AI policy
- Security, risk, compliance, and legal fully integrated
- Controls mapped to how AI is actually used today
- Infrastructure assessment to prevent rogue-AI incidents
- Built for Board, auditor, and regular scrutiny

1 AI Governance

Identify the AI Risk Owner, establish the Security Steering Committee, & define enterprise-level AI security prinicples.

2 Acceptable Use

Update AUP to explicitly address AI: approved tools, data handling, IP protection & enforcement alignment.

3 AI Training

Build staff capability on responsible AI use & threats: prompt injection, jailbreaks, & shadow AI.

4 Vendor Risk

Extend third-party risk assessments to AI vendors, integrated into the GRC platform with continuous monitoring.

Let's Put Governance Around Your AI Program

Your challenges are unique, our approach is tailored. Start a conversation, leave with a clear path forward.

Gui Cozzi, *Cybersecurity Risk & Compliance Director* | Visit Profile →

DEANDORTON

CYBERSECURITY RISK & COMPLIANCE

Structure, Accountability & Security

A complete program built for the Board, the auditor, and the regulator

DEANDORTON

CORE COMPONENTS

Six Components. One Defensible Program.

1 AI Governance

Formalize AI Risk Owner, activate Security Steering Committee & define AI security prinicples with clear decision rights.

2 Acceptable Use

Update AUP: approved vs. prohibited AI use, data handling, IP protection & enforcement alignment.

3 Vendor Risk Management

Extend third-party risk assessments to AI vendors. Evaluate model training data, privacy controls & security posture in GRC.

4 AI Risk Assessment

NIST AI-RMF assessments across governance, data, model & operational domains—tracked in your GRC platform.

5 LLM Penetration Testing

Offensive: Adversarial testing aligned to MITRE ATLAS—prompt injection, data leakage, model abuse & evasion techniques.

Configuration: Access controls, logging, model configuration & deployment architecture review.

6 Infrastructure & Controls Assessment

Assess whether AI agents operate within appropriate permission boundaries. Review least-privilege controls, confirmation guardrails, audit logging, backup architecure & environment segmentation to prevent autonomous destrucive actions.

AI RISK IN THE REAL WORLD

This Is Not an Industry-Specific Problem. It's a Leadership Problem

Any organization using AI faces these risks. The scenarios below are illustrative—the specifics change, the accountability gap does not.

HEALTHCARE EXAMPLE

Shadow AI exposes PHI: Staff use an unapproved LLM to summarize patient notes. No AI data-handling policy exists. PHI is retained in the vendor's model. OCR inquiry follows.

FINANCIAL SERVICES EXAMPLE

Agent with Unconstrained Access: An automated workflow agent is granted broad database permissions. A logic error triggers mass deletion of transaction records. Regulators demand an incident report.

ANY ORGANIZATION EXAMPLE

Vendor AI Trains on Your Data: A productiviy tool uses client inputs to improve its model. No contractual restriction exists. Client data becomes part of a commercial AI training dataset.

Leadership Will Gain: (1) clear ownership and governance of AI risk and (2) LLM pen test results mapped to MITRE ATLAS with remediation

ENGAGEMENT MODEL

Four Phases. One Cohesive Engagement.

Typical 12–16 week engagement, customized to your pace and organizational readiness.

1 Discovery & Scoping

Catalog AI tools & vendors, assess governance gaps & regulatory exposure.

2 Design & Build

Develop governance framework, update AUP, configure GRC, build training curriculum

3 Execute & Test

Deliver LLM pen testing, infrastructure review & vendor risk assessments

4 Report & Sustain

Executive report, MITRE ATLAS mapping, remediation roadmap & monitoring plan