AI Security & Governance

AI risk management and governance—before auditors and regulators ask

Overview

We embed AI security and governance into your stack: risk management, model governance, data protection, bias and explainability, audit readiness, and secure architecture patterns—so you deploy with confidence.

Business problems we address

- No structured AI risk management or policy frameworks—AI is going live without consistent governance.
- Model governance and supply chain assurance are unclear; third-party models and data flows lack oversight.
- Data protection and residency requirements block or complicate AI deployment across regions.
- Bias and explainability are not addressed; regulators and legal expect evidence of fairness and transparency.
- Audit readiness is missing—no documentation, runbooks, or evidence for auditors and procurement.

How we deliver

We start with an AI risk and governance assessment: map AI assets, data flows, and risk posture; define policy frameworks and ownership. We then design and implement secure architecture patterns: data protection and residency, model governance and supply chain assurance, bias and explainability (XAI) monitoring, and audit trails. We work alongside your security, legal, and engineering teams so controls are operational and documented for audit readiness.

Delivery steps

1. **AI risk & governance assessment**
   Map AI assets, data flows, and risk posture; define policy frameworks.
2. **Design secure architecture**
   Data protection, model governance, bias and explainability, audit trails.
3. **Implement & document**
   Deploy controls, runbooks, and evidence for auditors and procurement.

Security & governance

AI security and governance are the core of this service. We implement data protection and residency, model governance and supply chain assurance, bias detection and explainability (XAI), and policy frameworks aligned to EU AI Act, NIST AI RMF, and your internal standards. You get secure architecture patterns, documented controls, runbooks, and evidence for auditors and procurement—so you are audit-ready from day one.

Who should use this service

Ideal for CISOs, legal, compliance, and risk owners who need AI risk management, model governance, and audit readiness—and for engineering leaders who want to ship AI with policy frameworks and secure architecture patterns built in. We serve regulated industries and enterprises that cannot afford to retrofit security or governance after launch.

Example use case

A healthcare client required EU AI Act alignment and audit readiness before rolling out an AI-assisted triage system. We implemented AI risk management, model governance, data protection controls, bias and explainability monitoring, and policy frameworks—with secure architecture patterns and full documentation. Internal compliance and legal sign-off were achieved in 8 weeks.

Key capabilities

- AI risk management and policy frameworks
- Model governance and supply chain assurance
- Data protection and residency controls
- Bias detection and explainability (XAI)
- Audit readiness and documentation
- Secure architecture patterns