

How to Choose an AI Governance Platform

A practical way to pick AI governance software: know which layer you need, map it to your rules, then test inventory, risk workflow, evidence, integrations and security.

Updated October 29, 2026

AI governance software covers many products. Some keep a register of AI use cases. Some test models. Some block bad prompts. This guide helps you choose the right kind first, then the right vendor.

Step 1: Decide which layer you need

AI governance tools fall into four layers. Each does a different job.

Layer	What it does	Typical output
--- --- ---	Governance platform	Inventory, risk assessment, policy, approvals, evidence
---	AI register	risk records, audit reports
---	Guardrails	Screen inputs and outputs at run time
---	Blocked or redacted prompts and responses	Red teaming and evaluation
---	Attack and test a system before and after release	Findings, scores, test reports
---	Monitoring and observability	Watch models in production
---	Traces, drift alerts, performance metrics	

A governance platform does not stop a prompt injection. A guardrail does not prove compliance to an auditor. Most programs need more than one layer. Read [the glossary](#) if these terms are new.

Step 2: Map your rules

Write down the rules that apply to you. Then ask each vendor to show how its product supports them.

- **EU AI Act.** Main high-risk duties for Annex III systems apply from 2 December 2027 ([Article 113](#)). Transparency duties under Article 50 already apply. See [EU AI Act deadlines](#).
- **ISO/IEC 42001.** A management system standard for AI. It specifies requirements to establish, implement, maintain and continually improve an AI management system ([ISO](#)).
- **NIST AI RMF.** A voluntary framework with four functions: govern, map, measure and manage ([NIST AIRC](#)).
- **US banking.** SR 26-2 replaced SR 11-7 on 17 April 2026. Generative AI and agentic AI are outside its scope ([Federal Reserve](#), [OCC](#)).
- **Hiring in New York City.** Local Law 144 requires a bias audit of an automated employment decision tool within one year of use ([NYC DCWP](#)).

[ISO 42001 vs NIST AI RMF vs EU AI Act](#) compares the three.

Step 3: Score the core functions

Use this list in demos and pilots.

Inventory

- Can it find AI use you did not register, such as SaaS features and API calls?
- Does it record owner, purpose, data, vendor and risk tier for each use case?
- Does it cover third-party AI, not only models you build?

Risk and compliance workflow

- Can you edit the risk questions and approval steps?
- Does it map controls to more than one framework at once?
- Does it track who approved what, and when?

Evidence

- Can it export audit-ready reports?
- Does it pull test and monitoring results from other tools, or only from its own?
- Does it keep logs for as long as your rules need? Deployers of high-risk systems must keep logs for at least six months ([Article 26](#)).

Technical checks (if included)

- Which tests run on LLMs, and which on classic ML models?
- Does it cover the [OWASP Top 10 for LLM Applications](#)?
- Can it test for bias with your own protected attributes?

Step 4: Check fit and security

- **Integrations.** Your model registry, cloud ML service, ticketing tool, GRC tool and identity provider.
- **Deployment.** SaaS, private cloud or on-premise. Data residency for EU data.
- **Data access.** What data the tool reads, where it stores it and for how long.
- **Security.** SSO, role-based access, audit logs, recent third-party audit reports.
- **Price.** Most governance vendors do not publish prices. See [how AI governance software is priced](#).

Step 5: Run a pilot

1. Pick three real use cases: one high-risk, one LLM app, one vendor tool.
2. Load them into each shortlisted product.
3. Run one full risk review from intake to approval.
4. Export the evidence an auditor would ask for.
5. Ask the people who will use it every day to score it.

Use [questions to ask AI governance vendors](#) in your RFP.

Sources

- [Article 113](#) and [Article 26](#), [artificialintelligenceact.eu](#), accessed 2026-10-29

- [ISO/IEC 42001:2023](#), International Organization for Standardization, accessed 2026-10-29
- [AI RMF Core](#), NIST Trustworthy and Responsible AI Resource Center, accessed 2026-10-29
- [SR 26-2: Revised Guidance on Model Risk Management](#), Federal Reserve Board, 17 April 2026, accessed 2026-10-29
- [OCC Bulletin 2026-13: Model Risk Management: Revised Guidance](#), Office of the Comptroller of the Currency, accessed 2026-10-29
- [Automated Employment Decision Tools \(AEDT\)](#), NYC Department of Consumer and Worker Protection, accessed 2026-10-29
- [OWASP Top 10 for LLM Applications 2025](#), OWASP GenAI Security Project, accessed 2026-10-29

Frequently asked questions

Do I need one platform or several tools?

Often several. A governance platform keeps the inventory, risk assessments and evidence. Guardrails, red-teaming tools and monitoring tools do the technical checks. Choose tools that can send their results to the governance record.

Should the platform map to the EU AI Act?

If you place or use AI in the EU, yes. Ask the vendor to show how it classifies risk tiers, tracks deployer duties under Article 26 and keeps logs. Also check that the mapping reflects the 2026 Digital Omnibus dates.

How long should a pilot run?

Long enough to load real use cases and finish one full risk review. Use your own systems, not vendor demo data.