

Questions to Ask AI Governance Vendors

An RFP and security questionnaire for AI governance, guardrail, red-teaming and monitoring vendors. Each question says what a good answer looks like.

Updated October 29, 2026

A demo shows what a product does on a good day. Written answers show what the vendor will commit to. Use these questions in your RFP. Ask for evidence, not adjectives.

This guide is general information. It is not legal advice.

Scope and fit

1. Which layer does your product cover: inventory and risk workflow, guardrails, red teaming, or monitoring? What is out of scope?
2. Which AI types do you support: classic ML models, LLM applications, agents, third-party AI features?
3. Can you find AI use that is not registered yet? How?
4. Show three customers of our size and sector. Can we speak to one?

A good answer names clear limits. Be careful with a vendor that claims every layer equally well.

Regulatory mapping

1. Which frameworks do you map to out of the box? EU AI Act, ISO/IEC 42001, NIST AI RMF, others?
2. Does your EU AI Act content reflect Regulation (EU) 2026/1744 (the Digital Omnibus on AI)? For example, Annex III high-risk duties now apply from 2 December 2027 ([Article 113](#)).
3. How fast do you update mappings when a rule changes? Give a recent example.
4. How do you help a deployer meet Article 26 duties, such as human oversight, monitoring and log retention of at least six months ([Article 26](#))?
5. For US banks: how does your model risk content reflect SR 26-2, which replaced SR 11-7 on 17 April 2026 ([Federal Reserve](#))?
6. For hiring tools in New York City: do you support the Local Law 144 bias audit and notice duties ([NYC DCWP](#))?

A good answer shows dates and a change log. A mapping that still shows 2 August 2026 as the Annex III date is out of date.

Technical testing

1. Which risks do your tests cover? Map them to the [OWASP Top 10 for LLM Applications 2025](#), from LLM01 Prompt Injection to LLM10 Unbounded Consumption.

2. Can we add our own tests, attack prompts and evaluators?
3. How do you test bias? Which metrics? Can we use our own protected attributes?
4. What is the added latency of your run-time guardrails? Show it on our traffic, not a benchmark.
5. What is your false-positive rate on our sample data?

Evidence and audit

1. What does an auditor receive? Show a sample export.
2. Is every approval, change and override logged with user and time?
3. Can you import results from other tools, such as open-source test runs?
4. How long do you keep records? Can we set the period?

Security and data

1. What data do you collect from our models, prompts and users? Can we turn collection off?
2. Where is data stored and processed? Can we keep it in the EU?
3. Do you use our data to train any model? Put the answer in the contract.
4. Which LLMs run inside your product? Who hosts them?
5. Which third-party audit reports can you share?
6. Do you support SSO, role-based access and customer-managed keys?
7. How do you protect your own product against prompt injection and system prompt leakage (OWASP LLM01 and LLM07)?

The vendor's own AI governance

1. Do you run an AI management system in line with ISO/IEC 42001 ([ISO](#))? Is it certified?
2. Which framework do you use for your own AI risk? For example, the four NIST AI RMF functions: govern, map, measure, manage ([NIST AIRC](#)).

Commercial terms

1. What is the price at our size, and at twice our size? What is the unit: user, use case, model, trace or token?
2. Which features are only in the enterprise tier?
3. What happens to our data at the end of the contract?

Most governance vendors do not publish prices. See [how AI governance software is priced](#). To score the answers, use [how to choose an AI governance platform](#).

Sources

- [Article 113](#) and [Article 26](#), artificialintelligenceact.eu, accessed 2026-10-29
- [SR 26-2: Revised Guidance on Model Risk Management](#), Federal Reserve Board, 17 April 2026, 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
- [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

Frequently asked questions

Should I send a separate security questionnaire?

Yes, if your company has a standard one. Add the AI-specific questions in this guide to it. A governance tool often reads sensitive data about your models and your users.

Should I ask the vendor how it governs its own AI?

Yes. Many governance tools use LLMs inside the product. Ask which models they use, where data goes, and whether the vendor follows a framework such as ISO/IEC 42001 or the NIST AI RMF.

How many vendors should get the RFP?

Enough to compare, and few enough to run a real pilot with each finalist. Use the answers to cut the list before the pilot.