

Request for Proposal: AI Governance and Compliance Platform

Template from RFPhub.com. Replace everything in square brackets, adjust the Must-have/Nice-to-have priorities, and delete this note and the instructions section before you issue the RFP.

Instructions for the buyer (delete before issuing)

An AI governance platform sits between legal, risk, privacy, security, data science and procurement, so much of the early work is agreeing who owns which decision and which frameworks your program follows. The example schedule below runs about 13 weeks from kickoff to contract award; allow longer if your risk tiers and approval process aren't written down yet.

1. **Agree on scope and frameworks (weeks 1-2).** List the AI you need to govern: models you build, generative AI applications and agents, and AI features inside the software and services you buy. Name the frameworks your program follows, such as the EU AI Act, ISO/IEC 42001, the NIST AI RMF or sector model-risk rules, and have counsel confirm which obligations apply to you. Record the numbers vendors will price against: AI systems, users, legal entities and AI vendors to assess.
2. **Assemble the buying team.** Include an executive sponsor, the AI governance or risk lead, legal and privacy counsel, information security, data science and ML engineering, procurement and third-party risk, and internal audit. Decide early who owns the inventory and who approves higher-risk systems, because the platform's workflows will encode those decisions.
3. **Write down your process, and decide whether you need an RFI (weeks 2-4).** Draft your risk tiers, intake questions and approval steps before vendors shape them for you. If you're unsure whether you need a dedicated platform or an AI governance module in your existing GRC or ML platform, or your long list has more than five vendors, issue a short RFI first and use the answers to build a shortlist.
4. **Tailor and issue the RFP (weeks 4-7).** Delete requirements that don't apply, adjust priorities, and add your environment and framework list to

Section 1. Give vendors three weeks to respond, with a written question-and-answer window in the first week.

5. **Score the responses (weeks 7-8).** Have each evaluator score independently against the rubric before comparing notes, and set aside any vendor that misses a Must-have. For every framework a vendor says it supports, check who maintains the content, how it is reviewed, and whether it is generally available today or still on the roadmap.
6. **Run a proof of concept with two or three vendors (weeks 8-12).** Load a sample of real systems: one higher-risk model you built, one generative AI application or agent, and one vendor product with embedded AI. Take each through intake, classification, assessment, approval and evidence collection, connect at least one ML platform and your ticketing system, and time how long a reviewer takes to answer an auditor's question from the record. Call references of a similar size and regulatory scope.
7. **Negotiate and award (weeks 12-13).** Use the pricing table to compare like for like. Settle the pricing metric (AI systems, users, entities or assessments), framework content fees, renewal caps and full data export before you sign, not after.

Contents

1. Introduction and Background
2. Project Objectives
3. Scope of Work
4. Requirements
5. Vendor Questions
6. Evaluation Criteria and Scoring
7. Pricing Response
8. Submission Guidelines
9. Timeline

1. Introduction and Background

[Company Name] is seeking proposals for an AI governance and compliance platform that inventories the AI systems we build, buy and use inside vendor products, manages their risk assessments and approvals, maps our controls to the frameworks we follow, and keeps the documentation and evidence that show how

each system is governed. This RFP describes our requirements, the questions we need answered, and how proposals will be evaluated.

Organization background

- [Describe your organization: industry, size, locations]
- [Regulatory or contractual requirements that apply]
- [Current approach to AI inventory and approval: spreadsheets, GRC tooling, committee reviews]
- [Number of AI systems in production, in development and used through vendor products, and the number of AI vendors]
- [Frameworks and rules the program follows, for example the EU AI Act, ISO/IEC 42001, the NIST AI RMF or sector model-risk rules, as confirmed by counsel]
- [ML platforms, model registries, cloud AI services, GRC, third-party risk, data catalog and ticketing tools in use]
- [Governance roles in place, such as an AI committee, model risk, privacy, legal and internal audit]

2. Project Objectives

- Maintain one current inventory of AI systems, including AI inside vendor products, with risk classification and named owners
- Run consistent, proportionate intake, assessment and approval workflows across all business units
- Map our controls once to each framework we follow and track readiness by system
- Produce documentation and evidence for auditors, regulators and customers on request
- Connect governance to the ML platforms, GRC and ticketing tools our teams already use

3. Scope of Work

The selected vendor will be responsible for:

- Subscription to the proposed platform for [number] AI systems, [number] users and [number] entities

- Configuration of our risk tiers, assessments, approval workflows and the frameworks we follow
- Migration of our existing inventory, assessments, policies and evidence, in phases agreed with [Company Name]
- Integration with our ML platforms, GRC, third-party risk, data catalog and ticketing systems
- Training for governance reviewers, system owners and administrators
- Framework content updates, support and service management for the contract term

4. Requirements

For each requirement, enter a response code and a short explanation. Y = available today in the proposed edition; P = partially met or needs configuration/services (explain); R = on the roadmap (give a date); N = not supported. Must-have requirements are pass/fail; Nice-to-have requirements add points.

4.1 Functional

ID	Requirement	Priority	Response	Vendor comments
F1	A central inventory of AI systems, models, datasets, prompts, AI agents and use cases, with defined relationships between them, a named owner and accountable executive, business purpose and lifecycle status for each.	Must		
F2	Inventory records for AI used inside third-party products and services, including AI features enabled in software we already license, linked to the vendor and contract record.	Must		
F3	An intake workflow for proposed AI use cases, with configurable questions that route each request to the right reviewers (for example, legal, privacy, security or model risk) based on the answers.	Must		
F4	Risk classification against our own risk tiers and against the risk categories of the frameworks we follow (for example, the EU AI Act) [confirm with counsel], with the rationale recorded for every classification decision.	Must		
F5	Configurable risk and impact assessment templates with conditional logic, covering dimensions such as privacy, safety, fairness, security and effects on individuals,	Must		

ID	Requirement	Priority	Response	Vendor comments
	so lower-risk systems get a lighter review than higher-risk ones.			
F6	Inherent and residual risk scoring, with mitigations, accepted risks and exceptions recorded against each system. Exceptions must carry a named approver, an expiry date and a review reminder.	Must		
F7	Reassessment triggered on a schedule, by material changes (a new model version, new training data or a new use) and by monitoring alerts.	Must		
F8	A policy library with versioning, approval workflow and effective dates, where each policy statement maps to controls and to the AI systems in scope, and a policy change flags the systems it affects.	Must		
F9	Attestation campaigns in which staff acknowledge our AI acceptable-use policy, with completion tracking by team.	Nice		
F10	A control library that maps one control to requirements in each framework we name (for example, the EU AI Act, ISO/IEC 42001 and the NIST AI RMF) and to our own internal and sector frameworks, so evidence collected once counts toward each.	Must		
F11	Templates for model cards, system cards and dataset documentation, versioned with the model or system they describe, with role-based editing, review and sign-off, and export to PDF and structured formats.	Must		
F12	Automatic population of documentation fields (metadata, metrics and lineage) from connected ML platforms and model registries.	Nice		
F13	Storage of testing evidence against each system version, including bias and fairness, robustness and accuracy results and, for generative AI, evaluation and red-team results, with thresholds that can block an approval.	Must		
F14	Built-in bias, robustness or generative AI evaluation tests that run before release and on a schedule, in addition to importing results from our own tools.	Nice		
F15	Configurable approval gates at lifecycle stages (intake, before deployment, before material changes and at retirement), with sequential or parallel reviewers by role.	Must		

ID	Requirement	Priority	Response	Vendor comments
F16	Ingestion of monitoring signals for deployed systems (drift, performance, fairness and generative AI quality metrics) from our ML platforms and monitoring tools, with thresholds that open findings and notify the system owner.	Must		
F17	AI-specific third-party risk questionnaires and evidence requests, with tracking of supplier responses, provider documentation, attestations and reassessment dates.	Must		
F18	An AI incident workflow: intake from users, monitoring alerts and ticketing; severity levels we configure; links to affected systems and controls; corrective actions; and deadline tracking for any regulatory reporting our counsel identifies [confirm with counsel].	Must		
F19	Dashboards and on-demand reports by framework, AI system, business unit and period, including gap analysis per framework, overdue assessments, open findings and audit packs for auditors and customers.	Must		
F20	Discovery of unregistered AI use from sources such as single sign-on logs, SaaS and cloud accounts and code repositories, with a workflow to register, approve or block what is found.	Nice		
F21	Records specific to AI agents: the tools and data each agent can use, the identity and permissions it acts under, and the human oversight that applies, with review triggered when any of these change.	Nice		

4.2 Technical and architecture

ID	Requirement	Priority	Response	Vendor comments
T1	Multi-tenant SaaS with a choice of hosting region, or single-tenant or customer-cloud deployment. State which options you offer and any feature differences between them.	Must		
T2	A configurable data model for AI systems, models, datasets, agents, vendors, controls and evidence, with custom fields and record types we can add without vendor coding.	Must		
T3	Support for multiple legal entities, business units and regions, with separation of data and workflows between them where we	Must		

ID	Requirement	Priority	Response	Vendor comments
	require it.			
T4	Bulk import from spreadsheets and existing GRC records, registration of systems and evidence through the API, and full export of the inventory, assessments and evidence in documented formats.	Must		
T5	Support for [number] AI systems, [number] users and [number] vendors without loss of performance. State the largest volumes you have tested.	Must		
T6	Export of inventory records as an AI bill of materials in a standard format, such as CycloneDX (ML-BOM) or the SPDX 3.0 AI profile.	Nice		
T7	Built-in AI assistance, for example drafting assessment answers or documentation, or suggesting control mappings, with every AI-generated entry labeled and requiring human review. State which models are used and where they run.	Nice		

4.3 Integration

ID	Requirement	Priority	Response	Vendor comments
I1	Single sign-on with our identity provider via SAML 2.0 or OIDC, with user and group provisioning via SCIM.	Must		
I2	Connectors to our ML platforms, model registries and cloud AI services [list yours] that register models, pull metadata and metrics, and keep the inventory current. State which connectors are native and which need custom work.	Must		
I3	Two-way integration with our GRC platform, so controls, risks, issues and evidence are not maintained twice. State which objects sync, in which direction, and which system owns each record.	Must		
I4	Integration with our ticketing and IT service management platform to create remediation tasks, findings and incident tickets, with status synced back to the governance record.	Must		
I5	Integration with our third-party risk management and procurement tools, so an AI vendor assessment starts from the existing vendor record or purchase request.	Must		