

How AI Governance Software Is Priced

Most AI governance platforms do not publish prices. Some monitoring, guardrail and red-teaming tools do. Here are the pricing models and the public prices we could verify.

Updated October 29, 2026

AI governance software is sold in several ways. Some products charge per seat. Some charge per use case or per model. Guardrail and monitoring tools often charge by usage. This guide explains the models and lists the public prices we found.

The short answer

Most AI governance platforms do not publish prices. On the sites we checked, Credo AI, Holistic AI, Trustible and ModelOp offer a demo request, not a price list. OneTrust sends you to sales. Vanta offers "personalized pricing". You must ask these vendors for a quote.

Some vendors do publish prices. They are mostly monitoring, LLM tracing, guardrail and testing tools, plus one governance platform from IBM.

Common pricing models

| Model | What you pay for | Seen in | | --- | --- | --- | | Platform subscription | Instances, modules and concurrent users | IBM watsonx.governance | | Per use case | Number of AI use cases you govern | IBM on AWS (extra use cases in packs of 10) | | Per seat | Each named user | LangSmith | | Per usage unit | Traces, units, text units or tokens | Fiddler, Langfuse, AWS, Google | | Open source plus enterprise | Free core, paid features and support | promptfoo, Langfuse, Arthur | | Custom quote | Scoped by sales | Most governance platforms |

Public prices we verified

All prices are in US dollars and come from the vendor's own pricing page on 29 October 2026.

Governance platform

| Vendor and plan | Published price | What is included | | --- | --- | --- | | IBM watsonx.governance, Risk & Compliance Basic | Starting at \$3,500 (monthly pricing) | 1 basic instance, 1 module, 1 concurrent user, on IBM Cloud | | IBM watsonx.governance, Risk & Compliance Advanced | Starting at \$6,450 (monthly pricing) | 1 advanced instance, 1 module, 1 concurrent user, as enterprise software | | IBM watsonx.governance on AWS | Starting at \$42,000 | 1 basic instance, 5 AI use cases, 12,000 evaluations, 25 concurrent users |

IBM also offers a 14-day free trial. IBM says its prices are indicative and may vary by country ([IBM](#)).

Monitoring and LLM observability

| Vendor and plan | Published price | | --- | --- | | Arthur Free | \$0 per month | | Arthur Premium | \$60 per month, monitoring for up to 100 use cases | | Fiddler Developer | \$0.002 per trace | |

Langfuse Cloud Core | \$29 per month, 100k units included, then \$8 per 100k units | | Langfuse Cloud Pro | \$199 per month, same unit terms | | LangSmith Plus | \$39 per seat per month, then pay as you go |

Each of these vendors also has an enterprise plan with a custom price ([Arthur](#), [Fiddler](#), [Langfuse](#), [LangChain](#)).

Guardrails and red teaming

| Service | Published price | | --- | --- | | Amazon Bedrock Guardrails, content filters (text) | \$0.15 per 1,000 text units | | Amazon Bedrock Guardrails, sensitive information filters | \$0.10 per 1,000 text units | | Amazon Bedrock Guardrails, word filters and regex filters | Free | | Google Model Armor, standalone | Free up to 2 million tokens per month, then \$0.10 per million tokens | | promptfoo Community | Free, red teaming up to 10k probes per month |

On AWS, one text unit holds up to 1,000 characters ([AWS](#)). Google counts four characters as one token for Model Armor ([Google Cloud](#)). promptfoo Enterprise and On-Premise plans are custom ([promptfoo](#)).

A worked example

Say a chatbot screens 1,000,000 prompts a month on Amazon Bedrock Guardrails. Each prompt is under 1,000 characters, so each is one text unit. With content filters only, the list price is $1,000 \times \$0.15 = \$150$ a month. Each extra filter adds its own charge. If you also screen model responses, they count as text units too.

What drives your quote

Ask each vendor how these items change the price:

- Number of AI use cases, models or agents in scope
- Number of users, and whether they are named or concurrent
- Modules: inventory, risk assessment, policy, monitoring, evaluation
- Deployment: SaaS, private cloud, or on-premise
- Volume of traces, evaluations or screened tokens
- Support level, SSO, and data residency

Tips

1. Count your AI use cases before you ask for quotes.
2. Ask for the price at your size and at twice your size.
3. Ask which features are only in the enterprise plan.
4. Compare total cost against building it yourself. See [build vs buy](#).

Sources

- [watsonx.governance pricing](#), IBM, accessed 2026-10-29
- [Pricing](#), Arthur, accessed 2026-10-29
- [Pricing](#), Fiddler AI, accessed 2026-10-29

- [Pricing](#), Langfuse, accessed 2026-10-29
- [LangSmith pricing](#), LangChain, accessed 2026-10-29
- [Amazon Bedrock pricing](#), Amazon Web Services, accessed 2026-10-29
- [Security Command Center pricing \(Model Armor\)](#), Google Cloud, accessed 2026-10-29
- [Pricing](#), promptfoo, accessed 2026-10-29
- [Credo AI](#), [Holistic AI](#), [Trustible](#), [ModelOp](#) home pages, accessed 2026-10-29
- [Pricing](#), OneTrust, accessed 2026-10-29
- [Pricing](#), Vanta, accessed 2026-10-29

Frequently asked questions

Why do most AI governance vendors not publish prices?

Their price usually depends on the number of AI use cases, models, users and modules, and on the deployment type. Most of the governance platforms we checked ask you to request a demo or talk to sales.

Which vendors publish a price for an AI governance product?

IBM publishes starting prices for watsonx.governance. Several monitoring and LLM tools, such as Arthur, Fiddler, Langfuse and LangSmith, publish plan prices. Cloud guardrail services from AWS and Google publish usage prices.

Are the published prices final?

No. They are list or starting prices. IBM says its prices are indicative and may vary by country. Enterprise plans are usually custom.