EXECUTIVE BRIEF

What Is an AI Vendor Assessment? A Practical Guide to Third-Party AI Risk

Buying an AI tool does not transfer accountability to the vendor. An AI vendor assessment helps an organization understand what it is buying, what can go wrong, what evidence the vendor can provide, which safeguards belong in the contract, and whether the remaining risk is acceptable for the intended use.

IN THIS BRIEF

By F. Jay Hall and Stephan Pochet · GRC Careers · August 15, 2026
Source: **ai-governance-jobs.com**

EXECUTIVE SUMMARY

Organizations rarely build every artificial intelligence system themselves. They buy software with embedded AI, connect to general-purpose models through an API, license specialized models, use vendor data, adopt AI features inside systems they already own, and allow contractors to operate AI services on their behalf.

Each arrangement can create value. Each can also introduce risk that the customer cannot see directly. The model may change without warning. Training data may be poorly documented. Customer information may be retained or used to improve the product. Performance may look strong in a vendor demonstration but weaken in the customer's population or operating environment. A critical feature may depend on another provider several layers down the supply chain.

KEY TAKEAWAYS

- An AI vendor assessment evaluates the product, the provider, the intended use, and the customer's ability to govern the system after purchase.
- Traditional security and privacy reviews remain necessary, but they do not answer questions about model performance, bias, explainability, human oversight, training-data rights, or AI-specific change risk.
- The depth of review should follow the use case and potential impact, not the size or reputation of the vendor.
- Vendor claims are not evidence. Higher-risk uses require documentation, testing results, contract protections, and independent validation where appropriate.
- The assessment does not end at purchase. Monitoring, change notification, incident response, renewal, and exit planning are part of the same lifecycle.

WHAT THIS BRIEF COVERS

What is an AI vendor assessment?. An AI vendor assessment is a structured, risk-based evaluation of a third party that provides an AI system, model, data source, platform, component, or AI-enabled service.

Why an ordinary vendor review is not enough. Existing third-party risk programs give AI governance a valuable foundation.

What counts as an AI vendor?. The review should not be limited to companies that advertise themselves as AI vendors.

Start with the intended use, not the questionnaire. The same product can be low risk in one use and high risk in another.

Use risk tiers to determine the depth of review. Risk tiering should consider both the product and the proposed use.

The eight-step AI vendor assessment. Document the legal vendor, the product, the AI features, the underlying models, data providers, hosting environment, subprocessors, plug-ins, and critical integrations.

What evidence should the vendor provide?. No single document proves that a product is safe or compliant.

Core questions to ask an AI vendor.

Do not confuse the questionnaire with the assessment. A completed questionnaire is one source of information.

Test the system in the real workflow. Testing should reflect how the organization will actually use the product.

Contract protections that matter. Legal counsel should tailor contract language to the organization, jurisdiction, sector, and risk.

Approval outcomes and conditions. An assessment should end with a clear outcome:

Red flags that deserve attention. A red flag does not always require rejection.

Monitoring after purchase. Third-party AI risk changes over time. Monitoring should match the product's importance and risk. It may include:

Build the exit plan before you need it. A high-risk or operationally important AI service needs a contingency plan.

A practical approach for smaller organizations. A small nonprofit, school, public agency, or business may not have separate AI, model-risk, privacy, procurement, and security teams.

How the assessment fits leading frameworks and regulation. The NIST AI RMF treats acquisition and third-party resources as part of AI governance across the full lifecycle.

The bottom line. An AI vendor assessment is not designed to eliminate all third-party risk.
