Compliance · Essay

The Governance Gap: What AI Vendor Risk Actually Costs Mid-Market Companies

The data on AI vendor risk tells a less dramatic, more useful story than the sales pitch: a real coverage gap, real procurement friction, and very little evidence anyone is writing large new checks.

Jennifer McKinney · August 15, 2026 · 7 min read

AI vendor risk has acquired the familiar fog of a new enterprise category. The language is expansive. The product demonstrations are smooth. The market claims are frequently unhelpful.

What matters to a mid-market buyer is less glamorous: What is actually broken? What does it cost to address? What happens if we do nothing? And is the constraint a lack of tooling, a lack of people, or a lack of someone accountable for the decision?

The available data gives a more useful answer than the hype does. There is a real coverage gap. There is real procurement friction. There is also very little public evidence that mid-market companies are writing large, new checks specifically for AI vendor due diligence. Those facts can coexist. They should shape how we design the work.

Start with the gap. In Ncontracts' 2026 survey of financial institutions, 16% had not assessed vendor AI usage at all. Seventy-two percent were only partially aware of which vendors used AI, and no respondent was "extremely confident" in managing AI risk. Those are not edge cases. They describe an operating model in which vendor behavior is changing faster than the assessment process. Ncontracts

The finding is consistent across a different population. Panorays reported that 52% of CISOs still use general-purpose onboarding for AI vendors, while only 22% have a dedicated, documented AI-risk evaluation policy. A generic security questionnaire can be valuable. It usually cannot answer the questions a generative or decisioning system raises: what data crosses the boundary, what changes when a model updates, what authority does the system have, and how is a harmful output detected or corrected. Panorays

The easiest mistake is to interpret this as a software shortage. The data says otherwise. AuditBoard and Panterra Research found that only about two-thirds of surveyed GRC and internal-audit professionals conduct formal AI-specific assessment of third-party models or vendors. But the reported barriers were lack of clear ownership, insufficient internal expertise, and resource constraints. Fewer than 15% said lack of tools was the main problem. AuditBoard research coverage

That finding maps closely to what I see in program work. A new platform cannot decide whether legal, security, procurement, privacy, product, or the business owner has authority to accept a risk. It cannot establish a review threshold that people will actually use. It cannot resolve the mismatch between a sales timeline and an assessment that has no named decision-maker. Tools can make good process faster. They do not create an operating model.

Cost data needs the same discipline. I found no credible public survey that states the average or median amount a mid-market company spent on AI-specific vendor due diligence in the last year. That gap should be said plainly. We should not fill it with a made-up market average.

We do have useful anchors. Pivot Point Security publishes vendor-review pricing bands: $200 to $500 for an automated review, around $1,000 for an automated review with manual checking, $2,500 to $3,500 for a fully manual review, $6,000 to $7,000 for a complex high-risk vendor using a full SIG questionnaire, and $15,000 to $20,000 for a comprehensive onsite audit. These are published asking prices, not a survey of transactions. But they show the order of magnitude of the work that every new AI vendor can create. Pivot Point Security

The implication is practical. A company that treats every AI vendor as a bespoke legal and security investigation can create an expensive queue quickly. A company that treats every AI vendor as a standard SaaS tool can acquire risk it does not understand. The answer is tiering. Identify the few conditions that change the risk posture: sensitive data, automated decisions about people, customer-facing output, model training or retention terms, subprocessor dependence, cross-border use, and the ability to change behavior without meaningful notice. Then assign the depth of review to the tier, not to the loudness of the vendor’s AI marketing.

The broader third-party-risk stack has its own price. Mordor Intelligence says smaller buyers can face first-year platform, setup, and labor costs of $40,000 to $80,000, and characterizes implementation and integration cost as especially acute in SME markets. That is not an argument against formal risk management. It is an argument against treating a heavyweight platform as the only legitimate operating model. Mordor Intelligence

There is a regulatory nuance here that many sales narratives flatten. The EU AI Act’s Annex III high-risk obligations were postponed to December 2, 2027 as part of the Digital Omnibus changes. Article 50 transparency obligations did take effect on August 2, 2026. This means that

the regulatory picture is real but not uniform: disclosure and labeling requirements are active for in-scope systems, while the broader high-risk obligations that would drive conformity assessment and deeper documentation have a later timetable. Al Jazeera European Commission guidance

For a mid-market organization, that is not permission to wait. It is permission to be precise. Article 50 does not create a universal vendor-due-diligence regime. The better immediate question is whether an organization can identify the AI systems it deploys, understand when they interact with people, and substantiate the representations it makes about those interactions. The architecture matters because the evidence is much easier to produce when data flows, model boundaries, and decision authority are explicit.

The business pressure is often more immediate than the regulation. Levelpath's 2026 research on organizations purchasing AI software found that AI purchase cycles most commonly take 16 to 20 weeks, compared with 7 to 10 weeks for standard software purchases above $10,000. Fifty-eight percent of respondents identified security reviews as the leading cause of delay; 57% named vendor evaluations, and 52% named contract negotiations. Fifty-eight percent said an AI purchase involved at least seven people. Levelpath research via Businesswire

This is where the economics become less abstract. A 16-week delay is not merely a compliance inconvenience. It is a multi-team coordination failure: a product leader waiting on procurement, a procurement lead waiting on security, security waiting on vendor artifacts, legal waiting on someone to decide what terms are non-negotiable. The direct invoice for a review is visible. The cost of a stalled decision is distributed across calendars and is therefore easier to ignore.

There is no evidence that buyers are responding by opening unlimited new budgets. Ncontracts found that 64% of respondents expected flat third-party-risk budgets. Their behavior tells a similar story: collection of vendor AI documentation rose from 36% to 51%, while adding AI-usage language to contracts moved only from 35% to 41%. Companies are gathering artifacts faster than they are changing the terms or governance behind those artifacts. Ncontracts

That is why I would not position AI vendor risk as an invitation to buy another dashboard. The core product is a reliable decision path. It should establish ownership, classify the risk, identify the missing evidence, create a concrete escalation point, and produce an answer that procurement can use. In some organizations that will require a platform. In many, it will require a well-designed intake, a risk register, clear authority, and architecture that makes the relevant facts observable.

The governance gap is not imaginary. The spending story is simply less settled than the hype suggests. Mid-market companies are under-covered, short on expertise, and delayed by

fragmented review. The winning response will not be the loudest claim about compliance. It will be the system that makes a high-quality decision easier to reach, easier to defend, and materially faster to execute.

Sources

- Ncontracts
- Panorays
- AuditBoard research coverage
- Pivot Point Security
- Mordor Intelligence
- Al Jazeera
- European Commission guidance
- Levelpath research via Businesswire