

The AI You Can't See *third-party shadow AI*

A short guide to where vendor AI reaches your data and why standard vendor reviews tend to miss it.

SIGNED OFF BY YOUR TEAM

Approved copilot

Sanctioned LLM

Reviewed agent

BROUGHT IN THROUGH YOUR VENDORS

Browser AI extensions

Public chatbots

Vendor AI features

Unvetted agents

Nth-party models

WATERLINE

01 Blocking AI rarely works, and the problem reaches your vendors

Security leaders keep hearing the same question: how do we block AI? Blocking rarely holds. When an organization bans a tool people find genuinely useful, they keep using it out of sight, and the unvetted tools that result become invisible to the team responsible for managing them. The organizations handling this well give employees approved ways to use AI and then account for that use in their risk picture.

That approach covers the technology you manage directly. It does far less for the support agent at one of your vendors who is working through a queue of your customer tickets and pastes a batch into a public chatbot to summarize them faster. Nobody at your company approved that tool or assessed it, and nobody will hear about it. Your customer data has moved into an AI service outside your vendor inventory, through a relationship that looked healthy at onboarding.

That chatbot now holds enterprise data somewhere nobody can name, which makes it a third party in every sense that matters, even though it never went through review. The questions it raises are the ones any vendor assessment exists to answer, about what was shared and where it now lives, and today nobody can answer them.

43%

of security incidents involved unapproved AI tools, more than double the share a year earlier.

IBM Cost of a Data Breach Report 2026, conducted by Ponemon Institute

69%

of cybersecurity leaders have evidence or suspect that employees use prohibited public generative AI tools at work.

Gartner survey of 302 cybersecurity leaders

THE SUPPLIER SIDE

Your vendors are organizations too, and they come from the same population those surveys describe. Whatever share of companies are running unapproved AI tools today, your supplier list is drawn from that group, and nobody surveyed it on your behalf.

02 Four routes vendor AI takes to your data

Vendor AI reaches your data through more routes than procurement ever sees, and most of them involve no bad intent at all. Four come up again and again.

A

Tools your vendor's people use

A vendor's team works with browser AI extensions and public chatbots as part of their daily habits, often on data that belongs to you. None of it shows up in the vendor's questionnaire answers, because the questionnaire never asked.

B

AI features shipped in a product update

Your CRM adds a summarization feature for the notes your sales team writes every day. That feature calls a model API run by a company whose name appears nowhere in your vendor list, and it arrived through a product update rather than a procurement decision.

C

The providers behind the provider

A single meeting recording can pass through the app on your contract, a transcription service, a model provider and the cloud infrastructure underneath them. Four companies touch it, each with its own retention rules and jurisdiction, and your review covered the first.

D

Agents holding permissions

A marketing provider deploys an autonomous agent to run campaigns, with permission to change records in your CRM. The exposure now sits in what that agent is allowed to do, which a data-handling questionnaire is unlikely to capture.

A CLOSER LOOK · OPEN-WEIGHTS MODELS

Some vendors self-host open-weights models rather than calling an external API, including models released by developers in other jurisdictions. Keeping data inside the vendor's environment can reduce one kind of exposure, while the model's provenance and the way it was evaluated become the questions that matter.

In each case the vendor relationship itself looks unchanged. The contract, the security review and the risk rating all stay where they were, while the data path quietly extends to companies and systems the review never covered.

03 How organizations respond, and where it stalls

It usually starts at the top. The CEO or the board decides the company has to adopt AI or fall behind, and the organization answers by forming an AI committee. Membership varies widely from one company to the next, pulling in anyone from HR to R&D, and the CISO often ends up coordinating all of it. Without clear ownership of its own, the committee can become the place where requests go to wait.

Then a business owner asks to bring in an AI tool, or a long-standing vendor announces an AI feature, and someone has to decide whether it fits company policy. That one request lands on four desks, and each desk is asking something different.

WHAT EACH TEAM ASKS ABOUT AN AI REQUEST	
Business	Is the value worth the exposure, and who inside the company owns the relationship?
Technical	What can the AI reach, and what changes when the vendor updates the model?
Commercial legal	What does the contract say about AI subprocessors, and about notice when they change?
Privacy legal	Whose personal data enters the model, and is any of it kept or used for training?

WHERE THE PROCESS TENDS TO BREAK

Most stalls trace back to ownership. When accountability for AI risk is spread across a committee, a request can wait weeks while the business moves on without it, and the tools people wanted arrive anyway through a vendor's next product update. Programs that keep pace usually have one person accountable for AI risk and an inventory of which vendors use AI at all, since every other decision depends on that list.

04 Governance depends on seeing what runs where

GRC, TPRM and endpoint vendors all describe what they do as AI governance now. In practice, governance depends on instrumentation, meaning real visibility into which AI is running where. A questionnaire records what a vendor said at one moment, and a vendor that wasn't using AI last year can plug in a language model tomorrow without anyone updating a single document.

That gap gives continuous visibility more weight here than in most corners of third-party risk. Only 13% of CISOs in the 2026 Panorays CISO Survey track their nth-party vendors, the layers behind the layers where vendor AI usually runs, so most programs are working from an inventory that stops one hop short.

ONE WAY TO BUILD THAT VISIBILITY

Panorays AI Supplier Detection identifies which third parties use AI, drawing on external intelligence, certifications and discovery methods, even when a vendor hasn't disclosed it. Suppliers using AI are labeled and brought together in one report. Vendors that warrant a closer look can receive AI-focused questionnaires at increasing levels of depth, and AI exposure becomes its own factor in how vendors are tiered and monitored.

[Explore AI Supplier Detection →](#)