VendRisk360
ASSESS · MONITOR · STAY AHEAD

WHITEPAPER

Your Vendor's AI Is Your Risk: Assessing AI Suppliers and AI-Enabled Vendors

A practical assessment playbook for the AI now embedded in the products you buy: what to ask, what certification proves, how to find hidden model providers, and which obligations fall on you.

Published 27 September 2026

Written for Heads of TPRM, CISOs, procurement, model risk and compliance officers, internal audit

Online www.vendrisk360.com/resources/whitepapers/your-vendors-ai-is-your-risk/

Contents

A focused due-diligence playbook for AI suppliers and AI-enabled vendors: an evidence-based question set, how to uncover foundation model providers as fourth parties, what ISO/IEC 42001 and NIST AI RMF do and do not prove, EU AI Act deployer duties, shadow AI, and the contract clauses that matter.

Two or three years ago, "does this vendor use AI" was a niche question asked of a handful of analytics suppliers. Today it is the wrong question, because the answer is almost always yes. Generative AI has been added to the products you already buy: the support desk that now drafts replies, the code scanner that explains findings, the HR platform that summarizes candidates, the security tool that triages alerts. Many of these features are powered by a foundation model from a provider your organization has never assessed and, in some cases, your vendor did not prominently disclose. The AI your vendors have adopted is now part of your risk surface whether or not you have written it down.

This paper is a due-diligence playbook for that reality. It is deliberately narrower than a general treatment of AI governance: it assumes you have decided a vendor's AI is in scope and need to assess it well. It gives an evidence-based question set, a method for finding the model providers hiding behind your vendors, a clear account of what AI certifications do and do not prove, and the obligations that fall on you as the organization deploying someone else's AI. For the wider question of governing AI inside your own TPRM program, including using AI to review vendor evidence, see the companion paper, AI in Third-Party Risk Management. The guidance here is vendor-neutral; confirm specific obligations with counsel and your regulators.

KEY TAKEAWAYS

- Assume AI is present. The default assumption for new and renewing vendors should be that AI is somewhere in the product, and the assessment's job is to find where, on what data, and with what oversight.
- The foundation model provider is a fourth party. When a vendor's AI feature calls an external model, that model provider processes your data and belongs in your nth-party inventory as a concentration and data-flow question, not as an invisible dependency.
- Certification is evidence of a management system, not of a model's accuracy. ISO/IEC 42001 shows a vendor runs an AI management system; it does not prove any particular model is accurate, fair or safe on your use case. Read the certificate's scope.
- The two questions that matter most on data: is our data used to train or improve the vendor's or a provider's models, and where and how long are our prompts and outputs retained. Get both in writing.
- If you deploy a vendor's AI in a regulated use, obligations may fall on you directly. Under the EU AI Act, the organization using an AI system, the deployer, carries duties of its own, distinct from the provider's.
- Shadow AI is the near-term exposure. Staff paste sensitive data into consumer AI tools, and vendors switch on AI features by default. Discovery and contract terms matter more than a perfect model evaluation.

Start by finding the AI, not by scoring it

The first failure in AI vendor assessment is not a weak evaluation; it is a missed system. You cannot assess what you have not found. Before any questionnaire, establish where AI actually sits in the relationship:

- **Is AI in the product path you use,** or only in a module you have not enabled? A feature that is off by default is a different risk from one that processes every record.
- **Does it process your data,** or only the vendor's own operational data? An AI feature that never touches your regulated data is a lower priority than one that summarizes it.
- **Does it inform or make a decision that affects you or your customers?** A drafting aid a human edits is not the same as an automated eligibility or scoring decision.
- **Was it added recently, or by acquisition?** AI features arrive through product updates and through the vendor buying an AI startup. A prior assessment that predates the feature is stale.

Only once you know where the AI is and what it touches does scoring make sense. Scope the depth to that answer, in line with your criticality tiering.

The evidence-based question set

Questionnaires that ask "do you use AI responsibly" collect assurances, not evidence. The questions below are written to elicit artifacts a reviewer can verify, scaled to the materiality of the vendor's AI use.

AI SUPPLIER DUE-DILIGENCE QUESTIONS

Area	What to ask for (evidence, not assurance)
AI inventory	A list of AI features that process customer data or affect decisions, with what each does and whether it is on by default
Model provenance	Which models are used, built in-house or via a provider, and which foundation model and hosting providers are involved
Training on your data	A written statement of whether your data, prompts and outputs are used to train or improve any model, and the opt-out
Data flow and retention	Where prompts and outputs are processed, which providers see them, and retention periods
Human oversight	Where a person reviews or can override AI output, and where output is used without review
Evaluation and limits	Testing results, accuracy and error characteristics, known limitations, and instructions for safe use
Abuse and misuse defenses	How the feature resists manipulation through crafted inputs, and how data leakage between customers is prevented
Governance	ISO/IEC 42001 certificate (with scope), NIST AI RMF alignment, and internal AI policies
Change management	How model or prompt changes are tested and communicated to you before they take effect
Incident handling	How AI malfunctions, harmful outputs or data exposure are detected and disclosed to you

A vendor that can produce these artifacts is demonstrating a real program. A vendor that answers only in assurances is telling you where to look harder.

The foundation model provider is a fourth party

When a vendor's AI feature calls an external model, three things follow. Your data flows to that provider. Your resilience now depends on that provider's availability. And a small number of model providers sit behind a large share of AI features across the market, which makes them a concentration risk in the same way a hyperscale cloud region is.

Treat the model provider exactly as you would any material subservice organization. Record it in your nth-party inventory. Ask whether it is contractually flowed down as a subprocessor with equivalent obligations, and whether you get notice before it changes. Ask what the provider is permitted to do with the data it receives. A vendor that cannot name the model providers behind its AI features cannot answer the data questions that matter, because the answers are not theirs to give.

"WE USE A LEADING AI PROVIDER" IS NOT AN ANSWER

If a vendor will not name the model and hosting providers behind an AI feature that processes your data, you cannot assess data residency, retention, training use or resilience, because all of those depend on the specific provider and its terms. Name the provider or the feature is unassessed.

What AI certifications prove, and what they do not

Two references now appear in AI vendor conversations, and both are useful when read correctly.

ISO/IEC 42001, published in December 2023, is a management-system standard for AI, the AI equivalent of ISO/IEC 27001 for information security. A certificate is meaningful evidence that a vendor has a governed process for building and operating AI: roles, risk assessment, testing, oversight and improvement. Read it exactly as you read a SOC report. Check the scope: which parts of the organization and which systems it covers, the certification body, and the dates. A certificate scoped to one product line says nothing about the feature you are buying if that feature is out of scope. Crucially, it certifies the management system, not the accuracy, fairness or safety of any specific model.

The NIST AI Risk Management Framework, released in January 2023 with a Generative AI Profile added in 2024, is a voluntary framework, not a certification. A vendor that maps its practices to it is showing a structured approach. There is no NIST AI RMF certificate, so treat any claim of one with caution.

The practical stance: certifications and framework alignment are inputs that raise or lower how deeply you probe, not substitutes for the evidence in the question set above.

Obligations that land on you: the deployer

Assessing a vendor's AI is not only about the vendor. In several regimes, using someone else's AI creates duties for you. The clearest example is the EU AI Act, which entered into force on 1 August 2024 and applies in phases: prohibitions on certain practices from February 2025, obligations for general-purpose AI models from August 2025, and most high-risk system obligations from August 2026. The Act distinguishes the **provider** who builds an AI system from the **deployer** who uses it, and places obligations on both. If your organization deploys a vendor's AI in a use the Act treats as high-risk, you may carry duties such as human oversight, using the system per instructions, and monitoring, regardless of what the vendor has done.

The point for TPRM is that the assessment must record which AI systems you deploy, in what use, and therefore which obligations attach to you, not just to the vendor. That record is also what an examiner or regulator will ask to see.

WHAT AUDITORS AND EXAMINERS WILL ASK

- Show your inventory of vendor AI features that process regulated data or affect decisions, and how you found them.
- For a material AI feature, name the model and hosting providers behind it and where they sit in your nth-party inventory.
- Is our data used to train or improve any model, and where is the written statement?
- For AI you deploy in a regulated use, which obligations fall on you as the deployer, and how are they met?
- How are model or prompt changes by the vendor tested and communicated before they take effect?

Shadow AI: the exposure that is already here

While programs debate model evaluation, the immediate exposure is simpler and closer. Staff paste customer data, source code and confidential documents into consumer AI tools that were never assessed. Vendors enable AI features by default in products already in use, so data begins flowing to a model provider without a decision being made. Neither of these is a sophisticated model-safety problem; both are discovery and contract problems. The controls that help most are unglamorous: know which vendors have AI features and whether they are on, require disclosure and change notice in contracts, and give staff a sanctioned alternative so the incentive to use unsanctioned tools falls.

Contract clauses for vendor AI

Standard technology contracts rarely address AI. The following clauses, scaled to materiality, close the gaps the assessment surfaces:

- **AI disclosure and change notice:** the vendor discloses AI features that process your data or affect decisions, and notifies material changes before they take effect
- **No training on your data:** your data, prompts and outputs are not used to train or improve any model without written agreement
- **Model and hosting provider flow-down:** foundation model and hosting providers named as subprocessors with equivalent obligations and prior notice of changes
- **Data residency and retention:** where prompts and outputs are processed and how long any provider retains them
- **Human oversight and logging support:** features and information you need to oversee and log AI use
- **AI incident notification:** timely notice of malfunctions, harmful outputs or data exposure, aligned to your incident regime
- **Regulatory cooperation:** the vendor supports your obligations under applicable AI laws, including the EU AI Act where relevant
- **Evaluation and audit rights:** access to testing results, limitations and evidence about the AI feature, consistent with existing audit clauses

How VendRisk360 supports this

VendRisk360 lets you run the assessment this paper describes, and treats AI the same way it asks your vendors to: scoped narrowly, with people making the decisions.

Artifact-based assessments across 30+ control domains and configurable questionnaires let you add an AI-specific evidence request to any vendor whose product uses AI, scoped by criticality tier, with reassessment cadence set to your own policy. Because the questions ask for artifacts, the review turns on evidence rather than assurances. The vendor portal gives each vendor one-time-code access to only their own requests, and evidence is checked for scope and currency, not just filed. See Comprehensive Vendor Risk Assessment Services.

The foundation model provider problem is a concentration problem, and **[nth-party intelligence](/solutions/nth-party-intelligence/)** is built for it: relationships and subservice organizations discovered from SOC reports and your own records map the model and hosting providers sitting behind several vendors, so a provider that supports many of your AI features surfaces as a blast-radius and data-flow question rather than a hidden dependency. Where a vendor's own AI features

process your evidence, VendRisk360's optional AI is limited to completeness checks and key-date extraction and an AI-assisted first pass on SOC reports that an expert assessor verifies; every review, rating and sign-off is performed by a person, and the platform does not depend on AI. Our AI governance is aligned to the NIST AI RMF and ISO/IEC 42001; VendRisk360 is not certified to ISO/IEC 42001.

Findings, remediation and formal risk acceptance run through **multi-stage sign-off with segregation of duties** and a sign-off certificate, so a decision to accept an AI-related risk has an owner and a review date. When an examiner asks how you assessed a vendor's AI, the **examiner package export** bundles the assessment, evidence, sign-offs and audit trail. Customer data is protected by **row-level-security tenant isolation**, encryption in transit and at rest, SSO (SAML/OIDC), mandatory MFA, custom roles and a full audit trail. See Security and Board and executive reporting.

How to work with VendRisk360. On the Vendor Lifecycle Management Platform, your team manages vendors, sends AI-specific due diligence through the vendor portal, performs the review and signs off, with every step tracked. With Comprehensive Vendor Risk Assessment Services, you onboard the vendor and certified assessors collect the evidence, perform the review scaled to the vendor's tier with a second-expert quality review, and follow findings through remediation, while you keep final approval. Continuous Monitoring Services add the outside-in view between assessments, and Report-Specific Reviews are available standalone or alongside either. Assessors hold certifications such as CISSP, CISA, CISM, CRISC and ISO/IEC 27001 Lead Auditor and Lead Implementer.

To see an AI-supplier assessment, from evidence request to human sign-off, book a demo.

About VendRisk360

VendRisk360 is an independent company providing a third-party risk management platform and expert services for regulated organizations: banks, credit unions, fintech and payments, healthcare and SaaS. The platform runs vendor lifecycle management, risk-tiered assessments, continuous monitoring, nth-party intelligence and board-ready reporting from one governed record per vendor, combining point-in-time, evidence-based review with continuous outside-in monitoring. Customers run it themselves on the Vendor Lifecycle Management Platform, or add Comprehensive Vendor Risk Assessment Services, Continuous Monitoring Services and Report-Specific Reviews delivered by certified VendRisk360 assessors.

Learn more at **vendrisk360.com** or write to **info@vendrisk360.com**.

This paper is general guidance for practitioners, not legal advice. Regulatory requirements change and vary by jurisdiction and institution; confirm your obligations with counsel and your supervisors.