RSAA

RS ASSURANCE & ADVISORY

Vendor AI Due Diligence

Five questions to ask before a vendor's AI touches your data — and how to read the answers

WHITE PAPER · OCTOBER 2026

CONTENTS

What this paper covers

1	Why vendor AI is now your AI	3
2	How to use the addendum	3
3	The five questions	4
4	Reading the answers	7
5	After the review: keeping it current	7
6	How RS Assurance & Advisory supports vendor AI reviews	8
A	The AI addendum, ready to paste	9
B	Reviewer checklist	11
C	Framework mapping	12

EXECUTIVE SUMMARY

The short version

Most of the AI inside a small or midsize company did not arrive as an AI project. It arrived inside software the company already licenses, switched on in a product update, or adopted by one team without a review. That makes the vendor review the most practical place to govern it, and the security questionnaire most companies already send was written before anyone asked what a vendor's model does with their data.

Five questions close most of that gap: where your data goes, whether it is used for training, who is liable when the output is wrong, whether you can get your data deleted, and what independent evidence the vendor can produce. Each one is short to ask. Each one changes a decision when the answer is wrong.

This paper explains why each question matters, what a good answer looks like, and which answers should stop a purchase or send it to legal. Appendix A contains the addendum itself, written to paste into the questionnaire you

© 2026 RS Assurance & Advisory | rsassure.com

already use. A tool that passes goes into your AI inventory with a named owner and a data class, and the first two questions are asked again at every renewal, because vendors change their AI terms more often than their contracts.

SECTION 1

Why vendor AI is now your AI

A few years ago, the AI question in a vendor review was whether the vendor had any. Today the answer is almost always yes, and the AI is rarely a separate product. It is summarization in the ticketing system, drafting in the CRM, a note-taker in the meeting platform, and an assistant in the office suite. Vendors add these features in routine updates, often switched on by default, with no new contract and no new review.

That creates three problems for the buyer. Data that was shared under one set of assumptions may now flow to a third-party model provider. Output that used to come from a person now comes from a model, and it can be wrong in ways that reach customers. And the security questionnaire on file asks about encryption and access control, not about training, model subprocessors, or who answers for a bad output.

> **RESPONSIBILITY DOES NOT TRANSFER TO THE VENDOR**
>
> Courts and regulators have generally held the company that deploys an AI tool responsible for what it does, including tools a vendor built and runs. In *Moffatt v. Air Canada*, a tribunal rejected the airline's argument that its chatbot was responsible for its own statements. In *Mobley v. Workday*, a federal court allowed discrimination claims over AI applicant screening to proceed, reasoning that delegating a hiring function to software does not remove the obligations that come with it. Confirm the current status of any case with counsel before relying on it.

The practical conclusion is simple: if a vendor's AI touches your data or your customers, it is part of your AI program, and it needs to be reviewed like it.

SECTION 2

How to use the addendum

The addendum is not a new process. It is a section you add to the vendor review you already run. Paste Appendix A into your existing security questionnaire, alongside the questions on encryption, access control, and incident response.

© 2026 RS Assurance & Advisory | rsassure.com

When to send it

1. **New tool intake.** Any new software that includes AI features, whether or not the AI is the reason you are buying it.
2. **A vendor turns on AI in a tool you already approved.** Product announcements, release notes, and SaaS admin settings reviews are where these show up first.
3. **Renewal.** Re-ask questions 1 and 2 at minimum. Training and subprocessor terms are the ones vendors change most.
4. **Scope change.** When a tool starts touching a more sensitive data class, a new group of users, or customers directly.

Who does what

The business owner of the tool sends the addendum and collects the answers, because they own the relationship. The AI governance owner, or the security lead in a smaller company, reviews the answers against Appendix B. Legal reviews any red flag under question 3, and any contract where the vendor will not put training or deletion commitments in writing.

> **SCALE THE REVIEW TO THE USE**
>
> Not every tool deserves the same scrutiny. A drafting tool that never sees confidential data can move through a fast-track lane with short written answers. A tool that touches confidential or regulated data, makes or informs decisions about people, or acts on its own should provide evidence for all five questions before it goes live. If the review is heavier than the risk, people will route around it.

SECTION 3

The five questions

Each question below is followed by what a good answer looks like and the answers that should slow a purchase down. The vendor-facing wording, with sub-questions and evidence requests, is in Appendix A.

© 2026 RS Assurance & Advisory | rsassure.com

1 Where does our data go?

AI features frequently send data to a model provider the buyer has never heard of. The subprocessor list is where that shows up, and retention terms decide how long prompts, inputs, and outputs stay there.

A GOOD ANSWER

- A current subprocessor list that names model providers
- Defined retention for prompts, outputs, and logs
- Advance notice before a new AI subprocessor is added

RED FLAGS

- No subprocessor list, or model providers missing from it
- Indefinite retention of prompts and outputs
- "We may change subprocessors at any time" with no notice

2 Is our data used for training?

Once data is used to train or fine-tune a model, it is effectively impossible to take back. This is the question where a settings page and a contract can say different things, and only the contract counts.

A GOOD ANSWER

- No training on customer data, stated in the contract or DPA
- Applies to the vendor's model providers too
- Notice before any change to how data is used

RED FLAGS

- Training on by default, with an opt-out only in settings
- "We may use data to improve our services" with no carve-out
- Refusal to put the commitment in writing

3 Who is liable for bad output?

AI output can be inaccurate, biased, or infringing, and it can reach your customers under your name. The indemnification and limitation-of-liability clauses decide who pays when it does.

A GOOD ANSWER

- Indemnity that covers AI output, including intellectual property claims
- A liability cap proportionate to the use
- Described safeguards on output before it reaches you

RED FLAGS

- AI features carved out of indemnity
- A cap so low it is meaningless
- "Output is provided as is" for a customer-facing use

© 2026 RS Assurance & Advisory | rsassure.com

4 Can we get our data deleted?

Deletion is how you leave a vendor cleanly. With AI features, data often sits in more places than the core application: logs, caches, and the model provider's systems.

A GOOD ANSWER

- Deletion on request and at contract end, with a timeframe
- Deletion that reaches subprocessors and model providers
- Written confirmation of deletion

RED FLAGS

- Deletion that stops at the vendor's own systems
- No timeframe, or "within a commercially reasonable period"
- No written confirmation available

5 What certifications do you have?

Independent evidence is what separates a vendor's description of its controls from proof that they operate. The question is not whether a vendor has a report, but whether the report covers what you are buying.

A GOOD ANSWER

- A current SOC 2 Type II report whose scope includes the AI features
- ISO/IEC 27001, and ISO/IEC 42001 or a plan for it
- Reports and certificates provided, not linked

RED FLAGS

- A logo on a website and nothing else
- A report whose scope excludes the AI features
- A report period that ended long ago, with no bridge letter

© 2026 RS Assurance & Advisory | rsassure.com

SECTION 4

Reading the answers

Every review should end in one of three decisions, recorded with a reason. Most tools will land in the middle.

Decision	When it applies	What to record
Approve	All five answers are acceptable for the data class and use, with evidence where the risk calls for it.	Inventory entry with owner, purpose, data class, and approval tier; next review date.
Approve with conditions	The tool is acceptable if its use is limited, or a gap is closed before go-live.	The conditions, such as masked data only, a named user group, or the training opt-out added to the contract, and who verifies them.
Escalate or decline	A red flag that cannot be closed: training without a written opt-out, no indemnity for a customer-facing use, or no deletion from model providers.	The reason, the decision maker, and the alternative offered to the requesting team.

A LOGO IS NOT EVIDENCE

For question 5, read the report itself. In a SOC 2 report, confirm the system description includes the AI features you are buying, read every exception the auditor noted, and list the complementary user entity controls, because those are your responsibilities and they are easy to miss. For an ISO certificate, read the scope statement and confirm the certification body is accredited. For ISO/IEC 42001, ISO/IEC 42006 sets the requirements for bodies that audit AI management systems.

Declining a tool should come with an alternative. A team that asked for a tool had a real need. If the approved route does not meet it, the request will come back as shadow AI.

SECTION 5

After the review: keeping it current

A vendor review is a snapshot. Vendors change models, add subprocessors, and switch on new features between renewals. Three habits keep the answers from going stale.

1. **Record the tool in your AI inventory.** Tool, business owner, purpose, data class, approval tier, the date of the review, and the renewal date. The inventory is what every other governance artifact references.

© 2026 RS Assurance & Advisory | rsassure.com

2 **Re-ask questions 1 and 2 at every renewal.** Subprocessor and training terms are the ones most likely to have changed, and renewal is when you still have leverage to change them back.

3 **Watch for vendor-pushed features.** Review SaaS admin settings for newly enabled AI features. When an approved tool starts doing more, a short scope-change review decides whether to accept the new scope, roll it back, or monitor it with conditions.

These steps are also how vendor due diligence shows up in a formal program. Under ISO/IEC 42001, supplier review of AI vendors and hosted-model providers falls under control A.10.3, and the allocation of responsibilities between you and the vendor under A.10.2. In a SOC 2 examination, it is part of vendor and business partner risk management. Appendix C maps each question to the relevant controls.

SECTION 6

How RS Assurance & Advisory supports vendor AI reviews

RSAA helps organizations build vendor AI due diligence into the governance program they already run, and prove it to customers and auditors.

Vendor AI review support

Independent review of vendor answers and evidence against the five questions, with a clear approve, condition, or escalate recommendation.

SOC 2 report reviews

Reading the reports vendors provide: scope, exceptions, and the complementary user entity controls that become your responsibility.

AI inventory and governance program

A complete AI inventory, acceptable use policy, data-use matrix, and named ownership, built to scale with the company.

ISO/IEC 42001 readiness

Scoping, gap assessment, Statement of Applicability, and internal audit, so supplier controls like A.10.3 hold up at certification.

Our unified methodology maps controls across SOC 1, SOC 2, HIPAA, ISO 27001, ISO 42001, and CMMC, so evidence gathered once can serve multiple examinations. Vendor due diligence is one of the clearest examples: the same review supports your SOC 2, your ISO 27001 supplier controls, and your ISO 42001 program.

© 2026 RS Assurance & Advisory | rsassure.com

APPENDIX A

The AI addendum, ready to paste

Copy everything inside the box into your vendor security questionnaire as a new section. It is written to be sent to the vendor as is.

AI Addendum: Five Questions for Vendors Using AI

Please answer the following for any AI or machine learning features in your product or service, including features provided by third-party models or subprocessors. Attach supporting documents where requested.

1. Where does our data go?

a. Where is our data stored and processed? Please list countries and hosting providers.

b. Which subprocessors, including third-party AI model providers, receive or process our data?

c. How long do you retain our data, including prompts, inputs, outputs, and logs?

d. Will you notify us before adding a new AI subprocessor or changing where our data is processed?

Evidence requested: current subprocessor list and data retention schedule.

2. Is our data used for training?

a. Is our data used to train, fine-tune, or improve any AI model, yours or a third party's?

b. If yes, can we opt out, and is the opt-out on by default for our account?

c. Will you confirm in our contract that our data will not be used for training without our written consent?

d. Will you notify us before changing how our data is used?

Evidence requested: the contract or DPA clause covering training use.

3. Who is liable for bad output?

a. Who is responsible if your AI produces inaccurate, harmful, biased, or infringing output?

b. Do you indemnify customers for claims arising from AI output, including intellectual property claims?

c. How does your limitation of liability apply to AI features?

d. What human review or safeguards do you apply to AI output before it reaches us or our customers?

Evidence requested: the indemnification and limitation-of-liability clauses that apply to AI features.

4. Can we get our data deleted?

a. Can we request deletion of our data, including prompts, outputs, and logs, at any time?

b. How do you delete our data at the end of the contract, and how quickly?

c. Does deletion extend to your subprocessors and AI model providers?

© 2026 RS Assurance & Advisory | rsassure.com

d. Will you provide written confirmation of deletion?

Evidence requested: the deletion procedure and a sample deletion certificate.

5. What certifications do you have?

a. Do you hold a current SOC 2 Type II report? Does its scope include the systems that deliver your AI features?

b. Do you hold ISO/IEC 27001 certification?

c. Do you hold ISO/IEC 42001 (AI management system) certification, or are you working toward it?

d. Please provide the most recent reports or certificates, not a link or logo.

Evidence requested: SOC 2 report (under NDA if needed), ISO certificates with scope statements, and bridge letters if the report period has ended.

© 2026 RS Assurance & Advisory | rsassure.com

APPENDIX B

Reviewer checklist

For internal use only. Do not send this page to the vendor. Keep the evidence listed in the right-hand column with the review record.

#	Question	Good answer	Red flags	Evidence to keep
1	Where does our data go?	Named model providers on the subprocessor list; defined retention; notice of new subprocessors	No list; model providers missing; indefinite retention	Subprocessor list; retention schedule
2	Is our data used for training?	No training, in the contract or DPA, covering model providers	On by default; opt-out only in settings; no written commitment	Contract or DPA clause
3	Who is liable for bad output?	Indemnity covering AI output and IP claims; proportionate cap	AI carved out of indemnity; trivial cap; "as is" for customer-facing use	Indemnity and limitation-of-liability clauses; legal sign-off if flagged
4	Can we get our data deleted?	On request and at contract end, with a timeframe, reaching subprocessors	Stops at vendor systems; no timeframe; no confirmation	Deletion procedure; sample certificate
5	What certifications do you have?	Current SOC 2 Type II covering AI features; ISO 27001; ISO 42001 or plan	Logo only; AI features out of scope; stale report with no bridge letter	Reports and certificates; noted exceptions and complementary user entity controls

Decision

Field	Record
Decision	Approve / Approve with conditions / Escalate or decline
Conditions, if any	
Data class and approval tier	
Business owner	
Reviewer and date	
Next review (renewal) date	

© 2026 RS Assurance & Advisory | rsassure.com

APPENDIX C

Framework mapping

Where each question supports controls in common frameworks. Use this to reuse one vendor review as evidence across programs.

Question	ISO/IEC 42001	ISO/IEC 27001:2022	SOC 2 (TSC)	NIST AI RMF
1. Where data goes	A.10.3	A.5.19, A.5.21	CC9.2	GOVERN 6.1
2. Training use	A.10.2, A.10.3	A.5.20	CC9.2	GOVERN 6.1
3. Liability for output	A.10.2	A.5.20	CC9.2	GOVERN 6.1, 6.2
4. Deletion	A.10.3	A.5.20, A.8.10	CC9.2, C1.2	GOVERN 6.1
5. Certifications	A.10.3	A.5.22	CC9.2	GOVERN 6.1

Reference titles: ISO/IEC 42001 A.10.2 allocating responsibilities, A.10.3 suppliers. ISO/IEC 27001:2022 A.5.19 information security in supplier relationships, A.5.20 addressing security within supplier agreements, A.5.21 managing security in the ICT supply chain, A.5.22 monitoring, review, and change management of supplier services, A.8.10 information deletion. SOC 2 CC9.2 vendor and business partner risk management, C1.2 disposal of confidential information. NIST AI RMF GOVERN 6.1 policies for third-party AI risks, GOVERN 6.2 contingency processes for third-party failures. Titles are paraphrased; check them against your licensed copies of each standard.

© 2026 RS Assurance & Advisory | rsassure.com

ABOUT

RS Assurance & Advisory

RS Assurance & Advisory is an assurance and advisory firm headquartered in Southlake, Texas. We work exclusively with seasoned practitioners — no junior staff on engagements — and our unified methodology maps controls across frameworks so your team answers each evidence request once.

Service lines: SOC 1, SOC 2, and SOC 3 · SOX · ISO 27001 · ISO 42001 · CMMC 2.0 · HITRUST · HIPAA · GDPR · Nacha ACH Rules Compliance · PCI DSS readiness

Richard Stevenson

Managing Partner, RS Assurance & Advisory

richards@rsassure.com · (903) 229-0341 · rsassure.com

This paper is an educational starting point. It is not legal advice and does not replace review of vendor contracts by counsel. Framework references are paraphrased and should be checked against your licensed copies of each standard. Case law and regulatory positions change; confirm current requirements with counsel before relying on any statement here in a contract or public claim. *© 2026 RS Assurance & Advisory. All rights reserved.*