CHECKLIST · FREE RESOURCE

# Enterprise Security Review Readiness for AI Products

A checklist and Excel tracker for SaaS and AI companies facing enterprise security reviews: what questionnaires ask about data use, model providers, isolation, access and AI risks, and the evidence that closes each item.

By **Ratomir Jovanovic** · ratomir.com/resources/enterprise-ai-security-review-readiness

---

This checklist is for founders, product leaders and security leads at SaaS and AI companies that are starting to sell to larger organizations, or that keep losing weeks to security questionnaires once a deal reaches procurement. It is written from the seat of someone who leads product and also audits information security management systems against ISO/IEC 27001, so it treats a security review as what it is: a buyer deciding whether your controls are adequate for the data and the risk they are about to hand you.

The outcome is one owned view of every question you are likely to face, the honest answer, and the evidence behind it. With that in place, questionnaires become an assembly job rather than a fire drill, sales stops promising things engineering has not built, and the gaps you do have are known, prioritized and on a roadmap instead of discovered by a customer's risk team in week six of a deal.

## How to use this

1. Open the Excel tracker that comes with the PDF guide. Each row is one check, with columns for status, owner, evidence link and notes, plus the reason buyers ask and a priority.
2. Assign an owner per section, not per row. Data use and sub-processors usually sit with product and legal, isolation and encryption with engineering, incident response and continuity with whoever runs operations.
3. For every check, write the answer you would give a buyer today in one or two plain sentences. If you cannot write it without hedging, mark it as a gap.
4. Link the evidence itself, not a description of it: a policy, a configuration screenshot, a diagram, a test report, a ticket.
5. Close P1 gaps first. P1 items stop or stall deals, P2 items come up in most reviews, P3 items matter mainly for regulated or very large buyers.
6. Review the tracker every quarter and after any architecture change, new model provider or new sub-processor. A stale answer is worse than a known gap, because it is a statement you can be held to.

## What buyers are actually trying to learn

Questionnaires look long and arbitrary, but under the hundreds of questions a reviewer is trying to answer four things. What data of ours will you hold, and what will you do with it? Who else touches it? How would you know if something went wrong, and what would you do? Will you still be running next month? For AI products a fifth question now sits on top: can your model or agent be manipulated into doing something harmful with our data, and will our data end up training a model we do not control?

Reviewers are rarely trying to catch you out. Most work from a standard template, score your answers against their own risk appetite and escalate only the items that fall outside it. A short, specific answer with evidence closes an item. A long, vague answer opens follow-up questions, and follow-up rounds are where deals lose weeks. The goal of this checklist is to make your first-round answers good enough that there is no second round on most sections.

## The areas every review covers

The tracker groups checks into ten sections. The table below summarizes what each section is really asking and the kind of evidence that usually closes it.

| Area | What reviewers want to know | Evidence that closes it |
|---|---|---|
| Data use and model training | Whether customer data trains any model, yours or a provider's | Contract clause, provider terms, product setting, data flow diagram |
| Retention and deletion | How long data lives, including prompts, outputs and logs | Retention schedule, deletion procedure, sample deletion record |
| Sub-processors and model providers | Who else processes the data, where, and under what terms | Sub-processor list, DPAs, change notification process |
| Tenant isolation | Whether one customer could ever see another's data | Architecture diagram, isolation design notes, test results |
| Access control and SSO | Who can log in, with what, and who on your side can see data | SSO and SCIM support, role model, access review records |
| Audit logging | Whether actions are recorded and available to the customer | Log catalogue, retention, sample export, customer-facing audit log |
| Encryption | How data is protected in transit and at rest, and who holds keys | Configuration evidence, key management description |
| Incident response | How you detect, contain and notify | Incident plan, notification commitments, tabletop exercise record |
| Vulnerability management | How you find and fix weaknesses, and how fast | Scan reports, patch SLAs, penetration test summary |
| Business continuity | Backups, recovery targets and whether you have tested them | BCP, RPO and RTO, restore test record |

Every section has a P1 core. If you only have time for a short sprint before a large deal, fix the P1 items in data use, sub-processors, access control and incident response first. Those four are where reviewers most often stop the process rather than add a note.

## AI-specific questions in more detail

AI products get the standard questions plus a newer set that many reviewers are still writing as they go. Expect the wording to vary; the substance is fairly stable.

**Training and data use.** The first question is almost always whether customer content is used to train or fine-tune models. Answer it for three layers separately: your own models, any fine-tuned models you host per customer, and the third-party model providers you call. "We do not train on your data" is incomplete if your provider's terms allow it, or if your product analytics pipeline copies prompts into a dataset someone later uses for evaluation. Check the provider's current enterprise or API terms, not the consumer terms, and keep a dated copy.

**Provider retention.** Many model providers retain inputs and outputs for a period for abuse monitoring unless a zero or reduced retention arrangement applies. Know which arrangement you have, for which endpoints, and say so plainly. If you route to more than one provider, the answer has to hold for all of them, including fallbacks.

**Prompt injection.** Reviewers want to know that you understand that any text the model reads (documents, emails, web pages, tool results) can carry instructions, and that you have limited what a manipulated model can do. The credible answer is about blast radius rather than detection: the model acts with the user's permissions and no more, tools that write, send or delete require confirmation or are scoped narrowly, and retrieval respects the same access rules as the rest of the product.

**Output handling.** Model output is untrusted input to the next system. Reviewers ask whether outputs are rendered safely (no script execution, no automatic link following), whether outputs that reach other systems are validated, and whether users are told when content is generated.

**Human oversight and change.** Expect questions on who is accountable for model behaviour, how you evaluate a new model version before switching, and whether customers are told about material model changes. A short model change procedure with an evaluation step and a changelog answers most of this.

## How ISO/IEC 27001 and SOC 2 fit

Neither certification is required to pass a security review, but both change the conversation. ISO/IEC 27001 certifies that you run an information security management system: you assess risk, choose controls (the 2022 version has 93 in Annex A), document them in a Statement of Applicability, and improve them through internal audit and management review. SOC 2 is an attestation report from an independent auditor against the Trust Services Criteria. A Type I report covers control design at a point in time, a Type II report covers whether controls operated over a period, usually several months to a year.

In practice, buyers in North America ask for SOC 2 more often and buyers in Europe tend to ask for ISO/IEC 27001, but most large buyers accept either, plus a short questionnaire for what it does not cover. Two details matter more than which one you choose. First, scope: a certificate that covers your corporate IT but not the product the buyer is purchasing does little for them, and experienced reviewers check the scope statement. Second, neither framework goes deep on AI-specific risks. ISO/IEC 42001, the management system standard for AI, is starting to appear in questionnaires, and even without it you will still need the AI section of this checklist.

From the audit side, the most useful thing to know is that auditors and reviewers look for evidence that controls operate, not for well-written policies. An access review policy with no completed reviews is a nonconformity. A modest policy with twelve months of records is a strength. If certification is a year away, a maintained tracker, a small set of real policies and the operating records behind them will carry many mid-market reviews in the meantime, and they become the foundation of the certification work anyway.

## Building the evidence pack

Most of the questionnaire effort goes into finding and re-explaining the same material. Build it once as a pack you can share under NDA:

- A security overview of four to eight pages, written for a reviewer rather than a marketer, covering architecture, data flows, the AI components and the controls in each section of this checklist.
- A current architecture and data flow diagram that shows where customer data, prompts and outputs go, including every model provider.
- The sub-processor list with locations, purposes and a link to how customers are notified of changes.
- Your core policies: information security, access control, incident response, business continuity, vulnerability management, data retention and acceptable AI use.
- A penetration test summary or attestation letter from the last twelve months, with remediation status for high findings.
- The certificate and scope, or the SOC 2 report, or a dated roadmap to either.
- A completed copy of one widely used standard questionnaire, which answers a large share of any custom one.

Keep each document dated and owned in the tracker. When a reviewer asks for something not in the pack, add it after the deal so the next reviewer does not have to ask.

## Worked example

This example is hypothetical. A 40-person company sells an AI meeting summary product and reaches procurement with a large insurer. The questionnaire has a little over 200 questions. The first pass through the tracker shows three P1 gaps: the model provider's API retains inputs for abuse monitoring and the company had not checked whether a reduced retention option applied; SSO is available only on the top plan and has no SCIM provisioning; and the last penetration test is 16 months old.

The product lead takes ownership of the provider question, confirms the current terms, requests the reduced retention arrangement and updates the data flow diagram and sub-processor list to match. Engineering commits SCIM to the next quarter and the sales lead agrees SSO will be included for this customer's tier. A new test is booked, and in the meantime the company shares the previous report with remediation evidence plus the booking confirmation. The answers go back as short statements with links. The reviewer accepts the first two with a contractual commitment and records the third as a tracked risk with a date. The deal proceeds, and all three items stay in the tracker until the evidence of completion is attached.

## Common mistakes

- Answering "yes" to capabilities that exist only on the roadmap. Questionnaire answers often end up attached to the contract.
- Letting sales complete questionnaires alone. They know the customer, but not which answers are true this quarter.
- Describing intentions instead of controls. "We take security seriously" closes nothing; "access to production requires SSO with MFA and is reviewed quarterly, last review attached" closes the item.
- Forgetting the AI layer in the sub-processor list. Model providers, embedding services and evaluation tools all process customer content.
- Treating deletion as a database delete. Backups, logs, caches, vector indexes and provider retention all count.
- Over-answering. Extra detail invites extra questions; answer what was asked, with evidence.
- Letting the tracker go stale after the deal closes.

## When not to use this

If you sell only to individuals or small teams that never send a questionnaire, this is premature; a clear privacy policy and sensible defaults will serve you better. It is also not a substitute for a formal risk assessment, a certification audit or legal review of your contracts and data processing terms. And if a buyer operates under sector rules such as financial services or health regulation, expect additional requirements that this checklist names only at a general level; use it as the base layer and add their specific controls on top.

---

Editable template: Readiness tracker (Excel), available at ratomir.com/resources/enterprise-ai-security-review-readiness

© 2026 Ratomir Jovanovic. Free to use inside your team. Please link to ratomir.com/resources/enterprise-ai-security-review-readiness when you share it.