

AI Readiness & Security Checklist

A self-assessment for business leaders

This checklist is the companion to our AI Security for Business Leaders series. It gathers the questions each article asked and puts them in a form you can work through with your leadership team or hand to whoever manages your IT.

It is not a test, and there is no passing score. Most organizations answer “no” or “not sure” to a good portion of it the first time through. That is the normal starting point, not a warning sign. AI arrived in most businesses faster than anyone’s ability to govern it.

How to use it

Work through the seven sections in order. Each opens with a short note on why the section matters and who is best positioned to answer it. Some sections are leadership questions; others might need whoever handles your IT in the room. Answer honestly rather than aspirationally; a checklist filled out the way you wish things were is worth nothing.

Answer “not sure” freely. It is a real answer and, as the scoring section explains, often a more useful one than “no.” Scoring takes about a minute at the end. The whole thing should take 15 to 20 minutes.

1. Visibility

You cannot govern what you cannot see. Most organizations have a gap between the AI tools leadership has approved and the ones employees actually use day to day, and that gap is where nearly every other risk in this checklist originates. This section is usually the hardest one to answer honestly. That difficulty is itself the finding.

Best answered by: Leadership, together with whoever handles IT.

	Yes	No	Not sure
We can name the AI tools currently in use across the organization.			
We know which of those tools were formally approved and which were adopted independently.			
If a new AI tool appeared in a department’s workflow next month, someone would find out.			
If an employee leaves tomorrow, we would know what company data lives inside tools we don’t manage.			

2. Data Protection

Every prompt is a data transfer. Information an employee types into an AI tool leaves the organization the moment they hit enter, and what happens to it after that depends on terms most people have never read.

Best answered by: Leadership, with compliance or operations.

	Yes	No	Not sure
We have defined, in writing, which categories of information must never be entered into an AI tool.			
Employees have seen that list and know where to find it.			
We know whether any AI tool in use is currently receiving regulated data → client records, employee information, health information, or financial data subject to compliance requirements.			

3. Verification

AI produces confident output whether or not it is correct, and the tone and format of a response carry no information about its accuracy. AI being wrong occasionally isn't really the problem – people are wrong too. What matters is proportionality: the higher the consequence of being wrong, the more human verification is required.

Best answered by: Department heads, since the answers vary by function.

	Yes	No	Not sure
We have named the categories of work that require human review before use (i.e. client-facing content, compliance research, financial figures, regulated matters).			
For each of those categories, a specific person is responsible for the review.			
Employees understand that confident AI output is not the same as accurate AI output.			

4. Access and Permissions

An AI tool that drafts text is one level of risk. An AI tool that can send, update, delete, or approve is another entirely. Three questions govern the difference: what the tool is able to do, whose access it is using, and who approves before it acts.

Best answered by: Whoever handles your IT.

	Yes	No	Not sure
We know which AI tools can change something (send, update, delete, approve) rather than only produce text.			
Each of those tools has permissions scoped to its task, not to the person who set it up.			
AI tools that take action run under their own identity rather than on an employee's login.			
Actions that cannot be undone (anything deleted, anything sent to a customer, anything involving money) require human approval before they happen.			

5. Vendor Review

The AI capability inside a business application is often not built by the company selling it. A model provider, a cloud host, and a specialist vendor can all stand between your business and a working tool, and typically only one of them appears in the procurement conversation. Vendor review already exists in most organizations; this is about extending it to a category it was not built for.

Best answered by: Whoever owns vendor relationships and contract review.

	Yes	No	Not sure
For each AI-enabled tool, we know all the companies involved in delivering it			
We have reviewed what each vendor does with our data: what is retained, for how long, and where it is stored.			
Where required, data processing or business associate agreements are executed and explicitly cover the AI functionality.			
We know how we would get our data back, and in what format, if a vendor were acquired or discontinued the product.			

6. Policy and People

An unnamed policy is not the absence of a policy. It means employees are setting the policy themselves, one individual decision at a time, without knowing what is at stake. A policy that only lists what is prohibited leaves the harder question unanswered → “what should I use instead?”

Best answered by: Leadership.

	Yes	No	Not sure
We have a written AI acceptable use policy.			
That policy names which tools are approved, not only what is prohibited.			
Someone specific is accountable for AI oversight, and people in the organization know who that is.			

7. Monitoring and Response

Most of this section is about the distance between something going wrong and someone noticing. AI-related problems rarely announce themselves: there is no alert when a tool’s output quality drifts, and no error message when an agent does something it shouldn’t have.

Best answered by: Whoever handles your IT, with leadership.

	Yes	No	Not sure
AI tool usage and permissions are reviewed on a recurring schedule, not only at adoption.			
Where AI tools take action, activity is logged somewhere the tool itself does not control.			
We know what we would do if sensitive information were entered into an AI tool it should not have been.			
AI-related scenarios are included in our incident response planning.			

Scoring

Count one point for each “yes” while leaving “no” and “not sure” unscored for now. The section below explains why they are worth separating. There are 25 questions in total.

0–9 points — the tools are in use; the decisions haven’t been made yet

This is the most common starting point, and it is not a sign of negligence. It reflects how quickly AI entered most organizations relative to anyone’s ability to put structure around it. The highest-value next steps here are also the cheapest ones: find out what is actually in use and write down what should never be entered into it. Sections 1 and 2 are where to start, and almost everything else in this checklist gets easier once those two are answered.

10–18 points — some guardrails exist, usually where a problem surfaced

Organizations in this range have typically addressed the risks that announced themselves: a data policy written after someone asked a good question, a review step added after an error made it out the door. What is usually missing is the part nobody has yet been forced to think about: what the tools are actually able to do, and who would find out if something went wrong. Sections 4 and 7 are where the gaps tend to cluster.

19–25 points — the structure is there; the work now is keeping it current

At this level the failure mode changes. It is no longer about missing controls, it is about controls that describe conditions which no longer exist: a vendor review from eighteen months ago, an approved-tools list that predates three new adoptions, a permission grant nobody has revisited. Set a recurring date to review this and treat the checklist as the agenda for that conversation.

What your “not sure” answers are telling you

Count them separately, because they mean something different from a “no.” A “no” means somebody made a decision. A “not sure” means nobody in the room has the visibility to answer the question at all, which is frequently the more urgent of the two.

If your “not sure” answers cluster inside a single section, that section is where the next hour of effort belongs. If they are spread evenly across all seven, the underlying issue is usually visibility rather than any specific control, and Section 1 is the place to begin.

Where to go from here

Most of this checklist can be answered in a single leadership conversation, and the items that can’t are usually a short exercise for whoever manages your IT.

Even if you have an IT Department or existing IT provider, bringing this all together can be a significant feat. We’re always happy to have a conversation, learn more, and see how we could be of help in your AI journey. Reach out for a conversation about AI Governance & Enablement.

This checklist draws on guidance from the OWASP Top 10 for Large Language Model Applications, 2025, and accompanies the AI Security for Business Leaders series.