



CyberSilience

Cybersecurity | AI Governance | Training

FREE RESOURCE | AI ADOPTION CHECKLIST

10 Questions to Ask Before Deploying AI at Work

A security-first checklist for leaders, managers and IT teams who want AI to help without creating new risk.

Why security comes first

Most AI conversations start with what the tool can do. Ours starts with what it can expose. Every AI deployment connects your data, your people and your systems to a new technology and, usually, a new vendor. Getting the security and governance questions answered first is what lets you move faster later with confidence.

The ten questions below are ordered deliberately: data and vendors first, then access and attack paths, then legal, accountability, monitoring and people. Use them before a pilot, a purchase or the moment a team asks to switch a new AI feature on.

HOW TO USE THIS

1. Work through each question with the business owner, IT or security, and legal or privacy in the same room.
2. Score every question on the final page: Ready, Partly, or Not yet.
3. Any question that is not answered means the deployment is not ready. Resolve it or reduce the scope.

The ten at a glance

Data and vendors	1 to 3
Access and attack paths	4 to 6
Legal, accountability, monitoring and people	7 to 10

01

QUESTION

What data will this AI touch, and how sensitive is it?

Data exposure is the number one way AI deployments go wrong. If you have not classified the data first, you cannot judge any other risk.

ASK YOUR TEAM AND VENDOR

- Which data classes (customer, employee, financial, IP, regulated) will be entered or connected?
- Is any of it prohibited from leaving our environment today?
- Can we limit the tool to the minimum data it needs?

RED FLAG Nobody can list the data sources, or the answer is "everything in the shared drive."

02

QUESTION

Where does our data go, who can see it, and is it used to train models?

Prompts, files and outputs may be stored, logged, reviewed by staff or reused. Terms differ sharply between consumer, business and enterprise tiers.

ASK YOUR TEAM AND VENDOR

- Where is data processed and stored (region), and for how long is it retained?
- Is our data excluded from model training by contract, not just by a setting?
- Can we delete our data on request and get confirmation?

RED FLAG Retention or training terms are vague, or they change with a plan tier no one has checked.

03

QUESTION

Has the vendor passed a real security and third-party risk review?

An AI vendor becomes part of your supply chain and your attack surface the day you connect it.

ASK YOUR TEAM AND VENDOR

- Do they hold independent assurance (SOC 2 Type II, ISO 27001) and will they share reports?
- What is their breach notification commitment, and who are their sub-processors?
- What is our exit plan if they change terms, get breached or shut down (vendor lock-in)?

RED FLAG The tool was adopted with a credit card and a click-through agreement.

04

QUESTION

Who gets access, and how is that access controlled?

Unmanaged accounts and "shadow AI" put company data into tools that security cannot see or switch off.

ASK YOUR TEAM AND VENDOR

- Is access through single sign-on with MFA, and tied to joiner, mover and leaver processes?
- Are roles least-privilege, with admin rights tightly held?
- Do we have a way to find and address unsanctioned AI tools already in use?

RED FLAG Shared logins, personal accounts used for work, or no list of who is using what.

05

QUESTION

How could this system be attacked or manipulated?

AI introduces new attack paths: prompt injection, data leakage through outputs, poisoned content, and manipulated or spoofed AI-generated messages.

ASK YOUR TEAM AND VENDOR

- Has anyone threat-modelled this use case, including prompt injection from emails, documents and web pages?
- Have we tested it adversarially before go-live, and will we retest after changes?
- Could it be used to help attackers target us (deepfakes, convincing phishing, impersonation)?

RED FLAG The only testing done was checking that it gives good answers.

06

QUESTION

What can the AI actually do, not just say?

Once an assistant can send email, edit records, run code or call other systems, a bad output becomes a bad action.

ASK YOUR TEAM AND VENDOR

- Which connectors, plug-ins and permissions does it hold, and are they read-only where possible?
- Which actions require a human to approve before they run?
- Can we revoke its access instantly?

RED FLAG Broad write access granted "to make it more useful" with no approval step.

07

QUESTION

Which laws, contracts and policies apply?

Privacy law, sector regulation, client confidentiality, copyright and emerging AI rules all follow the data and the decision, not the tool.

ASK YOUR TEAM AND VENDOR

- Have legal, privacy and compliance reviewed the use case and any impact assessment required?
- Do client contracts or NDAs restrict use of third-party AI on their information?
- Who owns the outputs, and are there intellectual property or licensing risks?

RED FLAG Compliance is asked to sign off after the pilot is already live.

08

QUESTION

Who is accountable when the AI is wrong?

AI outputs can be confident, inaccurate and biased. Someone named must own the decision and the outcome.

ASK YOUR TEAM AND VENDOR

- Who is the business owner for this system, and who is the risk owner?
- Where is human review required, and what does the reviewer check?
- How are errors, bias and complaints reported and corrected?

RED FLAG "The AI said so" is treated as a sufficient explanation.

09

QUESTION

Can we monitor it, log it and respond if something goes wrong?

You cannot investigate what you did not record. AI needs to be inside your existing detection and incident response, not beside it.

ASK YOUR TEAM AND VENDOR

- Are prompts, outputs and admin actions logged and fed to our security monitoring?
- Does our incident response plan cover an AI data leak or misuse, with a tested kill switch?
- Who reviews usage and drift, and how often?

RED FLAG No logs, no owner for alerts, and no plan for an AI-specific incident.

10

QUESTION

Are our people trained, and is the value worth the risk?

The strongest control is an informed employee, and a deployment without a clear business case is risk with no return.

ASK YOUR TEAM AND VENDOR

- Is there an acceptable use policy in plain language, and has everyone been trained on it?
- Do staff know what never goes into an AI tool and how to report a concern?
- What measurable outcome justifies the risk, and when will we review it?

RED FLAG A rollout announced by email with no training and no success measure.

Readiness scorecard

Mark each question honestly. One "Not yet" in questions 1 to 6 is a stop sign.

#	Question	Ready	Partly	Not yet
01	Data sensitivity understood	☐	☐	☐
02	Data handling, retention and training	☐	☐	☐
03	Vendor security review passed	☐	☐	☐
04	Access and identity controlled	☐	☐	☐
05	Attack paths threat-modelled and tested	☐	☐	☐
06	AI permissions and human approvals limited	☐	☐	☐
07	Legal, privacy and contract review done	☐	☐	☐
08	Named accountability and human oversight	☐	☐	☐
09	Logging, monitoring and incident response	☐	☐	☐
10	Training, policy and business case	☐	☐	☐

How to read your result

8 to 10 Ready: proceed with a controlled pilot and keep monitoring.

5 to 7 Ready: close the gaps first, starting with data, vendor and access.

Under 5 Ready: pause. Reduce scope or fix the foundations before deploying.

Need help answering these?

CyberSilience trains teams and leaders to adopt AI safely, with practical security awareness, governance and risk training built for real workplaces.

Website: www.cybersilience.com

Email: info@cybersilience.com

This checklist is general guidance for awareness and discussion. It is not legal advice or a substitute for a formal risk assessment.