The AI security roadmap for small business.

Ninety days from "people are quietly using ChatGPT" to "we know what we run, we've decided what's allowed, and we'd notice if it went wrong." No new headcount. No platform purchase. Mostly conversations, one page of policy, and a handful of settings.

Written for the owner or manager of a 5–250 person firm who has read one too many AI-risk headlines and wants a plan measured in evenings, not quarters.

document	ai-security-roadmap
author	peter bassill
version	1.0 · 2026-07
locale	en_GB · UTC
licence	CC BY 4.0

§ 00 / THE SITUATION

Your business is already using AI. The only question is whether you know about it.

Surveys through 2025 and 2026 keep landing on the same number: roughly half of employees use AI tools their employer never approved, and around half of those have pasted non-public company information into them. Most small firms have no AI policy at all. If you have more than five staff, assume it is happening in your business this week.

That is not a moral failing on anyone's part. The tools are useful. Your bookkeeper found that ChatGPT drafts chase-up emails faster than she can. Your sales lead is summarising discovery calls with an AI notetaker he pays for on a personal card. Somebody in ops has wired an "agent" to their inbox because a YouTube video showed them how. They are trying to do their jobs well.

The risk is not "AI". The risk is the same three things it has always been, arriving through a new door:

- **Data leaving the building.** Customer records, pricing, contracts and staff data pasted into consumer tools, on consumer terms, where you cannot audit, delete, or account for it. If you have obligations under UK GDPR — and you do — those obligations did not pause because the destination was a chat window.
- **Attacks getting cheaper.** Phishing with perfect grammar, deepfaked voices on the phone to your finance team, fraud at a marginal cost near zero. The defences are procedural and they still work — but only if you have actually put them in place.
- **Software acting on your behalf.** AI agents that send emails, change records and move files are software with your credentials and their own judgement. Given broad permissions and no supervision, they are a new class of insider risk you configured yourself.

WHAT THIS ROADMAP IS NOT

It is not a compliance framework, an AI strategy, or a sales funnel for a platform. It is the sequence of small, boring, high-leverage actions I would take in a firm of 5–250 people, in the order I would take them. Every action states who does it and how long it takes. If a step needs money, it says so; most do not.

How to use it

Four phases across ninety days: **Look** (days 0–7), **Decide** (days 8–30), **Defend** (days 31–60), **Detect & drill** (days 61–90) — then a quarterly loop that takes an hour. Do them in order. The temptation is to jump to the technical controls in phase three; resist it. Controls chosen before you know what your people actually use will be the wrong controls, and your staff will route round them within the month.

Companion documents in this kit: the one-page staff cheat sheet (02), the AI usage policy template (03), the agents & MCP guide (04), the board briefing template (05), the self-assessment (06), and the DPIA guide (07). This roadmap tells you when to reach for each.

PHASE 1 · DAYS 0–7

Look — find out what you are actually running

effort: ~3 hours · cost: £0 · owner: you

You cannot defend what you cannot see. The good news: in a small firm, visibility is a conversation, not a procurement.

1.1 — Ask, without threatening

Send one message to all staff. The tone matters more than the wording; if it reads like the opening of a disciplinary, you will get silence and the shadow use will simply go quieter. Something like:

```
"We're putting a proper footing under how we use AI here — mostly so we
can say yes to more of it, safely. Reply with: which AI tools you use for
work (including personal accounts), what for, and roughly how often.
Nobody is in trouble. The tool you tell me about is a tool I can
make safe. The one I find out about later is the problem."
```

In my experience the amnesty framing roughly doubles what gets declared. You are trying to move use from shadow to daylight, and daylight has to be a safe place to stand.

1.2 — Check the places tools leave footprints

- **Expenses and card statements** — search for OpenAI, Anthropic, Otter, Fireflies, Jasper, Midjourney, Perplexity, Gemini, Copilot. Ten minutes with your bookkeeper.
- **The company Microsoft / Google admin console** — both list third-party apps staff have connected to their work accounts. This is where AI notetakers that joined your meetings show up.
- **Browser extensions** — if you manage browsers centrally, list installed extensions; AI writing assistants that read every page live here.
- **Your website and products** — did anyone add an AI chat widget to the site, or wire an AI feature into something customers touch? That one goes to the top of the list.

1.3 — Write the inventory. One page.

For every tool: *name · who uses it · what goes into it · personal or business account · does it touch customer or staff personal data — yes or no.* A spreadsheet with five columns. Do not build a register with thirty fields; you will never update it.

END OF WEEK ONE — YOU SHOULD BE ABLE TO SAY

- "These are the AI tools in use at this firm, near enough."
- "These two or three handle customer or staff personal data — they get attention first."
- "This many are on personal accounts on consumer terms." (Expect a bigger number than you hoped.)

PHASE 2 · DAYS 8–30

Decide — the policy, the approved list, the data rules

effort: ~4 hours · cost: £0–40/user/yr · owner: you + whoever runs IT

One page of policy beats forty pages of policy, because one page gets read. Document 03 in this kit is the template; this phase is about the three decisions you must make to fill it in.

2.1 — Decide the data rule first. It carries most of the weight.

Sort your information into three bins. This is the part of the policy people will actually remember:

BIN	WHAT'S IN IT	RULE
GREEN	Public or would-be-public material: published pages, blank templates, general research questions, text with nothing identifying in it.	Any approved tool, freely.
AMBER	Internal but not regulated: draft proposals, internal process docs, code without secrets.	Approved *business-tier* tools only — accounts the firm controls, with training-on-your-data switched off.
RED	Customer personal data, staff data, financials, contracts, credentials, anything under NDA, anything you'd owe the ICO a phone call about.	Does not go into AI tools. At all. Unless a named person has approved a named tool with a contract that covers it.

Staff do not need to remember GDPR. They need to remember three colours and one question: "*which bin is this in?*"

2.2 — Pick the approved list. Say yes to something.

A policy that only says no is a policy that gets ignored. Pick two or three tools that cover most declared use from phase one — typically one general assistant (ChatGPT, Claude or Copilot on a business plan), plus whatever is already inside your Microsoft or Google subscription. Put them on **business tiers, on company accounts**. The business tiers of the mainstream tools do not train on your data by default and give you central control of who has access; that single upgrade retires more risk than any other action in this roadmap, at £20–40 per user per month for the users who need it — and many need only the free bundled options.

2.3 — Decide who says yes to new tools

Name one person. New tool requests go to them; they check four things — *who owns the data · does it train on inputs · where is it processed · can we delete our data* — and answer within a week. If the answer takes longer than a week, people stop asking.

SHIP IT

Fill in the policy template (03), email it to everyone, and pin the cheat sheet (02) where people will see it. Then say the sentence out loud in a team meeting: "*the safe list exists so you can use this stuff — bring me anything you want added.*" The meeting matters more than the email.

PHASE 3 · DAYS 31–60

Defend — settings, suppliers, and the fraud calls

effort: ~5 hours · cost: mostly £0 · owner: you + IT

Now — and only now — the technical controls. Everything in this phase is a setting you already pay for or a procedure that costs nothing.

3.1 — Harden the accounts you just created

- **Single sign-on where the tier allows it,** so leavers lose AI access the day they leave, same as email. Where SSO isn't available, company email + MFA, enforced.
- **Turn off training-on-your-data** explicitly and screenshot the setting. Terms change; your screenshot is your evidence of the position you set.
- **Kill the orphans.** Personal-account use of the tools you have now licensed properly should end. Ask staff to export anything they need and delete work data from personal AI accounts — most mainstream tools have a delete-my-data control; link to it in the same email.
- **Restrict who can connect third-party apps** to your Microsoft/Google tenancy. This one admin setting is what stops the next AI notetaker joining your board meeting uninvited.

3.2 — The two fraud procedures AI made non-optional

Deepfaked voices and flawless phishing mean *recognising the sender* is no longer a control. Two procedures replace it:

- **Payment changes:** any new payee, any change to bank details, any "urgent" payment request — verified by a call to a number you *already had on file*, never one supplied in the request. No exceptions, including for the boss. Especially for the boss.
- **Out-of-band confirmation:** for anything irreversible requested by voice or video call, agree a callback or a pre-agreed challenge between the people who move money. Thirty seconds of friction against the year's most expensive category of fraud.

3.3 — Ask your suppliers what they are doing

Your accountant, your MSP, your CRM vendor and your payroll bureau are all "using AI" now. Send the important ones four questions — in writing, so the answers are on record:

```
1. Which AI tools or features touch OUR data in your service?
2. Is our data used to train anyone's models? Where is it processed?
3. Can you switch AI features off for our account if we ask?
4. Who do you notify, and how fast, if an AI-related incident touches
   our data?
```

You are not auditing them. You are creating a paper trail that you asked — which is precisely what your insurer, your regulator and your biggest customer will want to see if something goes wrong downstream.

PHASE 4 · DAYS 61–90

Detect & drill — noticing, and knowing what you'd do

effort: ~3 hours + 45-minute exercise · cost: £0 · owner: you

Perfect prevention is not on offer at any budget. What a small firm can genuinely have is *noticing quickly* and *a rehearsed first hour*.

4.1 — Three lightweight tripwires

- **The monthly question.** Add to your team meeting: "*anyone seen anything this month that felt AI-generated, or nearly pasted something they shouldn't have?*" Near-miss reporting is the cheapest detection system ever invented, and it only works if the answer to a confession is "thank you", not a lecture.
- **The admin console, quarterly.** Fifteen minutes in the Microsoft/Google third-party app list, looking for anything new. Diary it; it will not happen otherwise.
- **The expenses grep, quarterly.** Same search as phase one. New AI subscriptions on expenses = new shadow tools = a conversation, not a punishment.

4.2 — Write the one-page playbook

Four scenarios, each with three lines — *first action · who to tell · who decides what's next:*

SCENARIO	FIRST ACTION
Someone pasted RED-bin data into an AI tool	Note what, where, when. Use the tool's delete controls. Assess whether it's a personal-data breach — if it is, you may have 72 hours to tell the ICO. Say thank you for reporting.
Deepfake or AI-phish attempt on the firm	No blame for nearly falling for it. Preserve the message/recording. Warn the team same day — the second attempt usually follows the first within the week.
An AI tool or agent did something wrong on your systems	Revoke its credentials first, investigate second. Screenshot before you fix.
A supplier tells you their AI leaked your data	Get specifics in writing: what data, whose, when. Then treat as your own breach clock.

4.3 — Run one tabletop. Forty-five minutes.

Pick the deepfake-payment scenario. Get the three people who'd be involved in a room: "*the phone rings, it's my voice, I'm asking you to pay a new supplier £18,400 today. Walk me through what happens.*" The first run is always mildly embarrassing. That is the point of running it before it is real.

DAY 90 — WHERE YOU NOW STAND

You know what you run. Staff know the three bins and use tools you control. Payment fraud has a procedure that doesn't depend on spotting fakes. Suppliers are on notice, in writing. And an incident has a rehearsed first hour. That position is better than most firms fifty times your size — at a cost of roughly fifteen hours and the price of a few business licences.

§ 05 / AFTER DAY 90

The quarterly hour, and the regulatory weather

The loop — one hour, four times a year

MINUTES	DO
00–15	Re-run the inventory checks: expenses grep, admin console, one "what's new?" question to the team.
15–30	Review the approved list. Add what people are asking for (with the four vendor questions). Remove what nobody uses.
30–45	Re-read the policy. Terms of service drift; so do tools. Update the one page — version and date it.
45–60	Score yourself against the self-assessment (06). The trend matters more than the number.

The regulatory picture, briefly and without panic

United Kingdom. There is no general "AI law" to comply with. What binds you is what already bound you: **UK GDPR** applies in full whenever personal data goes into, or comes out of, an AI tool — the ICO has published AI-specific guidance and has been increasingly pointed about security expectations under Article 32. If you sell to larger firms, expect AI-usage questions in their supplier questionnaires now; your phase-two paperwork answers them.

European Union. The EU AI Act is in force and arriving in stages; prohibitions and AI-literacy duties already apply, transparency duties for AI-generated content land from late 2026, and the high-risk regime — after the 2026 postponements — bites from late 2027. If you merely *use* mainstream AI tools in ordinary business, your obligations are modest: know where AI touches customers, label synthetic content honestly, and keep a human accountable for decisions about people. If you *build* AI into a product sold in the EU, read further than this paragraph.

The honest summary. For a UK small business in 2026, regulatory risk from AI is overwhelmingly ordinary data-protection risk wearing a new hat. Handle the three bins properly and you have handled most of it.

WHEN TO SPEND ACTUAL MONEY

- **Business tiers** — phase two. Worth it from your first employee.
- **Cyber Essentials certification** — not AI-specific, but the baseline that makes everything here easier to stand on. If you don't hold it, put it on this quarter's list.
- **A DPIA** — the day AI touches personal data in anything beyond trivial ways, and certainly before client data goes near a tool. The kit's DPIA guide (07) gets you through the standard cases for £0; specialist legal help is for the hard ones it flags.

§ 06 / THE WHOLE THING ON ONE PAGE

Ninety days, twenty ticks

PHASE 1 · LOOK · DAYS 0–7	
☐	Amnesty message sent; declared tools recorded
☐	Expenses + card statements searched for AI subscriptions
☐	Microsoft/Google connected-apps list reviewed
☐	One-page inventory written; personal-data tools flagged
PHASE 2 · DECIDE · DAYS 8–30	
☐	Three-bin data rule agreed (green / amber / red)
☐	Approved tool list picked; business tiers purchased
☐	New-tool approver named; four vendor questions adopted
☐	Policy (03) issued; cheat sheet (02) pinned up; said out loud in a meeting
PHASE 3 · DEFEND · DAYS 31–60	
☐	SSO/MFA on AI accounts; training-on-data off, screenshotted
☐	Work data deleted from personal AI accounts
☐	Third-party app connections restricted in the tenancy
☐	Payment-change callback procedure live — no exceptions
☐	Four questions sent to key suppliers; answers filed
PHASE 4 · DETECT & DRILL · DAYS 61–90	
☐	Monthly AI question added to team meeting
☐	Quarterly console + expenses checks diarised
☐	One-page incident playbook written
☐	Deepfake-payment tabletop run once
ONGOING	
☐	Quarterly hour in the diary
☐	Self-assessment (06) scored; number written down
☐	Policy re-versioned at least twice a year

© Peter Bassill · peterbassill.com · Licensed CC BY 4.0 — use it, adapt it, share it, keep the attribution. This document is general guidance, not legal advice; when AI touches decisions about people, take proper advice.