
SME AI Governance Starter Kit

Quick Start Guide

A simple way to put sensible AI controls in place

This guide shows a small business how to use four practical governance documents without creating an enterprise-style compliance programme.

The aim is to help you answer five questions:

- What AI are we using?
- What information is it using?
- What is it allowed to do?
- What happens if it gets something wrong?
- Who is responsible?

Important: This pack is practical internal governance guidance. It is not legal, regulatory, cybersecurity, data-protection or other specialist professional advice.

1. WHAT THIS PACK IS FOR

AI governance does not need to mean committees, lengthy policies or stopping people from using useful tools.

For a small business, minimum viable governance is about visibility, ownership, proportionate controls and sensible human oversight. The objective is to make useful AI adoption safer and easier - not to create unnecessary administration.

The four documents

#	Document	Purpose
1	AI Acceptable Use Policy	Sets the basic rules employees follow when using AI for company work.
2	AI Systems & Use Case Register	Shows what AI is being considered, trialled or used, who owns it, what data is involved and why the business is using it.
3	AI Risk Assessment	Records how much authority the AI has and what happens if it is wrong, then derives the Green, Amber or Red level and the human oversight that goes with it.
4	AI Incident & Review Log	Records significant problems, near misses and reviews so the business can learn and adjust controls over time.

Who should own it?

Nominate one person at management level to act as the AI Policy Owner. In a small business this may simply be the Managing Director, Operations Director, Finance Director or another manager. They do not need to be an AI specialist; they need enough authority to coordinate the process and make sure important issues are not ignored.

Keep the objective simple

The business should be able to explain what important AI it uses, who owns it, what information it touches, what it is allowed to do, and how significant risks are controlled.

2. HOW THE FOUR DOCUMENTS FIT TOGETHER



1. The Policy sets the rules

Employees need straightforward guidance on approved tools, protecting company and personal information, checking important outputs, customer-facing use, the uses that need formal approval, and reporting problems.

2. The Register creates visibility

Create one row for each material AI use case. A single tool can appear more than once: using an AI assistant to draft marketing copy is a different use case from using the same tool to analyse employee information.

3. The Risk Assessment decides the level of control

Use it for material or uncertain use cases, particularly where customers, employees, sensitive information, money, external communications, business systems or autonomous actions are involved.

4. The Incident & Review Log closes the loop

Record significant incidents, near misses, scheduled reviews and material changes in an AI system. Update the policy, register or risk assessment when the evidence shows that something needs to change.

One tool, different risks

Do not classify an AI product as simply "safe" or "unsafe". Risk depends on what the business is using it to do, what information it accesses, who may be affected and how much authority it has.

The lifecycle



3. IMPLEMENT THE STARTER KIT IN SEVEN STEPS

1	Appoint an AI owner. Choose one manager/director to coordinate the process and act as the point of escalation.
2	Find the AI already in use. Ask staff and managers what AI they use for work. Include public AI assistants, AI built into existing software, trials, automations and tools being considered.
3	Complete the AI Register. Record the purpose, owner, data involved, the AI authority, the consequence if it goes wrong, the resulting risk level, human oversight, expected benefit and review date.
4	Assess material or uncertain uses. Use the AI Risk Assessment for anything that could materially affect people, customers, money, sensitive information, important decisions or business systems.
5	Adopt and communicate the policy. Tailor the Acceptable Use Policy, list approved tools, name the owner and make sure employees know the basic rules.
6	Put controls around Amber and Red uses. Define what must be checked, approved, monitored or escalated before the use goes live.
7	Review and learn. Record significant incidents and review use cases when their purpose, data, access or authority changes. Periodically confirm that they are still creating value.

Do not wait for perfect information

Start by finding the AI that matters most. You can improve the register and controls over time. The goal is useful visibility and proportionate action, not a perfect inventory on day one.

4. TWO QUESTIONS DECIDE HOW MUCH CONTROL A USE CASE NEEDS

The starter kit uses a simple internal classification to keep controls proportionate. It is not a legal or regulatory classification.

Two questions decide the level of control. Ask both. Neither one answers it on its own.

Input 1: how much authority does the AI have?

This is the furthest the AI is allowed to go without a person involved. Draft and Recommend leave someone in the loop before anything happens. Decide and Act do not.

DRAFT	RECOMMEND	DECIDE	ACT
Creates content or analysis for a person to review.	Suggests what a person should do.	Selects or determines an outcome.	Takes an action in a system or the real world.

Input 2: what happens if it goes wrong?

Judge the worst realistic outcome, covering both an incorrect result and information reaching somewhere it should not.

Consequence	What it means
LIMITED	An internal inconvenience, easily corrected, with nobody outside the business affected.
MATERIAL	A customer, client, employee or candidate is affected, or money, reputation or a commercial relationship is at stake.
SERIOUS	Significant financial loss, legal or regulatory exposure, confidential or personal information disclosed, or real detriment to an individual.

Read the two together

What the AI is doing	Limited	Material	Serious
Draft	Green	Amber	Amber
Recommend	Green	Amber	Amber
Decide	Amber	Amber	Red
Act	Amber	Red	Red

Two things follow from the grid, and both matter. A use case that only drafts can still need sign-off, because consequence carries it there. And an automated task with trivial consequences does not need formal approval, because autonomy on its own is not what makes something high risk.

What each band requires

Band	What it requires	Typical examples
GREEN	Use the output and own it. Normal professional care, no extra process.	Brainstorming; first drafts for your own use; summarising non-sensitive information; internal meeting notes.
AMBER	A named person checks and signs off before the output leaves the business or influences a decision. Owner recorded, controls documented, review date set.	Customer replies; proposals; document classification; finance analysis; routine workflow decisions; automated tasks with limited consequences.
RED	Formal approval required. Documented assessment, testing, defined controls, a working override and owner sign-off before it goes live.	Employment screening; significant customer decisions; moving money; changing permissions; committing the company contractually.

A useful rule

The more sensitive the information, the greater the consequence of an error, and the more authority the AI has to decide or act, the stronger the governance and human oversight should normally be.

5. KNOW WHEN TO REVIEW OR ESCALATE

Reassess an AI use when its job materially changes. The technology may be the same while the business risk becomes very different.

Purpose changes: The AI starts being used for a materially different task or decision.

Data changes: New personal, customer, employee, confidential or sensitive information becomes involved.

Authority increases: The AI moves from drafting or recommending to deciding or taking action.

System access increases: The AI gets access to additional systems, records, permissions or tools.

Repeated errors appear: People repeatedly correct the same problem, a complaint occurs, or a significant near miss is identified.

Provider or feature changes: A supplier materially changes the AI feature, settings, data handling or capabilities.

Consider specialist advice where appropriate

- The AI materially affects recruitment, employment or other significant outcomes for individuals.
- Particularly sensitive personal information is involved.
- The use relates to regulated activity, safety, healthcare, legal or other specialist areas.

- The AI can move money, commit the company, change permissions or take other consequential actions.
- A significant data, cybersecurity, contractual or regulatory incident may have occurred.

The starter kit does not replace specialist obligations

If a use case or incident raises legal, regulatory, data-protection, cybersecurity, employment, safety or other specialist issues, use the organisation's relevant formal procedures and obtain appropriate advice.

6. WHAT GOOD LOOKS LIKE - AND WHAT TO DO NEXT

You have a useful minimum viable governance system when the answers below are broadly "yes".

✓	A named manager/director owns AI governance.
✓	The business has reasonable visibility over material AI tools and use cases, including AI embedded in existing software.
✓	Employees have simple rules for approved use, sensitive information and checking important outputs.
✓	Material or uncertain use cases are assessed before they become routine.
✓	Human oversight is stronger where the consequences of an error are greater, and stronger again where the AI can decide or act without a person reading the output first.
✓	Important approvals, controls and review dates are recorded.
✓	Significant incidents and near misses are recorded and acted upon.
✓	AI is reassessed when its purpose, data, access or authority changes.
✓	The business also asks whether each AI use is producing enough value to justify continued use.

The four starter documents

- Document 1 - AI Acceptable Use Policy
- Document 2 - AI Systems & Use Case Register
- Document 3 - AI Risk Assessment
- Document 4 - AI Incident & Review Log

Need help putting this into practice?

A short external review can help identify what AI is already being used, assess the most important use cases, tailor the policy, train staff and identify practical AI opportunities as well as governance gaps.

Contact David Tongeman david@tongeman.co.uk or visit tongeman.co.uk

Free starter version - designed for small UK businesses.