



The CHART Method

A guide to safe adoption of AI for UK SMEs'



Cert No. 2921 ISO 9001



Cert No. 2921 ISO 27001

ramsac
the secure choice

Foreword

Your people are already using AI. The real question is whether anyone is guiding how they use it.

Across UK SMEs, staff are drafting emails in ChatGPT, summarising documents in Copilot and building small automations, usually without anyone having agreed the rules. It tends to happen quietly, well before the leadership team has discussed it, and it brings genuine productivity gains alongside genuine exposure. Most organisations can't see either clearly.



Good governance fixes that. It gives your people clear rules about which tools they can use and how, it protects the data you're responsible for, and it lets you adopt AI faster, because you're no longer second-guessing every new tool. In our experience, firms that see governance as a brake fall behind, while those that treat it as a set of guardrails tend to move more quickly than their more nervous competitors.

This guide sets out how we approach AI governance at ramsac. CHART is the five-step method our AI team uses when we work with clients, and we've kept it practical and in plain English so you can start using it straight away.

Dan May
CEO, ramsac

Who will find this guide useful


- CEOs and MDs who own strategic risk and set direction
- IT and operations leaders deploying or reviewing AI tools
- Finance, HR and compliance leaders handling data, policy and suppliers
- Any UK organisation of roughly 10 to 500 people using or considering AI

You can read this in ten minutes and act on most of it within thirty days. Come back to it every six months or so, because the landscape is changing quickly. And please share it with peers and suppliers: the more organisations that govern AI well,
the safer it is for everyone.

The CHART Method

CHART is the method we use with clients to take AI governance from a standing start to something that runs as part of normal business. It has five steps. You start by checking where you stand today, then work out which rules apply to you, take the next step that suits your position, put AI to work on a few safe, high-value tasks, and keep it all up to date as things change.

The steps form a cycle. Once you reach Track, you'll come back round to Check every six months or so, re-scoring your position as your use of AI grows and the rules move on.





CHECK

Check where you stand.
Two questions, five minutes.

Most maturity models assume that as you use more AI, your governance improves alongside it. That rarely happens. We regularly meet firms with heavy AI use and almost no controls, and others with excellent policies and very little use at all. How much AI you use is driven by ambition, while how well you govern it is driven by caution and obligation, and the two can move quite independently.

That's why Check comes down to two questions:

1. How much do you use AI?
2. How well do you govern it?

Each questions has its own short set of statements below.

Score your AI governance

Answer the questions below honestly, as this is for your benefit rather than anyone else's. Count your yes answers in each column, then plot yourself on the grid that follows.

How much you use AI...

This measures how widely and deeply AI is used across your organisation.

1. Staff use AI tools for real work, not just to experiment
2. More than one team relies on AI
3. AI is built into a live process or service
4. You automate work with AI, with limited human checking
5. You build or customise AI in-house

Score:

0 to 1 yes = Low use (You're not using AI widely yet. Focus on safe, practical first steps.)

2 to 3 = Growing (Momentum is building. Strengthen your governance as you scale.)

4 to 5 = Heavy use (AI is embedded and delivering value. Keep your governance strong.)

How well you govern it...

This measures the strength and maturity of your AI governance.

1. You have a written AI acceptable use policy that staff know about
2. You know what data your AI tools process, and have run a data protection impact assessment (DPIA) where needed
3. Your security and incident plans cover AI
4. You vet AI suppliers and check for a data processing agreement
5. Staff are trained and can report an AI concern
6. You review AI use and outputs, including for bias, on a regular cycle

Score:


0 to 2 yes = Low control (AI use isn't governed yet. Start with a policy, a tools register and a named owner)

3 to 4 = Partial (You have some controls, but there are gaps. Close them before your use grows.)

5 to 6 = Strong control (You govern AI well. Keep maintaining and improving your controls.)

Where you sit

Plot your two scores on the grid. If either score is in the middle band, Partial or Growing, you sit on one of the centre lines, so read both positions either side of it.



Read your two scores together. Heavy use with low control puts you in **Moving without a map**, while light use with strong control puts you in **Cautious to a fault**. Low on both means you're **Not started**, and the foundations in "Start with three safe wins" are the place to begin. High on both puts you in **Confident and in control**. Whatever your result, the gap between your two scores is usually where the real work lies.

Two positions deserve particular attention. **Moving without a map** is both the most common and the most costly: staff, and sometimes developers, are well ahead of your policies, so you're carrying exposure every day that you can't see. **Cautious to a fault** is quieter, but the cost is just as real. You govern tightly but use so little AI that competitors start to pull away. Wherever you start, the aim is to reach **Confident and in control**, where governance and adoption support each other.



H HANDLE

Handle what the rules ask.
Know what UK law expects of you now.

The UK doesn't have a dedicated AI Act, and government has signalled it isn't planning one soon. That doesn't mean AI is unregulated. Existing law already applies, and for most SMEs what you owe comes down to two things: data protection law and your own sector's rules. Everything else applies only if something specific about your business triggers it. Here's what matters, in proportion.

UK GDPR, the Data Protection Act 2018 and the Data (Use and Access) Act 2025

APPLIES TODAY

Any AI tool that processes personal data is covered. The Data (Use and Access) Act updated the rules on automated decision-making in February 2026. Decisions with legal or similarly significant effects can now be automated in more situations, but you must tell people, let them make representations, offer human intervention and let them challenge the outcome. You'll still need a DPIA for high-risk processing. This is a short, documented assessment of the privacy risks of a new tool or process and how you'll reduce them; the ICO provides a free template. You'll also need to check any AI supplier that handles personal data on your behalf, and have a data processing agreement in place with them.

In practice: keep a register of your AI tools, run a DPIA before anything touches personal data, and check that every supplier has signed a data processing agreement.

ICO guidance and the coming AI code

READ NOW

The ICO (Information Commissioner's Office) is now required by law to produce a statutory code of practice on AI and automated decision-making. Its updated guidance on automated decision-making is due to be finalised this winter, and it's gathering evidence on agentic AI to shape further guidance.

In practice: if you use AI in any process involving personal data, make the ICO's AI guidance required reading for whoever owns governance, and watch for the finalised code of practice.

Supply-chain and procurement questions

INCREASINGLY COMMON

Large customers and public bodies are increasingly asking SME suppliers for evidence of AI governance before they'll sign. Having a policy, a tools register and a record of your DPIAs ready to share can make the difference in a tender.

EU AI Act

ONLY IF IT APPLIES TO YOU

The EU AI Act can apply to a UK business that sells AI-enabled products or services into the EU, or whose AI produces outputs that are used there, for example in decisions about EU-based customers or staff. If neither describes you, it's unlikely to apply. Its stricter rules for high-risk uses such as recruitment and credit decisions have been pushed back to December 2027, though some transparency duties around AI-generated content already apply.

In practice: spend five minutes checking where your customers, staff and AI outputs sit. If none of them touches the EU, note that and move on. If they do, take advice.

Your sector regulator

TAKES PRECEDENCE

Financial services firms should follow FCA expectations, legal services SRA and BSB guidance, and healthcare organisations MHRA and NHS frameworks. Where your sector guidance and general advice differ, the sector guidance wins, so let it lead. We can help you work out what it means in practice.

Keep it in proportion. If you have a policy, a tools register and a DPIA for anything involving personal data, and you follow your sector's guidance, you've covered what applies to almost every UK SME. Treat everything else as a quick check rather than a cause for alarm.

What's coming

The picture is still moving. None of the following needs action today, but they're worth watching.

- **The Regulating for Growth Bill.** Announced in the May 2026 King's Speech in place of a dedicated AI bill, it would let government relax specific rules temporarily so new technology can be tested in controlled "sandboxes". Expect the UK to keep regulating AI through existing sector regulators rather than a single AI law.
- **The ICO's statutory AI code.** This will set out what good practice looks like for organisations using AI with personal data. Once published, it will be the benchmark you're measured against.
- **Product liability reform.** The Law Commission is consulting until January 2027 on treating software and AI systems as "products", which would make it easier to claim compensation for harm caused by defective AI. This matters most to organisations that build or customise AI systems, not just those that buy them.

Now turn to Act. Find the table for the position you scored in Check. Each one shows what that position typically looks like, the risk that matters most, your next best move and what you stand to gain.



A ACT

Act on your next move.
One clear step for your position.

Having decided where you're at, it's important to take action.

Not started Low use - Low control



What it looks like

Little deliberate AI use and no governance. Leadership has not formally discussed it. You may believe AI does not affect you yet.



The risk that matters

There are really two risks here. Staff are probably already using consumer AI tools without telling anyone, and competitors are learning faster than you. Standing still feels safe, but it has a cost.



Your next best move

Brief your leadership team on AI and this framework, run the five-minute check across departments, and pick one safe use case to pilot. Governance and a first win can start together.



What you gain

You enter the market on your own terms, with control in place before scale, rather than catching up after an incident.

Where ramsac helps: a short readiness session helps your leadership pick a safe, high-value first pilot.

Cautious to a fault Low use - High control



What it looks like

You take governance seriously and use AI sparingly, sometimes because leadership restricted or banned it. Policies are strong. Adoption is thin.



The risk that matters

The costs here are quieter. Competitors who adopt AI safely will pull ahead on productivity and service. And if you've banned AI outright, staff are probably using it anyway, out of sight, which recreates the shadow AI you were trying to prevent.



Your next best move

Use your governance strength to safely widen access. Approve two or three low-risk use cases, publish clear rules for them, and turn a ban into a managed yes.



What you gain

Your caution becomes an advantage, because you can adopt AI with a confidence others lack, backed by controls you already trust.

Where ramsac helps: We approve and govern new use cases with you, without widening risk.

Moving without a map High use - Low control



What it looks like

Staff, and sometimes developers, use AI across the business faster than your policies can keep up. Tools appear without review. Nobody owns the risk.



The risk that matters

Personal or client data flows into tools you have not checked, often with no data processing agreement. A single leak, or a wrong AI output you relied on, can become a GDPR or reputational incident with no recourse.



Your next best move

Find your shadow AI first. List every tool in use, issue an interim acceptable use notice, and run a DPIA on anything touching personal data. Name one owner this week.



What you gain

You keep your speed and remove the exposure that comes with it, turning an invisible liability into a managed advantage.

Where ramsac helps: An AI Readiness Assessment helps find shadow AI and drafts your policy.

Confident and in control High use - High control



What it looks like

AI runs across functions, governance is owned at board level, and controls are monitored, reviewed and enforced rather than merely written. You can deploy AI that acts on your behalf with proper oversight.



The risk that matters

Complacency as maturity settles, talent concentrated in a few key people, and agentic systems whose failure modes outpace yesterday's controls.



Your next best move

Formalise a review cycle, extend oversight to agentic and automated workflows, and make your governance a visible trust signal to customers and partners.



What you gain

Your governance starts to set you apart. It helps you win work that requires proof of control, and lets you scale AI faster because you're scaling it safely.

Where ramsac helps: Ongoing support keeps your controls current as you adopt agentic AI.

There's no fixed route

You don't have to work through the positions in order. An organisation that's **Moving without a map** can often reach **Confident and in control** with one focused piece of work, because the AI use is already there and only the governance needs to catch up. Wherever you sit, start with the next move for your position.

Situations that change your priorities

Sometimes a specific situation matters more than where you sit on the grid. If one of these applies to you, start here, whatever your scores. We see all four regularly.

You're regulated, or handle sensitive data

E.g. legal, wealth management, financial services

By regulated, we mean firms overseen by a sector regulator with its own rules, such as the FCA, SRA or CQC. By sensitive data, we mean information about people's health, finances, ethnicity, religion, sexuality or criminal records, data about children, and confidential client information. If either describes you, putting that data through an unchecked AI tool is a live risk today. Prioritise a DPIA, data processing agreements with your suppliers, and a clear rule that a person reviews any AI-assisted decision.

A customer has sent you an AI questionnaire

You need evidence, quickly

Large customers and public bodies increasingly send suppliers a due diligence questionnaire before they'll sign, asking which AI tools you use, what data they touch and how you control them. You don't need everything to be perfect. What you need is a policy, a tools register and a DPIA record you can show them. It's often the most urgent reason to act, and we can get you ready quickly.

You tried AI and it went wrong

How to try again safely

Perhaps an AI tool gave a confident but wrong answer, there was a scare over data, or a pilot didn't deliver. The cause is usually missing guardrails rather than the technology itself. Start again with one low-risk task, have a person check the output before it's used, and let a small, well-governed success rebuild confidence.

You've banned AI outright

Turning a ban into clear rules

Bans rarely stop people using AI. They just push it out of sight, where you can't see what data is going where. Good governance gives you a safe way to lift the ban: approve a short list of tools and uses, publish the rules, and give staff an approved route so they've no reason to work around you.



R RUN

Run safe, high-value wins.
Put AI to work without the risk.

Once the basics are in place, it's worth putting AI to work on something useful straight away. These three uses are a good place to begin. For most SMEs they're low-risk, they fit within a simple acceptable use policy, and they start saving time quickly.

WIN 01

Draft and summarise

Use an approved business AI assistant, covered by your data processing agreement, to summarise reports, recap meeting notes and draft emails. Keep sensitive data out, and have someone check anything before it goes to a client.

Payback: a few hours back each week for anyone who writes or reads a lot.

Effort: low. You can start this week.

WIN 02

Search your knowledge

Let staff ask questions of your own policies and procedures, and get answers in seconds. Check folder permissions first, as AI will surface anything a person technically has access to, including files they aren't meant to see.

Payback: save time hunting for information, faster onboarding of new starters.

Effort: medium, Start with a small pilot.

WIN 03

Automate a routine task

Pick one repetitive job, such as logging invoices, sorting a shared inbox or updating records after meetings, and automate it. Have a person sign off before anything is sent, paid or deleted. Start small, measure the time saved, then expand.

Payback: time saved every time that task comes round.

Effort: medium, one workflow at a time.

Your first thirty days

Wherever you sit on the map, these foundations will move you forward, and most of them take a morning rather than a project team. Tick off what you have in place and prioritise what you don't.

Own it and see it:

- ✓ Name one person accountable for AI governance
- ✓ List every AI tool in use or piloted, and the data each touches
- ✓ Brief leadership on this framework and your position

Protect and enable it:

- ✓ Publish a plain acceptable use policy and tell staff
- ✓ Run a DPIA on anything processing personal data
- ✓ Check each supplier for a data processing agreement
- ✓ Give staff one clear route to raise an AI concern

Govern AI that acts for you

Most AI tools answer questions. Agentic AI goes further and takes actions on your behalf, such as sending emails, moving data, running workflows and working with your other systems. It's now realistic for UK SMEs to use, but because an agent acts rather than just advises, its mistakes have real consequences. That means it needs stronger controls than a chat assistant.

Regulators are paying attention too. In October 2026 the ICO opened a call for evidence on agentic AI, and it has been clear that an agent's independence doesn't reduce your responsibility for how it handles personal data.

The risk's



How you control it



Agents take real actions. Their mistakes reach real people and systems.



A person approves anything that can't be undone. Critical tasks wait for sign-off, so a mistake is caught before it happens.



Mistakes can spread. One action can trigger another before anyone notices.



Every workflow can be paused and its actions reversed. Every action is recorded, with the reason.



Agents can reach a lot. They often have access to email, files and other systems at once.



Each agent gets access only to what its task needs. An agent that sorts invoices has no reason to read HR files.



Agents can be tricked. Instructions hidden in an email or document can hijack them (known as prompt injection).



Test with dummy data before going live, and keep its access tight, so a trick can't do much damage.

If you're already experimenting with agents, give each one a named person who's responsible for it, write down what it can access, and agree who signs off its higher-value actions. When ramsac builds agents for clients, we put all of these controls in place before anything goes live.



TRACK

Track, review and improve.
Keep it current as the ground
shifts, on an ongoing basis.

Governance isn't a one-off project, because the ground keeps shifting. New tools appear, staff come and go, suppliers change their terms, regulators publish new guidance, and as you grow you take on new kinds of risk. The Track step keeps the first four steps accurate six months from now, as well as on the day you wrote them.

What to track

- ✓ **Your AI tools register** - add anything new the moment it is adopted or piloted, and remove what has been retired.
- ✓ **Policy awareness** - do new starters know the acceptable use policy, and do existing staff still follow it?
- ✓ **Supplier changes** - has a supplier updated its terms, been acquired, or had a security incident?
- ✓ **Incidents and near-misses** - log every AI-related concern raised, even minor ones, and check the pattern over time.
- ✓ **Your position on the map** - are you drifting toward Moving without a map as use grows faster than control, or settling into Cautious to a fault as controls outpace adoption?
- ✓ **Regulatory movement** - the ICO's AI code and agentic AI guidance, updates from your sector regulator, the Regulating for Growth Bill, and the outcome of the Law Commission's product liability review.

When to review

Review	How often	What it covers
Light-touch check	Quarterly	Tools register up to date, incident log reviewed, any new starters briefed on the policy
Full governance review	Annually, or every six months if you are Moving without a map	Re-score both axes, re-read your position on the map, refresh the policy, re-check supplier agreements
Triggered review	Whenever it happens	A new AI tool or agent goes live, a customer sends a governance questionnaire, an incident occurs, you enter a new sector or market, or a regulator changes the rules

Working with ramsac

ramsac is an award-winning managed IT and cybersecurity partner based in Surrey. For 35 years we've helped organisations stay secure and productive, looking after their infrastructure, cloud, Microsoft 365, security and 24/7 support as one accountable partner. We also have a dedicated AI engineering team that builds and governs AI solutions in-house, so the advice in this guide comes from work we do every day. We're certified to ISO 9001, ISO 27001 and Cyber Essentials Plus.

We believe every SME and not-for-profit should have access to the same quality of AI governance and capability that large enterprises take for granted. You're welcome to use this framework on your own, and when you'd like support, here's how we can help.



How we can help you



AI Readiness Assessment

A clear read of where you stand on both axes, and the five moves that matter most for you.



Governance Framework Setup

Hands-on help to put policies, training, supplier review and controls in place.



Agentic AI Build

Design and build of AI that acts for you, with governance and rollback built in.



AI and Data Health Check

An independent review of your controls, benchmarked against UK good practice.



AI Policy Starter Kit

An acceptable use policy, a tools register and a GDPR checklist.



Ongoing Governance Support

Retained help with reviews, incidents, training and new regulation as it lands.

Start with a free Readiness Assessment

A short conversation will show you where you stand and what to do next, with no obligation.

01483 412040 · info@ramsac.com · www.ramsac.com

ramsac
the secure choice