

TRUE9

Who's controlling your AI?

A practical guide on how to take back control of AI in your business,
on your terms, at your pace.

Created: October 2026

Control 
Clarity gained

 Thrive through
technology



AI is already in your business. The question is, who's steering it?

Somewhere in your business today, someone is pasting a client list into a chatbot to tidy up the formatting. Someone else is asking a free AI tool to summarise a supplier contract. Nobody's being reckless, they're just busy. And AI is a brilliant tool to get through the daily to-do list faster.

Trouble is, when you multiply that by every team, in every function, every week, suddenly you have a picture most leaders are blind to. Meanwhile, the investment keeps flowing. This year, large companies doubled their AI spend from 0.8% to 1.7% of revenue¹.

The question is no longer whether you should use AI, or even what the 'best' use cases are. It's whether your organisation's use of AI matches how well you understand it, because AI isn't an IT project, it's a strategic business decision.

Where does your business sit?

Every business sits somewhere on a spectrum of AI risk:

In control

AI use is visible, governed and intentional. You know which tools are in play. Your policies get followed. Someone owns the outcome.

Mostly aware

You've got a reasonable handle on things, but there are gaps. A policy that isn't quite followed. An owner who wears too many hats. Tools nobody's mapped.

Blind spots emerging

Some parts of the business are AI-aware, while others are exposed. You need to get control of confidential data, ownership, clarity and confidence.

High exposure

Sensitive information and operational decisions are flowing through tools with little visibility or control.

Critical risk

Something is likely to go wrong before anyone notices, because you have sensitive data, no real ownership and little clarity on what regulation applies.

Not sure where you sit?

That's normal, and it's exactly the point. AI is evolving at such a rapid pace, it's hard to retain control over how it's used and governed within your organisation.

Use the [True9 AI Reality Check](#) to find out where your business currently stands.



You can't steer what you can't see

We've spent over a decade designing, building and improving the technology platforms that businesses rely on every day. When it comes to AI, we see the same challenges getting in the way of progress:

People are adopting AI faster than organisations can govern it

AI rarely arrives in a business through a board decision. It arrives through a browser tab, often via unsanctioned tools or personal accounts.

55%

of employees admit to using AI tools that haven't been approved.²

Robust controls around AI infrastructure are essential, but only 17% of organisations have governance built into their AI by design³. What's more alarming is that two-fifths of all AI interactions involve sensitive data⁴. And despite knowing the risks, 1 in 10 employees still put sensitive information into AI platforms⁵.

Perhaps unsurprisingly, nearly 9 in 10 businesses (86%) have experienced at least one AI-related incident in the past year⁶. And let's not forget that breaches driven by AI cost around \$1 million more than attacks without AI⁷.

Money is spent on the wrong tools

In the rush to keep up, AI budgets often go on tools and use cases that are chosen in isolation, because leaders lack the visibility to know what's really needed in the business. A subscription here. A pilot there. And a usage-based bill that creeps up month after month... after month.

It's uncontrolled growth that hits performance, too. When compared to businesses that have control over their AI estate, those with high AI tool sprawl experience:⁸

17%

lower productivity

22%

lower revenue

19%

lower operating profit

So why are some businesses able to achieve great ROI from their AI investments, while others struggle to realise its value?

The answer: ownership. When leaders take clear accountability for AI outcomes, 57% say AI delivers meaningful business value. Without that accountability, it drops to 21%.⁹

Shadow AI is already here

Different teams are trying different AI tools. Some approved. Some not. And some you have no knowledge of. Without complete visibility of your AI estate, it's impossible to know exactly what tools are being used, where your data is going, and what it means for compliance.

When it comes to AI, the danger is rarely in the tool itself. It's not knowing the tool is being used at all.

You're not alone; over two-thirds (68%) of chief information security officers (CISOs) see unsanctioned AI use in their business.¹⁰ Even at the largest US technology firms, which benefit from dedicated security teams and bigger budgets, nearly half (45%) have experienced a leak of sensitive data through unauthorised AI tools.¹¹

The real cost: decisions made in the dark

Put it all together and you're left facing an impossible choice: Should you invest more? Restrict access? Provide more training? Or just hope for the best?

Each could be the right answer. But you can't make a good choice when you can't see the full picture. And for every month you don't have that clear picture, the gap between what's going on and what you know gets a little wider.

So, what do you do?

^{2, 5} <https://securityjournaluk.com/55-british-employees-using-unapproved-ai/>

^{3, 6} <https://www.onetrust.com/resources/onetrust-2026-ai-ready-governance-report/>

⁴ <https://www.cyberhaven.com/blog/sensitive-data-flowing-into-ai-tools>

⁷ <https://www.helpnetsecurity.com/2026/07/30/ibm-cost-of-a-data-breach-2026/>

⁸ <https://www.ibm.com/downloads/documents/us-en/17f55bdd44c4ea17>

^{9, 12} <https://assets.kpmg.com/content/dam/kpmgsites/uk/pdf/2026/06/global-ai-pulse-q2-2026.pdf>

¹⁰ <https://www.okta.com/newsroom/articles/global-ciso-insights-2026/>

¹¹ https://www.ey.com/en_us/newsroom/2026/03/ey-survey-autonomous-ai-adoption-surges-at-tech-companies-as-oversight-falls-behind

Take control of AI in your business

The good news is every one of those problems has the same root cause, a lack of visibility, which means every problem can be fixed.

Taking control doesn't mean slamming on the brakes. You don't need to ban the tools your people now rely on or sink budget into a six-month transformation programme. Taking control simply means **finding out what's really going on, deciding what you want to happen, and putting someone in charge of making sure it does.**

Think back to when the finance director joined your business. The first thing they asked for wasn't a strategy. It was the numbers: Where's the money going? What are we committed to? Where are the surprises hiding? AI deserves the same treatment, because before you can decide whether to invest, restrict, train, or leave well alone, you need a clear, honest baseline you can trust.

With the baseline in place, the boardroom conversation changes from, **"Should we be doing something with AI?"**, to, **"Which of these opportunities do we go after first?"**

What control looks like in practice

You can see it

You know which AI tools are in use across the business (approved or not), who's using them, and where your data goes once it's typed into a prompt.

You know where the value is

You have a handful of realistic, specific use cases for how AI can be used to save your team time, win more work, or operate the business more efficiently.

You know what to worry about, and what not to

You have a prioritised shortlist of risks to focus on, because every risk, whether data, compliance, operational, or reputational, is sorted by urgency.

You know what happens next

A short policy your people will actually read, with a named owner and a clear plan for adopting AI safely, delivered at a pace that suits your business.

How to make AI work for you, not against you

Plenty of trends come and go. And while a few lead to permanent changes, most pass after their moment in the spotlight. What they all have in common is that technology alone rarely decides whether a business succeeds. It's the decisions around the technology that determine the outcome.

In every client engagement we apply the following three principles as guardrails to guide decisions:

Clarity first

Because adopting a tool before understanding the problem is just guessing.

Delivery-led

Because advice is only useful if you can act on it.

Human

Because the most useful insight usually comes from the people doing the work.

Every business is different

And every business will be at a different stage in the adoption journey. For example, a business that's never even thought about AI faces different challenges and opportunities to those who have already automated half the back office.

It's why we think about AI in three levels of maturity, with each stage representing how well a business understands and governs its use of AI, rather than being dictated by its size, sector, or budget.

Level 1: Aware	Level 2: Adopting	Level 3: Established
You know AI matters and you suspect your people are using it, but nothing official is in place. There's no policy, no owner and no clear picture.	AI is part of how work gets done. There may be guidance of some sort, but it isn't consistently followed. You know it's time to tighten things up.	You're well versed in AI and automation, with policies and ownership already in place. Your challenge now is keeping it that way as tools, people and regulation change.

The following three sections explore each maturity level in more depth, and share insights from our AI specialists, along with practical steps about how you can gain greater control of your AI estate.

Control

Level 1: Aware

You know AI matters and you suspect your people are using it, but nothing official is in place.

At the awareness stage, you've read the headlines, you've probably tried ChatGPT, and you likely have a nagging feeling that AI is moving faster than your business. But there's no AI policy, nobody's officially in charge, and if you're honest, you don't really know what tools your team is using or what they're putting into them. That's not a failing. It's where many businesses are right now, because decisions about AI are being made in your business every day. Just not by you.

It's time to get your house in order, which means you need to find out what AI tools are in use, what they cost, and where the risks sit. Your team also need to know what's expected of them when it comes to using AI. And you need a short, realistic list of use cases where AI could genuinely help.

The following actions can help you get there:

Find out what AI tools your business actually uses

Start with one simple question for every team: “Which AI tools have you used in the past month, and what for?” You want to identify both known and unknown tools, so make it clear there’s no wrong answer and nobody’s in trouble. You need honesty to achieve a full picture of what’s in your AI estate.

Then push the teams to think deeper. ChatGPT, Gemini, Claude are the obvious big names. But AI is now built into tools you already pay for, from email and meeting software to your CRM and design apps. Plus, browser extensions, transcription tools and free trials all count, too. Remember, you need a list of everything.

Full disclosure: there will be surprises. The gap between what’s been officially rolled out and what’s really in use is usually bigger than leaders expect, and that’s ok. You can’t govern what you can’t see, and after completing this exercise, you will see everything.

Work out what AI is costing you

AI spend has a habit of hiding. A £20 monthly subscription on a company card. A usage-based bill that doubles because one team discovered a new feature. A free trial that quietly became a paid plan.

You need to pull all this information into one place, so you can see every tool, who uses it, what it costs each month, and who’s paying. Check expense claims too, because that’s often where personal subscriptions turn up.

During this process, it’s common to find a tool you’re paying for twice, and another nobody has opened in months. But with it all written down in the same place, you can begin to identify overlaps, redundant tools, and opportunities to consolidate for better pricing, allowing you to get back in control of your AI spend.

Be honest about the return

AI ROI comes in two forms:

- **The risk you avoid.**
- **The value you unlock.**

Much of the AI conversation centres on exciting benefits. But at Level 1, it’s more important (and usually more urgent) to focus on risk, because every client list pasted into an unapproved tool is a data protection question you can’t answer. Simply closing that gap is a return in itself, even if it never shows up on a spreadsheet.

Also, achieving value from AI tools is often closer than you think, because tools like Microsoft 365 and Google Workspace already have AI built in, it’s just nobody has switched it on. These are quick, low-risk wins you can unlock immediately, without investing more budget or time creating the business case.

Spot your blind spots

Once you’ve mapped out what AI tools are in use, look for three warning signs:

Tools nobody can see: AI in use that nobody in leadership knew about.

No named owner: you ask, “Who’s responsible for AI here?” and the room goes quiet.

Uncertainty about the rules: nobody’s sure how compliance limits what you can do.

These signs tend to travel together, so if you spot one, the other two usually aren’t far behind. Whichever feels weakest, start there, because it’s probably the visible edge of a wider gap.

Fix what matters most, first

You don't need to fix everything this week. Instead, segment risks into three buckets:

Now

Any risks that could cause real, immediate harm to your business today. For example, client data going into a free AI tool with no data protection agreement.

Next

Risks that will be mitigated once the right foundations are in place. For example, a short policy, some basic training, or a simple way for people to ask, "Can I use this tool?"

Later

Risks that require formal governance, and if your clients or sector expect it, certification. For example, documented frameworks or ISO certification for supplier questionnaires.

Working in this structured way makes sure your biggest vulnerabilities are closed in days, with ample time to get the rest right.

Find where AI can make the biggest difference

Once you have control of the risks, you can turn your attention to the opportunity AI presents. Start by asking your team how they spend their time. For example, which tasks are repetitive, manual or frustrating? Where are people copying data from one system into another? Are there instances where the same email is being drafted for the tenth time this week?

Each answer is a small, realistic use case for where AI can deliver efficiencies in your business, and free people to work on the tasks where they add most value.

Start by selecting two or three options. Test them properly. Measure what changes. One small, proven win will do more for confidence (yours and your team's) than a long list of ideas that never get started.


Talk to your team, so they don't fear AI

Change is always a tricky subject, because while some employees are excited about what's coming, others are quietly worried about what AI means for their job.

Tackle the issue head on with clear, concise communications. Tell your team:

- **Why you're using AI.**
- **What it will be used for.**
- **How it will never be used to replace team members.**
- **The important role they play in using AI safely.**

When people know where they stand, they're far more likely to accept AI and use it properly. And remember, a well-run business before AI will still be a well-run business with AI, as long as someone is steering it. Good leadership doesn't change. You simply have a new tool to work with.

A man with short brown hair and a light beard is smiling and gesturing with his right hand, palm facing forward. He is wearing a blue denim button-down shirt with a small polo player logo on the left chest. The background consists of vertical wooden slats and a blue wall. The text is overlaid on the lower left of the image.

AI isn't an IT project,
it's a strategic
business decision.

Control

Level 2: Adopting

AI is part of how work gets done but any guidance isn't consistently followed.

At the adoption stage, AI has made its way into how your business works. Your marketing team drafts content with it. Your sales team summarises calls with it. Someone in finance has built a clever automation that nobody else quite understands.

The productivity gains are real, and nobody wants to give them up. But the guardrails haven't kept pace. There might be a policy somewhere, but it isn't enforced. People use a mix of personal and company accounts. And if a client asked you tomorrow exactly where their data goes when your team uses AI, you'd struggle to give a confident answer.

It's time to embed AI properly into your business, where it's used deliberately, in approved tools, under clear rules people actually follow, with one person accountable for making it work.

The following actions can help you get there:

Pause your riskiest AI use

Before you think about building anything new, stop and deal with the thing(s) keeping you up at night. It could be client data going into a free chatbot or automated records management nobody's checking. If you feel unsure, pause it now and work out how to fix the issue.

This isn't a crackdown to stop people using AI. It's a temporary, targeted step to close your biggest exposure(s) immediately, buying you time to fix things properly before you encounter catastrophic consequences.

Again, talk to your team. Explain why you've placed a temporary pause on the activity, what help you may need from them, and when you expect to come back with a timeline for getting back online.

Move everyone to company-controlled AI accounts

Personal AI accounts are one of the biggest blind spots, because when someone uses their own login, you can't see what they share or control which models they use. And when they leave, their chat history (and your data) leaves with them.

Company-controlled accounts instantly fix that. You can set spending limits, choose which models are allowed, switch off features you're not comfortable with, and see where your data goes. This is essential if you operate in highly regulated industries where data must stay in the UK.

Then make it a simple rule: anything involving your business, your clients or your projects happens in an approved company account. No excuses.

Keep sensitive data out of prompts

Firstly, before you approve any AI tool for sensitive work, check that it doesn't train the underlying model on your data.

Secondly, the easiest way to prevent a data leak is to make sure it never gets typed. Agree a simple, clear, specific rule for what can and cannot go into an AI tool. For example, "Never put client names, contract values or personal details into an AI tool that isn't on the approved list."

Finally, make data redaction a habit. For example, swapping "Company ABC £400k contract renewal" for "a large client's contract renewal". It's a small change, but it protects data privacy, keeps you aligned with UK GDPR, and honours the confidentiality your clients expect.

Decide where AI does and doesn't belong

Just because you can use AI for something, doesn't mean you should. Does a two-line reply to a colleague really need a chatbot? Should a difficult conversation with a client be drafted by a tool that doesn't know the history?

AI is one tool in a well-stocked toolbox. You should be the one to choose the tool, guide the work and take responsibility for the result.

As a business, agree where AI adds real value, like first drafts, meeting summaries, research, or tidying data. And agree where it doesn't belong, such as anything that depends on judgement, relationships, ethics, or hard-won professional experience.

Mind the human cost

There's a side of AI adoption that rarely makes the headlines: tool fatigue. Some people may start to feel worn down by the constant stream of new tools, while others slowly stop thinking for themselves. It's known as 'AI brain rot' and it's a real threat to the skills your business depends on.

To stop it affecting your business, keep an eye on how your team feels, not just the outcomes they produce. Encourage people to form their own view first before using AI to test it.

Write a policy people actually follow

Many policies are written and then filed away, never to see the light of day again. It happens because they're simply not fit-for-purpose. If you find your AI policy being overlooked or ignored, you need to find out why. Is it too long? Full of legal language? Did anyone ever tell the team it exists?


The best policies are short, sharp, and to the point. If you can condense it down into a single page, great. Write it in plain English, talk it through with your team in person, and revisit it often. If you need help writing one, simply drop us an email: hello@true9.co.uk

Put someone in charge

If you only perform one action from this guide, make it this one.

When leaders take clear accountability for AI outcomes, 57% say AI delivers meaningful business value. Without that accountability, it's just 21%.¹²

Your AI owner doesn't need to be technical. They simply need the authority to make decisions, the time to follow through, and the backing of the leadership team. Without that, adoption drifts, which is where cost and risk creep in. And if the right structure will take a while to put in place, appoint an interim owner this week to take control, while you identify the best fit long-term.

A man with a beard and light brown hair, wearing a light blue button-down shirt, is sitting at a wooden desk. He is smiling and looking down at a spiral-bound notebook, holding a pink pen over it. His left hand is resting on the desk. In the foreground, there is a black insulated tumbler cup. A computer keyboard is partially visible on the desk to the right. In the background, another person is blurred, and a grey office partition is visible. The scene is brightly lit, suggesting a modern office environment.

—
You don't need another
supplier who hands over a
report and disappears.

Control ⓘ

Level 3: Established

Your challenge is how to retain control of AI as tools, people and regulation change.

Early adopters got ahead early. AI is woven into how your business works, your policies are in place, and someone owns it. You're ahead of most businesses, and that's worth protecting. The risk now is complacency. New tools launch every week. Regulation is still taking shape. People join, suppliers change and, little by little, the picture you were confident about a year ago starts to drift.

It's time to change your mindset and treat AI as a sustained, long-term effort, with governance that keeps pace with change and the evidence to prove it.

The following actions can help you get there:

Treat AI as a habit, not a project

AI adoption isn't something you finish and file. Build regular check-ins into the calendar (quarterly works well for many businesses) and ask three questions each time:

- What's working?
- What isn't?
- What's changed?

Keep a short written record of your answers and your wider governance approach. It's useful for board reporting, tenders, and insurance renewals, and it gives you something concrete to point to when a client asks how you manage AI. Following every review, be sure to communicate updates with your team, so their knowledge remains up-to-date on how your business approaches AI.

Review your policy every 6 to 12 months

The AI policy you wrote last year was built for last year's tools. Since then, the tools you approved have added new features, new tools have arrived, and regulation has moved on. Set a fixed review date and treat it like any other compliance deadline, so your policy keeps pace with what's really happening, not just what you knew when you wrote it.

Watch the edges

Well-run businesses rarely lose control of AI in one dramatic moment. They lose it at the edges. A new hire brings their favourite AI tool from their last job. A supplier quietly adds AI features to software you already use. A team trials something 'just for a week' and it's still running six months later.

Build AI into your onboarding, your supplier checks, and your tool approval process, so the edges stay in view.

Keep your authentic voice

When everyone has the same AI writing tools, everything starts to sound the same. The same structure, the same phrases, the same polished but empty tone. Your clients will notice, even if they can't say why.

Whatever AI helps you produce, make sure it still sounds like you. At True9, we stand behind every idea and every sentence in the documents, reports and code we deliver, whether or not AI helped shape them. But when AI is involved, that matters even more.

Verify, don't just forward

AI is brilliantly confident. But sometimes, it's confidently wrong. It can hallucinate, invent a statistic, misread a contract or write code that works in testing and fails in the real world. So make it a rule: nothing AI-generated leaves the business without a person reviewing it first.

Assume it might be incomplete, inaccurate, generic, or overconfident. Check it, challenge it, correct it. If you wouldn't put your name to it, don't send it.

Get an outside view

The better your governance, the harder it becomes to spot your own blind spots. It's 'the curse of knowledge'. And with 78% of executives lacking strong confidence that they could pass an independent AI governance audit,¹³ being able to prove your position is a genuine advantage with clients, insurers and investors.

An independent assessment gives you that proof and catches anything worth tightening before it becomes a gap. Its cost is small compared to the cost of the incident it could help prevent.

¹³ <https://www.grantthornton.com/services/advisory-services/artificial-intelligence/2026-ai-impact-survey>

Your AI control checklist

Wherever you are on your journey, use this checklist to see how far you've come and what's next.

Level 1: Aware - **Get your house in order**

- ☐ We know which AI tools every team uses, approved or not.
- ☐ We know what AI is costing us, including expensed subscriptions.
- ☐ We've checked for our three warning signs: hidden tools, no owner, unclear rules.
- ☐ We've sorted what to fix into now, next and later.
- ☐ We've picked two or three realistic use cases to test.
- ☐ Our team knows why we're using AI and what it means for them.

Level 2: Adopting - **Embed AI properly**

- ☐ Our riskiest AI use is paused or fixed.
- ☐ All work happens in company-controlled AI accounts.
- ☐ Everyone knows exactly what data never goes into an AI tool.
- ☐ We've agreed where AI does and doesn't belong.
- ☐ We're keeping an eye on how our team is coping, not just what it produces.
- ☐ We have a one-page policy that people actually follow.
- ☐ One named person owns AI, with the time and authority to do it properly.

Level 3: Established - **Keep it sustainable**

- ☐ We review what's working with AI at least quarterly, and write it down.
- ☐ Our policy has a fixed review date every 6 to 12 months.
- ☐ New hires, suppliers and tools are checked for AI use.
- ☐ Everything we send sounds like us and has been checked by a person.
- ☐ We've had an independent view of where we stand.

/

We're in this together

You don't need another supplier who hands over a report and disappears. You need someone who'll help you understand where you stand, then roll up their sleeves and stick around while you put it right.

At True9, some of our client partnerships now span eight years or more, because we care about outcomes, not outputs.

Writing good code is a given

Our real value is in what happens around it. The questions we ask, the assumptions we challenge and the better ways we find to get things done.

We're intentionally small and senior

You work directly with experienced people who understand your business, not layers of account management.

Success means making a positive difference

Faster conversations, less bureaucracy and better continuity over the long term.

Every AI register, data flow map and policy we hand back is built by a True9 analyst who has sat down with your team, asked the awkward questions, and followed your data from one system to the next. None of it is generated by an AI agent. That's deliberate. Real governance takes real legwork, and you'll see it in what you get back.

We'll also be honest about who we're for. If you're a large enterprise with a dedicated AI governance team, we're probably not the right fit. We work best with established UK businesses that know AI matters, but don't yet have the visibility, policy or plan to govern it safely.

Of course, talk is easy. That's why we're always happy to start small, earn your confidence and grow from there.

Find out where you stand

Remember that client list being pasted into a chatbot?
It might be happening in your business right now. It might not.
Either way, you should know.

Take the [True9 AI Reality Check](#) to see where your business sits today, whether fully in control or facing critical risks. Then, when you're ready for the full picture, our AI Risk & Opportunity Assessment will show you exactly what to fix first and where AI can have the biggest impact.

- Answer 12 quick questions about how AI is used and managed.
- You'll get an instant Control Score out of 100 with practical next steps.
- It's free, and you don't need to sign up to see your score.

T9 | AI Reality Check

Question 12 of 12

Separately from governance, do you have people with the skills to actually build or deploy AI tools, rather than just approve them?

☐ Yes, dedicated

☐ Yes, as part of a wider role

☒ No, we'd rely on an outside partner

☐ No

Back

Your results:

19 out of 100

Risk factor: Critical Risk

With a control score of 19, something is likely to go wrong before anyone notices, because you have sensitive data, no real ownership and little clarity on what regulation applies.

What to do next:

Pause the riskiest use
Stop the one use you're least comfortable with, such as sensitive data going into public AI tools, while you work out the rest.

Name someone this week
An interim owner with real authority beats no owner while you build the proper structure.

Get an outside view
At this level, a proper assessment costs far less than dealing with an incident.

This needs attention now

Our AI Risk & Opportunity Assessment gives you a clear, prioritised picture within weeks: a register of every AI use in your business, each one risk rated, and a plan for what to fix first. It's delivered by our team, not generated by software.

[Arrange a full assessment](#)

Download as a PDF

Take control of AI in your business.

