KPMG

The C-suite's guide to building trust in AI

How to build trust in your AI initiatives – from first pilots to scaled integration

KPMG. Make the Difference.

How to build trust in AI for your Board, your people and you

Trust is essential to delivering the full value of AI in your organisation.

If your Board doesn't trust AI, you won't get the necessary investment. If your people don't trust AI, they won't use it and you won't see efficiency gains. And if you don't trust it then you risk falling short of your strategic objectives.

The UK is grappling with a complex AI trust issue as the technology increasingly integrates into our daily lives.

Research from KPMG and the University of Melbourne* found that just

42%

of the UK public are willing to trust AI

78%

are concerned about negative outcomes.

While most UK workers are now using AI – and experiencing the benefits – they're often doing so in ways that create risk for their organisations. They're relying on AI outputs without verifying the results. And they're putting company information into public AI.

Wherever you are on your AI journey, you'll be facing trust issues, whether it be:

- Employees concerned about what AI means for their roles
- Managers who don't trust that AI is being used securely, accurately or effectively
- Customers wanting reassurance as to how their information is being used
- Or your Board wanting evidence that your AI investment is delivering value.

We've drawn on our own experiences of integrating AI into KPMG and supporting clients with their AI programmes to create this short guide to building trust in AI.

*Gillespie, N., Lockey, S., Ward, T., Macdade, A., & Hassed, G. (2025). Trust, attitudes and use of artificial intelligence: A global study 2025. The University of Melbourne and KPMG. DOI 10.26188/28822919

Where are you on your AI journey?

We've identified three phases that organisations travel through to realise the full value of AI. We look at the trust issues you'll face at each stage, provide tips on how to tackle them, and list the key education, governance and controls you should have in place.

Which of these sounds like you?

Enable:

Initial adoption and piloting of AI tools to augment employee capabilities. You're rolling out tools like Microsoft Copilot to deliver productivity gains.

Go to page 4

Embed:

Integration of AI into specific processes and functions. You're implementing agentic AI to manage end-to-end business processes and functions, and may be looking to scale across your organisation.

Go to page 7

Evolve:

Full-scale AI-driven processes and workflows. Very few organisations have scaled agentic AI across their organisation. We consider what doing so looks like and the questions it poses.

Go to page 10

Phase 1: Enable

How to put in place the foundations for secure, safe and effective use of AI solutions.

Enable phase: the initial adoption of AI tools to augment employee capabilities.

At the enable stage, you're giving your people access to Gen AI tools that help them carry out tasks more effectively and efficiently. This is about improving productivity. At this stage, you may have begun implementing more basic customer-facing chatbots but we're not talking about agentic AI here. We're largely talking about the use of solutions like Microsoft Copilot as personal productivity tools across your organisation; or like GitHub Copilot – an AI-powered coding tool – in your tech teams.

The trust issues you're likely facing, include:

- Having clear guidelines, policies and standards so that users can start adopting AI technology safely and securely.
- Addressing employee concerns about the impact on their jobs.
- Spotting 'hallucinations' and inaccuracies in AI results.
- Stopping unauthorised use of external AI tools, leading to data privacy or intellectual property (IP) risks.
- Demonstrating to the Board that your investment in AI is delivering value.

Does this sound like you? Read on for our tips on building trust during the 'enable' phase.

> **At the start, you're building trust around a proof of concept, without the maturity. And you need to build corporate confidence in the value.**
>
> **– Douglas Dick**
>
> Head of Emerging Technology Risk
>
> KPMG in the UK

6 steps to building trust at the enable stage

1. Make your people part of the journey

You can roll out Microsoft Copilot to your people. But if they don't use it – or use it ineffectively – you're not going to see any benefits. So, how do you get them on board?

People don't like having things imposed on them. It can lead to distrust in the tools being implemented. They like to feel that they have a say and part to play. So, **make them part of your transformation journey from the start. Involve them in testing and validating of your AI systems and processes.** That will give them greater confidence in AI – that it delivers trustworthy results, saves them and their teams time and helps them do better work. It will also help to assuage any concerns that AI is going to replace their jobs.

2. Talk about AI in a way your people can relate to

It feels like a given, but being able to present a clear strategy around your use of AI has a huge part to play in building trust. That means 'clear' to your employees as well as you. When you're communicating your AI plans, don't focus in on how your business will benefit. **Use language and messages that will appeal to your people.**

What we're saying here is stop talking about 'productivity' and 'efficiencies'. Instead, say "it'll help you get mundane tasks done quicker", "it'll free your time to do work you're more interested in", "you know that spreadsheet that takes you hours to complete, AI can do that for you." Explain what's in it for them.

3. Break ingrained habits by encouraging your people to get hands-on

Convincing your people of the benefits is key to breaking ingrained habits, getting them using AI tools and delivering value from your investment. Your people are busy and under pressure. They don't feel like they have time to learn new ways of working. So, they fall back on the tried and trusted.

The best way of addressing that is to give them the time, space and support to get hands-on.

Give them access to AI and encourage them to use it. You can run organisation-wide events, where they can try out your AI with support. And you can gamify their learning – set up challenges around how many different day-to-day activities AI can help them with.

At KPMG, we applied this approach one January, renaming it 'Youcanuary'. We also established a group of 'AI ninjas' – employees who have volunteered to champion AI and support their colleagues in getting more from it. And every summer, we now run a 'Summer of Tech' where our people get insights and time to help them continue their learning on the use of emerging technology solutions.

4. Establish guardrails around the use of public AI

Almost two-fifths (39%) of UK employees have uploaded company information (for example, financial, sales or customer data) into a public AI tool*. Once your information has been entered into a public AI prompt, there's no way of retracting it. Uploading information inappropriately could be a breach of confidence, a breach of contract or in breach of the Data Protection laws – that's in addition to giving away your IP. How can you manage that risk?

For a start, you can implement enterprise versions of Gen AI, which prevent your information escaping into the public domain. However, these tools aren't yet as sophisticated as publicly available ones. And that means employees often seek out and use public Gen AI solutions. You can, of course, completely block access to public Gen AI tools. But then you're limiting what your people can do with AI. No one uses just one tool to do a job. It's about having the right blend and recognising the inherent limitations of internal and external tools and providing guidance on using them.

It's important your people are clear on what they can and can't do, so it's best to reinforce the message at the point of use. **Set up pop ups that appear when your employees try to access an external Gen AI tool to remind them of the terms of use.**

5. Encourage healthy scepticism in Gen AI answers

There's still a lack of trust in the results AI provides. We've all experienced 'hallucinations', where Gen AI has provided a false answer – for example, when you ask it to summarise a meeting it tells you people attended who were on leave. You can ask Gen AI the same question 100 times and it may never answer in the same way.

Your people are used to working with deterministic tools like Excel, where inputs and formulae produce predictable, verifiable outcomes. AI and particularly Gen AI, by contrast, is probabilistic – it generates outputs based on patterns in data, not fixed rules. This means results may be useful but not always accurate or explainable.

To get the most out of AI, **you need to encourage your teams to adopt a mindset of critical evaluation. They need to validate outputs, understand limitations and develop a healthy scepticism.** AI results can seem convincing and authoritative but they often need verifying by a human.

> **When you're working with Gen AI tools like Copilot, think of them as a very smart junior member of your team. They're helping you do your job better, not doing your job for you. And you have to check everything that they do.**
>
> **– Douglas Dick**
> Head of Emerging Technology Risk
> KPMG in the UK

6. Provide hands-on training that's relevant to role – and covers the risks

You get better results from AI if you ask better questions. With a well-composed, simple prompt, Gen AI can create lines of Python code for you.

Of course, being able to write good prompts takes training and experience. This is no one-size-fits-all textbook training exercise. **Identify the areas of your business where AI can have the biggest impact. Then roll out bespoke prompt training in those areas.**

Your AI training shouldn't just be about how to use the tools at your disposal. It also needs to cover how to use them safely and appropriately. **Build in your guidelines on appropriate use of IT – when you can share data, when you can't; what tools are permissible.** And don't just roll that training out once. Repeat it to reinforce messaging and to ensure it's up to date with new developments and risks.

*Gillespie, N., Lockey, S., Ward, T., Macdade, A., & Hassed, G. (2025). Trust, attitudes and use of artificial intelligence: A global study 2025. The University of Melbourne and KPMG. DOI 10.26188/28822919

Phase 2: Embed

How to build trust in process and business transformation enabled by AI

Embed phase: Integration of AI to handle whole processes or transform functions, and scaling of AI across the organisation.

At the embed stage, you're using AI to carry out a particular task – like answering customer queries or checking compliance. As you progress, AI is handling entire processes through an ecosystem of AI agents. Ultimately, you're scaling your AI solutions across your whole enterprise.

The trust issues you're likely facing, include:

- Upskilling, reskilling and resizing roles as AI changes how you operate.
- Managing security and data privacy risks as your number of AI agents increases.
- Ensuring that your AI tools operate effectively and deliver accurate results.
- Demonstrating to the Board that increased investment and scaling of AI will deliver a strong return on investment.
- Managing and mitigating third-party risk.

Does this sound like you? Read on for our tips on building trust during the 'embed' phase.

> **Many of our clients are embedding AI in pockets but struggling with scaling. How do you keep a handle on all those thousands of AI agents managing your processes?**
>
> **– Leanne Allen**
> Head of AI Advisory
> KPMG in the UK

6 steps to building trust at the embed stage

1. Rethink how you do things – don't just do the same thing but with AI

During the 'enable' phase, you're looking at how you can use readily available AI tools to deliver workplace efficiencies. But as you scale up, you'll be looking at the part AI can play across whole processes.

To make the biggest impact, don't just look at how AI can help you do parts of existing processes better. This is an opportunity to transform your business. The question shouldn't be, how can AI improve what I do now? It should be, how can I do this best and where does AI fit into that? And a step on from there, what new things can I do with AI?

That means how your people work and what they do is going to change. You need to think through what you're doing with your human capital. It's no longer just a question of how you upskill them to use AI effectively. It's now starting to be about how you upskill or reskill them to do a job that's changed, or take on an entirely new role.

Map out new career pathways for your people to demonstrate that AI means career growth not job loss.

2. Use a broad scorecard to measure return on investment

Your C-suite and Board need to trust that AI is making a difference. But we find that many organisations are struggling to measure and communicate the return on investment. They've deployed AI tools but don't really have a good view of how much time they're saving.

It's not straightforward. To understand the impact AI is having on productivity, you need an accurate baseline measure. That can be hard to gather. Any effort to measure employee productivity can lead to discontent and add to a sense of mistrust in AI. And if asked how long a task takes, employees who feel that their jobs are under threat may under – or overestimate – depending on what they perceive to be in their best interests.

That's why it's so important to communicate openly and clearly with employees. **Explain to them why you're measuring productivity and what it means for them and how they work.**

Don't just think in terms of productivity and efficiency when you're setting KPIs for your AI initiatives though. You should look at a broader scorecard. Think back to what we said about talking in employees' language about AI, and include softer people measures. What's the impact on levels of stress, on work-life balance, on employee engagement?

3. Think through the risk of becoming dependent on vendors

As AI becomes more embedded within your business processes and technical infrastructure, it becomes harder to unpick. Think through the risks of becoming increasingly dependent on providers of AI as a Service. How can you diversify your AI solutions so that you're not entirely wedded to one vendor and at risk if they're systems are compromised?

You should also consider the risks of working with a third party from a legal perspective. For example, transparency obligations under both AI laws and Data Protection laws require the deployer of AI to understand exactly what it does and how it uses data so it can inform individual users.

4. Keep a human in the loop – at least for now

With agentic AI, whole processes, such as 'Know Your Customer', can be handled by a chain of AI tools. And the calculations made by one AI can be checked by another, helping to cut out the risk of errors or hallucinations, without any human involvement.

But how much do you and your business trust AI? Most people are more comfortable if there's a human in the loop, checking results and providing a fail-safe. But that doesn't cut out all risk and provide 100% accuracy. We worked with one organisation that used manual corrections by the human in the loop as a metric for the accuracy of its AI system. The issue? When the number of overwrites drops, is that because the AI is more accurate or due to human complacency?

Ask yourself for each process, do we need 100% accuracy? At what point are we comfortable to take the human out of the loop? What does AI do completely and where does AI triage to a human?

> **Eventually, an ecosystem of AI agents will do the checking. Before you get an answer the data will have automatically looped through multiple AI, checking for errors.**
>
> **– Leanne Allen**
> Head of AI Advisory
> KPMG in the UK

5. Check the data feeding your AI

What you put in is what you get out. If your AI is drawing from inaccurate or incomplete data sets, you can't expect robust answers.

You'll need to consider the extent to which the personal data you hold – on customers and employees – can lawfully be used in conjunction with or to train AI. Data gathered for one purpose can't automatically be used for another. You should also think through how you'll manage objections, opt outs and subject access requests. The aim is to have the most comprehensive, clean data set that's compliant with data protection laws.

6. Grow your governance and controls with your AI

You don't necessarily need a wholescale policy change when you're first introducing AI at the enable stage. You'll already have existing policies and standards that cover, for example, appropriate use of technology and set out appropriate controls.

But that changes when you get to the point that you're embedding and scaling AI. That doesn't mean you need everything in place before you start. In fact, that's not practical anyway – by the time you've built a perfect policy everything will have moved on and you'll need to start again.

Create your governance and controls framework as you scale your AI.

Many of your existing policies and standards may stand. But the controls you'll need to secure and configure an AI agent will be different.

Phase 3: Evolve

What does an AI-first business look like?

Evolve phase: AI is fully integrated into processes and workflows and operates with minimal human intervention.

Most organisations we're supporting are still enabling AI or at the early stages of embedding and scaling AI. The evolve phase is where they're heading on their AI journeys.

During this phase, the AI is managing much of the risk for you. Trust is built in by design. So, the system stops people from using AI inappropriately. And checks and balances behind the scenes help ensure the accuracy of data.

Very few organisations have embedded end-to-end agentic AI. So, our thinking around exactly what the 'evolve' phase looks like is likely to change as organisations build greater trust in AI. But we do know it's going to change how you approach your work, what you do and how you generate value. It will involve re-engineering your value streams and functions with an AI-first mindset.

> **When we get to the evolve phase, I think the question flips and becomes, what do you value in terms of human skills? What needs to stay human and what can be handed off to AI?**
>
> **– Leanne Allen**
> Head of AI Advisory
> KPMG in the UK

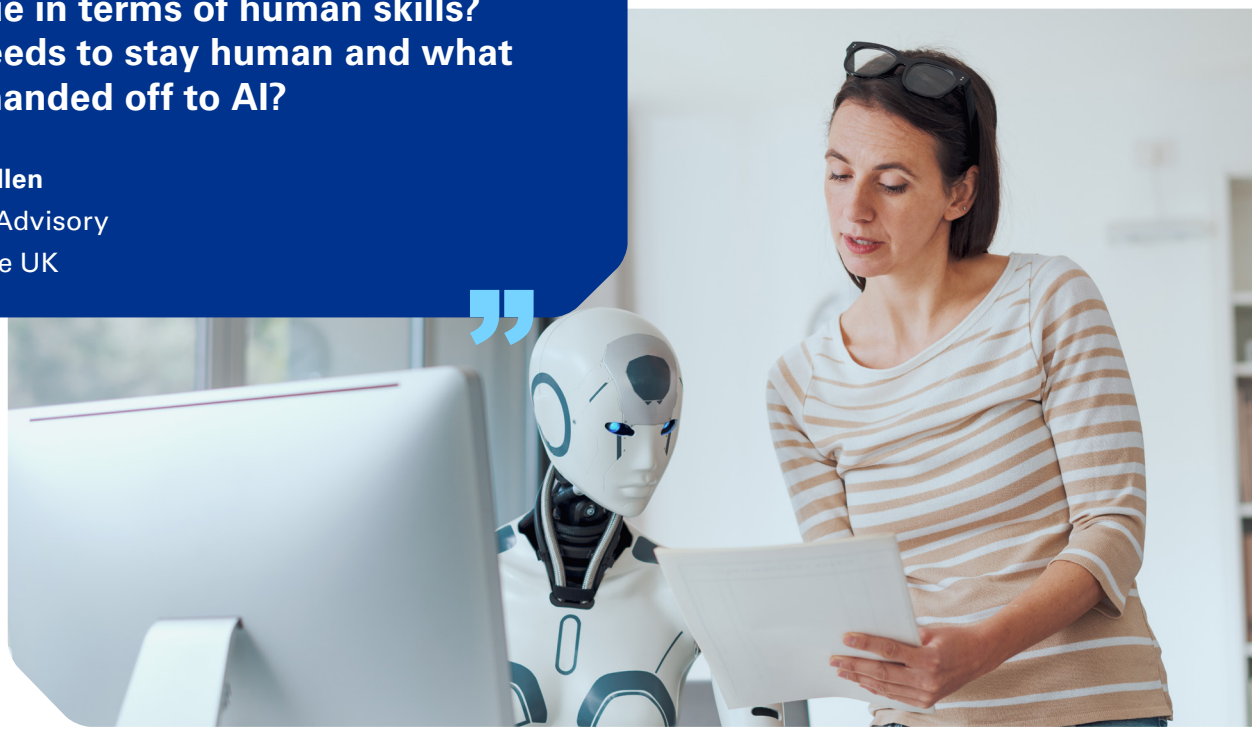

You can build trust and deliver value from AI. You can with KPMG.

AI has the power to transform your business. But only when your leaders trust the output from AI, your customers believe what AI is telling them, and your employees feel confident with the tools they're using.

We combine deep industry experience with technical skills to help you build AI systems that have ethics baked in. And we can provide legal advice too.

Our approach is underpinned by our **Trusted AI Framework**. The framework is based on ten ethical pillars and includes 150 controls to address critical risks across the AI lifecycle.

Developed in 2017, it has helped us establish a long-standing reputation for building responsible AI solutions. We integrate it across the end-to-end AI lifecycle, from strategy and design to ongoing operational excellence.

Find out more about our Trusted AI Services

Want to chat about your AI programme? Contact our KPMG UK experts:

Leanne Allen
Head of AI Advisory
KPMG UK
leanne.allen@kpmg.co.uk

Douglas Dick
Head of Emerging Technology Risk
KPMG UK
douglas.dick@kpmg.co.uk

James Cassidy
Director, Data Protection
KPMG Law
james.cassidy@kpmg.co.uk

The KPMG Trusted AI ethical pillars

Fairness

AI solutions should be designed to reduce or eliminate bias against individuals, communities and groups.

Explainability

AI solutions should be developed and delivered in a way that answers the questions of how and why a conclusion was drawn from the solution.

Security

Robust and resilient practices should be implemented to safeguard AI solutions against bad actors, misinformation or adverse events.

Sustainability

AI solutions should be designed to be energy efficient, reduce carbon emissions and support a cleaner environment.

Transparency

AI solutions should include responsible disclosure to provide stakeholders with a clear understanding of what is happening in each solution across the AI lifecycle.

Accountability

Human oversight and responsibility should be embedded across the AI lifecycle to manage risk and comply with applicable laws and regulations.

Safety

AI solutions should be designed and implemented to safeguard against harm to people, businesses and property.

Data integrity

Data used in AI solutions should be acquired in compliance with applicable laws and regulations and assessed for accuracy, completeness, appropriateness and quality to drive trusted decisions.

Reliability

AI solutions should consistently operate in accordance with their intended purpose and scope and at the desired level of precision.

Privacy

AI solutions should be designed to comply with applicable privacy and data protection laws and regulations.

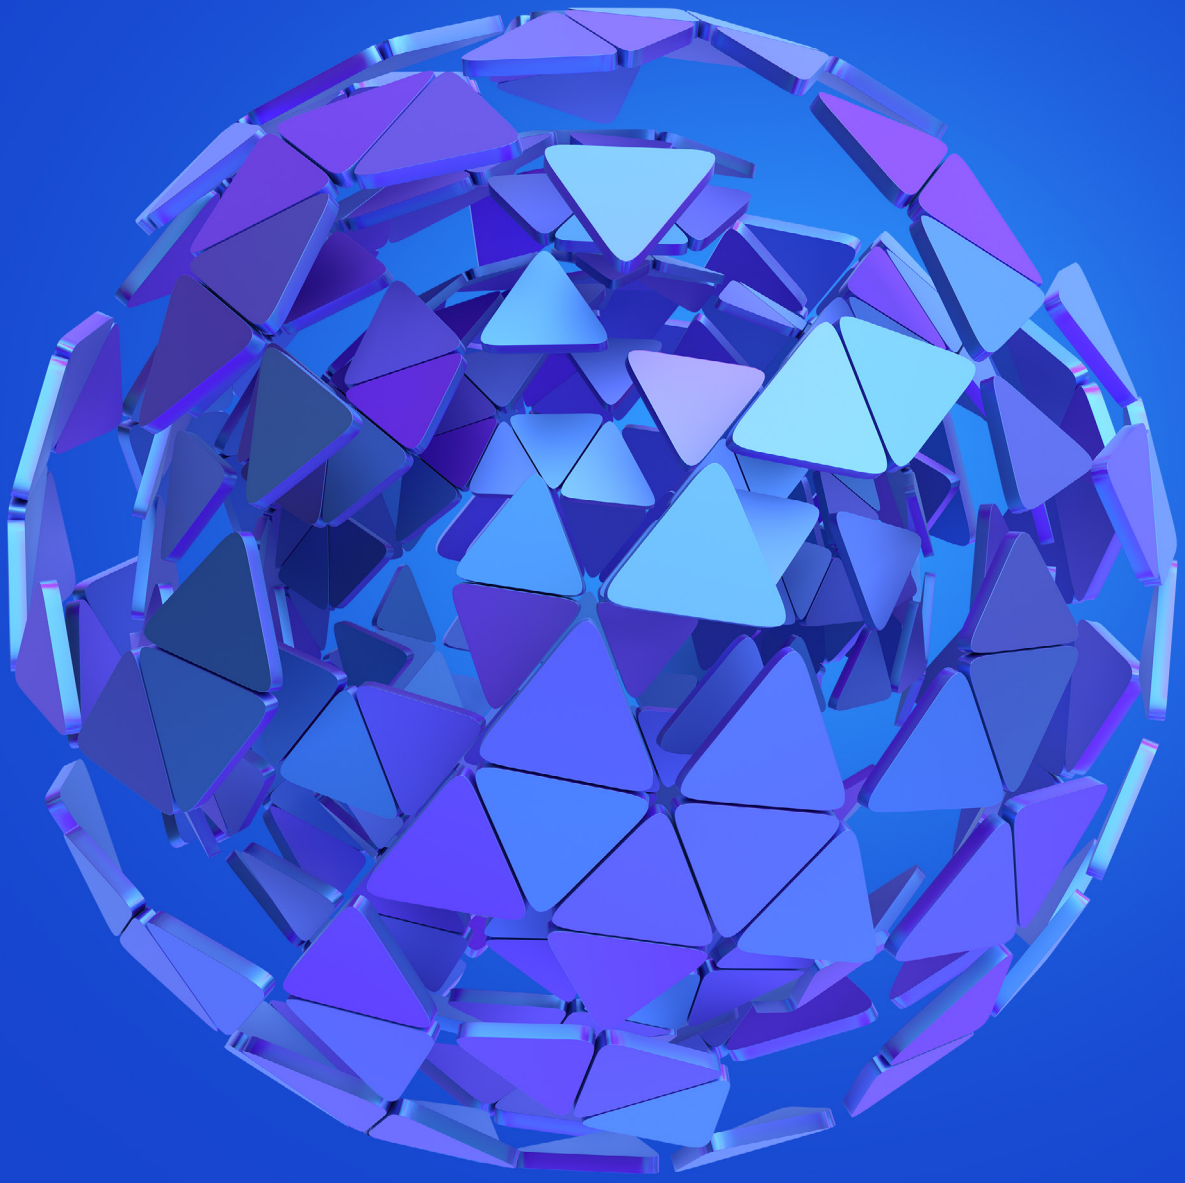

Some or all of the services described herein may not be permissible for KPMG audited entities and their affiliates or related entities.

kpmg.com/uk

CREATE. | CRT162841A | August 2025

Document Classification: KPMG Public