MIND THE GOVERNANCE GAP

How **AI Adoption** is **Outpacing Governance** and How to **Get Back on Track.**

transparity

CONTENTS:

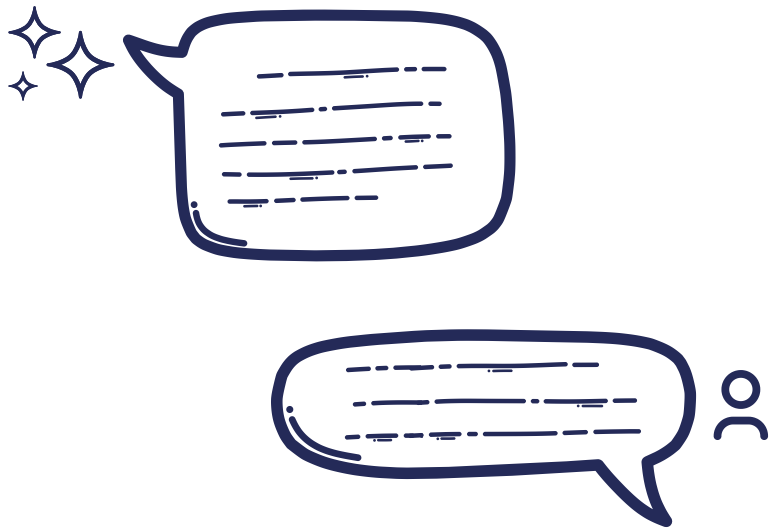

AI HAS MOVED INTO THE OPERATING MODEL. NOT HYPOTHETICALLY, NOT EVENTUALLY. NOW.

The organisations that win with AI over the next two years won't necessarily be the fastest adopters.

They'll be the ones that can scale AI with confidence, backed by operational governance, strong cyber foundations and the ability to demonstrate control when it matters most.

THE QUESTION ISN'T

"SHOULD WE BE ADOPTING AI?"

IT'S

METHODOLOGY:

THE WHO, WHAT AND WHY OF OUR RESEARCH

We wanted to go beyond the existing data and understand the real story on AI governance and security, from IT leaders across industries.

What's keeping them up at night, how they are juggling competing priorities between AI adoption and keeping their organisations secure and how are they really balancing innovation with control?

Not what the policy documents say. Not the aspirational governance frameworks sitting in SharePoint. What is actually happening on the ground and what separates the organisations scaling AI confidently from those struggling with visibility, risk and compliance?

We partnered with YouGov to survey **over 500 IT leaders and decision-makers** from organisations with **100-1,000 UK-based employees**. We heard from leaders across a range of industries including Financial Services (FSI), Professional Services, Retail, Nonprofits, Health and the Public Sector.

The results told a clear story. AI adoption is well underway. Governance is not keeping pace. And the gap between the two is creating a confidence crisis that most organisations have not yet fully reckoned with.

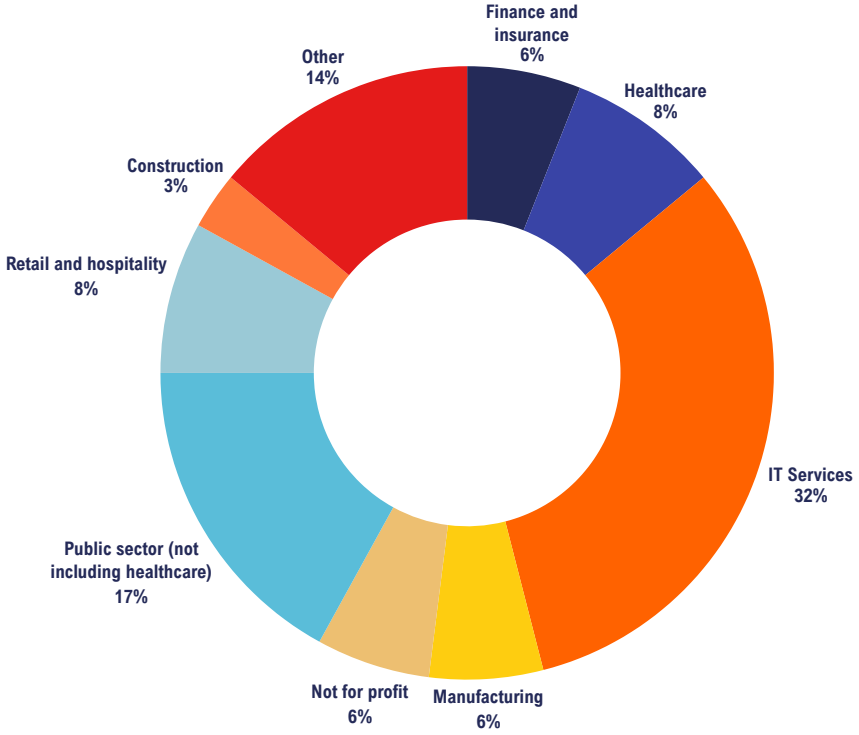

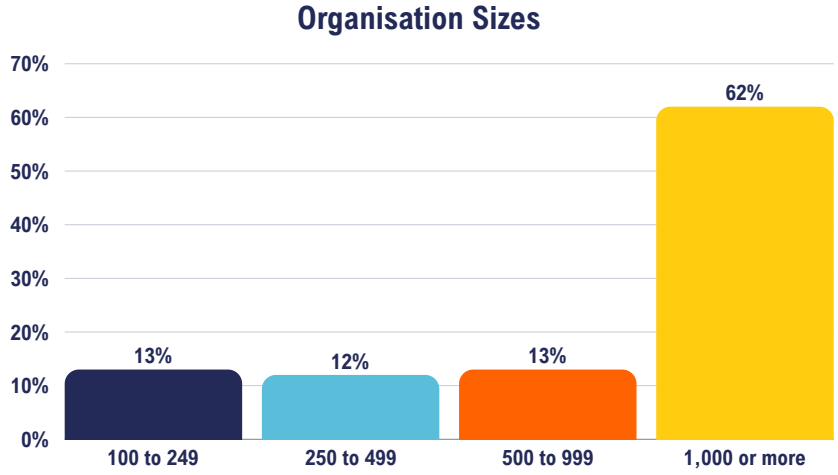

THE RUNAWAY TRAIN OF AI ADOPTION

AI IS NO LONGER A PILOT.

For the majority of UK organisations, AI has already moved into the way work gets done; into workflows, decisions, customer interactions and business-critical processes.

THE EXPERIMENT PHASE IS LARGELY BEHIND US AND THE AI TRAIN HAS LEFT THE STATION.

The data makes this clear.

69%

of organisations are beyond limited AI experimentation, actively piloting, embedding or deploying AI in core operations.

And the ambition is only growing.

75%

of organisations say AI is important or critical to their business goals over the next two years. That is not a marginal interest. That is a strategic priority, backed by investment and expectation at board level.

The deployment landscape is equally telling.

65%

of organisations are using or planning to use AI assistants: tools like Microsoft Copilot that sit inside everyday productivity environments, accessing emails, meetings, documents and data.

33%

are already using or planning to implement AI agents: systems that go further, capable of taking autonomous action across workflows, systems and data sources.

AI IS IN THE OPERATING MODEL

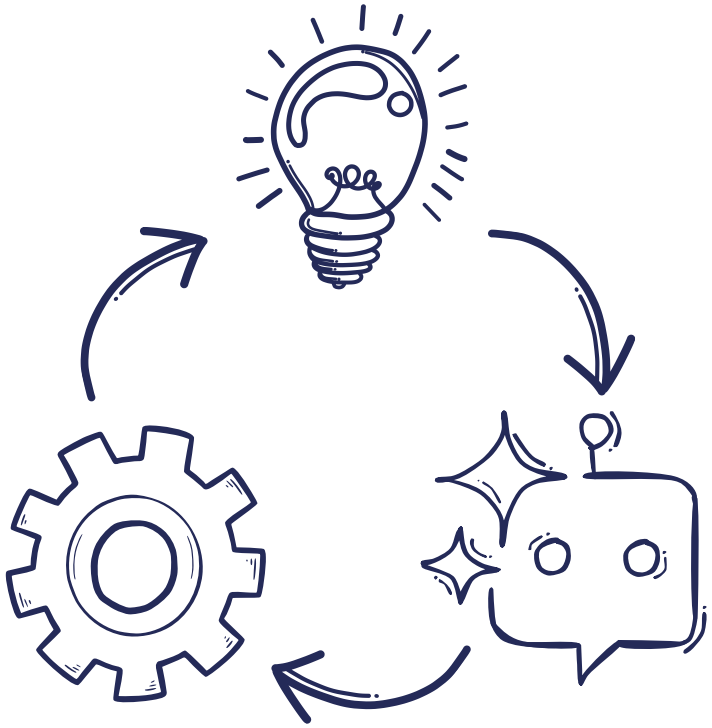

WHEN WE LOOK AT WHERE AI IS BEING APPLIED, THE OPERATIONAL REALITY BECOMES EVEN SHARPER.

43% of organisations are using AI-driven analytics, drawing on machine intelligence to inform business decisions.

35% are running automated workflows, removing human touchpoints from processes that previously required them.

27% are using custom or third-party AI tools, extending AI capability beyond the major platforms and into specialist use cases.

THIS IS NOT A WAVE THAT IS COMING. IT'S ALREADY HERE.

When AI becomes part of the operating model, it also becomes part of the security model.

Every assistant, every automated workflow, every agent and every third-party tool expands the attack surface, increasing data exposure risk and raising the bar for what operational control requires.

THE GOVERNANCE GAP

ONLY 12.5%

OF ORGANISATIONS SAY THEIR AI GOVERNANCE IS AHEAD OF THEIR AI ADOPTION.

ONE IN EIGHT.

In a landscape where **69%** of organisations are actively deploying AI in core operations, only a small minority have built the governance foundations to match.

More broadly, **51%** of organisations either say AI adoption is outpacing their governance or admit they have not fully considered AI governance at all. For more than half of UK organisations surveyed, the governance gap is growing.

WHEN WE DIG DEEPER, THE DETAIL BEHIND THOSE NUMBERS IS ILLUMINATING.

Only 28% of organisations have ongoing governance with continuous oversight; the kind of always-on capability that matches the always-on nature of AI deployment. At the other end of the scale, **15%** have no formal governance in place whatsoever and **47%** are relying on guidelines only. They have no specific AI controls, or are simply unsure what controls exist within their organisation.

SO HOW DID WE GET HERE?

Simply put, AI adoption moved faster than most governance functions were designed to handle. From AI assistants deployed at speed to the democratisation of AI Agents and the proliferation of unmanaged third-party AI tools, the looming challenge of shadow AI has grown faster than the frameworks designed to govern it.

The central problem isn't that organisations are ignoring governance entirely.
It's that governance hasn't kept pace with how deeply adoption has been built into everyday practice.

AI is not creating entirely new governance problems. It's exposing the ones organisations have been carrying for years. Weak data classification, over-permissive access, inconsistent labelling and unclear ownership have always created risk. The difference now is that tools like Copilot can surface that risk in seconds.

THE CONFIDENCE COLLAPSE.

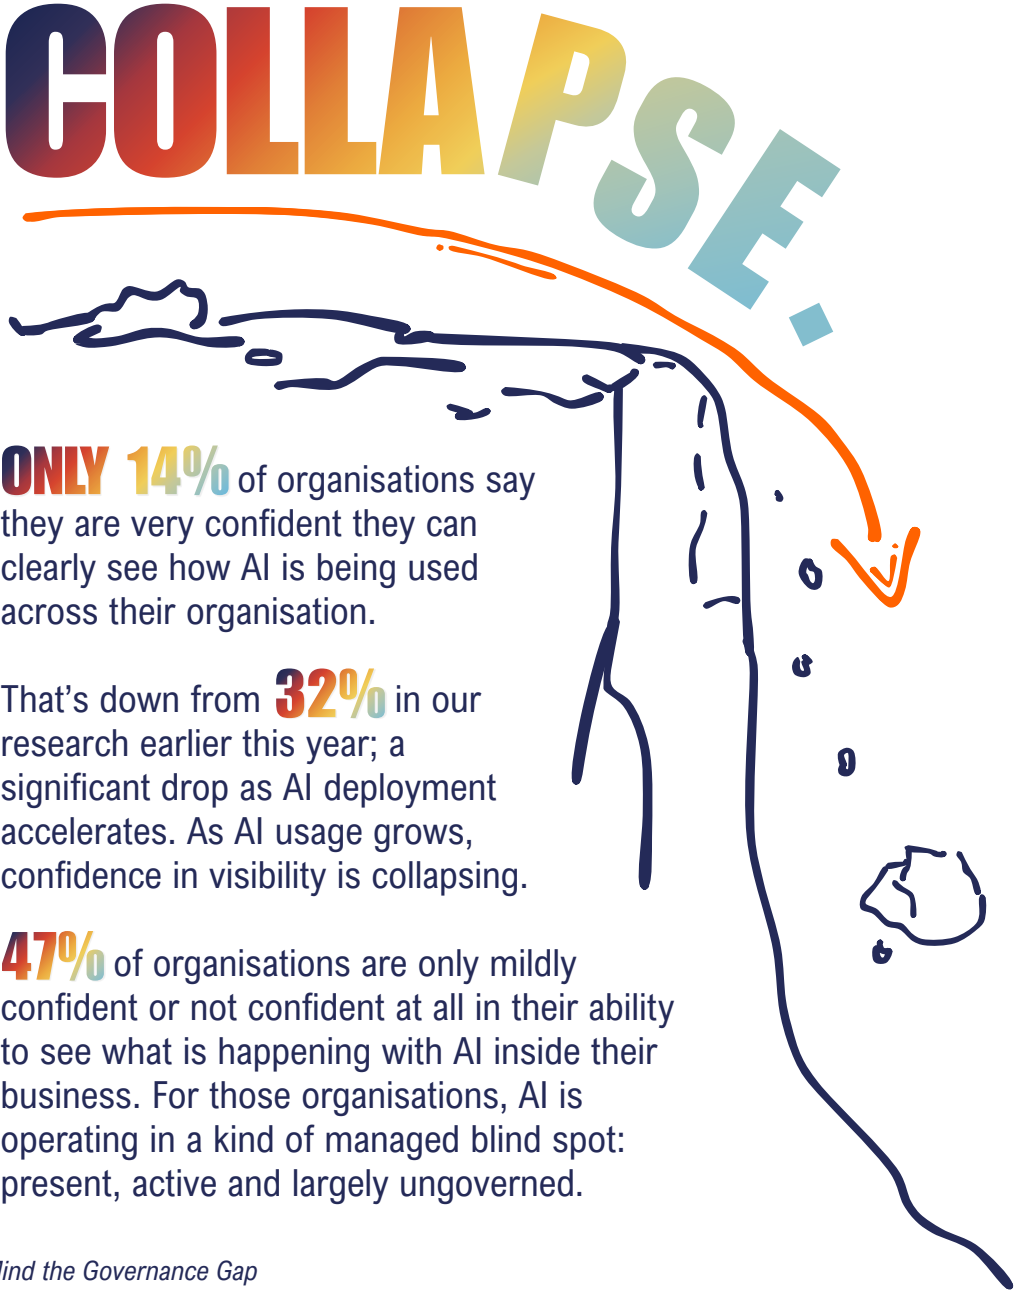

ONLY 14% of organisations say they are very confident they can clearly see how AI is being used across their organisation.

That's down from **32%** in our research earlier this year; a significant drop as AI deployment accelerates. As AI usage grows, confidence in visibility is collapsing.

47% of organisations are only mildly confident or not confident at all in their ability to see what is happening with AI inside their business. For those organisations, AI is operating in a kind of managed blind spot: present, active and largely ungoverned.

ONLY 12% of organisations are very confident that sensitive data is consistently protected across AI tools and use cases while **53%** lack strong confidence in how their sensitive data is being handled across the AI tools their people are using every day.

Operationally, that means AI assistants accessing emails, documents and internal databases. Automated workflows ingesting customer data, financial records, HR information. Custom and third-party AI tools with varying levels of data governance.

It's no wonder the majority of organisations can't say with confidence that this data is being consistently protected.

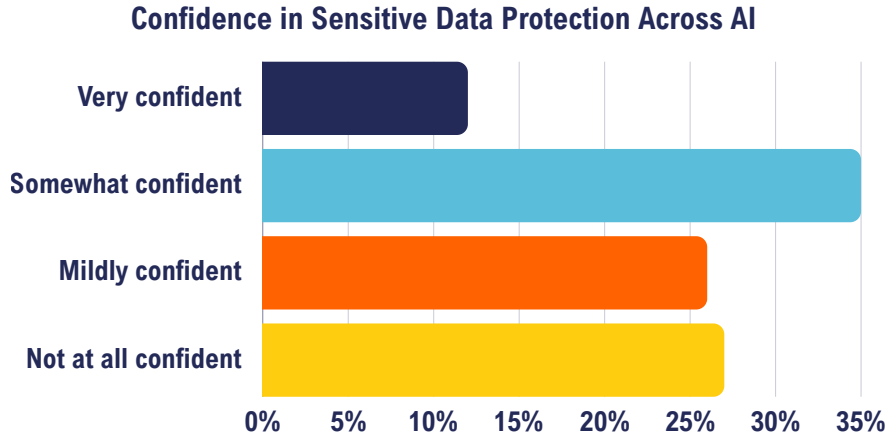

AUDIT ANXIETY

AUDIT READINESS IS MORE LIKE AUDIT ANXIETY FOR THE MAJORITY OF OUR RESPONDENTS.

Just **15% of organisations** say they are very confident they could provide evidence of compliant AI usage within 72 hours if asked by leadership or auditors while **49% are not very** confident or not confident at all that they could meet that bar.

WHY DOES 72 HOURS MATTER?

ICO reporting requirements make that window a concrete test of real-world governance readiness.

Organisations across every industry may need to report personal data breaches within 72 hours.

Regulated sectors face additional scrutiny from sector-specific compliance obligations, auditors and supervisory bodies.

AI usage and controls must be evidenced quickly. If leaders can't explain where AI is being used, what data is exposed, who owns the risk and what controls are operating, the gap between policy intent and operational reality becomes a measurable compliance exposure.

And the commercial impact is real: reputational damage can weaken customer trust, stall deals and turn AI from a growth story into a board-level risk.

ARE YOU CONFIDENT YOUR ORGANISATION COULD ANSWER THAT CHALLENGE TOMORROW?

For most organisations, the honest answer is **NO**.

That's not a governance problem alone.
It is a security operations problem and it requires a security operations answer.

UNDERCOVER AI AGENTS

AI AGENTS ARE CHANGING THE GOVERNANCE GAME.

33% of organisations are already using or planning to implement AI agents. That figure will only grow. And it matters to security and governance pros everywhere because agents don't just support work, they take action.

The shift from generating outputs to taking autonomous actions fundamentally changes the governance question. And the data among agent-adopting organisations reveals just how much pressure that creates.

Among organisations already using AI agents, **57%** are also running automated workflows, **64%** are using AI-driven analytics and **42%** rely on custom or third-party AI tools.

These are not organisations dipping a toe into AI.

They are running interconnected, AI-driven operating environments and the governance challenge across all of those systems compounds quickly.

The confidence data among this group is a stark wake-up call.

Even in organisations that have committed to agent-driven AI, fewer than one in four are very confident they can clearly see how AI is being used across their organisation.

53% of agent-using organisations say governance is slowing AI initiatives moderately or significantly.

THE TOP BLOCKERS?

Unclear policies **(36%)**, lack of visibility into risk **(36%)**, regulatory concerns **(36%)** and a lack of skilled resource **(35%)**. These are not niche operational problems. They are structural blockers affecting the majority of agent adopters.

AND YET THE DIRECTION OF TRAVEL IS UNDENIABLE.

83% of agent-adopting organisations say AI governance must become a continuous operational capability.

AGENTIC ERA

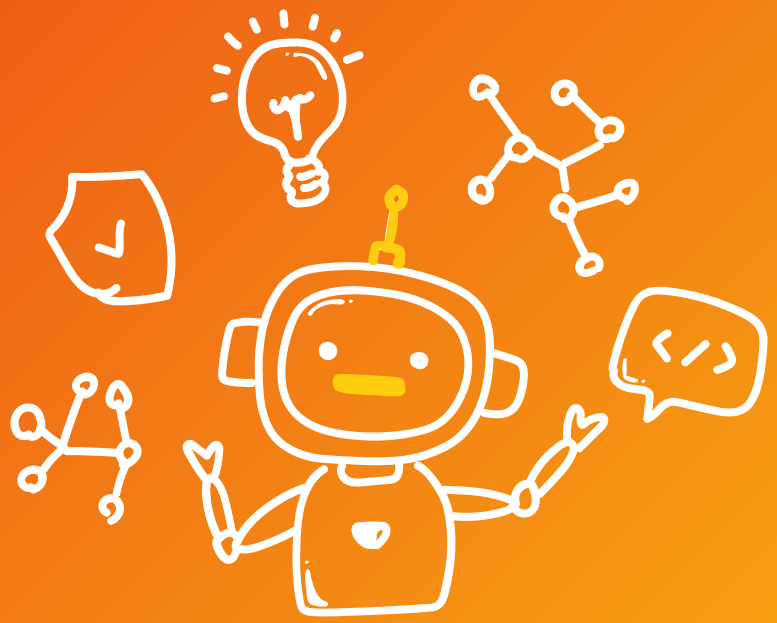

SO, WHAT MAKES AGENTS SO SIGNIFICANT?

Unlike AI assistants, agents can take action without a human in the loop at every step.

AGENTS CAN INHERIT USER PERMISSIONS.

They can access sensitive data repositories. They can trigger workflows across connected systems. They can interact with third-party applications. They can operate continuously, across multiple processes, without the natural circuit breakers that human involvement provides.

And in many organisations, there is currently no mechanism to discover which agents exist, what permissions they hold, what data they can reach, or how agent-related activity would be investigated if something went wrong.

Prompt injection attacks, where malicious instructions are embedded in content that an agent processes, are already documented in the wild.

Jailbreak attempts target AI systems to bypass their guardrails. Data exfiltration through AI interfaces is an emerging and under monitored risk.

And the same AI capabilities that make agents powerful for productivity also make them interesting to attackers looking to accelerate their own tooling or exploit ungoverned access points.

AI agents represent a structural shift in how AI interacts with the organisation and creates new demands on how security teams keep up.

THAT IS NOT A FUTURE REQUIREMENT.

FOR THE 33% of organisations already deploying agents, it is a present one.

CYBER FUNDAMENTALS SET THE PACE

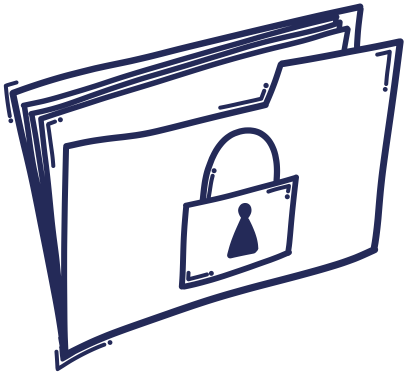

HERE IS THE REALITY:

AI governance can't be operationalised without the cyber foundations to support it.

Now, more than ever, security and risk teams must find a way to enable their organisations at speed. Saying 'no' isn't an option and the 'lock down and hope for the best' approach will be upturned by the proliferation of AI.

YOU CAN'T GOVERN WHAT YOU CAN'T SEE

You can't protect data you haven't classified. You can't manage identity risk without strong identity controls. And you can't do any of this at scale without the operational muscle to keep it running continuously.

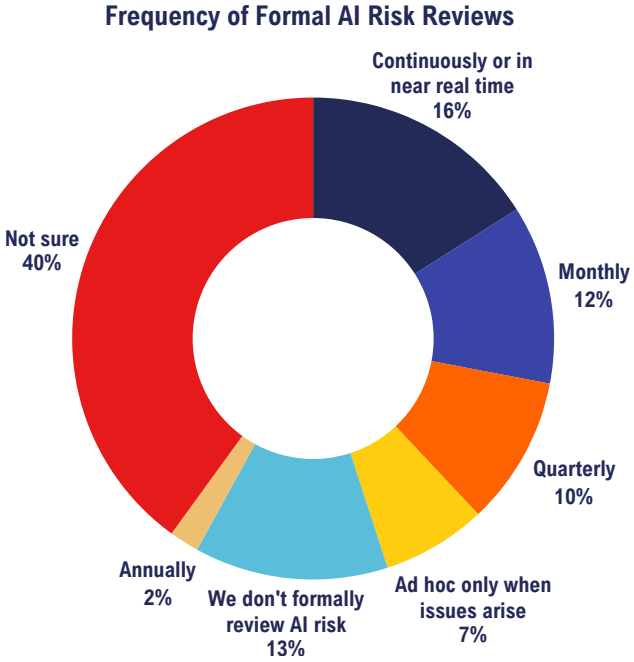

OUR RESEARCH MAKES THIS CLEAR.

Only **22%** of organisations have consistent controls and monitoring in place. **47%** are relying on guidelines only, have no specific AI controls, or are unsure what controls even exist.

30% say responsibility is unclear, not formally defined, or they simply do not know who owns AI governance.

And **53%** either do not know how often AI risk is reviewed or say it is not formally reviewed at all, with only **16%** reviewing AI risk continuously or in near real-time. The cyber fundamentals that make AI governance possible aren't new. They're the building blocks of a mature security posture and they need to be in place before you can innovate at scale.

IT'S BACK TO BASICS:

- **Cyber hygiene** from patching to permissions
- **Risk management** that comes all the way from the top
- **Data and identity** as core foundations
- **Cyber awareness** and education for every user
- **Knowing what really counts** and where it is, to ringfence your most valuable data

YOUR TICKET TO SCALE WITH TRUST

WHEN WE LOOK AT RESPONDENTS WITH CONTINUOUS AI OVERSIGHT, THE RESULTS BECOME CRYSTAL CLEAR.

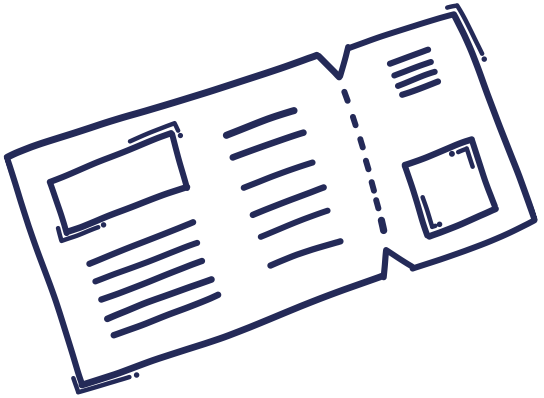

54% have AI embedded in multiple business processes or made it core to how work is designed, compared to a fraction of that figure in the broader sample.

60% say AI is important or critical to their business goals compared to **37%** overall. They are not just more confident. They are more ambitious, because governance gives them the capability to move.

ON THE METRICS THAT MATTER MOST:

31% are very confident in their AI visibility, more than double the overall figure of **14%**.

These organisations can evidence compliant AI usage within **72 hours**, again more than double the overall figure of **15%**.

79% are using or planning AI assistants and **54%** are using or planning AI agents; both well above the overall average.

88% of this group agree that AI governance must be a continuous operational capability.

For IT leaders with continuous AI oversight, governance reduces friction rather than creating it:

77% avoid policy-related delays

72% have clearer visibility of AI-related risk

69% aren't slowed down by approvals

THE MESSAGE COULDN'T BE CLEARER:

Governance does not slow AI down.
Poorly operationalised governance does.

GOVERN OPERATIONALLY, MOVE FASTER

71% of all organisations surveyed agree that AI governance must become a continuous operational capability, not a periodic compliance exercise. That's a supermajority of UK IT leaders.

38% cite continuous monitoring of AI risk as a top capability gap and say they lack the skilled resources to manage AI governance effectively. The confidence gap in visibility, data protection and 72-hour audit readiness confirms that policy alone is not closing those gaps.

As the pace of AI adoption accelerates and governance capabilities lag, the window for getting ahead of this runaway train is narrowing.

SO, WHAT DOES THE PATH FORWARD LOOK LIKE?

SEEING AI USAGE CLEARLY: Discovering what tools are in use, which identities are accessing them, what data is being exposed and where ungoverned AI activity is happening outside sanctioned environments.

UNDERSTANDING YOUR RISK: Not in the abstract, but operationally. What is the organisation's risk tolerance? Where does current AI usage sit relative to that tolerance? What are the highest-priority control gaps?

EVIDENCING CONTROL: Implementing the monitoring, documentation and reporting capability that lets leaders answer hard questions quickly.

RESPOND QUICKLY: Building the capability to investigate incidents, contain exposure, report accurately and maintain stakeholder confidence.

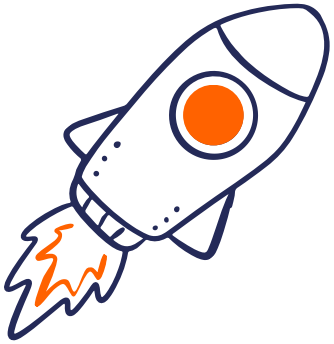

AI scales faster when governance and security give organisations the confidence to move. The risk is not in adopting AI. The risk is in adopting it without the foundations to see it, secure it, govern it and evidence it.

The organisations that build those foundations now will not just be more compliant. They will be more capable, more resilient and better positioned to take AI further, faster, than the organisations still working out what controls they have in place. That means better ROI on AI initiatives, competitive advantage and insulation from the commercial impact of a breach.

HOW TRANSPARITY HELPS ORGANISATIONS SCALE AI SAFELY

TRANSPARITY WORKS WITH UK ENTERPRISES TO TURN AI GOVERNANCE INTENT INTO OPERATIONAL REALITY.

We combine managed Governance, Risk and Compliance (GRC) with a broader cybersecurity capability that addresses the full picture: identity, access, data protection, posture management, threat detection, incident response and managed security operations.

Whether you are working to understand your current AI risk exposure, strengthen the cyber foundations that governance depends on, or build the continuous monitoring capability that lets you evidence control at pace, we can help you move from where you are to where you need to be.

READY TO SAFELY SCALE AI?

Book a strategic consultation with a **Transparity expert** today. Let's bridge the governance gap.