

The AI control gap: what UK boards can't see, and what it will cost them

AI adoption in UK business has outrun the ability to see it, police it and evidence it. This paper sets out the size of the gap, what it costs, what regulators now expect, and the operating discipline that closes it.

Executive summary

In one paragraph

AI is in daily use across UK business, and most of it arrived without a decision being taken. 71% of UK employees have used unapproved AI tools at work (Microsoft UK, 2025), yet only 31% of UK employers have a formal generative AI policy (CIPD, 2026) and 63% of UK organisations lack AI access controls (IBM, 2025). That gap between use and oversight is the AI control gap. It is a data exposure: over a third of what staff paste into AI tools is sensitive (Cyberhaven, 2025) and high shadow-AI use adds £498,000 to a UK breach (IBM, 2025). It is a regulatory exposure: the Data (Use and Access) Act is in force, a statutory ICO code on AI is coming and ISO 42001 is appearing in tenders. And it is a board exposure, because the people accountable cannot see it. Bans do not close the gap and governance platforms were not built for a 300-person business. What closes it is an operating discipline: audit, rules policed by a usage log, watch, measure. The aim is not to slow AI down; it is to be able to say yes to more of it, with confidence. The paper ends with ten questions a board can answer today.

1. AI walked in through the front door

Most boards think of AI adoption as a project: scoped, budgeted, procured and reported on. For most UK businesses it has been an event that happened to them. It arrived on personal accounts, in browser tabs, and through features switched on by vendors the company already paid. Nobody signed it off because there was nothing to sign.

The official statistics understate this. Around 35% of UK businesses with ten or more employees now use AI, up from roughly 12% in September 2023, and only 10% use it "extensively" (ONS, 2026). Active AI use among SMEs is 54%, up from 35% a year earlier (British Chambers of Commerce, 2026). Those are figures as reported by the business: what the company believes it has decided to do.

The employee-level data tells a different story. Microsoft's survey of 2,003 UK employees found that 71% had used unapproved consumer AI tools at work, 51% weekly, and 28% said their employer provided no approved alternative (Microsoft UK, 2025). MIT found personal AI tools in use at more than 90% of companies, against 40% with an official subscription (MIT NANDA, 2025). The average organisation now runs eight generative AI applications, with 47% of users on personal accounts and 223 AI data-policy violations a month (Netskope, 2026). 80% of enterprise AI tools operate without IT oversight (Zluri, 2025).

The company believes it has 35% adoption. Its staff report 71%. The difference is the part nobody owns.

When a director asks "how are we using AI?", the honest answer in most UK businesses is "we do not know". Finance holds a handful of subscriptions. IT holds a list of approved tools. Neither holds the marketing team that has moved its drafting into a free chatbot, or the

customer-service lead who pastes complaint letters into an assistant to summarise them. The estate exists. It is simply unmapped.

This is not a criticism of staff, who use these tools because they work and because the business gave them nothing better. It is a statement about visibility. A business cannot control, cost or evidence what it cannot see.

2. Policy without policing

The instinctive response is to write a policy. The CIPD found that 31% of UK employers now have a formal generative AI policy, up from 16%, and that only 35% provide any AI training (CIPD, 2026). Put the other way round, seven in ten UK employers have no formal policy and two in three give no training on tools most staff already use.

Having a policy is the easy half. Only 31% of UK organisations have AI governance policies and 63% lack AI access controls (IBM, 2025). A policy that says "do not enter customer data into public AI tools" changes nothing if no control prevents it, no log records it and nobody is looking. It is a document, not a system.

What does "policed" mean in practice? Three things working together. First, a usage log: which AI tools were used, by whom, when, and with what kind of data, across browser tools, licensed products and any API spend the technology team has switched on. Second, business rules applied to that log: this team may use these tools for these purposes; this category of data may never leave; this use case needs human sign-off before output reaches a customer. Third, someone who reviews exceptions and can show what happened when asked.

This is the discipline every finance director applies to expenses. Nobody relies on the policy alone; there is a system that records the claim, rules on what may be claimed, approval thresholds and an audit trail. The document sets the expectation, the system provides the evidence.

Policy is a document. Policing is a system. The ICO will ask for the second one.

Training matters too. The 35% of employers who train staff (CIPD, 2026) are addressing a skills gap that 31% of UK firms admit to (Lloyds Business Barometer, 2026). Training that explains which tools are approved, what must never be pasted into a chatbot and how to check an answer before acting on it is a control in its own right. Training nobody can evidence is not.

3. What is in the data going in and out

This matters to the board, not only to IT, because of what is inside the traffic. Cyberhaven found that 34.8% of data entering AI tools is sensitive, and that 71.7% of the AI tools in use

are rated high or critical risk (Cyberhaven, 2025). In a mid-market business that means customer records, pricing, contracts, HR matters and source code, entered into tools whose terms most users have never read.

The cost when it goes wrong is measurable. The average UK data breach costs £3.78 million, and breaches at organisations with high shadow AI use cost £498,000 more (IBM, 2025). 92% of organisations suffering an AI-related breach had no AI-specific access controls (IBM, 2026). Gartner forecasts that 40% of organisations will suffer security or compliance incidents from unauthorised AI by 2030 (Gartner, 2025).

The best-known example remains Samsung, which banned generative AI in May 2023 after engineers uploaded internal source code to a public chatbot (CNBC, 2023). The lesson is not that Samsung's staff were unusual. It is that a global technology company with a mature security function found out only after the data had left.

The regulator is not passive. The ICO's largest fines of 2025 went to Capita (£14 million), Advanced Computer Software (£3.07 million) and 23andMe (£2.31 million) (Trowers & Hamlins, 2026). None was an AI case as such. All were failures of control over personal data, which is exactly what an unmapped AI estate is. The ICO needs no AI-specific power to act on a customer file pasted into a public tool; existing law already covers it.

Table 1. The control gap in numbers (UK unless stated)

Measure	Figure	Source
UK employees who have used unapproved AI tools at work	71%	Microsoft UK, 2025
UK employers with a formal generative AI policy	31%	CIPD, 2026
UK organisations lacking AI access controls	63%	IBM, 2025
Share of data entering AI tools that is sensitive (global)	34.8%	Cyberhaven, 2025
Extra cost of a UK breach where shadow AI use is high	£498,000	IBM, 2025
AI-related breach victims with no AI access controls (global)	92%	IBM, 2026
UK businesses using AI whose cyber practices address AI risk	24%	DSIT, 2026
Boards with limited or no AI knowledge (global)	66%	Deloitte, 2025

Two further figures complete the picture. Only 24% of UK businesses using or considering AI have cyber practices that address AI risk (DSIT, 2026). And 66% of boards have limited or no knowledge of AI; at 31% it is not on the agenda at all (Deloitte, 2025). The data is going out. The people accountable are mostly not looking.

4. The regulatory clock

The UK has no standalone AI Act and, after the King's Speech in May 2026, is not getting one soon (IAPP, 2026). UK regulation of AI is principles-based and runs through existing law: data protection, consumer protection, sector rules and cyber-security duties. That is sometimes read as a reason to wait. It is the opposite. Principles-based regulation does not tell you what to do; it makes you responsible for deciding, and for showing your reasoning afterwards. The clock is running on several fronts.

Data (Use and Access) Act 2025

The main data-protection provisions of the Data (Use and Access) Act commenced on 5 February 2026, including reformed rules on automated decision-making, recognised legitimate interests, and PECR fines aligned with UK GDPR levels (DLA Piper, 2026). Any business using AI to make or shape decisions about customers or staff needs to know where that is happening and what safeguards apply.

A statutory ICO code on AI

Under a statutory instrument in force from 12 May 2026, the ICO must produce a code of practice on AI and automated decision-making (Bratby Law, 2026). Its strategy shows final guidance on automated decision-making due in winter 2026, alongside a technology report on agentic AI published in January 2026 (ICO, 2026). When the code lands, "we did not know staff were using it" will be a poor defence.

EU AI Act, if you sell into Europe

UK businesses placing AI-enabled products or services on the EU market fall within the EU AI Act. General-purpose AI obligations have applied since 2 August 2025; Article 50 transparency obligations from 2 August 2026; Annex III high-risk obligations were deferred to 2 December 2027 and Annex I to 2 August 2028 (Gibson Dunn / Pinsent Masons, 2026). The deferral bought time. It did not remove the need to know which uses would count as high-risk.

Cyber Security and Resilience Bill

The Cyber Security and Resilience Bill has been in the House of Lords since June 2026, with Royal Assent expected late in 2026 and reported fines of up to £17 million or 4% of turnover for in-scope organisations (to be confirmed against the Bill as enacted). For businesses in critical supply chains, an unmapped AI estate is a cyber-security gap as well as a data-protection one.

ISO 42001 as a tender question

ISO/IEC 42001, the standard for AI management systems, had roughly 350 certified organisations worldwide by April 2026, and BSI became the first UKAS-accredited certifier in January 2026 (UKAS, 2026). A small number, which is why it is appearing in procurement questionnaires as a differentiator. A mid-market supplier need not be certified tomorrow, but it must be able to answer "how do you manage AI risk?" with something better than a PDF.

NCSC guidance on AI agents

Finally, the National Cyber Security Centre published guidance in August 2026 on autonomous AI agents, calling for scoped credentials, dedicated identities, logging, human oversight and a kill-switch (NCSC, 2026). As AI moves from answering questions to taking actions in business systems, attributable and reversible actions become the baseline.

Principles-based regulation does not tell you what to do. It makes you responsible for deciding, and for showing your reasoning afterwards.

5. Why bans and platforms both fail

Boards reach for one of two answers. The first is to ban it. 25% of UK employers prohibit generative AI outright, while 61% allow it (CIPD, 2025). Set that 25% against the 71% of employees who have used unapproved tools anyway (Microsoft UK, 2025) and the problem is obvious. A ban does not stop use; it moves use onto personal devices and accounts, where the business has even less sight of it. It also forfeits the upside. The 28% who say their employer offers no approved alternative (Microsoft UK, 2025) do not stop; they go elsewhere.

The second is to buy an AI governance platform. Gartner published its first Magic Quadrant for these in June 2026, sizing the market at \$65 million in 2024 and \$1.4 billion by 2030 (Gartner, 2026). The leading products are sold to risk teams at large enterprises, at annual costs from tens of thousands of pounds upwards, and are built around model inventories, regulatory mapping and bias auditing. Good tools for a bank with a model risk function. They assume staffing, and questions, that a 300-person business does not have.

Nor do vendor add-ons close the gap alone. Microsoft's Purview tooling sees Microsoft's world, and needs device onboarding and browser extensions to see more. It does not see a team's spend with another provider, the vendor who quietly added an AI feature to a product already in use, or whether any output was correct.

Both responses fail the same way. They treat AI as an exception, to be excluded or fenced off in a specialist tool. Businesses that get this right treat AI as a normal operation, with the routines already applied to money, people and data.

6. An operating discipline for AI

The discipline has four parts. None is exotic.

Audit. Find every AI tool in use, licensed or not, and map what data goes in and out, what it costs and who owns it. A fixed piece of work that produces a document the board can read.

Rules. Turn the audit into a policy written for a UK business, with the controls to enforce it: approved tools by team, data that may never leave, use cases needing human sign-off, and training. Align it to UK GDPR and the Data (Use and Access) Act, and to ISO 42001 and SOC 2 where those are, or will become, customer questions.

Watch. Put one hub in the middle that every AI tool, model, agent and vendor connects through, so the usage log runs against the rules and exceptions are visible when they happen, not a year later. This is the part most businesses skip, and the part that turns policy into evidence and lets the business say yes to more.

Measure. Report to the board monthly on what AI is in use, what it costs, whether it is behaving and whether it is delivering. One page.

None of this needs a large team. It needs a decision that AI is an operation the business runs, not something that happens to it. The table below gives ten questions a board can put to the executive team this quarter. Each has a yes-or-no answer, and each "no" is a control gap.

Table 2. AI control health check: ten questions the board can answer today

#	Question	Yes / No
1	Do we have a complete list of every AI tool in use across the business, including free and personal-account tools?	
2	Do we know what categories of data are going into each of those tools, and which of it is personal or commercially sensitive?	
3	Do we have a written AI policy that names approved tools, prohibited data and required sign-offs?	
4	Is that policy enforced by a control or a log, rather than relying on staff having read it?	
5	Have all staff received AI training in the last twelve months, and can we evidence who attended?	
6	Do we know where AI is making or shaping decisions about customers or employees, and what human review applies?	
7	Can we show a regulator, within a week, who used which AI tool with what data on a given date?	

#	Question	Yes / No
8	Do we know our total AI spend across licences, cloud usage, vendor add-ons and expenses?	
9	Have we assessed whether any of our AI use falls under the EU AI Act or sector-specific rules?	
10	Does the board receive a regular report on AI use, cost, risk and results?	

A business answering "yes" to all ten has closed the control gap. Most UK mid-market businesses today would manage fewer than three. That is a reason to start, not for alarm; the starting point is cheap and quick.

What to do next

Sequence matters. Businesses that start by buying tools, or writing rules for an estate they have not seen, spend money on the wrong problem. The order that works is the order above, and it is how GenWA works with clients.

Start with an Independent AI Audit. In three weeks, an independent review establishes every AI tool in use, what data goes in and out, what it costs and where the exposures are, delivered as a board-ready report with a risk register and a 90-day plan. It answers questions one, two, six, eight and nine in one piece of work.

Then set the rules with AI Policy & Controls. A policy written for a UK business, with access controls, sign-off rules, training and an evidence pack to show it is enforced, aligned to UK GDPR, ISO 42001 and SOC 2. This covers questions three, four, five and seven.

Then connect it all through Business Watch. Business Watch is the central hub for all your AI systems: the one place every AI tool, model, agent and vendor connects through, gets its rules, is costed and is measured. It shows usage against your rules, data in and out, spend by tool and team and behaviour alerts, and it produces a monthly board report with risk and return side by side. This is what turns policy into policing, answers question ten every month, and gives the board the confidence to approve more AI, not less.

Build with control only when the first three are in place. When the business needs AI to do real work, in customer communications, case handling or the contact centre, it should be built auditable, with humans in the loop by design.

If the board can answer fewer than five of the ten questions with confidence, the next step is a 30-minute conversation about the audit. Write to hello@genwa.co.uk.

Sources

1. Microsoft UK / Censuswide, "Rise in shadow AI tools raising security concerns for UK", November 2025.
<https://ukstories.microsoft.com/features/rise-in-shadow-ai-tools-raising-security-concerns-for-uk/>
2. Office for National Statistics, "Artificial intelligence in UK businesses: 2023 to 2026", July 2026.
<https://www.ons.gov.uk/businessindustryandtrade/business/businessservices/articles/artificialintelligenceinukbusinesses/2023to2026>
3. British Chambers of Commerce, "Half of SMEs using AI with limited headcount impact so far", March 2026.
<https://www.britishchambers.org.uk/news/2026/03/half-of-smes-using-ai-with-limited-headcount-impact-so-far/>
4. MIT NANDA, "The GenAI Divide", August 2025 (via Virtualization Review).
<https://virtualizationreview.com/articles/2025/08/19/mit-report-finds-most-ai-business-investments-fail-reveals-genai-divide.aspx>
5. Netskope, Cloud and Threat Report 2026, January 2026. <https://www.netskope.com/resources/cloud-and-threat-reports/cloud-and-threat-report-2026>
6. Zluri, enterprise AI tools without IT oversight, June 2025.
<https://www.businesswire.com/news/home/20250618738949/en/>
7. CIPD, Labour Market Outlook Autumn 2025 and blog "Can AI deliver the productivity boost the UK needs?", January 2026. <https://www.cipd.org/en/about/blogs/can-ai-deliver-productivity-boost-uk-needs/>
8. IBM, Cost of a Data Breach Report 2025, UK findings, July 2025. <https://uk.newsroom.ibm.com/2025-cost-of-data-breach-UK>
9. Lloyds Bank Business Barometer, August 2026.
<https://www.lloydsbankinggroup.com/media/press-releases/2026/lloyds/more-than-half-of-uk-businesses-say-ai-has-created-new-jobs.html>
10. Cyberhaven, "Majority of corporate AI tools present critical data security risks", April 2025.
<https://www.cyberhaven.com/press-releases/cyberhaven-report-majority-of-corporate-ai-tools-present-critical-data-security-risks>
11. IBM, Cost of a Data Breach Report 2026, global findings, August 2026 (via Northdoor).
<https://www.northdoor.co.uk/insight/blog/cost-of-a-data-breach-2026-global-findings/>
12. Gartner, "40% of organisations will suffer incidents from unauthorised AI by 2030", November 2025 (via Infosecurity Magazine). <https://www.infosecurity-magazine.com/news/gartner-40-firms-hit-shadow-ai/>
13. CNBC, "Samsung bans use of AI like ChatGPT for staff after misuse of chatbot", May 2023.
<https://www.cnbc.com/2023/05/02/samsung-bans-use-of-ai-like-chatgpt-for-staff-after-misuse-of-chatbot.html>
14. Trowers & Hamlin, "A round-up of ICO enforcement in 2025", February 2026.
<https://www.trowers.com/insights/2026/february/a-round-up-of-ico-enforcement-in-2025>
15. Department for Science, Innovation and Technology, Cyber Security Breaches Survey 2025/26, April 2026.
<https://www.gov.uk/government/statistics/cyber-security-breaches-survey-20252026/cyber-security-breaches-survey-20252026>
16. Deloitte, Governance of AI: a critical imperative for today's boards, 2nd edition, 2025.
<https://www.deloitte.com/global/en/issues/trust/progress-on-ai-in-the-boardroom-but-room-to-accelerate.html>
17. IAPP, "King's Speech signals diffuse UK digital policy agenda but no AI bill", May 2026.
<https://iapp.org/news/a/king-s-speech-signals-diffuse-uk-digital-policy-agenda-but-no-ai-bill>
18. DLA Piper, "UK: commencement of the data protection provisions in the Data (Use and Access) Act", February 2026. <https://privacymatters.dlapiper.com/2026/02/uk-commencement-of-the-data-protection-provisions-in-the-data-use-and-access-act/>

19. Bratby Law, "ICO AI and biometrics code of practice" (SI 2026/425), 2026. <https://bratby.law/ico-ai-biometrics-code-of-practice/>
20. Information Commissioner's Office, AI and biometrics strategy update, March 2026. <https://ico.org.uk/about-the-ico/our-information/our-strategies-and-plans/artificial-intelligence-and-biometrics-strategy/ai-and-biometrics-strategy-update-march-2026/>
21. Gibson Dunn / Pinsent Masons, "EU AI Act omnibus agreement: postponed high-risk deadlines and other key changes", May 2026. <https://www.gibsondunn.com/eu-ai-act-omnibus-agreement-postponed-high-risk-deadlines-and-other-key-changes/>
22. UK Parliament, Cyber Security and Resilience Bill, 2026 (secondary reporting; confirm fine levels on [parliament.uk](https://bills.parliament.uk)). <https://bills.parliament.uk/>
23. UKAS, "UKAS grants first AIMS (ISO/IEC 42001) accreditation", January 2026. <https://www.ukas.com/resources/latest-news/ukas-grants-first-aims-accreditation/>
24. National Cyber Security Centre, guidelines for secure AI system development, including autonomous AI agents, August 2026. <https://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development>
25. Gartner, inaugural Magic Quadrant for AI Governance Platforms, June 2026 (via Sanjeev Mohan, Medium). <https://sanjmo.medium.com/inside-gartners-first-ai-governance-platform-magic-quadrant-fa1f182f75e9>

About GenWA

All the upside of AI. None of the surprises.

GenWA helps UK mid-market businesses take the upside of AI with confidence. We show you what AI is doing in your organisation, put business rules around it and make it pay. Our services are the Independent AI Audit, AI Policy & Controls, Business Watch (the central hub for all your AI systems), Build with control and Managed AI Ops. GenWA is an AWS Marketplace seller and our work is aligned to UK GDPR, ISO 42001 and SOC 2; we prepare you for certification.

John White, Founder & CEO. 30+ years driving digital transformation across financial services and fintech. Former CTO at Vanquis Bank (FTSE 250), Chetwood Bank, and Arrow Global/Intrum. Head of Solutions at Capital One UK, establishing their first international operation. Led major transformation initiatives at Barclays/Barclaycard. Previously CEO (SMF1) at Shoreditch fintech, achieving regulatory approval and launching credit products from concept to market in nine months.

Lukasz Piotrowski, CTO & Co-founder. 20+ years building and scaling engineering organisations across financial services and capital markets. Former Head of Engineering at Intrum UK and Chetwood Financial, delivering AWS-native contact centre, collections, and cloud-native banking platforms at scale. Deep capital markets experience from Numis Securities.

Contact: hello@genwa.co.uk · genwa.co.uk · linkedin.com/company/genwa