

WHITE PAPER · COMPLIANCE

Rules as controls

What UK regulators now expect of a business that uses AI, why a policy on the intranet does not meet it, and how one gateway does.

John White, GenWA

SUMMARY

A policy is a document. A control is a system.

Most UK businesses that have an AI policy have a page on the intranet. The regulators, the courts and the insurers have moved on from that. What they now ask is whether the rule was applied at the point of use, and whether you can show it.

This paper sets out the rules already in force for a UK business using AI on people's data and in decisions, in plain English; what the regulators say they expect in practice, from the ICO's finding that "clicking approve on an AI-generated ranking" is not human involvement to the High Court's statement that a lawyer owns what AI drafts; why the gap between policy and control is where the risk sits; and how a gateway, one door between the business and every AI provider, turns each rule into something that happens on every request and leaves a record.

IN SHORT

Only 31% of UK employers have a formal policy on generative AI (CIPD, 2025) and 63% of UK organisations lack AI access controls (IBM, 2025). The Data (Use and Access) Act's safeguards for automated decisions have been in force since 5 February 2026; the ICO's statutory AI code is required from 12 May 2026; the CMA has said consumer law applies "regardless of whether that is through people or AI systems"; the SRA says AI "has no separate legal personality"; and professional indemnity insurers now ask for an AI policy, framework, training and monitoring evidence at renewal. A rule that exists only as text satisfies none of them. A rule enforced on every request, with a log, satisfies all of them at once.

31%

of UK employers have a formal policy on generative AI.

CIPD, 2025

63%

of UK organisations lack AI access controls.

IBM, 2025 · vendor research

16

employers written to by the ICO in March 2026 over automated recruitment decisions.

ICO, Recruitment Rewired

01 · THE RULES IN FORCE

Already law, already in force



UK GDPR and the Data (Use and Access) Act

2025. The Act's main provisions commenced on 5 February 2026. Section 80 widens the scope for automated decision-making on non-special-category data, but attaches mandatory safeguards: the person must be told, must be able to contest the decision, and must be able to obtain meaningful human intervention. Any business using AI to make or shape decisions about individuals, in recruitment, credit, claims, pricing or service, is already inside these rules.

The ICO's statutory code. A statutory Code of Practice on AI and automated decision-making is required by SI 2026/425, in force from 12 May 2026, with final guidance due in winter 2026. The ICO's enforcement posture is set: eight UK GDPR penalties totalling £32.4 million in 2025/26, including £14 million against Capita; investigations into a foundation-model developer; engagement with eleven others (ICO Annual Report 2025/26). The average fine rose to £1.45 million (URM Consulting, January 2026).

ISO/IEC 42001. The management-system standard for AI is becoming the buyer's and insurer's question. About 350 organisations were certified worldwide by April 2026; BSI became the first UKAS-accredited certifier in January 2026.

Consumer law. "UK consumer law applies whether decisions are made by people or by AI... businesses are responsible... regardless of whether that is through people or AI systems" (CMA, Agentic AI and consumers, March 2026). A tribunal reached the same conclusion about Air Canada's chatbot in 2024.

Sector regulators. The PRA's model risk standard SS1/23, updated April 2026, applies to AI and machine learning and puts responsibility with a named senior manager. The FCA will publish Consumer Duty guidance on AI by the end of 2026. The SRA's warning notice of August 2026 says solicitors "remain accountable for their work and outputs, regardless of how that work has been prepared". The FRC published the first audit-regulator guidance on generative and agentic AI in March 2026. The CQC says AI "can enhance, but not replace human decision making". The MHRA has ruled that ambient voice tools that support diagnosis are medical devices.

Employment. The Fair Work Agency has had powers to inspect electronic records since April 2026; a consultation on workplace surveillance technology closes this month.

The EU AI Act applies to any UK business serving EU customers; high-risk obligations begin in August 2026, and a deployer who substantially modifies a high-risk system becomes its provider (Article 25).

02 · WHAT REGULATORS EXPECT

What “meaningful human involvement” actually means


The regulators have stopped talking in principles. They are describing behaviour, and the behaviour they describe is a control, not a document.

WHO	WHAT THEY SAID	WHAT IT REQUIRES IN PRACTICE
ICO	“Clicking approve on an AI-generated ranking without independent analysis” is not meaningful human involvement. Many employers using automated recruitment “are likely relying on solely automated decisions” (Recruitment Rewired, March 2026).	Decisions about people are held for a named reviewer who sees the inputs, not just the output, and whose review is logged.
ICO	Organisations “remain responsible for their use of agentic AI systems”; hallucinations could “cascade across tools, databases and other agents”; suggests “a standalone monitoring system to monitor logs, interpret them, and intervene” (Tech Futures, January 2026).	Every agent has an identity, a scope, a log and a kill-switch, monitored by something that is not the agent.
NCSC	Scoped credentials, dedicated identities, logging, human oversight and a kill-switch for AI agents (August 2026).	Agents run on the same control layer as people, with narrower permissions.
High Court	“Those who use artificial intelligence to conduct legal research... have a professional duty therefore to check the accuracy of such research by reference to authoritative sources” (Ayinde v Haringey, June 2025).	Evidence produced with AI carries a verification step before release, and the verifier’s name.
SRA	Putting client material into an open tool “is to place this information on the internet in the public domain”; privilege may be permanently waived (August 2026).	Client material, personal data and privileged content are detected and stopped before a request leaves the business.
PRA	Model risk management applies to AI and machine learning; a named senior manager is responsible (SS1/23, April 2026).	An approved model list, an owner, and a change process when a model is retired, repriced or updated.
CMA	Businesses are responsible for what AI says and does to consumers (March 2026).	Customer-facing agents constrained to approved sources; money, eligibility and deadlines held for a human.
Insurers	An AI policy, risk framework, training and monitoring evidence required at renewal; AI-related payouts capped by some London-market carriers (Lockton, March 2026; Fenwick, June 2026).	Evidence generated from the record, not assembled the week before renewal.

Only 28% of financial-services compliance functions that use AI test or validate its outputs; 24% have policies for third-party AI (ACA Group/NSCP, October 2025). Only 14% of chief information officers have defined at executive level who is accountable for AI results (Logicalis, 2026, secondary).

03 · THE GAP

Where the risk sits



CIPD, 2025; Microsoft UK, 2025 (vendor survey); IBM, 2025 (vendor research, inverse of 63% lacking controls); DSIT Cyber Security Breaches Survey, April 2026; Cloud Security Alliance, April 2026.

The gap is not between businesses that have thought about AI and those that have not. It is between the rule and the request. A business can have a two-page policy, a training module and a signed acknowledgement from every employee, and still have source code in a consumer chatbot, a candidate ranked by an algorithm nobody reviewed, and an agent with production access and no budget. The policy is true. It is simply not in the path.

The incidents bear this out. Samsung’s engineers broke no rule; there was none in the way. Air Canada’s chatbot was not malicious; it was unconstrained. The solicitor in AI-Haroun did not intend to mislead the court; nothing in his workflow required a citation to resolve to a source. Uber’s engineers did what they were told; nothing stopped at a limit. In each case the control that would have caught it is mundane, cheap, and absent.

Two-thirds of boards have limited or no AI knowledge, and at a third AI is not on the agenda at all.

Deloitte, 2025.

The regulators know this, which is why their language has shifted from principles to behaviour. The ICO did not ask the sixteen employers whether they had a policy. It asked whether a person had looked at the inputs before approving the output. The High Court did not ask whether the firm had guidance on AI. It asked whether anyone had checked the citations. Gartner expects that by 2027, 75% of regulated organisations relying on manual AI compliance processes will face fines exceeding 5% of global revenue (Gartner, March 2026). Whether or not that figure proves right, the direction is clear: the question is what happened on the request, and whether you can show it.

One door. Every request. Your rules.



A gateway is a single control point between everything in the business that uses AI, people, agents and systems, and every provider of it, the labs, open-source models on your own cloud, and AI features inside software you already own. Every request passes through it. On each one the gateway **identifies** who or what is asking and on which account; **checks** the content against your rules, personal data, client material, special-category data, source code, privileged content; **chooses** the model that satisfies the task and the policy at the lowest cost; **records** the request, the decision and the outcome; and **stops** what should not happen, saying why.

The rules are written about tasks, data and risk, not about a vendor. That is what makes them durable. When a model is retired, repriced or changed, the rule stands and the routing changes, under a change process with a test and a record, which is what the PRA's model risk standard and the EU AI Act's "substantial modification" provision both describe.

SIX RULES, AS CONTROLS

THE RULE	THE CONTROL
No customer personal data in consumer AI tools	Detected and blocked or redacted before the request leaves
No client, privileged or source material in open tools	Classified content stopped at the gateway; approved tool on a company account offered instead
Decisions about individuals need a human	Outputs tagged "decision" held for a named reviewer who sees the inputs; review logged
Evidence must be verified before release	Citations resolve to a source; the verifier's name in the record
Customer-facing agents state only approved policy	Constrained to an approved source; money, eligibility and deadlines held for a human
Agents act within a budget and a scope	Own identity, scoped credentials, spend cap, action log, kill-switch

From the record to the regulator

The same log that enforces the rules is the evidence that you did. Generated from the record, mapped to the clause, produced on request.

WHAT THE RECORD HOLDS

For every request: who asked, on which account, what class of data it contained, which rule applied, what the gateway did, which model answered, what it cost, and, where a human reviewed it, who and when. For every agent: its identity, its scope, its budget, every action it took and every action it was stopped from taking. For every model: when it was approved, by whom, when it changed, what was re-tested and with what result.

WHAT THE EVIDENCE PACK MAPS IT TO

FRAMEWORK	WHAT THE PACK ANSWERS
UK GDPR and DUAA s.80	Lawful basis, DPIA generated from the register and rules, the safeguards for automated decisions and the human-intervention log
ICO AI and ADM code	Transparency, contestability and meaningful human involvement, evidenced per decision
ISO/IEC 42001	The AI management system: inventory, risk assessment, controls, monitoring, change, with clause-by-clause evidence
PRA SS1/23	Model inventory, ownership, validation and change control for AI and ML models
NCSC agent guidance	Identity, scope, logging, oversight and kill-switch, per agent
Sector codes: SRA, FRC, CQC, PCRT	Verification steps and named reviewers for AI-assisted professional work

WHO ASKS, AND WHAT FOR

The regulator. The ICO’s first question in every AI matter so far has been whether the risk was assessed before the system was switched on (Snap, 2023) and whether a human was involved in a way that means something (Recruitment Rewired, 2026). A DPIA generated from the register answers the first; the review log answers the second.

The insurer. Renewal now asks for the policy, the framework, the training record and the monitoring evidence. Beazley and QBE have capped AI-related payouts at around 10% of limits; the answers determine terms.

The customer. Public buyers have required suppliers to declare AI use since February 2025 (PPN 017); 43% of financial-services firms do enhanced due diligence on AI vendors (ACA Group, 2025). ISO 42001 is becoming the tender question.

The court. Air Canada lost because it could not show what its chatbot had been told to say. Al-Haroun’s solicitor was referred because nothing showed the citations had been checked. The record is the defence.

£32.4m

in ICO penalties in 2025/26, across eight cases.

ICO Annual Report 2025/26

From policy to control in one quarter

1

Weeks 1–3 · Know what is running.

The register: every tool, account, agent and AI feature, with an owner and the data it touches. 82% of enterprises have agents they do not know about (CSA, 2026). You cannot assess what you have not listed.

2

Weeks 4–5 · Write the rules about tasks, data and risk.

Two pages at most. What may be done, with which class of data, with what human check, and where AI must not be used. Align to DUAA s.80, the ICO code and the sector regulator; note where ISO 42001 will ask for more.

3

Weeks 6–8 · Put the rules in the path.

The gateway in front of every tool, the six controls on page 6 switched on, company accounts for everyone, personal accounts migrated. A policy that lives only on the intranet is a document.

4

Weeks 9–10 · Name the humans.

A reviewer for every class of decision about people, a verifier for every class of evidence, an owner for every agent. Their names in the record.

5

Weeks 11–12 · Generate the evidence.

The first DPIA from the register, the first evidence pack mapped to the frameworks, the first board report. Then ask: if the ICO, the insurer or a customer wrote tomorrow, what would we send them? The answer should already exist.

WHERE YOUR BUSINESS STANDS

The AI health check scores your business on seven questions, four of them about control and evidence. **Ten minutes online, 48 hours with your usage data, or half a day on site.** No obligation.

genwa.co.uk/#/healthcheck

75%

of regulated organisations relying on manual AI compliance processes will face fines above 5% of revenue by 2027, Gartner predicts.

Gartner, March 2026

Related reading: The AI Control Playbook; Case notes: eight AI incidents; The AI Control Gap.

Sources

Official statistics, regulators and courts first; vendor-commissioned surveys and vendor telemetry are marked as such.

1. CIPD, Labour Market Outlook, Autumn 2025.
2. IBM, Cost of a Data Breach Report UK, 2025 (vendor research).
3. ICO, Recruitment Rewired, 31 March 2026.
4. Data (Use and Access) Act 2025, s.80; DLA Piper commencement note, 2026.
5. SI 2026/425; ICO, AI and automated decision-making code of practice, 2026.
6. ICO, Annual Report 2025/26; URM Consulting, ICO enforcement review, January 2026.
7. UKAS, ISO/IEC 42001 accreditation, January 2026.
8. Competition and Markets Authority, "Agentic AI and consumers", 9 March 2026.
9. *Moffatt v Air Canada*, 2024 BCCRT 149.
10. Bank of England / PRA, SS1/23 Model Risk Management, updated April 2026; PRA roundtable note, November 2025.
11. House of Commons Treasury Committee, 20 January 2026; FCA, 2026.
12. Solicitors Regulation Authority, "Misuse of AI" warning notice, 17 August 2026.
13. Financial Reporting Council, guidance on generative and agentic AI in audit, 30 March 2026.
14. Care Quality Commission, statement on AI, 21 May 2026; MHRA and NHS England, 29 July 2026.
15. Lewis Silkin, Employment Rights Act 2025 note, September 2026 (secondary).
1. EU AI Act, Articles 14 and 25.
2. ICO, Tech Futures: Agentic AI, January 2026.
3. NCSC, guidance on AI agents, August 2026.
4. *Ayinde v London Borough of Haringey*; *Al-Haroun v Qatar National Bank* [2025] EWHC 1383 (Admin), 6 June 2025.
5. Legal Futures / Lockton, 27 March 2026; Fenwick, June 2026; Insurance Business, July 2026 (secondary).
6. ACA Group / NSCP, AI Benchmarking Survey, 29 October 2025.
7. Logicalis, CIO Report 2026 (secondary).
8. Microsoft UK / Censuswide, 2025 (vendor survey).
9. DSIT, Cyber Security Breaches Survey 2025/26, April 2026.
10. Cloud Security Alliance / Zenity, 16 April 2026; CSA, 21 April 2026.
11. Deloitte, Board practices and AI, 2025.
12. Gartner, 17 March 2026.
13. ICO, conclusion of Snap "My AI" investigation, May 2024.
14. Cabinet Office, PPN 017, February 2025.
15. CNBC, 2 May 2023; Bloomberg, 2 June 2026.

© GenWA Ltd 2026. This white paper is general guidance and not legal advice. Figures are attributed to their sources and were correct to the best of our knowledge at publication. Readers should check current guidance from the ICO, the FCA, the PRA, the SRA, the FRC, the CQC and the NCSC before relying on regulatory statements.



All the upside of AI. None of the surprises.

[GENWA.CO.UK](https://genwa.co.uk)