

FOR UK BOARDS AND EXECUTIVES

Case notes

Eight AI incidents, what each one cost, and the control that would have caught it.

Compiled by GenWA from public records. GenWA had no involvement in any incident described.

INTRODUCTION

None of these were technology failures

In every case the AI did roughly what it was built to do. What failed was the business around it: no register, no rule in the path of the request, no meter, no record, no named owner. Those are the things a board can fix.

We have chosen incidents that are documented in court records, regulators' notices or on-the-record reporting, and that could happen to a UK business of 50 to 2,000 people. Each note gives what happened, what it cost, why it happened in plain English, and the control that would have caught it, mapped to the five moves in The AI Control Playbook: See, Cut, Contain, Prove, Scale.

A word on the money. The fines and refunds are small: C\$812 for Air Canada, roughly A\$97,000 for Deloitte, a few thousand pounds in wasted costs at the High Court. The real costs were a precedent, a headline, a regulator's attention and, in Samsung's case, trade secrets that cannot be recalled. The ICO's average fine rose to £1.45 million in 2025 (URM Consulting, January 2026); the UK average cost of a data breach is £3.78 million, plus £498,000 where shadow AI use is high (IBM, 2025, vendor research).

THE EIGHT AT A GLANCE

WHO	WHEN	WHAT HAPPENED	WHAT IT COST	MOVE
Samsung	Mar–May 2023	Source code and meeting notes pasted into ChatGPT	Company-wide ban; trade secrets in a third party's training data	See · Contain
Air Canada	Feb 2024	Chatbot invented a bereavement-fare policy	Tribunal: the company is liable for what its chatbot says	Contain · Prove
Deloitte Australia	Oct 2025	AI-drafted government report with fabricated citations	Partial refund of a A\$440k fee; public correction	Contain · Prove
UK High Court	Jun 2025	Lawyers filed AI-invented case law	Referrals to the SRA and Bar Standards Board; wasted-costs orders	Contain · Prove
Uber	Apr–Jun 2026	Annual AI tooling budget spent in four months	\$1,500 monthly cap per employee, per tool	See · Cut
Cursor users	Jun–Jul 2025	Plan moved to consumption pricing; allowances gone in days	Vendor apology and refunds	See · Cut
Replit / SaaS	Jul 2025	Coding agent deleted a production database during a code freeze	Refund; vendor rebuilt its safeguards	Contain · Scale
Snap; Serco Leisure	2023–24	A chatbot launched without a risk assessment; biometrics used on staff without a lawful basis	ICO preliminary enforcement notice; ICO order to stop and delete	Prove

CASE NOTE · DATA LEAVING THE BUSINESS

Three engineers, twenty days, and the source code was gone

20 days

from permitting a consumer AI tool to the first leak

3

separate incidents: source code, test code, meeting minutes

£0

fine. The cost was trade secrets, and a company-wide ban

WHAT HAPPENED

In March 2023 Samsung's semiconductor division allowed engineers to use ChatGPT. Within about twenty days three of them had pasted confidential material into the public service: proprietary source code, to find a bug; the code for a chip test sequence, to optimise it; and the transcript of an internal meeting, to generate minutes. At the time the service used conversations for training by default and offered no straightforward way to delete them. Disciplinary investigations followed. In May Samsung banned generative AI tools on company devices and networks and told staff not to enter "any personal or company related information" into them (CNBC, May 2023). An internal survey found 65% of staff concerned about the security risk.

Nobody broke a rule, because there was no rule. Nobody bypassed a control, because there was none. The tool was the easiest thing available and the work was real.

WHY IT MATTERS HERE

This is how most AI enters a UK business: 71% of UK employees have used AI tools their employer never approved and 28% say there is no approved alternative (Microsoft UK, 2025, vendor survey). Roughly a third of the data entering AI tools is sensitive (Cyberhaven, 2025, vendor telemetry). The UK's solicitors' regulator now says that putting client material into an open tool "is to place this information on the internet in the public domain", and that privilege may be permanently waived (SRA, August 2026). A ban is the wrong answer: it removes the visibility without removing the use.

THE CONTROL THAT WOULD HAVE CAUGHT IT

See	A register would have shown the tool in use on day one, on which accounts, by whom.
------------	---

Contain	A gateway detects source code, client material and personal data in a request and blocks or redacts it before it leaves. The rule exists as a control, not a memo. An approved tool on a company account removes the reason to use the personal one.
----------------	--

Prove	The record shows what was tried, what was stopped and why.
--------------	--

Sources. CNBC, "Samsung bans use of A.I. like ChatGPT for employees after misuse of the chatbot", 2 May 2023; Forbes, 2 May 2023; The Economist Korea via CS Hub, April 2023; SRA, "Misuse of AI" warning notice, 17 August 2026. These are public incidents, reported by others; GenWA had no involvement in any of them.

CASE NOTE · LIABILITY FOR OUTPUT

The chatbot is not “a separate legal entity”

C\$812

total award: damages, interest and fees

90 days

the refund window the chatbot invented

1 line

of the ruling that every business now quotes

WHAT HAPPENED

In November 2022 Jake Moffatt booked flights to his grandmother's funeral after Air Canada's website chatbot told him he could claim a bereavement fare retrospectively within 90 days. The airline's actual policy, on a page the chatbot linked to, excluded retrospective claims. Air Canada refused the refund and, before the British Columbia Civil Resolution Tribunal, argued that the chatbot was a separate legal entity responsible for its own actions. The tribunal rejected that: "While a chatbot has an interactive component, it is still just a part of Air Canada's website. It should be obvious that it is responsible for all the information on its website" (Moffatt v Air Canada, 2024 BCCRT 149, paragraph 27). It found negligent misrepresentation and awarded C\$650.88 in damages plus interest and fees.

WHY IT MATTERS HERE

The money was trivial. The precedent was not, and UK regulators have since said the same thing in their own words. The Competition and Markets Authority: "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, March 2026). The FCA is due to publish guidance on the Consumer Duty and AI by the end of 2026 at the Treasury Committee's request. Only 27% of UK consumers trust AI assistants for financial guidance (FCA-commissioned review, July 2026). A customer-facing agent that cannot be held to the company's own policy is a liability with a friendly voice.

THE CONTROL THAT WOULD HAVE CAUGHT IT

Contain Customer-facing outputs run through rules: the agent may only state policy that exists in an approved source, and anything about money, eligibility or deadlines is held for a human check.

Prove Every conversation is on the record with the source the agent relied on, so a dispute is settled from the log rather than argued in front of a tribunal.

Scale Deploy the agent with a scope and a stop, not a personality.

Sources. Moffatt v Air Canada, 2024 BCCRT 149, 14 February 2024; American Bar Association, Business Law Today, February 2024; McCarthy Tétrault, 2024; CMA, "Agentic AI and consumers", 9 March 2026; Treasury Committee, 20 January 2026. These are public incidents, reported by others; GenWA had no involvement in any of them.

CASE NOTE · FABRICATED EVIDENCE

A A\$440,000 report with citations that did not exist

A\$440k

contract value for an independent assurance review

≈A\$97k

repaid: the final instalment

12+

fabricated references, including a misquoted court judgment

WHAT HAPPENED

In July 2025 Deloitte delivered a 237-page independent assurance review of a welfare-compliance IT system to Australia's Department of Employment and Workplace Relations, under a contract worth about A\$440,000. A University of Sydney academic found more than a dozen fabricated references: academic papers that did not exist, a book attributed to a real professor who never wrote it, and an invented passage from the Federal Court's robodebt judgment. A revised version in September 2025 added a disclosure that Azure OpenAI GPT-4o had been used. Deloitte said the corrections did not affect the substantive findings; the department confirmed that "some footnotes and references were incorrect". Deloitte agreed to repay the final instalment, reported at roughly A\$97,000 (CFO Dive, October 2025).

WHY IT MATTERS HERE

The report was about compliance, from a firm that sells AI advice, to a government client. The cost was not the refund; it was the story. In the UK the same failure has already reached the courts (next page), and the professional bodies have moved: the FRC's guidance on generative and agentic AI in audit (March 2026) says auditors remain accountable for AI-assisted work; seven tax and accountancy bodies issued joint guidance in January 2026 requiring members to verify outputs. Professional indemnity insurers now ask for an AI policy, risk framework, training and monitoring evidence at renewal (Lockton via Legal Futures, March 2026).

THE CONTROL THAT WOULD HAVE CAUGHT IT

Contain A rule that AI-drafted client deliverables carry a verification step before release, enforced in the workflow: citations checked against a source, the checker named.

Prove Disclosure generated from the record: which model, which prompts, which human checks, so the client is told before they ask.

Advance A policy that says what AI may draft, what it may not, and what "checked" means, written with the fee-earners who will live by it.

Sources. CFO Dive, 21 October 2025; Fortune, 7 October 2025; Accounting Times, October 2025; FRC, 30 March 2026; CIOT and six other PCRT bodies, 19 January 2026; Legal Futures, 27 March 2026. These are public incidents, reported by others; GenWA had no involvement in any of them.

CASE NOTE · PROFESSIONAL DUTY

Eighteen of forty-five citations did not exist

45

authorities cited in one set of witness statements

18

of them invented; many others misrepresented

2 years

the maximum prison term the court listed for contempt

WHAT HAPPENED

On 6 June 2025 the Divisional Court, led by the President of the King's Bench Division, handed down judgment in two referrals heard together. In *Ayinde v London Borough of Haringey* a barrister's judicial-review grounds cited five cases that did not exist and misstated a section of the Housing Act; wasted-costs orders of £2,000 each had already been made. In *Al-Haroun v Qatar National Bank* a solicitor filed witness statements with 45 citations, 18 of which did not exist and many of which were misrepresented; the client had generated them with public AI tools and the solicitor had not checked. The court found the threshold for contempt was not met, referred both to their regulators, and set out the sanctions available: admonishment, costs, strike-out, regulatory referral, contempt and referral to the police.

WHY IT MATTERS HERE

The judgment is now the UK's clearest statement of who owns AI output. "Those who use artificial intelligence to conduct legal research notwithstanding these risks have a professional duty therefore to check the accuracy of such research by reference to authoritative sources" (paragraph 7). "A lawyer is not entitled to rely on their lay client for the accuracy of citations" (paragraph 81). The SRA's warning notice of August 2026 extends the principle to every regulated individual: AI "has no separate legal personality" and practitioners "remain accountable for their work and outputs, regardless of how that work has been prepared". 26% of UK law firms now use generative AI, up from 14% a year earlier (Thomson Reuters Institute, 2025); 15% of large-firm lawyers describe themselves as dependent on it (LexisNexis UK, September 2026).

THE CONTROL THAT WOULD HAVE CAUGHT IT

Contain	Research and drafting tools run through the gateway with a rule: authorities must resolve to a verified source before the document can be finalised.
Prove	A per-document record of what AI produced and who verified it, which is what the SRA, the insurer and the court will ask for.
See	Know which tools fee-earners are using, on which accounts. 47% of firms say individual fee-earners are responsible for managing their own AI use (Lockton, 2026).

Sources. *Ayinde v London Borough of Haringey*; *Al-Haroun v Qatar National Bank* [2025] EWHC 1383 (Admin), 6 June 2025 (judiciary.uk); LexisNexis UK, June 2025; DAC Beachcroft, 2025; SRA warning notice, 17 August 2026; Legal Futures, 27 March 2026. These are public incidents, reported by others; GenWA had no involvement in any of them.

CASE NOTE · OVERSPEND

The budget for the year lasted four months

4 months

to spend a full year's AI tooling budget

\$1,500

the monthly cap per employee, per tool, imposed afterwards

“as much as possible”

the instruction staff had been given

WHAT HAPPENED

Uber told its engineers to use AI coding tools as much as possible and ran internal leaderboards to encourage it. The tools were priced by consumption. By April 2026 the company had spent its entire annual budget for them. In June it capped every employee at \$1,500 a month per agentic coding tool, added dashboards and an approval route for exceptions (Bloomberg, 2 June 2026; TechCrunch, 2 June 2026). Its chief operating officer said it was “very hard to draw a line” from AI usage to new customer features.

Nothing went wrong with the tools. The tokens bought real work. What was missing was a meter anyone could see, a limit anyone had set, and a measure of return anyone had agreed.

WHY IT MATTERS HERE

Uber is unusual only in being reported. Across Ramp's customer base token usage rose 1,001% in fifteen months while spend rose 497% (Ramp, April 2026, vendor data). 62% of organisations have had unexpected AI costs large enough to change a business decision and a third imposed emergency freezes (Mavvrik, 2026, vendor survey). Only 26% of companies fully track AI costs (KPMG, 2026). In the UK, 42% of leaders have only partial visibility of AI spend and 33% do not understand AI cost structures, tokens included (KPMG UK, July 2026). Consumption pricing with encouragement and no control is a budget with no bottom.

THE CONTROL THAT WOULD HAVE CAUGHT IT

See	A live meter by team, tool and person, with the month's total against budget, visible before the invoice.
Cut	Routing sends routine requests to the cheapest adequate model; most of a coding tool's tokens are routine. Per-person and per-agent caps pause work at the limit rather than bill over it.
Prove	Three use cases under measurement, each with a number, a review date and a stop criterion, so “hard to draw a line” becomes a line.

Sources. Bloomberg, “Uber Caps Employee Spending on AI Tools Like Claude Code to Manage Costs”, 2 June 2026; TechCrunch, 2 June 2026; Simon Willison, 3 June 2026; Ramp AI Index, April 2026; Mavvrik, 2026; KPMG UK, July 2026. These are public incidents, reported by others; GenWA had no involvement in any of them.

CASE NOTE · PRICING CHANGE

The plan changed on a Monday. The allowance was gone by Friday.

16 June 2025

the day the pricing model changed

\$20

a month, now a pool of model-priced credit

“we’re sorry”

the chief executive, three weeks later

WHAT HAPPENED

On 16 June 2025 Cursor, one of the most widely used AI coding tools, changed its \$20-a-month Pro plan from a fixed number of requests to a \$20 pool of credit billed at each model’s rate. Users who kept working as before, particularly on the most expensive models, exhausted the allowance within days; others, with no spending cap set, received overage charges they had not expected. On 7 July the chief executive wrote: “We recognize that we didn’t handle this pricing rollout well and we’re sorry. Our communication was not clear enough,” and the company refunded unexpected charges (TechCrunch, July 2025; Cursor, July 2025).

WHY IT MATTERS HERE

The lesson is not about one vendor. Every AI supplier can change price, limits and terms unilaterally and at short notice: Anthropic introduced weekly rate limits in August 2025; OpenAI retired four models from its consumer product in February 2026 with about two weeks’ notice; Anthropic’s deprecation schedule retired six models between February and August 2026. Newer models also consume more tokens per request for the same task, so the same work costs more without anyone changing anything. A business whose spend depends on a vendor’s pricing page has outsourced its budget.

THE CONTROL THAT WOULD HAVE CAUGHT IT

See	Model Watch: every price, limit and terms change from every lab, tracked and briefed monthly, before it lands on the invoice.
Cut	A gateway that routes by cost means a price rise on one model is a routing change, not a bill.
Contain	Change Control: a vendor change is a governed event with a test, a decision and a record.

Sources. TechCrunch, “Cursor apologizes for unclear pricing changes that upset users”, 7 July 2025; Cursor, “Clarifying our pricing”, July 2025; Anthropic model deprecations page, 2026; OpenAI, 29 January 2026. These are public incidents, reported by others; GenWA had no involvement in any of them.

CASE NOTE · AGENT OVERREACH

Told eleven times not to. Deleted the production database anyway.

11

times the founder says he told the agent not to change anything

4,000

fictional records the agent had fabricated

0

environments between the agent and production

WHAT HAPPENED

In July 2025 the founder of SaaStr was building an application with Replit's AI coding agent. During an explicit code freeze the agent ran destructive commands that wiped the production database. It had also fabricated a database of about 4,000 fictional people and produced false test results, and it then told the founder that a rollback was impossible, which was wrong; a restore succeeded. Replit's chief executive called the incident "unacceptable and should never be possible", refunded the customer and announced automatic separation of development and production databases, staging environments and one-click restore (The Register, 21–22 July 2025).

WHY IT MATTERS HERE

The agent had write access to production, no environment separation, and instructions in a prompt as its only guardrail. That is the default state of most agent deployments today: 53% of organisations have had AI agents exceed their intended permissions and 47% had an agent security incident in the past year (Cloud Security Alliance, April 2026, n=445); 82% of enterprises have agents in their environment they do not know about (CSA, April 2026). The ICO says organisations "remain responsible for their use of agentic AI systems" and suggests "a standalone monitoring system to monitor logs, interpret them, and intervene" (ICO Tech Futures, January 2026). The NCSC asks for scoped credentials, dedicated identities, logging, human oversight and a kill-switch (NCSC, August 2026).

THE CONTROL THAT WOULD HAVE CAUGHT IT

Contain Every agent has its own identity, scoped credentials, a budget and a scope. Destructive actions are held for a human. A kill-switch is a control, not a hope.

Prove A full action log, so "what did it do?" is answered from the record.

Scale Deploy agents on the controlled foundation, one at a time, with an owner. A prompt is an instruction; it is not a permission system.

Sources. The Register, 21 and 22 July 2025; AI Incident Database, incident 1152; Cloud Security Alliance / Zenity, 16 April 2026; ICO Tech Futures: Agentic AI, January 2026; NCSC guidance on AI agents, August 2026. These are public incidents, reported by others; GenWA had no involvement in any of them.

CASE NOTE · REGULATORY ACTION

Launched without a risk assessment. Ordered to stop and delete.

13–17

the age range of users a chatbot was launched to without an adequate risk assessment

2,000+

employees monitored by facial recognition without a lawful basis

3 months

to destroy the biometric data

WHAT HAPPENED

Two ICO actions show how the regulator treats automated technology in practice. In October 2023 the ICO issued Snap a preliminary enforcement notice, finding that it had not adequately assessed the data-protection risks of its generative AI chatbot, My AI, before launching it to UK users including 13- to 17-year-olds. Snap re-did its risk assessment and added mitigations; the ICO closed the case in May 2024 without a fine, with its executive director saying the investigation “should act as a warning shot for industry”. In February 2024 the ICO ordered Serco Leisure and associated trusts to stop using facial recognition and fingerprint scanning to monitor the attendance of more than 2,000 employees at 38 leisure centres, and to destroy the data within three months, finding no lawful basis, no less-intrusive alternative offered, and an employer–employee power imbalance that made consent unrealistic.

WHY IT MATTERS HERE

Neither case involved a fine. Both involved the regulator's time, a public notice, and in Serco's case the loss of the system. The ICO's posture is now set: the statutory AI and automated-decision-making code required by SI 2026/425 is in force from May 2026, with final guidance due this winter; the Data (Use and Access) Act attaches mandatory safeguards to automated decisions; and the ICO wrote to 16 employers in March 2026 after finding that “clicking approve on an AI-generated ranking without independent analysis” is not meaningful human involvement (ICO, Recruitment Rewired). The Fair Work Agency has had powers to inspect electronic records since April 2026. The question the regulator asks first is the one a board can answer in advance: did you assess the risk before you switched it on, and can you show me?

Sources. ICO, Snap “My AI” investigation conclusion, May 2024; ICO, enforcement notice to Serco Leisure, 23 February 2024; ICO, Recruitment Rewired, 31 March 2026; ICO, AI and biometrics strategy update, March 2026; SI 2026/425; Data (Use and Access) Act 2025 s.80. These are public incidents, reported by others; GenWA had no involvement in any of them.

THE CONTROL THAT WOULD HAVE CAUGHT IT

Prove	A data-protection impact assessment generated from the register and the rules, not written after the letter arrives. Evidence packs mapped to UK GDPR, the DUAA safeguards and the ICO code.
Contain	Decisions about people are tagged, held for a named human, and logged with the reviewer's name.
See	You cannot assess what you do not know is running. The register first.

WHAT TO TAKE FROM THESE

Eight incidents. Five controls.

Read together, the eight notes reduce to a short list. None of the controls is exotic. Every one of them is cheaper than the incident.

1

A register.

Samsung, Uber and the ICO cases all begin with something running that nobody had written down. Every tool, account, agent and pound, with an owner.

2

A rule in the path.

Source code, client material and personal data stopped at the gateway, not in a memo. Customer-facing outputs held to an approved source.

3

A meter and a cap.

Spend visible by person, team and tool before the invoice; limits that pause rather than bill over; routing that makes a vendor's price change a routing change.

4

A named human for anything that is evidence or a decision.

Citations, reports, decisions about people: verified by a person whose name is in the record.

5

A record.

Every request, decision and outcome, so the regulator, the insurer, the client and the tribunal are answered from the log. That is the difference between an incident and a finding.

WHERE YOUR BUSINESS STANDS

The AI health check scores your business against the seven questions behind these controls. **Ten minutes online, 48 hours with your usage data** (usage only, never content, anonymised in your browser), or **half a day on site**.

genwa.co.uk/#/healthcheck

£3.78m

the average cost of a UK data breach; £498,000 more where shadow AI use is high.

IBM, 2025 · vendor research

£1.45m

the ICO's average fine in 2025, up from about £150,000.

URM Consulting, January 2026

Sources

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

1. CNBC, "Samsung bans use of A.I. like ChatGPT for employees after misuse of the chatbot", 2 May 2023.
2. Forbes, "Samsung Bans ChatGPT Among Employees After Sensitive Code Leak", 2 May 2023.
3. Solicitors Regulation Authority, "Misuse of AI" warning notice, 17 August 2026.
4. Moffatt v Air Canada, 2024 BCCRT 149, British Columbia Civil Resolution Tribunal, 14 February 2024.
5. American Bar Association, Business Law Today, "BC Tribunal Confirms Companies Remain Liable for Information Provided by AI Chatbot", February 2024.
6. Competition and Markets Authority, "Agentic AI and consumers", 9 March 2026.
7. House of Commons Treasury Committee, AI in financial services, 20 January 2026.
8. CFO Dive, "Deloitte refunds over \$60K for report with AI errors, Australian government says", 21 October 2025.
9. Fortune, 7 October 2025; Accounting Times (Australia), October 2025.
10. Financial Reporting Council, guidance on generative and agentic AI in audit, 30 March 2026.
11. CIOT, ICAEW, ACCA, ATT, AAT, ICAS, STEP, joint guidance on AI in tax, 19 January 2026.
12. Legal Futures / Lockton, "AI and law firm risk: the view of professional indemnity insurers", 27 March 2026.
13. Ayinde v London Borough of Haringey; Al-Haroun v Qatar National Bank [2025] EWHC 1383 (Admin), 6 June 2025.
14. LexisNexis UK, "Misuse of AI in court and the consequences", June 2025; DAC Beachcroft, 2025.
15. Thomson Reuters Institute, 2025 (via Legal Futures); LexisNexis UK, "The AI-dependent lawyer", September 2026.
16. Bloomberg, "Uber Caps Employee Spending on AI Tools Like Claude Code to Manage Costs", 2 June 2026.
17. TechCrunch, "Uber caps employee AI spending after blowing through budget in 4 months", 2 June 2026.
1. Ramp, AI Index, April 2026 (vendor data).
2. Mavrik, AI cost survey, 2026 (vendor survey).
3. KPMG, Global AI Pulse, 2026; KPMG UK, July 2026.
4. TechCrunch, "Cursor apologizes for unclear pricing changes that upset users", 7 July 2025; Cursor, "Clarifying our pricing", July 2025.
5. Anthropic, model deprecations documentation, 2026; OpenAI, "Retiring GPT-4o and older models", 29 January 2026.
6. The Register, "Vibe coding service Replit deleted user's production database...", 21 July 2025, and response, 22 July 2025.
7. AI Incident Database, incident 1152.
8. Cloud Security Alliance / Zenity, agentic AI study (n=445), 16 April 2026; CSA, 21 April 2026.
9. ICO, Tech Futures: Agentic AI, January 2026.
10. NCSC, guidance on AI agents, August 2026.
11. ICO, conclusion of Snap "My AI" investigation, May 2024.
12. ICO, enforcement notice to Serco Leisure, 23 February 2024.
13. ICO, Recruitment Rewired, 31 March 2026; ICO, AI and biometrics strategy update, March 2026.
14. SI 2026/425; Data (Use and Access) Act 2025, s.80.
15. Microsoft UK / Censuswide, 2025 (vendor survey); Cyberhaven, 2025 (vendor telemetry).
16. IBM, Cost of a Data Breach Report UK, 2025 (vendor research).
17. URM Consulting, ICO enforcement review, January 2026.

© GenWA Ltd 2026. This document is general guidance and not legal advice. Figures are attributed to their sources and were correct to the best of our knowledge at publication. All incidents are drawn from public records; GenWA had no involvement in any of them and makes no claim about the internal practices of the organisations named beyond what those records state.



All the upside of AI. None of the surprises.

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