BRIEFING · SEPTEMBER 2026

Artificial Intelligence in *UK Legal Services*

What the evidence actually says, what the regulator expects, and what can safely be done now.

Neil Catton

Independent Technology Strategist · Technology, Ethics & Human Purpose

NEILCATTON.COM

Solicitors in private practice, England and Wales. Every figure is traced to a regulator, a court, an ombudsman, published legislation, a trade body, or an operator speaking on its own record.

ONE

A note on the numbers

The most consequential thing in this evidence is not a statistic. It is what a court asked a regulator to look at.

In a housing case in the Administrative Court, a pupil barrister filed grounds for judicial review citing five authorities. None of the five existed. A statutory provision was materially misstated alongside them. In a second case heard at the same time, a schedule prepared by a judicial assistant identified forty-five problematic citations, of which eighteen were to cases that had never been decided. The Divisional Court's summary of that second set is one sentence: "The vast majority of the authorities are made up or misunderstood."

The pupil was referred to the Bar Standards Board. What the court asked the regulator to consider, though, was not only the false citations. It was the pupil's supervision during pupillage, how work had been allocated to them, and how their competence had been assessed. The instructing solicitor was referred to the Solicitors Regulation Authority over the steps he had taken to satisfy himself that the pupil had sufficient experience and competence. In the second case a solicitor was referred for having relied on his own lay client to do the legal research.

Three referrals, and not one of them is really about a chatbot. They are about who was supposed to check the work. That is the argument of this briefing, and it is not an opinion imposed on the evidence — it is what the primary sources say when they are read in full rather than summarised.

What this briefing includes

Only primary sources. The SRA's own guidance, warning notices and data. Judgments read in full rather than through commentary. Legal Ombudsman complaint data. HM Land Registry's annual report. Trade body market data, with the genre named at every use. Two law firms' own published test results on their own tools.

What it leaves out, and why

Almost every published statistic about AI in the legal market is produced by a company selling AI into the legal market. Nine sources were located, read, and excluded. Four are vendor adoption surveys, one of which is carried on a trade body's own website as paid partner content — which makes it more likely to be repeated, not more reliable. One is a consultancy projection of productivity gains with no inspectable method. One is a paywalled commercial market estimate. Five are second-hand figures repeated by the regulator itself in a 2023 publication, including an adoption percentage sourced to a legal publisher, a ninety-per-cent case-prediction claim from an unnamed provider, and a discrimination statistic derived from a different country's mortgage market.

One further exclusion is worth recording separately, because it would have been the most quotable number here. A trade body report states that 82% of lawyers are adopting or planning to use AI tools. The underlying survey is not named. A trade body repeating an unattributed adoption figure is still an unattributed adoption figure.

Where the record is thin, and said to be

No UK primary source was found on what legal AI tools were trained on, on model security in legal practice, or on what AI is doing to trainee numbers. Each was searched for and each is reported here as an absence rather than passed over in silence. One document — HM Land Registry's annual report — gives two fraud figures that it does not reconcile; both are given here, with that fact attached. Two sources could not be retrieved at all, and that is recorded as not obtainable rather than as evidence of absence.

Sources: SRA; Courts and Tribunals Judiciary; Legal Ombudsman; HM Land Registry; TheCityUK; Ashurst; Linklaters. Ayinde v London Borough of Haringey; Al-Haroun v Qatar National Bank [2025] EWHC 1383 (Admin), 6 June 2025, at [33]–[37], [70], [72], [74] and [81].

TWO

The wider market

This is a strong sector, and the strength is not confined to London.

UK legal services contributed £38bn in gross value added in 2024, about 1.5% of total UK GVA, growing just over 3% on the year. Sector revenue was £52.3bn, up 11%, of which the hundred largest firms account for £37bn. The trade surplus was £8.9bn — the highest of any UK professional services sector. The sector employs around 364,000 people, and almost two-thirds of them work outside London.

£38bn
Gross value added, 2024 — 1.5% of UK GVA

£52.3bn
Sector revenue, 2024, up 11% on the year

364,000
People employed, almost two-thirds outside London

The genre of those figures matters and is stated here rather than buried. They come from a trade body's fourteenth annual report on the sector, published with a commercial partner. They are not regulator statistics. They are the best traceable numbers available for the market's size, and they are cited as what they are. A paywalled commercial estimate putting the market above £55bn was excluded because its method cannot be inspected.

Against that strength, the structure of the profession is moving. In the twelve months to the end of July 2026, 371 firms opened. In the twelve months to the end of May 2026, 520 closed — a figure that now includes firms ceasing to practise, mergers, splits and changes of status. The two periods are not co-terminous and should not be netted against each other. The direction is nonetheless consistent with a market in which the smallest firms are under the most pressure, and it is the smallest firms with the least supervisory capacity to spare.

TheCityUK, ‘UK legal services 2025: Legal excellence, internationally renowned’, fourteenth annual report, published with Barclays, 9 December 2025; figures for 2024. SRA regulated population statistics, firms opened and closed.

THREE

The sector's pressure point

More people complained, the same number were answered, and the redress did not move.

Someone selling a house waits three weeks for an answer to an email. They ring, and are told the person handling the file is away. They write to complain, and hear nothing back. Eventually they complain to the ombudsman instead. That is not an anecdote chosen for effect — it is the modal complaint in the largest category of complaints in the sector, and the numbers behind it moved sharply in 2025/26.

The Legal Ombudsman received 14,259 complaints in the year to 31 March 2026, against 10,447 the year before — a rise of 37%. It accepted 8,412 for investigation, up 16%. It resolved 8,199, which is 1% fewer than the 8,270 it resolved the previous year. Redress ordered was £3,700,986, against £3,708,674 the year before. Essentially flat.

Complaints received	14,259, up 37% from 10,447
Accepted for investigation	8,412, up 16% from 7,203
Resolved	8,199, down 1% from 8,270
Redress ordered	£3,700,986, against £3,708,674
No final response from the firm	25% of consumers

Read those five lines together and the pressure point states itself. Demand for redress is rising considerably faster than the sector's capacity to answer it, and the amount of money changing hands at the end of the process has not moved at all. A quarter of the people who complained to the ombudsman had never received a final response from their firm in the first place.

THE COMMERCIAL READING

The gap between complaints made and complaints answered is where the case for AI in a law firm is being made. It is a real gap. It is not a manufactured one, and this briefing does not treat it as one.

Legal Ombudsman, annual complaints data and insight, 1 April 2025 to 31 March 2026. No publication date is stated on the dataset; the period is given instead.

FOUR

What the sector is facing

Two drivers account for most of the complaints, and neither of them is about getting the law wrong.

Poor communication is 24% of what consumers complain about. Delay or failure to progress is 22%. Together they are 46% of all complaint causes, and they remain the top two among complaints that are upheld. Failure to advise adds a further 19%. Costs, which is what most people outside the profession would guess, is 9%.

That is an operational failure pattern, not a legal one. It describes firms whose work is probably sound and whose clients cannot find out what is happening.

It concentrates in one place

Residential conveyancing produced 3,039 accepted complaints — 36% of the total — and rose 41% in a year. Poor service was evidenced in 78% of the conveyancing complaints investigated. Wills and probate follows at 1,147 complaints, with poor service evidenced in 79% of those investigated, and personal injury at 1,140. Immigration and asylum rose 49%, employment 37% and litigation 13%. Property work fell 13%.

Residential conveyancing	3,039 accepted — 36% of all accepted complaints, up 41%. Poor service evidenced in 78% of those investigated.
Wills and probate	1,147 accepted — 14%. Poor service evidenced in 79% of those investigated, the highest rate in the data.
Personal injury	1,140 accepted — 14%. Poor service evidenced in 74%.
Fastest rises	Immigration and asylum up 49%; employment up 37%; litigation up 13%.

And it shows in how firms handle being told

Poor service was evidenced in 72% of all investigated complaints. Unreasonable complaint handling was found in 46% of investigative outcomes, three percentage points better than the year before. Within the residual 'other' category, complaints about unreasonable complaint handling rose 84% in a single year, to 990.

The consequence lands on individuals in small amounts that matter to them. Compensation for emotional impact was awarded in 86% of investigated complaints, averaging £429. Fee refunds, reductions or waivers were ordered in 29% of cases, averaging £3,520. Compensation for direct financial loss was ordered in 14%, averaging £4,553.

Legal Ombudsman, annual complaints data and insight, 1 April 2025 to 31 March 2026.

FIVE

The challenges of using AI

Two firms have tested their own tools and published the results. Both found less than the market claims, and both said so.

In October 2023 a City firm set fifty English law questions across ten practice areas and had its own specialists mark the answers out of ten — five marks for substance, three for citations, two for clarity. The best model scored 4.4. Across the exercise, 47 of 150 answers contained hallucinated authorities.

In February 2025 the same firm repeated the exercise against a newer generation of models. OpenAI o1 scored 6.4 out of 10 and Gemini 2.0 scored 6.0. Hallucinated authorities fell to 9% of answers. That is a substantial improvement across roughly sixteen months, and it is the honest number to quote — the 2023 score still circulates widely without its date attached, and using it today would flatter the argument rather than test it.

What did not improve was the firm's conclusion. Its recommendation after the second round reads: “despite the significant improvement in Gemini 2.0 and OpenAI o1, we recommend they should not be used for English law legal advice without expert human supervision.”

> THE FINDING
>
> *The tools got materially better across two model generations and the supervision requirement did not change. That is the whole question for a managing partner, answered by a firm testing its own software.*

What a large trial actually measured

A second firm ran three concurrent trials between November 2023 and March 2024 across 411 participants in 23 offices and 14 countries, and published the results in June 2024. In a blind study, generative AI output was marked for accuracy at an average of 3.0 out of 5 against human work at 3.5. Half of the AI outputs were mistaken for human work or judged indistinguishable from it. None of the human work was mistaken for AI. Time savings by task type were 80% on UK corporate filings, 59% on research reports and 45% on first-draft briefings. Of the participants, 77% agreed it got them to a first draft faster and 61% that it helped them manage their workload.

The firm's own caveat is the most useful sentence located anywhere in this research: “Given the sample size, much of the data cannot be simply extrapolated to get to definitive answers on the actual time and costs savings that the firm might achieve.” It added that the industry “is still some years away from being able to determine or implement any necessary changes to law firm business models” as a result of adopting the technology.

The measurement problem, stated directly

Four firms were approached for self-published evidence. Two had published substantive measured results; two had not. A trade body's own financial benchmarking survey could not be retrieved and is recorded as not obtainable rather than absent. The honest output of that sweep is not a ranking of who has adopted most. It is that the only two UK firms to have published measured results both found less than the vendor consensus and both attached explicit warnings against extrapolation. A board benchmarking itself against a competitor's published AI figure is, in this sector, almost always benchmarking against a disclosure decision rather than a performance.

Linklaters, LinksAI English law benchmark, version 1 (October 2023) and version 2 (10 February 2025). Ashurst, ‘Vox PopulAI’, 10 June 2024 — 411 participants, 23 offices, 14 countries, trials November 2023 to March 2024. The firm now trades as Ashurst Perkins Coie.

SIX

The threat side

The threat model does not transfer from other sectors, and here it runs in two directions at once.

In insurance, the AI threat to a firm was fraud committed against it. In social housing, it was legitimate but better-equipped use by claimants and their advisers. Neither framing fits legal services, and adopting one of them without checking would have produced a briefing about the wrong risk.

Inward: against the firm's own client checks

In April 2026 the SRA passed on an alert from the National Economic Crime Centre. Its wording is specific: “AI has been detected in various parts of the CDD process, including to manipulate or create identity documents and produce deepfakes to deceive liveness detection tests.” And, on whether this is hypothetical: “Most attempts are unsuccessful, but some have been able to bypass controls, confirming this as a current threat not a theoretical one.”

The regulator's sectoral risk assessment, published in June 2026 and updated in September, carries the same risk forward: deepfakes “may increase the risk of identity fraud, impersonation and the circumvention of customer due diligence processes, particularly where firms rely on remote onboarding”, and the exposure is heightened “where firms rely on remote verification methods or digital onboarding processes”. The practical point for a firm is that this lands on conveyancing and on client account — which is also where the complaint volume is concentrated.

Outward: from the firm's own desks

Between July 2025 and July 2026 the SRA received 42 reports of potential AI misuse and has investigations open. It says those investigations concern inaccurate legal citations, supervision and confidentiality. Two of those three are about how a firm is run rather than what a person typed.

On 4 September 2026 the Solicitors Disciplinary Tribunal decided its first case involving AI-generated false citations. A registered foreign lawyer was struck off. The generated citations had appeared in his response to the regulator's own allegations; when the false quotations were identified, the corrected response contained further fabricated material, also generated. The tribunal found culpability “very high” and a “persistent lack of insight”, holding that he had “embarked on the use of AI without giving any proper thought to the consequences” and had “failed to take proper steps to ensure his work output was accurate”. No costs were ordered; the regulator accepted he could not pay.

That case is the sanction end of the range the Divisional Court had described fifteen months earlier, and it arrived quickly. It is also the one case in this evidence where the individual framing fits better than the systemic one — a person acting alone, with no supervisory structure implicated. The pattern across the rest of the cases is different, and this briefing does not bend this one to match it.

> WHERE THE DIRECTION OF THE THREAT MATTERS
>
> The best-documented AI harm in UK legal services to date is produced by regulated professionals and counted by their own regulator. That is not true of any other sector in this series, and it changes what a firm's first control should be.

The fraud figures, and what they do not say

HM Land Registry's annual report for 2025-26 states that suspected fraudulent registration applications were cancelled against approximately £60 million worth of property. The same section of the same report refers to preventing fraudulent property applications worth more than £63 million. The report does not reconcile the two figures, and both are given here for that reason. In April 2025 the principal offender in one case was convicted of 76 fraud and invalid identity offences and sentenced to ten years. Neither figure separates AI-enabled fraud from any other kind, and the report does not mention deepfakes at all — the registry's own fraud reporting does not yet distinguish it.

SRA update on AI and identity checks, 22 April 2026 (NECC-sourced). SRA sectoral risk assessment, 8 June 2026, updated 4 September 2026. SRA news release, 17 August 2026. Solicitors Disciplinary Tribunal decision of 4 September 2026 — reported via trade press and firm commentary; the tribunal's own judgment was not obtained. HM Land Registry, Annual Report and Accounts 2025-26.

SEVEN

The UK legal position

There is no AI rulebook for law firms, and the striking thing is that the profession has just been asked whether it wants one and said no.

The SRA regulates outcomes rather than tools. Its position, unchanged since 2023, is that “we do not specify the technologies that firms should or should not adopt”. Its warning notice of 17 August 2026 sets out what does not change: “The use of AI does not diminish or transfer your professional responsibilities. You remain accountable for all work and advice delivered to clients, whether or not AI has been used in its preparation.”

What that notice cites is the interesting part. Alongside the obvious paragraphs on not misleading the court and only advancing what is properly arguable, it cites the Code for Solicitors on competent service and on accountability and effective supervision, and the Code for Firms on governance, systems and controls, and on supervision and competence systems. The obligations a firm is most likely to be found against are the ones about how it is organised.

Accountability	Principles 1, 2, 4–5 and 7; Code for Solicitors 1.4 and 2.4 — do not mislead, and only advance what is properly arguable.
Competence and supervision	Code for Solicitors 3.2 and 3.5; Code for Firms 4.3–4.4 — a competent service, and systems that supervise it.
Governance	Code for Firms 2.1 — effective governance, systems and controls. The paragraph most firms have not mapped to their AI tools.
Confidentiality	Code for Solicitors and Code for Firms 6.3.
Automated decisions	Data (Use and Access) Act 2025 s.80, in force 5 February 2026 — inform, represent, intervene, contest.
The statutory code	SI 2026/425 requires an ICO code on AI and automated decision-making. No deadline is set for it.

The duty the court set out

The Divisional Court was direct about the tools: freely available generative AI trained on a large language model “are not capable of conducting reliable legal research”. Anyone using them anyway has “a professional duty therefore to check the accuracy of such research by reference to authoritative sources, before using it in the course of their professional work”. It named those sources — the Government's legislation database, the National Archives database of court judgments, the official Law Reports, and the databases of reputable legal publishers. Most are free. The verification infrastructure this sector needs already exists and costs nothing.

The court then set the standard that matters commercially. The duty “rests on lawyers who use artificial intelligence to conduct research themselves or rely on the work of others who have done so”, and it is “no different from the responsibility of a lawyer who relies on the work of a trainee solicitor or a pupil barrister”. It put the obligation on firm leadership by name — heads of chambers and managing partners — requiring “practical and effective measures”.

The question nobody has answered

In November 2025 the Upper Tribunal said something in passing that a firm now has to have a position on. Putting client letters and Home Office decision letters into an open source AI tool such as ChatGPT, it said, “is to place this information on the internet in the public domain, and thus to breach client confidentiality and waive legal privilege”, and any regulated professional or firm doing so “would…be advised to consult with the Information Commissioner's Office”.

HOW TO HOLD THIS

This is not settled law, and treating it as settled would be wrong. The observation was not necessary to the outcome, it has not been tested on appeal, and whether a commercial AI service places material "in the public domain" depends entirely on the product, the contract and the configuration. What is true is narrower and still significant: a judge has said it once, in writing, about a real firm's use of a named tool. A firm can no longer assume the question will not be asked — of its tools, its contracts, and what its people put into them.

The rule change that was consulted on, and declined

On 17 February 2026 the Civil Justice Council published an interim report and consultation on the use of AI for preparing court documents. The consultation closed on 14 April 2026 and a findings update followed on 30 June. That update records "a high degree of alignment across respondents that, for legal professionals preparing pleadings, advocacy documents, skeleton arguments and similar documents, no additional formal requirements relating specifically to AI are currently necessary".

The proposed transparency measures land elsewhere: an amendment to Practice Direction 35 requiring expert witnesses to "identify and explain" AI use beyond administrative functions, and a declaration for trial witness statements under PD57AC that "AI has not been used for the purposes of generating its content (including by way of altering, embellishing, strengthening, diluting or rephrasing the witness's evidence)". A final report is expected later in 2026 and had not been published as at 12 August 2026.

An independent judicial advisory body asked the profession whether AI needed its own rules for lawyers' documents, and the answer was that the existing duties already cover it. That is this briefing's argument reached from the other direction. The gap is not in the rulebook. It is in supervision.

SRA warning notice and news release on the misuse of AI, 17 August 2026. SRA Risk Outlook: the use of artificial intelligence in the legal market, 20 November 2023 — three years old, and its statistics are second-hand and excluded here. Ayinde v Haringey; Al-Haroun v QNB [2025] EWHC 1383 (Admin) at [6]–[9]. UK v Secretary of State for the Home Department [2026] UKUT 81 (IAC), 17 November 2025, at [4] and [21], reported by ICLR. Civil Justice Council interim report and consultation, 17 February 2026; consultation findings update, 30 June 2026. Data (Use and Access) Act 2025 s.80; SI 2026/425.

EIGHT

Beyond England and Wales

Three positions a firm with work outside this jurisdiction has to hold at the same time.

The rest of the United Kingdom

The SRA regulates England and Wales. It does not regulate Scotland or Northern Ireland, and neither the Law Society of Scotland nor the Law Society of Northern Ireland has adopted its guidance by reference — each has issued its own AI guidance for its own solicitors. A firm operating across the United Kingdom is answering to three sets of professional guidance on this question, not one. The assumption that a single UK-wide regulator covers all four nations is a common and expensive error.

The European Union

The EU AI Act classifies AI used in the administration of justice as high risk, and places obligations on both providers and deployers of such systems. The UK has no equivalent statute and has chosen to regulate through existing professional duties instead — which is why this briefing spends its regulatory section on the Codes of Conduct rather than on AI-specific law. For cross-border work, the stricter regime sets the operating floor, and a firm advising on EU matters cannot rely on the UK position alone.

The litigation signal

Fabricated citations are now a cross-jurisdiction pattern rather than an English curiosity. The same failure has reached courts in the United States and elsewhere on materially the same facts: unverified output filed under a professional's name. No jurisdiction that has examined it has found the technology at fault. Every one of them has located the failure in the duty to check.

> WHAT THIS MEANS PRACTICALLY
>
> A firm's AI position needs to be written once and then tested against each jurisdiction it practises in, rather than written for England and Wales and assumed to travel. The supervision requirement is the part that travels; the specific rules are the part that does not.

Law Society of Scotland and Law Society of Northern Ireland AI guidance. Regulation (EU) 2024/1689. Cross-jurisdiction pattern as recorded in the judgments read for this briefing.

NINE

Today

Seven things a firm can do now. All seven sit inside duties that already exist, and none needs a new rule to permit it.

- → **Put verification into the file-closing checklist, not the training deck.** The court named the authoritative sources and most of them are free. Verification that depends on someone remembering a training session is not a control.
- → **Write down which tools may touch client material, on what contractual terms.** Then answer the privilege question before a judge asks it. The answer will differ by product and by contract, which is precisely why it has to be written down rather than assumed.
- → **Make the review step visible.** Every referral in the case law turned on whether supervision could be shown to have happened. A review that leaves no record is, evidentially, a review that did not occur.
- → **Use AI where the complaint data actually is.** Acknowledgement, progress updates and chasing account for the 46% of complaints that are communication and delay. This is the lowest-risk, highest-value application in a law firm and it is not the one being sold hardest.
- → **Re-test client due diligence against deepfake impersonation.** Particularly where onboarding is remote and the matter is conveyancing. The regulator has said some attempts are getting through.
- → **Separate assistive use from anything that reaches a client or a court.** Nothing goes out without a named person having read it in full. That is a workflow rule, not a policy statement, and it is enforceable.
- → **Record what the firm excluded and why.** The list of tools considered and rejected, and the reasons, is the governance artefact a regulator will ask to see. Most firms document what they bought and not what they refused.

None of these is expensive. Six of the seven are process changes rather than purchases. The one that costs money — re-testing identity verification — is the one with a named regulator alert behind it.

Derived from the obligations set out in the SRA warning notice of 17 August 2026 and the duties stated in Ayinde v Haringey at [6]–[9]. Complaint proportions from Legal Ombudsman 2025/26.

TEN

Tomorrow

Three levels, defined by where accountability sits rather than by how capable the system is.

The useful distinction for a board is not technical. It is about who answers for the outcome, and it produces three levels.

Assistive	Prepares, summarises, surfaces. A person decides. Accountability is unchanged. Almost all defensible value in a law firm sits here today, and the benchmark evidence in section five is the reason.
Augmentive	Recommends within bounds, with reasoning and confidence attached. A person confirms or overrides, and the override is recorded. The review must be demonstrably meaningful — which is the exact point every referral in the case law turned on.
Adaptive	Decides within a defined envelope. Oversight moves to the envelope, and accountability to whoever sets and monitors it. This is a reasoned projection, not an observation: nothing in the evidence read shows it operating in a UK firm today.

Near, medium and long term

0 to 12 months	The Civil Justice Council's final report on AI in court documents is expected. Its June findings already point away from new rules for lawyers' own filings and towards transparency on expert evidence. Firms settle their privilege position, their tool inventory and their supervision record.
12 to 36 months	The statutory ICO code required by SI 2026/425 arrives, with no deadline currently set for it. Disciplinary outcomes accumulate and establish the sanction range. Augmentive use in drafting and disclosure becomes ordinary, and the review record becomes the thing that is audited.
36 months and beyond	Client expectations, rather than regulation, reset what a firm is paid for. If routine drafting carries no premium, the commercial question moves to judgement and accountability — the parts a firm cannot delegate to a tool.

The first band is evidenced. The second is reasoned from instruments already made. The third is a projection and is offered as one.

Civil Justice Council interim report and consultation, 17 February 2026; consultation findings update, 30 June 2026. SI 2026/425. The three levels are this series' own framing; the augmentive standard is drawn from Ayinde v Haringey at [8].

ELEVEN

Benefit realisation

A published saving is a disclosure decision before it is a performance.

The worked example is the trial described in section five. A firm ran it across 411 people in 23 offices and published time savings of 80% on UK corporate filings, 59% on research reports and 45% on first-draft briefings. Those are real measurements of real tasks. Then the same firm published this: “Given the sample size, much of the data cannot be simply extrapolated to get to definitive answers on the actual time and costs savings that the firm might achieve.”

That sentence is the most commercially useful thing in this briefing. A task-level saving in a trial is not a firm-level saving in a year. It does not survive contact with matter mix, with review time, with rework, or with the fact that a large proportion of the work was never the bottleneck. A board that benchmarks against a competitor's published AI number is benchmarking against that competitor's decision about what to publish.

What to measure instead

→ **Time to first substantive response to a client**, and the proportion of matters with no response inside the firm's own service standard. This is the axis the ombudsman judges on.

→ **Review time per output, recorded.** Not to police it, but because the record is the evidence that supervision happened. This is the single measure with regulatory as well as commercial value.

→ **Rework rate** — outputs returned by the reviewer, and what proportion required a source to be corrected rather than the prose improved. A rising source-correction rate is an early warning.

→ **Complaint causes mapped to the ombudsman's own categories**, so the firm measures itself on the same axis it will be judged on rather than on one of its own devising.

Three of those four are already captured somewhere in most firms' practice management systems and are simply not reported to the board. The fourth requires a mapping exercise that takes a day.

Ashurst, ‘Vox PopulAI’, 10 June 2024. Complaint categories from Legal Ombudsman, annual complaints data 2025/26.

TWELVE

A way in

Eight workstreams. The first three form a natural first step.

01	**Readiness assessment**	What is actually in use across the firm, including what nobody authorised, and where each use sits against the Codes.	4–6 weeks
02	**Use case identification and triage**	Ranked against the complaint data rather than against vendor categories, so effort goes where the service failure actually is.	3–4 weeks
03	**Regulatory and legal position paper**	The firm's written answer on confidentiality, privilege and supervision, defensible to the SRA and to a judge.	3–4 weeks
04	**Board and executive governance**	Who decides, who reviews, what reaches the partnership, and how Code for Firms 2.1 is evidenced.	4 weeks
05	**Threat and integrity response**	Client due diligence re-tested against deepfake impersonation, and the internal reporting route for AI errors.	4 weeks
06	**Benefit realisation framework**	The four measures in section eleven, instrumented and reported.	3 weeks, parallel
07	**Build, buy or partner evaluation**	Including the contractual terms that determine the privilege answer, which is usually decided by procurement without knowing it.	4–6 weeks
08	**Twelve-month mobilisation roadmap**	Sequenced against the Civil Justice Council's final report and the arrival of the statutory ICO code.	2 weeks

The first three — readiness, triage and the position paper — form a natural first step of six to eight weeks. They produce the three things a firm is most likely to be asked for and least likely to have: what is in use, why, and on whose authority.

THE CLOCK

There is no statutory deadline in this sector. What there is instead is 42 open reports and live investigations at the regulator, a first strike-off already decided, and a Civil Justice Council final report expected before the year is out. Six to eight weeks fits comfortably inside that, and a firm's position is better settled before the report lands than after.

Contact: nc@neilcatton.com · neilcatton.com · writing.neilcatton.com

THIRTEEN

Sources

Every figure traces to one of the following. Where a source could not be obtained, that is said.

Regulator

- SRA, Warning notice: misuse of AI — 17 August 2026. The current regulatory position.
- SRA, news release on the safe and responsible use of AI — 17 August 2026. Source of the 42-report figure, July 2025 to July 2026.
- SRA, Be aware of clients using AI to bypass identity checks — 22 April 2026, relaying an NECC alert.
- SRA, Sectoral risk assessment (anti-money laundering and sanctions) — 8 June 2026, updated 4 September 2026.
- SRA, Risk Outlook: the use of artificial intelligence in the legal market — 20 November 2023. Cited only for the SRA's statement of its own approach; its statistics are second-hand and excluded.
- SRA, Regulated population statistics — firms opened and closed. Population totals could not be retrieved.

Courts and tribunals

- Ayinde v London Borough of Haringey; Al-Haroun v Qatar National Bank QPSC [2025] EWHC 1383 (Admin), Divisional Court, 6 June 2025.
- UK v Secretary of State for the Home Department [2026] UKUT 81 (IAC), 17 November 2025, reported by ICLR.
- Cork & Anor v Smith [2026] EWHC 1199 (Ch). Cited only for what the SRA's warning notice says about it; the judgment was not obtained.
- Solicitors Disciplinary Tribunal, 4 September 2026 — the first SDT case involving AI-generated citations. Reported via trade press; the tribunal's judgment was not obtained.
- Courts and Tribunals Judiciary, AI Guidance for Judicial Office Holders — 31 October 2025.
- Civil Justice Council, Use of AI for Preparing Court Documents — interim report and consultation, 17 February 2026; closed 14 April; findings update 30 June 2026.

Ombudsman, registry and market

- Legal Ombudsman, annual complaints data and insight 2025/26 — 1 April 2025 to 31 March 2026.
- HM Land Registry, Annual Report and Accounts 2025-26.
- TheCityUK, UK legal services 2025 — fourteenth annual report, published with Barclays, 9 December 2025. A trade body report produced with a commercial partner.
- Law Society of Scotland and Law Society of Northern Ireland AI guidance.

Operators on their own published record

- Linklaters, LinksAI English law benchmark — version 1, October 2023; version 2, 10 February 2025.
- Ashurst, Vox PopulAI — 10 June 2024. The firm now trades as Ashurst Perkins Coie.

Excluded, and why

Four vendor adoption surveys, one carried on a trade body's website as paid partner content. A consultancy productivity projection with no traceable method. A paywalled market estimate. Five second-hand figures repeated in the SRA's 2023 Risk Outlook, including an unnamed provider's 90% case-prediction claim. One trade body adoption figure of 82% whose underlying survey is not named. Aggregated trackers of hallucination cases compiled by commentators — useful for locating cases, not citable for a count. *Not obtained rather than absent:* a trade body financial benchmarking survey, the SRA's release on the first AI-driven law firm authorisation, and SRA population totals.

FOURTEEN

About the author

Neil Catton is an independent technology strategist, working across a fractional and advisory portfolio — fractional CTO and CIO roles, AI board briefings, AI readiness reviews, technology due diligence, data strategy, and programme recovery. More than thirty-nine years in the technology industry, across more than twenty sectors including public sector, justice, policing, health, finance, insurance, telecoms, retail, utilities and construction.

The justice and public safety strand of that portfolio is the relevant one here. It includes advisory work to a group founded by senior former police officers, covering an AI and data mesh programme for UK custody data and advice on the ethical use of AI in evidential settings. The questions a law firm now faces about supervision, provenance, and what a record has to be able to show are the same questions, met earlier and under sharper scrutiny, in a setting where the consequences of an unverified assertion are measured in liberty rather than costs.

He is the author of *The Next Evolution*, *The Cognitive Crucible* and *The Shadow System*, and of the ebook *Working as Designed: Why systems that work still fail the people inside them*. He writes *The Next Evolution* on Substack, on technology, ethics and human purpose.

How this briefing was produced

Every figure is traced to a regulator, a court, an ombudsman, published legislation, a trade body, or an operator speaking on its own published record. Where a figure could not be traced to such a source, it is excluded, and the exclusion is named in section thirteen rather than left implicit. Where a source could not be obtained, that is recorded as not obtained rather than treated as evidence of absence. An absence is searched for before it is asserted. Publication dates are checked on every source, because a regulator's page can read as current and be three years old.

Notice

This briefing is not legal, regulatory, financial or professional advice, and should not be relied on as such. It states a position as at September 2026; the regulatory position, the case law and the market data may since have moved, and the Civil Justice Council's final report was outstanding at the date of writing. Third-party sources are cited in good faith and should be verified before being relied upon. Named organisations appear only from regulator, ombudsman or court findings and are used illustratively; none has been engaged or approached, and individuals named in published judgments are not named here. No advisory or client relationship is created by receipt of this document.

Independence. The author holds no vendor affiliations, reseller agreements, referral arrangements, commission interests or equity positions in any product, platform or firm referenced in this briefing, and has received no payment from any organisation named in it. That independence is what makes the exclusion discipline in section one credible rather than merely stated.

 nc@neilcatton.com · neilcatton.com