Lawyery

AI tool risk assessment for a regulated law firm

A working template for satisfying yourself before client data enters an AI system. Version 1.0, August 2026.

A working template for satisfying yourself, as the SRA's warning notice on the misuse of AI (17 August 2026) requires, before client data enters an AI system.

The questions are drawn from three places: the requirements of the notice itself, the obligations we already owe under data protection law, and the question set in ISO/IEC 42001:2023, used as a source of questions rather than as a certification target. This is not a certification scheme and it is not a substitute for a data protection impact assessment. Where you are processing special category data under Article 9 UK GDPR, which covers most immigration and family work, a DPIA under Article 35 is likely to be a legal requirement in its own right. Run that first, or run this alongside it and cross-refer. The structure below borrows from ISO/IEC 42001:2023, in particular its idea of assessing impact on individuals rather than only on the organisation, without pretending to be an implementation of it.

Complete one assessment per tool, not per matter. Assign a permitted tier at the end. Review at least annually, and immediately on any of the triggers in section 8.

Tiering

Decide the tier before you answer anything else, because it sets how hard the rest of the questions bite.

Tier	Material	Examples
1	No client data	Legal research on published sources, drafting a knowhow note, summarising a reported judgment
2	Client data, no special category, no court-facing output	Correspondence chasing, internal file summaries, routine company documents

Tier	Material	Examples
3	Special category or criminal offence data, safety-relevant material, commercially sensitive material, or anything going to a court, tribunal or the Home Office	Sponsor licence compliance files and suspension or revocation correspondence, right to work audit findings, modern slavery and trafficking matters, domestic violence and child arrangement files, fitness to practise evidence, applications resting on a relationship or an identity the client cannot safely disclose, grounds of appeal and judicial review papers

A tool cleared for tier 1 is not thereby cleared for tier 3. Most consumer chatbots will not clear tier 2.

1. The supplier and the contract

Who is the supplier, what is the legal entity, and where is it established?

Which product tier are we on: consumer, business or enterprise? Consumer terms and business terms for the same brand frequently differ on training and retention.

Is there a written data processing agreement compliant with Article 28 UK GDPR, and have we read it rather than accepted it?

Does the supplier train models on our content? If the default is yes, how is it disabled, who in our firm can re-enable it, and is that change logged?

What is the retention period for prompts and outputs after a session ends? Can we require deletion, by what mechanism, and within what timescale?

Who are the sub-processors, where do they hold data, and are we notified before the list changes?

Are there international transfers, and on what lawful basis?

What security certifications does the supplier hold, and do they cover this AI service specifically or only the corporate estate? This distinction is missed more often than any other on this list.

What are the supplier's incident notification obligations, and within what period?

What happens to our data on termination, and what happens to anything already ingested into a model?

Is the supplier connected to the firm or to any of its managers? If so, has this assessment been documented as though a stranger were reading it?

2. The data

What categories of client data will realistically reach this tool, as opposed to what the policy permits?

Does that include special category data under Article 9, or criminal offence data under Article 10?

Does it include material about third parties who have given no consent: witnesses, family members, children, opponents?

Would disclosure of this material endanger anyone, here or abroad? Answer it about the client's family and about third parties named in the file, not only about the client. In trafficking, modern slavery and domestic violence work this question is the assessment.

Is the material commercially sensitive or price sensitive? A sponsor licence file can contain an employer's whole payroll, its regulatory failings, and a restructuring or change of ownership it has not announced. Treat that as tier 3 whether or not any special category data is present.

Is the material privileged, and have we considered whether ingestion could compromise the confidentiality on which the privilege rests?

Is the material subject to an undertaking, a court order, reporting restrictions, or a confidentiality ring?

If redaction or pseudonymisation is our answer, is it realistic at volume and under deadline, and who checks that it happened?

What do we know about the provenance and quality of the data the model itself was trained on, and does the supplier make any representation about it? An unsatisfactory answer here is common and is not by itself disqualifying, but it should be recorded rather than skipped.

3. The use

What task is the tool performing: drafting, summarising, research, translation, transcription, review, or something else?

Will the output, or anything derived from it, go to a court, a tribunal, the Home Office, or a client?

Where authorities are involved, who applies the four tests in the notice, namely that the authority is genuine, relevant, verifiably cited and actually advances the argument being made?

Does that person have the domain expertise to apply tests two and four? The first and third can be done by anyone. The second and fourth cannot.

What is the failure mode if the output is wrong, who finds out, and how quickly?

Is there a non-AI route to the same result, and how much slower is it? If the permitted route is materially slower than the prohibited one, expect the prohibited one to be used.

What is the intended use, and separately, what is the reasonably foreseeable misuse? The tool the supplier demonstrates and the tool a fee earner reaches for at half past five are not always the same tool.

What is the level of automation, and at exactly which step does a human decide something rather than confirm something already decided?

How was this validated before we deployed it, how is it monitored now, and who would notice if the quality of its output changed?

4. The people

Who is authorised to use this tool, and at which tiers?

What training have they had, when, and how is that recorded?

Who supervises the output, and does that person satisfy the notice's expectation of experience proportionate to the complexity and risk of the work?

Does anyone signing hold less than three years' post-qualification practice, and if so, how does that sit with the supervision requirements in the Authorisation of Firms rules?

Who signs, and does the signature record acceptance of professional responsibility rather than mere approval?

5. The record

What evidence exists that the review actually happened, as distinct from that someone clicked approve?

Could we produce that evidence eighteen months from now for an insurer at renewal, a regulator responding to a complaint, or a court asking how a citation reached it?

Are prompts and outputs retained on the matter file, and if not, why not?

Does the record show what was changed and what was rejected, not only what was accepted?

6. The client

Does our retainer or engagement letter address AI use in terms a client would understand?

Would this client be surprised to learn how their material was processed? If yes, that is the answer to whether they should be told.

Has the client, a funder, or a panel agreement imposed any restriction on AI use?

For legal aid or third party funded work, is there any restriction we have not checked?

7. Impact beyond the firm

Borrowed from the AI impact assessment idea in ISO/IEC 42001:2023, and the section most firms will skip.

Who bears the harm if this goes wrong, and is it us or the client?

Is there a foreseeable disadvantage to a particular group of clients from using this tool, for example where performance is weaker in a language, an accent, or a name format common among our client base?

If the tool is used in translation or transcription, what happens when it is confidently wrong about a word that carries the case?

Would we be comfortable explaining this use to the client, in writing, in advance?

Who are the interested parties in this deployment: the client, the court or tribunal, the insurer, the regulator, the supplier, our own staff? What does each of them need to be told, when, and by whom? Working out what you would say to a judge about how a document was produced is a quick way of finding what you are not comfortable with.

8. Ownership and review

Who owns this assessment by name?

When was it completed, and when is it next due?

Which of the following triggers an immediate reassessment: a change of model or version; a change to the supplier's terms, training default or retention period; a change of sub-processor; a security incident at the supplier; a new matter type being brought within scope; a near miss on one of our own files?

Permitted tier assigned, and by whom.

Conditions or restrictions attached to that permission.

Prepared as a practical aid, not as legal advice, and not as a compliance guarantee. Completing it does not discharge any duty under the SRA Standards and Regulations or under data protection law. References to the warning notice are to the version published on 17 August 2026; check the current version before relying on it. References to ISO/IEC 42001:2023 are to the standard's structure and concepts in general terms; the standard itself is the authoritative text.