

AI selection and use checklist

For Legal Practitioners



A joint initiative by Queensland Law Society in collaboration with:

Law Institute of Victoria
Law Society of South Australia
Law Society of Western Australia
Law Society of the Australian Capital Territory
Law Society Northern Territory
Law Society of Tasmania



February 2026

Risk assessment overview for Law Practices investigating an AI tool supplied by a third party. (Small / Mid-tier firm)

This checklist is designed to assist small and mid-tier firms to select and adopt AI tools built by someone else. It is intended to complement the companion materials listed below, but as a shortened summary it does not capture all the content of those materials.

The assessment process focuses on the following main points:

- Can you satisfy your confidentiality obligations to clients?
- Will the legal work product produced (AI Tool + procedural guardrails + human oversight) be good enough that you will be confident putting your firm's name to it, and your name as the responsible practitioner?

It is likely that the answers to these questions do not produce a "yes" or "no" answer. In addition to the go/no go decision you will need to supply guardrails to ensure that the way a tool is used stays within appropriate boundaries. This document therefore moves beyond a checklist for selection and suggests some of the conditions and controls that might be appropriate.

Ideally, a legal practice will employ a third-party expert consultancy and establish a specialist project team before using AI in a legal context. There are many other relevant considerations beyond the two basic ones listed above. However, if resources do not permit this a law practice must still undertake a basic risk assessment.

Arrangements between AI companies and providers adding AI functionality into their own software can be complex. This potentially adds new layers of risk and data sharing that does not arise from use of the core product.

Companion materials:

Source	Details
Law Council of Australia central website resource	Artificial Intelligence and the Legal Profession - Law Council of Australia
Queensland Law Society	QLS Guidance Statement 37 – Outlines the ethical obligations of AI use in law firms. QLS AI Companion Guide – A summary of AI technologies and what they can be used for in law firms. Template: Warning to Clients – Use of AI Tools Template AI use policy – Ensure staff know that they can only use approved AI systems, and only within the terms of that approval.
Law Institute of Victoria	Ethical and Responsible Use of Artificial Intelligence Statement on the use of artificial intelligence in Australian legal practice VLSB+C

Source	Details
Law Society of Western Australia	Law Society Artificial Intelligence Hub Statement on the use of artificial intelligence in Australian legal practice LPBWA
Law Society Northern Territory	Policy in development.
Law Society of NSW	A Solicitor's Guide to Responsible Use of Artificial Intelligence Statement on the use of artificial intelligence in Australian legal practice VLSB+C
Law Society of the ACT	Guidance on the Use of Generative AI in Legal Practice
Court directions *not exhaustive, practitioners should ensure they are aware of all relevant directions in their respective jurisdictions	QLD: Supreme Court of Queensland, PD 5/25 Accuracy of References in Submissions (*replicated in other jurisdictions including QCAT), Supreme Court of Queensland, PD 14/24 Expert Evidence in Criminal Proceedings (Other than sentences) Victoria: Guidelines for litigants, responsible use of artificial intelligence in litigation NSW: SC Gen 23 SA: Guidelines concerning the use of Generative artificial intelligence in litigation in South Australian Courts WA: Supreme Court of WA Guidelines for the Use of Generative Artificial Intelligence NT: No current Court directions.
Office of the Australian Information Commissioner	Guidance on privacy and the use of commercially available AI products

1. Understanding the AI Tool and proposed usage

- **Purpose:** Stakeholders¹ must be clear what the AI tool *can* be used for, what it *should* be used for and how the firm intends to ensure that use is within guardrails. Are any of the proposed uses directly covered by relevant Practice Directions?
- **Confidentiality:** How much access will be given to confidential data?² This is a foundation question and will inform the scope of subsequent risk assessment and management.³
- **Private or Public?:** Can the vendor provide a clear explanation of how the various parts of the system interact and who controls them? Is an isolated version of the base model exclusively licensed to the vendor?⁴ Are additional third party systems not controlled by the vendor used? If so, how will their data access and usage be governed? (See **Annexure A** for a discussion on the difference between Private and Public AI systems).
- **Terms and conditions:** Are all contractual terms (both from the primary provider and any others) available to you? Free or trial services may not provide clear terms, and in many cases the trade-off for a free service is wide data access. This is usually inappropriate for professionals.
- **Limits and Capabilities:** All AI has limitations. Are these reasonably apparent and is a manual system to identify and correct errors realistic? Will the error checking be so onerous that all time savings become illusory?
- **Reliability:** Does the vendor have any research data showing the baseline accuracy of the system?⁵
- **Integration options:** Can the tool be integrated into your standard software suite (such as Copilot) or Practice Management Software? There are pros and cons.⁶ Dealing with a trusted brand that has a good knowledge of the legal regulatory environment is reassuring. However, integration necessarily entails significant data access – by definition boring a large hole in your firm’s confidentiality wall. This means your Privacy Impact Assessment will need to be more rigorous. The more organised your data is the easier it will be to ensure the tool can only look at appropriate material.

¹ “Stakeholders” must include the firm’s senior management, AI working group (however named), Risk Committee, work group supervisors, front line users, (possibly) clients.

² Factors: limitations inherent in the tool, the firm’s risk appetite and resources available for deployment, mandatory usage restrictions (E.g. Court Practice Notes/Directions, Ethics constraints.)

³ If you propose to skip the data use assessment on the basis that no access to confidential data will be given, consider at least a preliminary look to see if it would be possible to expand use later.

⁴ See Attachment 1 for an explanation of a *Public vs Private (also known as Open vs Closed) AI model* licensing arrangements and the implications.

⁵ Note: such information must be approached with caution, and not only because it comes from the sales team. AI systems can be quirky, producing very different results at different times.

⁶ Integration can be powerful. As an existing provider already has access to client data anyway you might think it has no additional privacy impact. Notwithstanding, you would need to carefully assess confidentiality and work quality implications despite the provider being one you have dealt with for some time.

2. Privacy and Confidentiality

A **Privacy Impact Assessment** (“PIA”) must be conducted⁷ before deploying the system. This is a statutory obligation for Privacy Act regulated entities, and a professional duty for others.

The main points to be considered are:

- **Data Handling:** What does *each* of the parties who will be given access to your data (“Processors” or “Subprocessors”) say about how it will be used? Can the data be resold (anonymised or not) or used for AI model training? (Eg: “May be used to improve system performance”). See **Annexure A** for a discussion on some of the considerations in model selection. Be careful which version of the privacy policy and terms of service you use. Applicable documents can vary depending on whether you are dealing directly with the AI provider (vs using a version licensed by a third party), and which subscription tier you are on.
- **Data Security:** Do the Processors/Subprocessors have a meaningful data security accreditation, such as SOC 2, ISO 27001 or ISO 27017?⁸
- **Storage and Retention:** Where and how will data be stored, and for how long?
- **Jurisdiction:** Is client data being stored or processed offshore? If yes, in which jurisdiction?⁹ Does such use conform to your obligations concerning Personal Information (“PI”) [under APP 8¹⁰](#) [*Privacy Act 1988* (Cth)] or your duty of confidentiality under rule 9 of the ASCR?
- **Consent:** Do you have (or need)¹¹ informed client consent to use *or* disclose client information? Is the use consistent with your obligations under the *Harman* principles?¹² In most cases clients should be informed of AI data sharing and given access to any relevant documentation. For less sophisticated clients, analysis or guidance on the privacy regime should be supplied, potentially prior to or as part of the client retainer agreement.
- **Privilege:** The impact that using AI to process privileged communication is not clear. Arguably, using such material to train an AI system is incompatible with the intention to keep it confidential and therefore privileged. AI prompts (like Google searches) may not be privileged and may provide significant insights into the weaknesses in your client’s case. This consideration applies to client use of AI as well. (See **Annexure B**).
- **Anonymous data:** Anonymising data before it is supplied to an AI system (or ensuring particularly sensitive information is not processed) is a useful way of reducing risk. For example, if AI research or timeline summaries are prompted using facts from which client/party names are replaced with “Party A” / “Party B”, some of the immediate confidentiality risk can be avoided. Some tools can automate that, either before or after it is uploaded away from your system. Anonymisation is not a panacea, though. For more detail see Chapter 11 of the [QLS AI Companion Guide](#).

⁷ For further information about what this is and how to do it [see the OAIC website](#).

⁸ There is a difference between an audited certification standard, such as ISO and a data security framework such as NIST or Essential 8. The latter are best practice guidelines, but compliance is not audited.

⁹ It is easier to avoid the whole issue of international PII transfer by selecting a vendor that processes all data in Australia. While there is no “whitelist” of jurisdictions appropriate for foreign data transfers, a recipient which complies with the GDPR *or* is in a jurisdiction subject to GDPR [“adequacy decision”](#) can also reduce the scope of assessment required. Note: rapid changes in US government policy means that confidentiality assurances may no longer be credible.

¹⁰ Noting that from July 2026 the Small Business Exemption for firms with less than \$3 Million in turnover may not apply to legal practices supplying Designated Services under the AML/CTF regime.

¹¹ A client imparts confidential information to the firm for the purpose of obtaining legal services. The firm is not allowed to use that for their own benefit or disclose it to third parties unless R.9 of the ASCR permits.

¹² Harman principles: obligation to use material obtained under court process (discovery or subpoena) *strictly* for the purpose for which it was supplied; *Hearne v Street* [2008] HCA 36.

- **Information in court proceedings:** Information obtained by subpoena or compulsory court process is subject to an implied obligation to use it only for the purpose of the litigation.¹³ It is arguable – depending on a number of factors – that using such material to train an AI model contravenes your obligations under these principles. Even if you are satisfied that model training will not take place and information is appropriately secure, it may be prudent to discuss any use of an AI tool for the review of such material with opposing solicitors in order that they can raise any objections in advance.¹⁴
- **Sensitive and proprietary data:** Some digital know-how costs a lot of money to create. Allowing an AI to train on it can allow the AI (or a data broker) to profit from that investment, even if no breach of confidentiality as it is traditionally understood occurs.

Ensuring a complex tool only has access to a limited part of your firm's information can require expert assistance. Copilot, for example, by default has access to anything a user can access on the network.

3. Vendor and Contractual Considerations

- **Vendor Reputation:** Is the vendor reputable, and do they have a track record of working with law firms?
- **Cost:** Is the pricing on an introductory or honeymoon basis? This goes beyond “free for three months” style arrangements. Many AI projects and companies are running at a loss. Can or will the vendor give a prediction of likely pricing once the introduction to market phase is over?
- **Indemnity:** Does the contract include indemnities for errors caused by the AI tool? (Note: it probably won't, but you can ask.) Be wary of warranties or indemnities flowing the other way, which may not be covered by your PI insurance. For example, you may “warrant that no Personally Identifiable Information” is uploaded to the system, even if such information is essential to the functioning of the product.
- **Who owns output information and documents?** This should be expressly addressed in the terms of use, and if not, in a written agreement with the vendor.
- **Termination:** Can the firm easily terminate the use of the tool, recover appropriate records and ensure proper deletion of any stored client data?
- **Longevity:** A larger provider is likely to have more resources to ensure security and reliability, and greater stability as a commercial partner.

¹³ QLS Website: Can I use documents obtained by discovery or the court process for a purpose that is collateral to the proceedings? Located here: [Harman Undertaking Note](#); see also paragraph 9 of the [Supreme Court of WA Guidelines for the Use of Generative Artificial Intelligence](#).

¹⁴ As at November 2025, the [NSW Supreme Court Practice Note SC Gen 23](#) prohibits the entry of documents the subject of the implied undertaking or suppression orders into any Gen AI program (as defined in that Practice Note) unless the practitioner is satisfied of certain matters. Other jurisdictions do not explicitly state this (as at April 2025), but the NSW Practice Note provides helpful guidance as to prudent practice and paragraph 10 of the [Supreme Court of WA Guidelines](#) cautions against this. Consideration should be given to disclosing any intended use of such tools as part of processes required by the court (for example under the Supreme Court of Queensland's [Practice Direction 18 of 2018](#)).

4. Professional and Statutory Obligations

- **Competence:** Does the firm have the necessary skills and understanding to use the AI tool effectively?¹⁵ Training may be needed to address skill gaps.
- **Supervision:** Supervision is in two layers: first, someone taking personal responsibility for AI generated output, and checking it carefully. Second, the usual firm supervision systems¹⁶ to ensure that those checks are actually carried out as required by the policy.
- **Transparency with clients and other stakeholders:** Are clients informed of the use of AI in their matter where appropriate?¹⁷
- **Limits on use:** Is the proposed use in accordance with any Court Practice Directions or guardrails binding on the client?¹⁸

5. Risk Mitigation

- **Audit and Monitoring:** How will you monitor the AI tool's performance and outputs? Is the person(s) doing that checking (a) sufficiently versed in the law and cognisant of the individual facts to spot errors? (b) given enough time to conduct a proper check?¹⁹
- **Error Management:** What is the protocol for identifying and correcting errors made by the AI tool? Are prompts, logs and copies of drafts preserved?
- **Backup Systems:** Are there alternative methods or systems in place if the AI tool fails or is suddenly withdrawn?

6. Decision-Making and Accountability

- **Responsibility:** Has the firm clarified accountability for decisions involving AI outputs and overall effectiveness of the system? This goes a bit further than just having someone checking for errors in each document. It requires a high-level supervision of all the working parts to make sure (a) use is not creeping outside approvals (b) potential for greater savings are noted when appropriate (c) the overall use of the system remains appropriate.
- **Record Keeping:** What records of the system use (including prompts and responses) is possible? Can this be easily incorporated into the client file or supplied to a Court if asked to explain the provenance of AI generated material?²⁰

¹⁵ Competence is in two domains: (1) using the tool and (2) ensuring that any errors can be identified. While a tool that assists with research, precedents and drafting might appear to be an attractive way to break into new areas of law, the need for line-by-line checking of drafts and outputs means that this is unrealistic while using this first generation of AI systems. Even where the vendor offers an expert check, caution is required. Firstly, the expert has not spoken to your client and is reviewing the output in a vacuum. Secondly, the cost of providing that service means that vendors may not be able to keep providing it in the longer term.

¹⁶ See *R (Ayinde) v London Borough of Haringey* [2025] EWHC 1383 for discussion of the fact that supervision obligations rest on the firm hierarchy, and solicitors instructing counsel to check output.

¹⁷ "Where appropriate" is nuanced. AI features in software is becoming ubiquitous. Even data handling and storage systems without an overt AI element may be provided under terms which allow unacceptable data use. Your privacy impact analysis should extend to all software in use. The cost of this means that dealing with fewer vendors may be appropriate.

¹⁸ Paragraph 8 of the [Supreme Court of WA Guidelines for the Use of Generative Artificial Intelligence](#) imposes personal responsibility for the accuracy of content of any court documents on the litigant or practitioner who signs or certifies a document, files a document with the court or otherwise relies on a document's contents in proceedings.

¹⁹ Time spent carefully checking automated output must be reflected in the oversight party's KPIs to avoid it becoming a tick-n-flick exercise.

²⁰ For an example, see para 22(b) of the NSW Supreme Court Practice Note SC Gen 23 on record keeping for experts using Gen AI tools.

7. Training and Awareness

- **Staff Training:** Have all relevant staff been trained to understand the AI tool's functions, risks, and the limits on approved use? It is likely that ongoing training is required.
- **Policy Development:** Have you updated internal policies and procedures to incorporate AI usage?

8. Ongoing Assessment

Do you have a process for periodically reviewing the tool's compliance with ethical and professional obligations? In particular,

- Will you be periodically monitoring for any changes to data access permissions and Terms of Service? (Note: task based AI systems can be set to warn you when documents change); and
- Will you be monitoring how AI is used by your firm to ensure that it stays within any guardrails that have been set? Risk mitigation will need to evolve with the technology and as the firm gains confidence and proficiency with the system.

9. Final Steps:

- **Document:** The decision-making process for adopting the AI tool, including risks, mitigations, and approvals. Map review intervals.
- **Implement:** Design and implement a training plan.

AI usage & verification statement: This document was prepared with the assistance of CoPilot, Chat GPT 4o, 5 and Perplexity.ai running Claude 3.5 Sonnet. Uses included: (1) a document review to identify SME self-assessment best practice, (2) a summary of common self-assessment elements (3) re-drafting assistance. Initial verification: David Bowles, QLS Ethics. Approval: AI committees/working groups/relevant staff attached to each Law Society / Institute that has adopted this guide.

Annexure A: Public vs Private AI models.

An AI model is the part of an AI service that provides most of the basic functionality. A model can be extremely expensive, costing millions of dollars to create. Many sellers of an AI service will not have trained their own model but will be accessing someone else's, adding additional functionality and re-selling the result. (Alternate terminology: Open / Closed). There are two basic access options:

- **Public:** Accessing an AI model hosted by a provider and made available for general use, such as ChatGPT. The functionality a re-seller gets is largely similar to that obtained by a person who subscribes directly, although pricing and terms of service might be very different.
- **Private:** Using a licensed copy of an AI model that is exclusive to the reseller (or supplied directly), customised or deployed in a way that ensures greater exclusivity and control, often incorporated into the service provider's infrastructure.

A person using, say, Claude, Chat GPT or Gemini will for the most part be accessing the Public version of a model.

A Private model licensed to a re-seller can operate in a variety of ways. It must be hosted on infrastructure, which may or not be controlled by the licensee entity. It may interact with other models and/or have a number of user interface, data retrieval and search options. How the overall system is implemented will affect the degree of control any single party has. Any promises the licensee/vendor makes about confidentiality need to be examined in light of the other providers involved.

Public AI Model limitations

- **Accessibility:** Public AI models are hosted on shared platforms and are accessible via APIs or web interfaces.
- **Data Handling:** User input and interactions may be processed on shared infrastructure, potentially exposing data to security or privacy risks. Wide range of privacy terms.
- **Cost:** Often low-cost or free, with paid plans for higher usage or advanced features.
- **Flexibility:** Suitable for general-purpose tasks and experimentation but may not meet specialised business needs.
- **Examples:** ChatGPT (OpenAI), DALL·E, or APIs from cloud providers like AWS or Google.

Pros:

Easy to access and implement. No need for significant infrastructure investment. Benefit from cutting-edge updates and community-driven improvements.

Cons:

Limited data privacy and control. The model provider will often use interactions to train or improve its model or even be at liberty to sell the interaction data to other companies. Model outputs are generalised, not tailored for specific industries or needs. Risk of data exposure if sensitive data is processed.

Conclusion: for the most part it is not appropriate for confidential information to be uploaded to a Public model, either directly or via a re-seller. If conducting research, structuring documents etc. no identifying particulars should be included in prompts or uploaded.

Feature	Public AI Models	Privately Licensed AI Models
Access	Open or shared via API	Restricted to licensee or organisation
Cost	Lower or usage-based	Higher upfront and operational costs
Data Privacy	Limited guarantees	High, with strict isolation
Customisation	Limited	Fully customisable
Infrastructure	Hosted by provider	Hosted privately or in a secure vendor environment
Use Case	General-purpose	Specialised
Compliance	May not meet strict regulatory needs	Designed to meet specific compliance standards

Annexure B: example AI disclosure / opt-out consent.

We use <<name of tool>> for limited tasks such as summarising material, drafting outlines and preparing chronologies <<ensure this list is accurate, and updated if firm use expands.>>.

Unless expressly stated otherwise, all AI-assisted content is reviewed by a solicitor, and any citations are independently verified. <<if applicable: We do not use AI to generate a witness's evidence or for any other use prohibited by a court direction.>>

We will not enter confidential information into AI tools unless it is kept within a controlled environment and not used to train any external AI models. Our firm's AI may be trained on all output created while working on your matter, including data used to create such output.

Please tell us if you prefer we do not use AI on your matter. If so directed, a revised cost or time estimate may apply to the work.

It is extremely important that you discuss AI use by you, your staff or consultants (such as an Accountant) when preparing material, briefing us or researching legal issues. This is essential to protect the integrity of the evidence or information which your matter rests and to ensure that the protections which apply to legal communication between client and solicitor are not lost.