Law Firm AI Governance Diagnostic

Summit Guard

A practical diagnostic worksheet for law firms using or considering ChatGPT, Claude, Microsoft Copilot, Gemini, legal AI tools, transcription tools, and AI-enabled practice software.

Use this diagnostic to support a leadership discussion about AI visibility, client-data boundaries, Microsoft 365 exposure, human review, vendor settings, and the evidence the firm should be able to produce.

> **Governance-only note:** This resource provides practical AI governance guidance. It does not replace legal, professional, or client-specific advice.

How to use this diagnostic

For each section, mark the current position:

- **Green:** clear owner, documented rule, and evidence available.
- **Amber:** partly defined, but evidence or ownership is incomplete.
- **Red:** unclear, informal, or not yet reviewed.

Record the top five actions at the end. Keep the first pass practical rather than perfect.

1. AI tool visibility

- ■ The firm has a current list of AI tools and AI-enabled software in use.
- ■ The list includes ChatGPT, Claude, Copilot, Gemini, legal research tools, transcription tools, document tools, practice platforms, and browser extensions.
- ■ Personal accounts are distinguished from firm-managed accounts.
- ■ Each approved or restricted tool has a business owner.
- ■ Each tool has an approved-use status and review date.
- ■ Staff know how to request approval for a new AI tool or feature.

Notes / gaps:

2. Client and matter data boundaries

- ■ Staff have plain-English rules for what client or matter information must not be entered into open or unapproved AI tools.
- ■ The rules cover privileged or confidential material, draft advice, correspondence, evidence, negotiation material, personal information, and commercially sensitive information.
- ■ Client-specific restrictions can be checked before AI is used.
- ■ Approved use cases explain when information must be de-identified, summarised, or replaced with synthetic examples.
- ■ Staff know who to contact when a proposed use case is unclear.

Notes / gaps:

3. Approved and restricted use cases

- ■ Low-risk approved use cases are documented.
- ■ Restricted use cases are documented, including summarising matter documents, drafting client-facing content, analysing client material, and using meeting transcripts.
- ■ Prohibited use cases are clear enough for staff to follow.
- ■ Restricted use cases have an approval path.
- ■ Exceptions are recorded.

Notes / gaps:

4. Microsoft Copilot and Microsoft 365 exposure

- ■ Copilot rollout scope is documented by users, teams, and practice groups.
- ■ High-risk SharePoint sites, Teams spaces, OneDrive folders, mailboxes, and shared mailboxes are identified.
- ■ Broad groups, inherited permissions, guest users, and old sharing links have been reviewed for high-risk locations.
- ■ Sensitive internal repositories such as HR, finance, partnership papers, pricing, strategy, and templates are included in the review.
- ■ Staff understand that Copilot can work with content a user can already access.
- ■ Permission review actions are tracked before pilot expansion.

Notes / gaps:

5. Human review before client use

- ■ AI-generated material is checked by an accountable person before use in client work.
- ■ Review expectations are stronger for higher-risk work.
- ■ Reviewers check accuracy, missing context, assumptions, confidentiality, sources, and client-specific restrictions.
- ■ Important outputs have review evidence that can be inspected later.
- ■ Staff are reminded that AI outputs can be incomplete, outdated, biased, or wrong.

Notes / gaps:

6. Vendor, logging, and retention questions

For each approved or restricted AI tool, the firm has considered:

- ■ What data is sent to the vendor or model provider.
- ■ Whether firm or client data is used to train or improve models.
- ■ Whether training or product-improvement use can be switched off for firm data.
- ■ Where prompts, files, outputs, transcripts, and logs are stored.

- ■ How long information is retained.
- ■ Which administrator controls, access logs, export options, and user-management features are available.
- ■ How access is removed when a user leaves or changes role.
- ■ How material feature changes are communicated.

Notes / gaps:

7. Incident and exception handling

- ■ Staff know how to report accidental or unclear AI use.
- ■ The firm records which tool was used and what information may have been entered.
- ■ Client or personal information issues have a clear escalation path.
- ■ Exceptions record who approved the use and why.
- ■ Lessons from exceptions or incidents are fed back into staff guidance.

Notes / gaps:

8. Leadership evidence pack

Leadership should be able to inspect:

- ■ AI tool register.
- ■ Approved, restricted, and prohibited use cases.
- ■ Client-data boundary guidance.
- ■ Vendor setting notes.
- ■ Copilot permission review notes.
- ■ Staff briefing material.
- ■ Human review examples for important outputs.
- ■ Exception and incident log.
- ■ 30/60/90-day action tracker.

Notes / gaps:

Top five actions

1.
2.
3.
4.
5.

Suggested 30/60/90-day plan

First 30 days: confirm tool visibility, write client-data boundaries, and pause unclear high-risk use cases.

Next 60 days: review Copilot exposure, document vendor settings, and define review evidence for client-facing outputs.

Next 90 days: brief staff, update workflows, review exceptions, and present leadership with a current risk-and-action view.

Need a practical starting point?

Summit Guard helps law firms turn AI usage into visible governance, clear data boundaries, human review points, and leadership-ready evidence.

Contact Summit Guard: https://summitguard.com.au/contact?source=/insights/law-firm-ai-policy-quick-start