Law Society of Ireland

SHADOW AI IN LEGAL PRACTICE

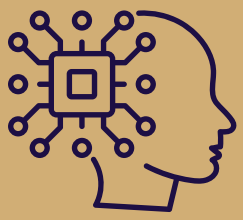

What it is, why it matters, and what to do about it

Law Society of Ireland, Blackhall Place, Dublin 7

Tel: +353 1 881 5768 | Email: solicitorservices@lawsociety.ie | www.lawsociety.ie

SHADOW AI IN LEGAL PRACTICE

What Is Shadow AI?

Shadow AI is the use of artificial intelligence tools in your practice without formal approval, governance, or management visibility. It is not a niche problem. Multiple workplace surveys suggest that a significant proportion of professionals who use AI at work do so without formal disclosure to their employer, and legal practice is unlikely to be an exception.

Shadow AI in a law firm typically looks like this:

- A solicitor pasting a client letter into a consumer chatbot to get a first draft.
- A trainee using a free AI summarisation tool to condense a lengthy document.
- AI features switching on automatically inside tools your firm already subscribes to — Microsoft 365 Copilot, document management systems, legal research platforms — without a deliberate procurement decision.
- A personal device being used to access AI tools that would not be approved on a firm device.

In each case, client data or privileged communications may have been processed by a third party under terms that were not designed to provide the confidentiality protections, data residency commitments, or restrictions on model training that legal practice requires.

Why This Is a Professional Obligation Issue

Shadow AI is not primarily an IT risk — it is a professional conduct risk. Three obligations are directly engaged:

Obligation	Relevance to AI use in practice
Client confidentiality	Entering client data into a third-party tool may constitute disclosure of confidential information. Consumer AI tools are not typically offered with a Data Processing Agreement or confidentiality undertaking adequate for professional practice. The obligation to protect client confidences is stricter than GDPR — it is a personal professional duty that cannot be delegated.
Legal professional privilege	Sharing privileged material with an AI vendor — even inadvertently — may constitute a waiver of privilege. Consumer tools are generally not structured to provide adequate safeguards against this risk.
Duty of competence	Using an AI tool whose limitations, data handling practices, and terms of service you have not reviewed may engage your duty to ensure that the work you produce meets the required standard.

> The Law Society's Guidelines for the Use of Generative Artificial Intelligence by Solicitors (2025) are clear: professional obligations apply to AI use in the same way as to any other tool or process in practice. The technology does not change the obligation.

What You Need to Do

The starting point is visibility. You cannot manage an AI risk you don't know exists. Every practice — regardless of size — should be able to answer three questions:

- What AI tools are being used across this practice, including by individual staff members and embedded in existing software?
- What data are those tools processing, and under what contractual terms?
- Which of those tools presents an unacceptable risk and needs to stop, and which can be formalised with the right governance in place?

A structured audit does not need to be complicated. For a small firm, a direct conversation with staff and a review of subscriptions and browser extensions will surface most of what you need to know. For larger practices, a formal survey and IT review is more appropriate. The important thing is to document what you find and what you decide — the record of a governance decision is itself evidence of competent practice.

Four Outcomes for Every Tool You Find

Every AI tool identified should receive one of four action classifications, applied consistently across your practice:

BLOCK	FORMALISE	ADOPT	MONITOR
Cease use. High risk, no viable governance path.	Value confirmed. Governance required before continued use.	Low risk, high value. Sanction with appropriate guardrails.	Borderline. Review within 30 days and reclassify.

Examples:

BLOCK — A fee earner has been using a free-tier AI tool to draft client correspondence. No DPA, inputs retained for model training, data processed outside the EEA. Client names and matter details have been entered. Use stops immediately.

FORMALISE — The firm is running an AI writing assistant embedded in its productivity suite. The vendor offers a paid enterprise tier with a DPA and EEA data residency, and the tool has clear value. But no one has signed the enterprise agreement and there is no policy governing its use. The tool can stay — once the agreement is signed and a usage policy is in place.

ADOPT — The firm uses an AI-assisted legal research platform under a signed enterprise agreement: DPA in place, EEA processing confirmed, no training on customer inputs. It is used for research queries only — no client-identifying information entered, outputs always verified. Risk is low, governance is sound, use is sanctioned.

MONITOR — A procured tool has quietly added an AI-assisted feature the firm did not knowingly adopt. The vendor has been asked about the data handling terms but has not yet responded. No client data is currently entering the feature. Classification deferred pending a written answer from the vendor within 30 days.

The goal is to move from unmanaged exposure to governed, productive use. Informal AI adoption is often a signal that staff are identifying real process gaps — workload pressure, template shortfalls, research friction. The classification system captures the risk while preserving that signal. Blocking a tool without providing an alternative is likely to drive continued informal use rather than eliminate it.

What Good Governance Looks Like

A practice with sound AI governance should be able to confirm all of the following:

Every AI tool in use is known to practice management and has been assessed for data protection and professional obligations compliance.
Consumer AI tools are not used with client data, privileged material, or court documents without an enterprise agreement that includes appropriate data protection and confidentiality terms.
Tools that process client data operate under a written Data Processing Agreement (GDPR Article 28) with a confirmed data residency position.
Staff understand which tools are approved, which are restricted, and why. Clear guidance reduces informal workarounds.
The firm's AI tool inventory is reviewed at least annually, and whenever a new tool is introduced or a vendor's terms change.

> **EU AI Act — Current Position**
>
> Prohibited AI practices have been banned since February 2025. General-purpose AI obligations have applied since August 2025. The main obligations for high-risk AI systems are currently set to apply from 2 August 2026. The European Commission's Digital Omnibus proposal — not yet enacted — proposes deferring these to December 2027 at the earliest. The current statutory deadline remains 2 August 2026.

This document does not constitute legal advice. It is a practical guide to help solicitors and practice managers understand and manage shadow AI risk. Outputs of any audit undertaken should be reviewed against your firm's specific circumstances and current Law Society guidance.

Law Society of Ireland

AI GOVERNANCE TOOLKIT

Tool 1.1

Shadow AI Audit — Summary Report Template

Know What's in Your Practice

Solicitors in private practice · In-house solicitors · Practice managers · Managing partners

Date completed: _____

This document does not constitute legal advice. It is a structured framework to help solicitors and practice managers assess their current AI exposure and governance position. Outputs should be reviewed against your firm's specific circumstances and current Law Society guidance.

About This Document

This template guides you through a structured audit of AI tools used in your practice. It complements Tool 1.1 in the Excel workbook, which contains detailed compliance checklists and risk matrices. Use this summary report to capture key findings and develop a governance roadmap.

SCOPE: This audit covers all generative AI tools and services used in your practice — both consumer tools (ChatGPT, Copilot, etc.) and enterprise solutions. It examines data exposure, compliance with professional obligations, GDPR compliance, and training data practices.

Section 1: Practice Overview

Firm name	
Number of solicitors	
Total staff	
Practice areas	
Audit conducted by	
Audit date	
Audit methods	

Section 2: Executive Summary

[Write a 3–5 sentence summary of audit scope, key findings, and recommended risk mitigation priority.]

BLOCK	FORMALISE	ADOPT	MONITOR
[count]	[count]	[count]	[count]

Section 3: Professional Obligations Assessment

Obligation	Source	Tools That Engage It	Risk Level & Mitigation
Client confidentiality	S.41		
Legal professional privilege	Common law / S.11 Legal Services Regulation Act 2015		
Duty of competence	S.64		
Duty of supervision	S.66		
Duty to the court	S.67		
Client communication	S.68		
Conflicts of interest	S.76		

Section 4: Data Protection Assessment

AI Tool	Lawful Basis (Art.6)	DPA (Art.28)	Transfer Mech.	DPIA Required?	Status/Notes

Add additional rows as needed for other AI tools.

Section 5: Detailed Findings

AI Tool	Data Exposure	Agreement	EEA Transfer	Trains on Input?	Risk Score	Classification

Risk Score: 1–10 scale (1 = low, 10 = critical). Classification: BLOCK (stop use immediately) | FORMALISE (implement safeguards) | ADOPT (approved for use) | MONITOR (periodic review).

Section 6: Key Risks Identified

[Summarise key risks arising from the audit. Examples of risk categories:

Consumer tools (ChatGPT, Copilot, Gemini, Claude) — no DPA, training on input, no EEA safeguards. Privilege waiver risk — disclosing privileged communications to third-party AI vendors. Court submissions — potential exclusion of AI-generated content if not properly reviewed. EEA data transfers — requirement for Standard Contractual Clauses or Adequacy Decision. Training on input data — vendor models may be improved using your client data. Personal devices — shadow IT and uncontrolled access to AI tools. Conflicts of interest — inadequate conflict checking in multi-matter environments.

> For detailed guidance on professional obligations and AI risk, refer to the Law Society of Ireland's AI Governance Framework (December 2025) and current Practice Notes.

Section 7: Recommendations

Immediate (within 2 weeks)

Block classified tools: Cease use immediately. Remove access. Notify relevant staff.

Short-term (within 1 quarter)

Formalise classified tools: Implement Data Processing Agreements, limit access, configure safeguards, log usage.

Ongoing governance

Quarterly review: Update the AI Tools Register (Tool 2.1). Monitor vendor updates, regulatory changes, and incident reporting.

Section 8: Next Steps

Action	Owner	Deadline	Prof. Obligation

> Use Tool 2.1 (AI Tools Register) to maintain a living record of all approved and monitored tools. Use Tool 2.3 (Data Processing Agreements) to document vendor contracts and DPA terms. Attend Law Society workshops on AI governance and professional obligations.

Sign-Off

Role	Name	Signature	Date
Prepared by			
Reviewed by			
Approved by			

This resource was developed by Acuity AI Advisory in consultation with the Law Society of Ireland.

Acuity AI Advisory
Buttermilk Lane,
Downings North
Prosperous,
Kildare W91 K2N7
E hello@acuityai.co
W www.acuityai.co

Law Society of Ireland,
Blackhall Place,
Dublin 7,
D07 VY24

E solicitorservices@lawsociety.ie
W www.lawsociety.ie