

FIVE WAYS LAW FIRMS LEAK CLIENT DATA INTO AI (WITHOUT KNOWING IT)

*A confidentiality audit for managing partners and firm administrators —
findable in one afternoon, fixable in one quarter*

Your firm almost certainly has an AI confidentiality problem. Not because anyone is reckless — because AI arrived *inside* tools you already trusted, and through habits nobody has updated. Privilege doesn't care how the information got out. Here are the five leaks, in the order you're most likely to find them.

Leak #1 — The personal chatbot habit

An associate pastes a contract into a free AI chatbot to summarize it. A paralegal pastes a client email thread to "clean up the timeline." On personal or free-tier accounts, that text leaves the firm, lands under consumer terms that may permit training, and exits every confidentiality control you have. **You can't even prove what was pasted.**

Find it: anonymous survey, one question: "Have you used a personal AI account for work text this month?" Brace yourself. **Fix it:** provide an approved, enterprise-terms alternative *first*, then prohibit the rest. A ban with no alternative just drives the habit underground.

Leak #2 — The AI your existing software quietly switched on

Your document management, email, research, and billing platforms all added AI features in the past 18 months — often enabled by default, processing firm content through the vendor's AI pipeline under terms nobody at the firm has read since the feature appeared.

Find it: list your top ten tools; check each one's release notes for "AI," "assistant," or "copilot"; then find where that processing happens and under what terms. **Fix it:** disable what you haven't vetted. Re-enable deliberately, in writing, tool by tool.

Leak #3 — The vendor's vendor

The e-discovery platform, the deposition-summary service, the intake chatbot on your website — each may pass client data to *their* AI subprocessors. Your engagement letter promises confidentiality; their subprocessor list is where that promise goes to be tested.

Find it: ask every legal-tech vendor two questions in writing: "Which AI subprocessors receive our data?" and "Is our data used to train any model?" **Fix it:** no written answer, no client data. This is the same discipline you'd apply to a contract attorney; apply it to software.

Leak #4 — The intake and marketing funnel

AI note-takers joining consult calls, chatbots collecting matter details from prospects, AI-transcribed voicemails — often before a conflicts check, often stored in marketing tools with the weakest terms in your stack. Prospective-client duties attach anyway.

Find it: map one prospective client's path from first contact to engagement letter. List every system that touched their story. **Fix it:** the intake path gets the same confidentiality standard as the matter file — or it collects less.

Leak #5 — The "anonymized" workaround

Someone conscientious strips the client name before pasting the fact pattern into a chatbot. But a fact pattern *is* the confidence: an acquisition of this size, in this industry, with this regulator circling, is identifiable to anyone who knows the market — and now it sits in an unlogged consumer system.

Fix it: the rule isn't "remove the name." The rule is **client matter content goes only into systems the firm controls, under enterprise terms, with a log.**

THE PATTERN — AND THE FIX THAT HOLDS

All five leaks share one root cause: **the AI lives outside the firm, so client information has to travel to it.** The durable fix inverts the geometry — AI that works *inside* systems the firm owns: reading only what it's permitted, logging every action, keeping matter data exactly where your engagement letters promise it stays. Then the productivity gain and the confidentiality duty stop being at war.

Firms that get this right won't just avoid the grievance. They'll be the ones who can tell clients, in a pitch, precisely how their matters are — and aren't — touched by AI. That answer is about to be a differentiator.

Slingr builds custom software for law firms and legal operations where the AI works inside the walls: firm-owned systems, matter-level permissions, every action logged. slingr.io