

WHITEPAPER

The AI answer on your renewal is a written representation about your firm. Decide now whether it is true.

For managing partners and firm administrators at Pennsylvania and
New Jersey law firms

Executive summary

Your firm already uses AI. Individual use among legal professionals went from 31% to 69% in a single year, while 43% of firms have no formal AI policy at all and only 9% have one that is written and enforced.¹ You are not deciding whether to adopt. You are deciding whether the use inside your firm is documented or undocumented.

That decision has acquired a date, and it is not one you set. Carriers have started asking about AI at renewal: CNA has been sending supplemental AI questionnaires since 2025 covering tools in use, governance policy, attorney training, work-product verification, and incident response.² Seven of thirteen major LPL carriers reported AI-related claims increases in the same survey.

The obligation is older than the questionnaire. ABA Formal Opinion 512, issued July 29, 2024, holds that firm leadership must establish written policy governing generative AI use and supervise lawyers and staff in that use, under Rules 5.1 and 5.3.³ That duty exists today at your firm whether or not anyone has written anything down.

Our recommendation: build the policy set and the evidence underneath it before the application arrives, not after. That is **AI Policies, \$4,900, fixed fee:** acceptable use, the tool-and-model inventory and vendor risk, mapped to NIST CSF 2.0 and CIS Controls v8.1, every line reviewed by a vCISO who signs an attestation. The signature is the deliverable. Maintaining it afterwards is **\$495 a month**, month to month.

The situation

Four things happened, all of them external to your firm, and all of them dated.

The carriers moved. Per a 2026 EPIC carrier survey reported in specialty insurance trade press, eight of thirteen major lawyers professional liability carriers reported overall claim-frequency increases (the first rise in five years), and CNA has sent supplemental AI questionnaires since 2025.² This is trade-press sourced; your broker can tell you exactly what your carrier is asking this cycle.

The courts started pricing it. At least \$145,000 in sanctions for fabricated AI citations in the first quarter of 2026 alone. *Whiting v. City of Athens*, 6th Cir. No. 25-5425, March 13, 2026: two attorneys sanctioned \$15,000 each plus appellee fees and double costs over 24 or more fabricated or incorrect citations, reported as the largest federal appellate AI-hallucination sanction on record. The Oregon Court of Appeals has effectively adopted a schedule: \$500 per false citation, \$1,000 per fabricated quote.⁴

Your clients started flowing the obligation down. The Association of Corporate Counsel publishes a sample AI guidelines template for general counsel to insert into outside counsel guidelines.⁵ The clause one client sent you is likely to be substantially the clause the next four send. No survey quantifies how many firms have received one, and we are not going to invent that number.

Your own practice management vendor started selling it to you. Clio launched Clio Work as a standalone AI product for solo, small and mid-size firms on April 21, 2026, at \$199 per user per month bundled with Clio Manage.⁶ Adoption is no longer a decision anyone is asking your permission for.

Four bodies in three states have already said the same thing the ABA said: verify output, protect confidentiality, disclose and obtain consent, supervise.⁷ None of them is waiting for legislation.

What most firms miss

Most firms respond to this by buying a policy. That is the cheap half, and it is the half nobody checks.

The unconsidered need is the evidence layer. A carrier supplemental does not ask whether you have a policy. It asks what tools are in use, who was trained and when, what happens before AI-assisted work product leaves the firm, and what incidents you have had. Every one of those is answered from an artifact, not from a document: an inventory, a sign-in sheet, a completed verification record, a log. A firm can hold a well-drafted policy and be unable to answer four of the five sections.

Here is the part firms have not costed. **The answer on the application is not a form. It is a written representation about your controls, made to a party with a financial interest in checking, and it sits in the file for the life of the policy.** The honest "not yet, here is our timeline" is cheap; most underwriters and most general counsel will take it. The optimistic "yes" is what creates the exposure, and it is an exposure the firm manufactured itself, in writing, under signature.

The precedent worth knowing sits on the cyber side, and we will state it precisely because the precision is the point. In 2022 Travelers and International Control Services jointly stipulated to void an active cyber policy from inception after ICS misrepresented its use of multi-factor authentication on the application.⁸ That was a stipulation between the parties, not a decision on the merits, and it was cyber rather than professional liability. **There is no documented, named case of an LPL claim denied specifically over an AI-related application misstatement, and we will not imply one exists.** The narrow lesson holds anyway: insurers act on application misstatements.

The second thing firms miss is that the ban does not work. It is the instinctive partner response and it is the worst available answer, because it does not reduce use: it converts documented use into undocumented use, which is precisely the gap between 69% and 9%.¹ A prohibition nobody can enforce produces the same behaviour with none of the record.

What good looks like

A standard, not a pitch. A firm of 15 to 75 attorneys can hold all of this without hiring anyone, and a firm that never calls us should still be able to work from this list.

1. A policy with seven parts. Scope and definitions, including AI features embedded in software you already own. An enumerated tool list by named product and permission tier: "public chatbots are discouraged" is not enforceable; "consumer ChatGPT is prohibited for any client information, the firm's paid business accounts are approved for non-client research" is. Permitted and prohibited uses by work type, because a deposition summary and a filed brief sit differently. The client consent rule, naming which matter types require consent before information enters a self-learning tool, who obtains it, and where it is recorded. The verification requirement. One named owner with the authority to switch a tool off, and an exceptions process. Incident response. Underneath all seven, an acknowledgement register: who signed, on what date, against which version.

2. A tool-and-model inventory. One table, dated, owned. Product, vendor, underlying model where disclosed, how it was acquired, whether the vendor's terms permit training on your inputs, where data is processed and how long it is retained, what firm data it can reach, who approved it, date of last review. Building the first version takes about a week and is uncomfortable: firms find individually expensed subscriptions, features their vendor switched on by default, and a browser extension somebody installed in 2024. That discomfort is the point.

3. A verification step that names a person. Every cited authority pulled from the citation as written. Every citation read for the proposition it is offered for: the common failure is no longer the invented case, it is the real case cited for something it does not say. Every quotation matched character for character. Every pin cite confirmed. Subsequent history checked. The verifier named on the record, and where the drafter is not the supervising attorney, the verifier is someone other than the drafter. The supervising attorney signs having seen the completed record. Eight lines. It fits on a card.

4. A near-miss log. When a verifier catches a bad citation before it goes out, that gets recorded, not to discipline anyone, but so that "no, we have had no incidents" is a real answer rather than an absence of information.

5. One crosswalk, three questionnaires. The carrier supplemental, the client security questionnaire and the OCG AI clause ask for the same five things. Map your policy sections and artifacts once to NIST CSF 2.0 and CIS Controls v8.¹⁹ and you answer all three by lookup instead of by composition. Do not chase ISO/IEC 42001 this year; it is an emerging procurement requirement, not settled practice at your size.

How we do it

ENGAGEMENT	DETAILS
------------	---------

AI Policies: \$4,900, fixed fee.

About three weeks.

We interview the people who actually use the tools, not just the people who approve them. We build the inventory from expense reports, tenant admin data and interviews, which is where the surprises are. We draft the acceptable use policy against your firm: your matter types, your system of record, your named owner, and the verification step AI-assisted work passes before it leaves the firm. We review each vendor's terms: whether it trains on your inputs, and where your data is kept. We produce the crosswalk to NIST CSF 2.0 and CIS Controls v8.1 so one answer set serves the carrier, the client and the OCG. Then our vCISO reads every line and signs an attestation.

WHAT THIS DOES NOT INCLUDE.

It is not legal advice and it is not a bar ethics opinion; we do not give either, and your outside ethics counsel should read the result. It does not cover the rest of your security program; that is Security Policies at \$4,900, or the Full Program at \$8,900, which includes the AI Policies. It does not deploy or configure any AI tool; that is Claude Setup, from \$2,900. It does not build an agent. It does not include ongoing maintenance: a policy set is only true on the day it is signed, and keeping it current is the **GRC Platform at \$495 a month**, month to month: policies in a portal, an annual review timed to your renewal rather than a random anniversary, regulatory monitoring, and questionnaire support when the next one lands. If a contract term is the only thing keeping you, we have already lost.

Two ways we bill and only two: fixed fee for known shapes, a monthly block for longer engagements. Pass-throughs bill at cost. We do not mark up model or API consumption.

Questions

QUESTION	ANSWER
"Why not our existing IT provider?"	Ask them who signs. Most MSPs will configure tools and decline to attest to a control. Our deliverable is an attestation from a named vCISO against a named framework, the same process we use to attest security policy for a law firm client. If your provider will sign, use them.
"Can't we just download a free template?"	You can, and it will describe a firm that is not yours. That is the failure mode, not the saving: a document asserting practices you do not operate, handed to a party with audit rights. The \$4,900 buys the interviews, the inventory, the vendor review and the signature.
"What if our attorneys aren't using AI?"	Thirty percent of lawyers at firms with 10–49 lawyers reported using AI in the ABA's 2024 survey, ¹⁰ and individual use across the profession hit 69% by early 2026. ¹ Run the expense-report check in question five of our assessment before you answer that on an application.
"Shouldn't we wait for the bar to issue a rule?"	ABA 512 issued July 29, 2024. Pennsylvania and Philadelphia issued Formal Opinion 2024–200 on June 18, 2024. New Jersey issued January 25, 2024. New York's task force reported April 6, 2024. ^{3,7} The guidance arrived two years ago. What arrived this year is somebody asking to see your compliance with it.

Signet, a division of Circle Square Consulting. AI that holds up. Radnor, Pennsylvania.

Footnotes

- 1 8am Legal Industry Report, published March 5, 2026: individual use 31%→69%; firm-level implementation 46%; 43% no formal AI policy; 9% written and enforced; 54% no training and no plans to add it.
- 2 2026 EPIC carrier survey reported in specialty insurance trade press (actuary.info, current to May 2026): 7 of 13 major LPL carriers reported AI-related claims increases; 8 of 13 reported overall claim-frequency increases; CNA supplemental AI questionnaires deployed from 2025. Trade-press sourced; confirm with your broker.
- 3 ABA Formal Opinion 512, July 29, 2024 (Rules 1.1, 1.4, 1.5, 1.6, 5.1, 5.3).
- 4 ComplexDiscovery, April 6, 2026, Q1 2026 aggregate of at least \$145,000. *Whiting v. City of Athens*, 6th Cir. No. 25-5425, March 13, 2026. *Couvrette v. Wisnovsky* (D. Or.), \$15,500. Oregon Court of Appeals: *Watson*, December 2025, \$2,000; *Ghiorso*, March 2026, \$10,000.
- 5 Association of Corporate Counsel, "Sample Artificial Intelligence (AI) Guidelines for Outside Counsel," acc.com resource library, verified August 2026.
- 6 Clio Work launched as a standalone product April 21, 2026, \$199/user/month bundled with Clio Manage.
- 7 Joint PBA/Philadelphia Bar Formal Opinion 2024-200, June 18, 2024; NJ Supreme Court Committee on AI Preliminary Guidelines, January 25, 2024; NYSBA Task Force on AI report approved April 6, 2024.
- 8 Insurance Journal, August 30, 2022; stipulated August 26, 2022, dismissed with prejudice.
- 9 NIST Cybersecurity Framework 2.0, February 2024 (no 3.0 exists). CIS Controls v8.1, June 2024 (no v9 exists); CIS publishes an official crosswalk of v8.1 to NIST CSF 2.0.
- 10 ABA 2024 Legal Technology Survey Report, published March 7, 2025: 30% of lawyers use AI overall; 30% at firms with 10-49 lawyers.