




OUTSMART THE BREACH

EXECUTIVE BRIEF | LEGAL SECTOR

Confidential by Default

Closing the AI Governance Gap Before It Becomes a Malpractice Problem

PUBLISHED BY

Katherine Majette, [nextGEM Academy](#)

July 2026

[cyreaux.com](#)

TABLE OF CONTENTS

01 EXECUTIVE SUMMARY4

02 THE GOVERNANCE GAP.....4

03 WHY NOW4

04 THE COMPOSITE CASE 5

05 THE READINESS GAP 6

06 WHAT CLOSING THE GAP LOOKS LIKE..... 7

07 ADDRESSING THE COMMON OBJECTIONS..... 8

08 THE NEXT STEP 9

09 ABOUT CYREAUX..... 9

10 WORKS CITED10

01 EXECUTIVE SUMMARY

Generative AI is already inside small law firms, drafting memos, summarizing depositions, researching case law, and in some cases recording client conversations, all well ahead of any written policy governing that use. This brief maps the readiness gaps CyReauX sees repeatedly across firms at this stage, ties them to the two named professional-responsibility opinions that already govern the space, and lays out a right-sized governance path for firms that will never staff an in-house security team.

*Every month without a policy is a month of undocumented exposure.
Not a neutral holding pattern.*

02 THE GOVERNANCE GAP

Generative AI has moved into small law firms faster than firm policy has. Attorneys and staff are already using AI tools to draft memos, summarize depositions, research case law, and communicate with clients, often without a single conversation about whether that use is permitted, what data it touches, or who signed off on it.

This is not a hypothetical risk. It is current behavior inside firms with no dedicated IT, security, or AI governance staff to manage it. And it is no longer a gray area under professional responsibility rules.

ABA Formal Opinion 512, issued in July 2024, sets national expectations for competence, confidentiality, client communication, and supervision when lawyers use generative AI, including securing informed client consent before their confidential information goes into an AI tool. State and local bars are layering more specific requirements on top: **NYC Bar Formal Opinion 2025-6**, issued December 2025, requires notice and consent before an AI tool records, transcribes, or summarizes a client conversation.

Treat the ABA opinion as the floor. Confirm what your jurisdiction has added on top of it.

03 WHY NOW

Both governing opinions are recent, which is exactly why most firms haven't caught up. The ABA opinion is two years old. The NYC opinion is months old. More is coming.



"We'll deal with it later" is not a fixed target. It's a moving one. Firms that build a governance habit now, rather than reacting to the next opinion, spend less time and carry less exposure in the interim.

ON JURISDICTION

The ABA opinion is the national floor and applies to any ABA-governed attorney. State and local bars are starting to add specifics. NYC 2025-6 is one early example, not a national rule. Build a policy anchored to the ABA baseline, then confirm what your own state or local bar has issued, or is likely to issue, on top of it.

04 THE COMPOSITE CASE

The pattern below is drawn from CyReauX advisory casework. It's illustrative, not a specific engagement, but it reflects what we see repeatedly.

CASE PROFILE

A small law firm. **Fewer than 15 attorneys**, no dedicated IT or security governance staff, no written AI use policy.

Attorneys across the firm have adopted generative AI on their own, using it for research, drafting, summarizing, and client communication. Each attorney made that choice **independently**, so there's no firm-wide visibility into which tools are in use or what data is flowing through them.

Some attorneys use AI notetaking tools on client calls **without a consistent process for telling the client beforehand**.

05 THE READINESS GAP

Four gaps show up consistently in firms at this stage. Each carries a distinct consequence.

01

No Written AI Use Policy

Attorneys and staff have no guidance on which use cases are permitted, restricted, or prohibited. Without that baseline, "competent supervision" under ABA Opinion 512 has nothing to point to.

02

No Shadow AI Visibility

No inventory of tools actually in use. Shadow AI is the default state, not the exception. A firm cannot vet, secure, or account for a tool it doesn't know exists.

- 03

Unmanaged Confidentiality & Privilege Exposure

Without policy or technical controls, nothing prevents confidential client information from being entered into public, unvetted AI tools. This is the exact risk ABA Opinion 512's confidentiality guidance is built to prevent.
- 04

No Client Consent Procedure

ABA Opinion 512 already recommends securing informed client consent before their confidential information enters an AI tool, and specifies that boilerplate engagement-letter consent isn't adequate. Most firms have no consistent, documented process for obtaining that consent at all.

Individually, each gap is manageable. Together, they leave a firm unable to answer the question a regulator, client, or malpractice carrier might reasonably ask: what AI tools are you using, and how do you know client confidentiality and consent are protected?

06 WHAT CLOSING THE GAP LOOKS LIKE

Closing this gap doesn't require hiring an in-house security or compliance team. Most firms this size never will. It requires a governance partner who can build the policy, vet the tools, and keep the program current as AI use and professional guidance keep evolving.

CyReauX's **PRAXIS Tier III** engagement is built for exactly this situation: an ongoing relationship, not a one-time audit. Five components map directly to the four gaps above:

COMPONENT	WHAT IT INCLUDES
Written AI Use Policy	Defines permitted, restricted, and prohibited use cases in plain language attorneys will actually read.
Shadow AI Inventory	A living record of tools in use across the firm, refreshed on a set cadence. Not a one-time survey.

COMPONENT	WHAT IT INCLUDES
Vendor & Tool Vetting	Each tool evaluated on data handling, access, and vendor certification before it's described as vetted. Never a blanket approval.
Client Consent Protocol	A consistent, documented process for notifying clients and obtaining consent before AI records, transcribes, or summarizes a conversation.
Ongoing Review	Scheduled reassessment as new tools, new attorneys, and new bar guidance change the picture. Governance as a habit, not a one-time project.

WHY VENDOR-NEUTRAL MATTERS

CyReauX doesn't implement tools. We assess your environment, produce a custom sanctioned AI stack, and give you the implementation blueprint. The judgment stays independent from the sale of any product, which is exactly what a firm bringing in AI on client matters needs from its advisor.

Any tool a firm is already piloting is evaluated on a conditional basis: what data it touches, what access it requires, and whether the vendor can demonstrate the certifications a law firm needs. No tool is endorsed until that review is complete.

07 ADDRESSING THE COMMON OBJECTIONS

WHAT WE HEAR	WHAT'S ACTUALLY TRUE
<i>"We're too small for this to matter."</i>	Size doesn't change the ethics rules. ABA Opinion 512 applies nationally to any lawyer using these tools, regardless of firm size. And state bars are adding specifics on top.
<i>"No client has asked us about it yet."</i>	The obligation exists independent of whether a client has raised it. The first time it surfaces is often in a malpractice claim or bar complaint. Not a routine conversation.
<i>"We'll write a policy when we have time."</i>	Shadow AI use is already the default in firms without a policy. Every month without one is a month of undocumented exposure. Not a neutral holding pattern.

08 THE NEXT STEP

The firms that get ahead of this aren't waiting for an incident or a bar complaint to force the conversation. They're starting with a confidential Executive Advisory conversation: a chance to map current AI use against ABA Opinion 512 and whatever their jurisdiction has added on top, and to build a governance plan sized to the firm's actual practice.

REQUEST YOUR FREE AI + CYBER READINESS ASSESSMENT

cyreaux.com

Confidential · No obligation · Right-sized for your firm

09 ABOUT CYREAUX

CyReauX helps law firms adopt AI without betting the practice on a vendor's roadmap. We assess your environment, identify which AI tools can be sanctioned for privileged work, and hand you the implementation blueprint. We take no referral fees and hold no vendor partnerships, so the recommendation you get is the one your risk profile supports.

Through PRAXIS, our diagnostic-first engagement model, we serve firms with no IT team or a stretched one that still carry every obligation of a firm that has both: client confidentiality, the duty of technology competence under ABA Model Rule 1.1 Comment 8, and whatever your malpractice carrier and largest clients decide to ask next year. Assessments are grounded in NIST AI RMF, NIST CSF 2.0, and the OWASP LLM Top 10, which matters if your firm is building its own client-facing chatbot or intake tool rather than just buying one. The ongoing governance tier gives firms a standing answer when a partner wants to run a matter through a new tool and no one knows whether that is allowed.

OutSmart the Breach. | cyreaux.com

10 WORKS CITED

1. American Bar Association, Standing Committee on Ethics and Professional Responsibility. (2024, July 29). *Formal Opinion 512: Generative Artificial Intelligence Tools*. Ethics Opinion.
2. New York City Bar Association, Committee on Professional Ethics. (2025, December 22). *Formal Opinion 2025-6: Ethical Issues Affecting Use of AI to Record, Transcribe, and Summarize Conversations with Clients*. Ethics Opinion.
3. CyReauX Advisory Casework. (2026). *Composite advisory patterns across small law firms*. Internal advisory synthesis.