Unsanctioned AI is in your firm. Privilege, liability, and productivity are on the line.

Most law firms have procured AI legal platform like Legora or Harvey, yet a survey conducted by Axiom Law found that 83% of legal teams are using AI tools like ChatGPT, Claude, and Gemini that were not provided by their company. Your firm remains accountable to clients, courts, and the bar for every output. Can you prove it was used under counsel direction, on appropriate matters, with client information protected?

The 3 Ungoverned AI Risks Firms Can't Ignore

01

Shadow AI That Compounds Malpractice and Privilege-waiver Risk

A written prohibition alone won't satisfy your Model Rule 1.6 and 5.1 duties when a court asks whether the attorney took reasonable steps to preserve the privilege. Policy must be paired with provisioned tooling, technical blocks on consumer AI, and documented training. Otherwise the next privilege fight is malpractice exposure the firm answers for.

02

Productivity That Reaches the Bottom Line

Attorneys using AI well are 30 to 40 percent more productive on research, drafting, and review. Attorneys using AI poorly create rework, write-offs, and realization loss. The difference is whether the firm see and steer the data shared with guardrails.

03

Hallucinations That Disrupt the Top Line

To date, 1398 legal decisions involved generative AI producing hallucinated content. When Sullivan & Cromwell filed AI-hallucinated citations in a federal bankruptcy court in April 2026, the apology came from the firm, not the associate. Reputational damage compounds without a defensible record of AI usage. One hallucinated brief could lead to embarrassing sanctions or pro hac vice revocations that puts your client relationships at risk.

WHAT COURTS ARE ALREADY DECIDING

United States v. Heppner (S.D.N.Y., Feb. 2026): exchanges with public AI are not privileged absent counsel direction and enterprise safeguards.

Warner v. Gilbarco, Inc. (E.D. Mich. Feb. 10, 2026): AI-generated content from sanctioned Enterprise platforms is not discoverable.

Wadsworth v. Walmart (D. Wyo., Feb. 2025): three attorneys sanctioned, one with pro hac vice revoked, for AI-fabricated citations.

MARO

AI GOVERNANCE FOR LEGAL TEAMS

Maro: AI oversight built for the way attorneys actually work.

Maro works in the browser, at the prompt, where AI use is actually happening. It preserves privilege and work-product protections by keeping client information out of unsanctioned tools, enforces firm policy in real-time, and produces a defensible record of every interaction. Attorneys get to stay productive in the tools they prefer. The firm gets visibility, control, and proof.

Without Maro	With Maro
Restrictive policies push attorneys toward shadow AI. Productivity gains are uneven and unmeasurable.	**Attorneys work 30 to 40 percent faster when using their tools of choice with guardrails.** Time recovered shows up in realization, not write-offs.
Prompts to public AI are discoverable in litigation. Once submitted, prompts can be subpoenaed and privilege protection collapses.	**Privileged data stays inside firm-licensed AI.** Automagically strike firm-defined data from spilling in real-time.
Hallucinated citations slip past manual review. Sanctions and bar complaints land on the attorney *and* the firm.	**Every AI output traceable to its prompt.** Defensible record of what was asked, what was produced, and what was filed.
Clients are writing AI clauses into OCGs. The firm has policy to point to, not enforcement.	**Auditable and provable AI governance.** Aligned with ABA Formal Opinion 512 and state bar guidance, ready for OCG review.

See what a defensible AI record looks like. seekmaro.com →

How Maro Works

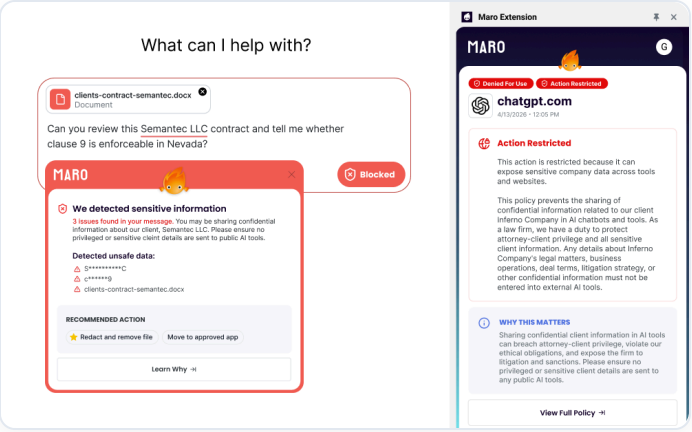

REAL-TIME COACHING

Stop Client Data from Reaching Public AI

When an attorney is about to expose privileged information, Maro pauses the prompt, names the violation, and offers safe paths forward: redact sensitive data or redirect to an approved tool.

FIRM-SPECIFIC POLICY

Tune What Counts as Sensitive for Your Firm

Client names, matter types, and OCG terms become enforceable rules at the prompt. A confidentiality clause for one client enforces differently than another, automatically.

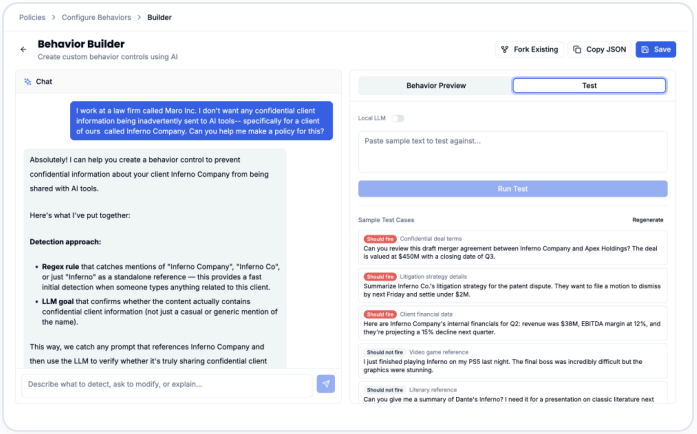

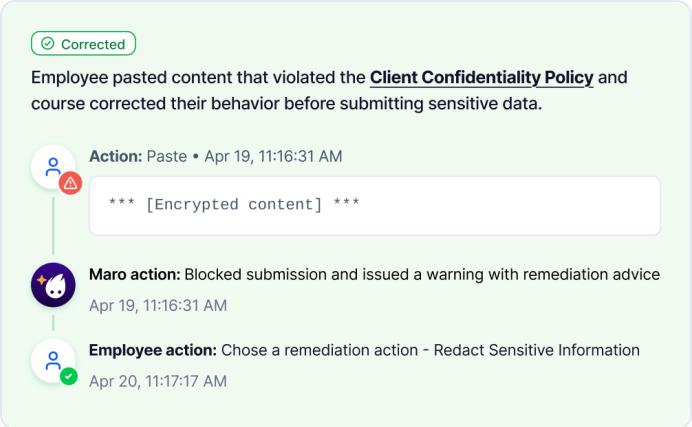

DEFENSIBLE RECORD

Every AI Interaction, on the Record

Who used what, when, and what was prevented. Queryable, exportable, and structured for discovery, bar inquiries, and OCG audits.