[FIRM NAME] Artificial Intelligence Use Policy

From: John A. Kelly

Date: August 4, 2026

Re: Model Law Firm AI Use Policy (research draft for adaptation)

Research draft prepared on a professional courtesy basis. This is a model for adaptation by the firm and its own counsel; it is not legal advice to the firm or any client. Bracketed items require firm-specific decisions.

1. Purpose and Adoption

This policy governs the use of artificial intelligence tools by all personnel of [FIRM NAME] in connection with client work and firm administration. It implements the firm's obligations under the California Rules of Professional Conduct and the State Bar Act, as explained in the State Bar of California Standing Committee on Professional Responsibility and Conduct (COPRAC), Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law (2026), and ABA Formal Opinion 512 (2024). Under Rules 5.1 and 5.3, managerial lawyers must establish clear policies on AI use, provide training, and update the policy as the technology and the rules evolve. This document is that policy. Adoption also has a practical driver outside the Rules: cyber liability carriers increasingly ask on their applications whether the firm has a written AI use policy.

Adopted by [the Managing Partner / the Executive Committee] on [DATE]. Questions go to the AI Policy Administrator designated in Section 3.

2. Scope and Definitions

Covered persons. All attorneys, paralegals, legal assistants, staff, contractors, and vendors performing work for the firm.

Covered tools. Any computer system that generates, selects, summarizes, or acts upon information in response to inputs, including:

1. **Generative AI**: systems that create new content (text, images, audio, code, analysis) in response to prompts, such as ChatGPT, Claude, Gemini, and Copilot.
2. **Embedded AI**: AI features inside tools not usually thought of as AI, including legal research platforms (Westlaw Precision AI, Lexis+ AI), document review systems, practice management software, and word processing assistants.
3. **Agentic AI**: systems that can plan, select tools, and execute multistep tasks with limited or no real-time human direction, including autonomous drafting pipelines, intake bots, and research agents.

COPRAC's 2026 guidance treats all of these as within a lawyer's professional responsibility obligations, and imposes heightened supervision duties as system autonomy increases.

3. Governance

AI Policy Administrator. [NAME], responsible for maintaining the Approved Tool List, vetting new tools, coordinating training, and reviewing this policy at least annually and upon any material change in the rules or the technology (a COPRAC-recognized duty of periodic reassessment).

Approved Tool List. Personnel may use only tools listed on the firm's Approved Tool List (Appendix A) for client work, in the approved configuration. Use of an unlisted tool, or of a listed tool through a personal account, for client work is prohibited. So is enabling the AI features embedded in software not on the Tool List (videoconference assistants, PDF tools, scheduling apps); an embedded AI feature is a tool, and it stays off for client work until vetted. Anyone may propose a tool for approval to the Administrator.

Vendor vetting. Before a tool is added to the Tool List, the Administrator confirms, at minimum:

1. Terms of use and privacy policy reviewed; understand that use alone may constitute assent.
2. Training on firm inputs is disabled or contractually excluded (enterprise or professional tier; zero data retention terms where available).
3. Memory and personalization features identified, and disabled or matter-scoped in the approved configuration. If the tool retains anything across sessions, know how it is partitioned between matters and whether it honors the firm's ethical walls (see Section 4).
4. Data retention period, storage location, and deletion rights identified. Note that some vendors retain inputs for a fixed period even under zero data retention terms for trust and safety reasons; know the period.
5. For add-ins and plugins (a Word add-in, a document management integration): which account the add-in is bound to. An add-in inherits the data posture of that account, not of the desktop software it runs inside.
6. Pricing model (per seat or per token) and usage caps. Token-metered plans can be exhausted before a deadline; know the exhaustion behavior and any overage terms, and note that some vendors charge for staff seats as well as attorney seats.
7. Security posture (SOC 2 or equivalent; encryption in transit and at rest).
8. For agentic tools: what systems and data the agent can access, what actions it can take without approval, and what audit logging exists.
9. Consultation with qualified IT or security professionals where the tool will receive confidential client information (a COPRAC "reasonable efforts" measure).

4. Confidentiality (Rule 1.6; Bus. & Prof. Code § 6068(e))

The core rule. Do not input confidential client information into any AI tool that presents material risk to confidentiality or security, absent the client's informed consent. Reliance on a vendor's marketing assurances is not enough; the Tool List reflects the firm's actual vetting.

Input tiers.

- **Tier 1 (any approved tool):** public information, general legal questions, anonymized hypotheticals with client identifiers, case numbers, and distinctive facts removed or replaced with placeholders.
- **Tier 2 (approved tools in enterprise configuration with training disabled):** client documents and matter facts, where the vetting in Section 3 is complete and no client instruction forbids it.
- **Tier 3 (prohibited absent informed client consent and Administrator approval):** material subject to a protective order or sealing order, information the client has designated as restricted, information subject to statutory confidentiality regimes (e.g., medical, financial, juvenile), and any input into a consumer-tier or unvetted tool.

Client instructions control. Where a client's engagement terms or outside counsel guidelines restrict AI use, those restrictions override this policy for that client's matters. The responsible attorney must flag such restrictions to the matter team at intake.

Agentic access limits. Agentic systems with persistent access to firm email, document management, or client files raise acute risks, including cross-matter disclosure. Such access requires Administrator approval, matter-level or need-to-know scoping, audit logging, and a prohibition on autonomous external transmission of client information (no automated filings, emails to third parties, or data transfers without human review).

Memory and cross-matter contamination. A tool's memory or personalization feature can carry one client's information into a session about another client (a risk ABA Formal Opinion 512 identifies), and can defeat an ethical wall: a screened matter must not surface through the AI layer in an unscreened user's session. Accordingly: (1) memory features on general-purpose tools are disabled for client work unless the Tool List records a matter-scoped configuration; (2) tools integrated with the firm's document or case management system must honor the access restrictions and ethical walls configured there, verified at vetting; (3) continuity comes from the matter file, not from vendor memory. Standing instructions, matter context summaries, and session logs live in the matter file and are loaded into a session deliberately, so each session starts clean and the lawyer controls exactly what the tool sees. (The firm's "Where the Data Lives" appendix maps retention, memory, and zero data retention across the three working modalities: web chat, document add-ins, and agentic harnesses.)

Work product. Prompts, research queries, and AI analytical sessions reflect counsel's mental impressions and strategy and are treated as attorney work product (Code Civ. Proc. § 2018.030). Store session records within the matter file, not in personal accounts.

5. Verification and Court Submissions (Rules 3.1, 3.3; Bus. & Prof. Code § 6068(d); Code Civ. Proc. § 128.7)

No delegation of candor. The lawyer signing a filing remains personally responsible for every citation, quotation, and factual and legal contention, regardless of whether AI assisted. Several California superior courts now say this expressly in local rules (see the firm's County Compliance Matrix), and pending SB 574 would write a personal read-and-verify duty into section 128.7 statewide.

The four-step citation ladder. Before any AI-assisted document leaves the firm:

1. **Existence.** Pull every citation on Westlaw or Lexis. If it does not appear, it was fabricated.
2. **Quotation.** Search the opinion's full text for every quotation attributed to it.
3. **Proposition.** Confirm through headnotes and the opinion's actual language that the case supports the stated proposition.
4. **Currency.** Run KeyCite or Shepard's on every case, checking negative treatment on the specific point cited.

Statutes and rules. Verify all statutory and rule citations against official or maintained text (e.g., leginfo.ca.gov), not against the model's memory.

Grounding practice. Prefer closed-corpus workflows: provide the current statutory text, the relevant treatise chapter, or the authority file, and instruct the tool to cite only from the provided material.

Stress test. Before filing any substantive brief, run it through an approved tool asking for weaknesses, contrary authority, and omissions. Assume opposing counsel will do the same the day it arrives.

AI in client documents. Opening a client document in a session with an AI add-in enabled sends that document to the account the add-in is bound to; the input tiers above apply to the whole document, not just the prompt. When an AI tool edits a document, it must leave a reviewable record of every change: tracked changes, or sequential saved versions compared before acceptance (Word's Compare feature produces the redline from two versions). Silent in-place editing is prohibited. The reviewing attorney examines the changes, not just the final text, before accepting them.

Agentic prohibition. No AI system may file a document, sign anything, communicate with a court, or communicate legal advice to a client without attorney review and approval of the specific output. Autonomous filing is prohibited without exception.

6. Supervision and Training (Rules 5.1, 5.2, 5.3)

1. All covered persons receive training on this policy at onboarding and at least [annually / on a recurring internal schedule, such as a monthly half-hour technology session], covering confidentiality settings, the verification ladder, and current sanctions case law (e.g., Mata v. Avianca, Inc. (S.D.N.Y. 2023) 678 F.Supp.3d 443; Noland v. Land of the Free, L.P. (2025) 114 Cal.App.5th 426).
2. Supervising attorneys are responsible for AI use by those they supervise, including nonlawyer staff and outside vendors.
3. A subordinate lawyer may not use AI in a manner that violates that lawyer's own professional obligations, even at a supervisor's direction, and must exercise independent judgment on whether a given use is proper.
4. No "don't ask, don't tell": undisclosed use of unapproved tools or personal accounts for client work is a policy violation. The remedy for a gap in the Tool List is a proposal to the Administrator, not a workaround.

7. Client Communication (Rules 1.2, 1.4)

The responsible attorney must consider, at intake and as the representation evolves, whether AI use should be disclosed to the client, including how the technology will be used and its benefits and risks. Disclosure is appropriate where AI use materially affects decision-making or where the client has expressed interest or concern. The firm's standard engagement letter includes [an AI disclosure clause; see Appendix / to be drafted]. A firm adopting task-unit pricing under Section 8 also states the per-task fee schedule in the engagement agreement as the basis of compensation (Bus. & Prof. Code § 6148). Boilerplate consent is not a substitute for informed consent where Tier 3 inputs are proposed (ABA Formal Op. 512).

8. Billing (Rule 1.5; Bus. & Prof. Code §§ 6147, 6148)

1. Bill actual time spent (prompting, refining, reviewing, editing), not time the tool saved.
2. General-purpose AI subscriptions are overhead, absorbed in rates, like the library or the network.
3. Matter-specific AI costs (per-use fees for specialized tools, custom implementations) may be passed through only if the fee agreement discloses them, the charge reflects actual cost, and any markup has informed written consent.
4. Do not bill clients for time spent learning general AI tools; time learning a tool the client specifically requests may be billable with disclosure.
5. **Task-unit pricing (optional).** The firm may instead price defined tasks at fixed charges, in effect a disclosed schedule of small flat fees (for example, [0.1 / 0.2] of the standard hourly rate per defined task), stated in the fee agreement as the basis of compensation (Bus. & Prof. Code § 6148). The unit of analysis is then the professional task (drafting, review, verification, judgment), not the minutes inside it, and efficiency from AI use stays with the firm by design rather than by rounding. The total fee remains subject to Rule 1.5's prohibition on unconscionable fees. Where billing remains time-based, rounding up to a disclosed minimum increment is permitted (ABA Formal Op. 93-379), but entries may not systematically exceed the time actually expended; minimum increments stacked across many short tasks are the classic fee-audit target. The documentation practice in Section 10 makes task-unit billing administrable: the same session records that document verification generate the granular task entries.
6. **Court-reviewed fees.** Where a court or a third party will review fees on an actual-time basis (fee-shifting

motions, probate matters, insurance billing guidelines), bill actual time. Minimum-increment billing draws percentage reductions in that setting (Welch v. Metropolitan Life Ins. Co. (9th Cir. 2007) 480 F.3d 942 (reduction affirmed where quarter-hour minimums for brief calls and emails inflated hours)).

9. Court Rules Compliance and Disclosure

Pre-filing check. Before filing in any court, the responsible attorney checks the local rules, the department's standing orders, and the judge's posted preferences for AI provisions. The firm maintains a County Compliance Matrix for the counties where it regularly files; the matrix is a starting point, not a substitute for checking current rules.

Mandatory disclosures. Where a rule requires disclosure (e.g., Santa Clara Superior Court Local General Rule 10(B)(5), requiring disclosure of AI or machine learning use in ex parte materials, effective July 1, 2026), compliance is not optional.

Voluntary disclosure posture. In matters where the firm's AI use is material, the responsible attorney should consider the firm's standard voluntary disclosure instruments (the CMC statement attachment and the AI Disclosure Certificate in the firm's Court Disclosure Kit). Disclosed, verified use is the firm's public posture; it names the accountable human, demonstrates the verification protocol, and positions the firm favorably if an error ever surfaces despite the protocol.

10. Documentation

For each matter in which AI materially assists work product, the matter file should reflect: the tool and configuration used, the general tasks performed, and the verification completed. The disclosure certificate format in the Court Disclosure Kit doubles as an internal documentation template even when nothing is filed.

11. Bias and Other Jurisdictions (Rules 8.4.1, 8.5)

AI outputs can encode bias; this risk compounds when multiple agents interact. Do not use AI for screening decisions about employment, clients, or candidates without human review calibrated to anti-discrimination law. Lawyers licensed or practicing in other jurisdictions must check those jurisdictions' AI rules; this policy implements California requirements only.

12. Violations and Review

Violations are reported to the Administrator and handled under the firm's ordinary professional standards. This policy is reviewed [annually] and upon: adoption of new State Bar rules (COPRAC's proposed amendments to Rules 1.1, 1.4, 1.6, and 3.3 are pending), enactment of SB 574 or successor legislation, new local rules in counties where the firm files, or adoption of materially new tools or agentic capabilities.

Appendix A: Approved Tool List (template)

[Example: Claude, enterprise]. Tier / account type: Enterprise | **Training disabled:** Yes | **Memory off / scoped:** Off | **Retention terms:** [state period] | **Permitted input tier:** Tier 2 | **Agentic capabilities:** [none / list] | **Approved date:** [date] | **Notes:** [note pricing model; note any Word add-in and the account it binds to]

[Example: Lexis+ AI]. Tier / account type: Firm subscription | **Training disabled:** Contractual | **Memory off / scoped:** [state] | **Retention terms:** [state] | **Permitted input tier:** Tier 2 | **Agentic capabilities:** none | **Approved date:** [date] | **Notes:** Closed corpus

Appendix B: Principal Authorities

- COPRAC, Practical Guidance for the Use of Generative AI in the Practice of Law (2026) (replacing the

2023 guidance; adds agentic AI at the request of the California Supreme Court)

- ABA Formal Opinion 512 (July 29, 2024)
- Billing: ABA Formal Opinion 93-379 (Dec. 6, 1993) (time-based billing must reflect actual time; rounding to disclosed minimum increments permitted); Welch v. Metropolitan Life Ins. Co. (9th Cir. 2007) 480 F.3d 942 (quarter-hour minimum increments reduced on lodestar review)
- Cal. Rules of Prof. Conduct, rules 1.1, 1.2, 1.4, 1.5, 1.6, 3.1, 3.3, 5.1, 5.2, 5.3, 8.4.1, 8.5; Bus. & Prof. Code §§ 6068(d), (e), 6147, 6148
- Code Civ. Proc. §§ 128.7, 2018.030
- Cal. Rules of Court, rule 10.430 and Standard 10.80 (court-side AI policies; structural model, not attorney-binding)
- Local rules: see County Compliance Matrix
- Sanctions canon: Mata v. Avianca, Inc. (S.D.N.Y. 2023) 678 F.Supp.3d 443; Noland v. Land of the Free, L.P. (2025) 114 Cal.App.5th 426; United States v. Farris (6th Cir. 2026) 171 F.4th 920; Fivehouse v. U.S. Dept. of Defense (E.D.N.C. 2026) No. 2:25-cv-00041
- Pending: SB 574 (Umberg); COPRAC proposed rule comment amendments (2026)