

Generative AI Use Policy — [Firm name]

Eight clauses covering what the ethics opinions and the courts now expect of a firm that uses generative AI. Adopt, edit or delete each one to fit your jurisdiction, your practice areas and your insurer’s requirements; the line under each clause says where the duty comes from.

1. **Approved tools and tiers.** [Firm name] permits generative AI only through the tools and account tiers on the approved list maintained by the policy owner, and a tool goes on that list only after someone has read its Terms of Use and privacy policy and recorded who can access inputs, whether they train the model, and how long they are retained.

Source: ABA Formal Opinion 512, “Generative Artificial Intelligence Tools” (29 July 2024), pp. 7 and 11 — read the terms, or be briefed by a colleague or an outside expert who has.
2. **What may never be entered.** No client name or matter identifier, privileged analysis, sealed or protective-order material, or health, financial or other sensitive record is entered into a general-purpose tool. [Optional stricter rule: none at any tier, approved business accounts included.]

Source: ABA Model Rule 1.6; Texas Rule 1.05 and Texas Opinion 705 (February 2025) — the lawyer “should be reasonably satisfied that the program will not reveal confidential information to others.”
3. **The anonymize-first rule.** Where an outside model has to work on a client document, the identifiers are removed before the text is submitted, and the result is checked rather than assumed.

Source: Alaska Ethics Op. 2025-1 — lawyers “must fully anonymize their inputs” to a self-learning tool outside a closed system unless the client consents; NYC Bar Formal Op. 2024-5 — consent is “not needed if no confidential client information is shared, for example through anonymization of client information”; D.C. Bar Op. 388 cautions that a growing dataset may still “connect the dots.”
4. **Verification before filing or advice.** Every authority, quotation and figure produced with AI assistance is confirmed against the primary source by the responsible lawyer, and no AI output reaches a court, a client or a counterparty unread.

Source: ABA Formal Opinion 512, p. 10 (candor to the tribunal; review the output before it is filed) — and the sanctions record: Mata v. Avianca (S.D.N.Y., June 2023), Johnson v. Dunn (N.D. Ala., July 2025), Whiting v. City of Athens (6th Cir., March 2026), Ibach v. Stewart (Ala., 24 April 2026).
5. **Client disclosure and consent.** The firm obtains and records the client’s informed consent before information relating to a representation goes into a tool that could disclose it, and discloses AI use when the client asks or the engagement terms require.

Source: ABA Formal Opinion 512, p. 7 — informed consent for self-learning tools; a boilerplate engagement-letter clause is not enough. The Florida Bar’s Proposed Advisory Opinion 24-1 recommended consent “prior to utilizing a third-party generative AI program if the utilization would involve the disclosure of any confidential information” (proposed text, as quoted in the NYSBA Task Force on Artificial Intelligence report, nysba.org, April 2024; final opinion adopted 19 January 2024).
6. **Court disclosure.** Before filing, the responsible lawyer checks the court’s rule and the judge’s standing order on AI use and files any certification they require.

Source: New York, 22 NYCRR Part 161 — AI use in preparing papers submitted to a court “should not be prohibited, as long as such use is in accordance with the duties and responsibilities that apply to individuals who submit papers to a court”; other courts require a certification that the cited authorities exist.
7. **Supervision (Rules 5.1 and 5.3).** A named partner owns this policy, every lawyer and nonlawyer is trained before using an approved tool, and no vendor is approved without a documented check of its security and confidentiality terms.

Source: ABA Formal Opinion 512, p. 10 — managerial lawyers “must establish clear policies regarding the law firm’s permissible use of GAI,” and supervisory lawyers must make reasonable efforts to ensure that the firm’s lawyers and nonlawyers comply; p. 11 — confirm the tool “is configured to preserve the confidentiality and security of information.”
8. **Incidents, training and review cadence.** A bad paste, a fabricated citation or any other AI incident is reported to the policy owner the same day and logged; training is repeated at least annually; and this policy is reviewed on the date below.

Source: ABA Formal Opinion 512, p. 10 (training on the tools’ limits and on secure data handling); NYSBA Task Force on Artificial Intelligence report (nysba.org, April 2024) — “you must take precautions to protect sensitive client data and ensure that no Tool compromises confidentiality.”

ADOPTION

Firm	Adopted on	Signature
Policy owner (Rule 5.1)	Next review	Name and title

Approved tools and tiers are listed in Schedule A, maintained by the policy owner.