CNA

PROFESSIONAL COUNSEL®

Advice and Insight into the Practice of Law®

Building a Safe and Practical Law Firm Artificial Intelligence Policy: A Risk Management Playbook

Artificial intelligence [AI] is no longer a "future tech" issue for the legal profession. From search and drafting assistants to e discovery, AI now touches our daily work product as well as client expectations. But without a firmwide, AI policy, lawyers and staff can easily stumble into ethical, confidentiality, billing, advertising, employment, and cross border regulatory traps. This article distills what lawyers need to know and offers a practical blueprint to set up (and continuously improve) a policy that protects you, your clients, your firm, and your reputation.

Why Having a Firm Wide AI Policy is Important

To Comply with Legal and Regulatory Standards[1]

AI technologies can be subject to specific laws and regulations. This includes laws regarding data privacy, intellectual property rights, and consumer protection. An AI policy provides guidance as to your firm's use of AI, making the firm aware of relevant laws and regulations, helping to mitigate the risk of legal issues or penalties. For example, the Federal Trade Commission is policing "AI washing" and deceptive claims (e.g., "robot lawyer")[2]; the Equal Employment Opportunity Commission warns that AI used in recruiting/HR can create disparate impact and employers must assess selection tools to avoid Title VII violations.[3] Globally, the EU AI Act phases in obligations and fines, and transparency/due diligence obligations for general purpose AI and high risk systems through 2026–2027.[4] Even non EU firms may be impacted if tools or outputs target EU users or are used in EU matters.

Guide Firm Employees on the Ethical and Responsible Use of AI Technology on Client Matters

Ethical obligations already apply. American Bar Association [ABA] Formal Opinion 512 is the ABA's first comprehensive ethics guidance on generative AI. It emphasizes that existing duties such as competence (Rule 1.1), confidentiality (Rule 1.6), communication (Rule 1.4), supervision (Rules 5.1/5.3), candor to the tribunal (Rule 3.3), meritorious claims (Rule 3.1), and reasonable fees (Rule 1.5), govern your use of AI.[5] Lawyers must understand the benefits/risks of the tools they use, verify outputs, and protect client information.

Further, state-level guidance is converging. The State Bar of California's Standing Committee on Professional Responsibility and Conduct published *Practical Guidance on the Use of Generative Artificial Intelligence in the Practice of Law*,[6] warning lawyers not to put client confidences into tools lacking adequate safeguards and urging anonymization unless informed consent is obtained. It also highlights competence, confidentiality, supervision, communication, candor, and fees. The New York State Bar Association[7] and the New York City Bar[8] have issued guidance encouraging cautious adoption, transparency to clients, and education over legislation. Their reports include sample disclosure language for engagement letters which lawyers may want to consider for their practice setting. The Illinois Supreme Court recently adopted a policy stating AI use by litigants, lawyers, and judges "should not be discouraged, and is authorized provided it complies with legal and ethical standards," while cautioning on accuracy, bias, confidentiality, and non disclosure requirements.[9]

1 The content of this article focuses on internal AI policies from a legal services perspective and does not reflect wider AI use guidance in other firm operations contexts such as employment/HR, etc.
2 Federal Trace Commission, *FTC Finalizes Order with DoNotPay That Prohibits Deceptive 'AI Lawyer' Claims, Imposes Monetary Relief, and Requires Notice to Past Subscribers*.
3 U.S. Equal Employment Opportunity Commission, Select Issues: *Assessing Adverse Impact in Software, Algorithms, and Artificial Intelligence Used in Employment Selection Procedures Under Title VII of the Civil Rights Act of 1964*.
4 Glocert International, *EU AI Act Timeline & When Obligations Kick In*.
5 American Bar Association, *Formal Opinion 512 Generative Artificial Intelligence Tools*.
6 The State Bar of California Standing Committee on Professional Responsibility and Conduct, *Practical Guidance for the Use of Generative Artificial Intelligence in the Practice of Law Executive Summary*.
7 New York State Bar Association, *Report and Recommendations on Artificial Intelligence and Access to Justice in 2025*.
8 New York City Bar, *Formal Opinion 2025-6: Ethical Issues Affecting Use of AI to Record, Transcribe, and Summarize Conversations with Clients*.
9 Illinois Supreme Court, *Policy on Artificial Intelligence*.

Identify Ways to Protect Clients' Confidential Information and the Firm's Internal Data

Generative AI can learn from sensitive and confidential data, and this data can be exposed to hackers or malicious actors. A law firm AI policy helps mitigate this risk by establishing stringent data protection protocols, e.g., use of only the firm's closed loop AI system for confidential information. It outlines how data will be collected, stored, and used to protect the privacy of your clients and employees, e.g., no extraction, deletion, or return of client information that has been input into the AI platform.

Provide Guidance on How to Prevent Data Bias and Discrimination

AI language models, like sponges, absorb the characteristics of their training data. If the input data is biased, it is likely the AI output will be as well. A law firm AI policy can use bias reviews and audits, ensuring AI-generated content does not discriminate based on race, gender, age, or other traits.

5 Ideas for Creating a Solid Internal AI Usage Policy

Step 1: Assess Your Organization's AI Needs and Goals

Before embarking on any AI projects, it is important to understand your firm's needs and goals. This step allows you to identify where AI can add value and helps you choose the most suitable AI technology for your law firm. At the same time, it allows you to evaluate potential risks linked to AI implementation.

Begin by examining your current systems and processes to identify where AI can improve operations or resolve issues. This might involve finding repetitive tasks that could be automated, bottlenecks that could be smoothed out with predictive analysis, or large datasets that could benefit from machine learning to uncover insights.

After identifying these potential areas of improvement, you can start comparing different AI technologies to find the most suitable platform and vendor. For example, if your law firm wants to scale and streamline the production of written content, a generative AI platform might be useful for grammar and plagiarism checks, analysis, or generating initial drafts.

Step 2: Create a Governance Framework for AI Usage

A governance framework for AI usage on client matters is essentially a set of rules for how AI should be used and managed within your law firm. This not only includes the day-to-day handling of AI but also the broader strategy for its development, deployment, and maintenance.

In your governance framework, you may want to include detailed processes that outline how to use AI ethically and responsibly on client matters, with a strong focus on data privacy, security, and transparency. You should also consider measures for preventing bias and ensuring compliance with all relevant regulations. For example, your firm's policy might include the following sections:

- **Purpose, Scope, and Definitions.** Consider a section that explains the purpose of the framework, establishes its scope (which is likely firm-wide, applicable to all personnel, as well as contractors and vendors), and defines terms like "AI tool," "generative AI," "confidential information," "client personal data," and "firm confidential information." You may want to keep definitions technology neutral to avoid constant rewrites.
- **Ethics first.** Consider anchoring the policy in your jurisdiction's rules of professional conduct and require lawyers to independently verify AI outputs and to avoid outsourcing legal judgment to tools. ABA Formal Op. 512 stresses verification and ongoing technological competence.
- **Client centric confidentiality.** The policy may address prohibiting entering client identifying or privileged content into any AI tool unless (i) the tool meets firm security standards (contractual, technical, and operational), and (ii) the client has consented when appropriate. California's guidance is explicit on avoiding confidential input into non secure tools and anonymizing where feasible. The policy might require matter by matter assessment of whether disclosure to the client or informed consent from the client is appropriate.
- **Transparency with clients when material.** Consider mandating disclosure if AI materially affects services (e.g., novel risks, reliance on third party processors, or significant cost or process changes).
- **Training and Certification.** The policy might require annual CLE style training on AI ethics, confidentiality, and tool use for all employees; new hire onboarding module; matter type micro trainings.
- **Permitted vs. Prohibited Use.** For example, the policy could provide guidance on permitted and prohibited use. Permitted use might include brainstorming, first draft outlines, style edits, idea generation, brief summaries of public materials, quality checks (e.g., issue spotting). Prohibited use might include uploading client identifying or privileged material into any non-approved AI tool; relying on AI outputs as legal authority without human verification and citation to real sources; marketing claims that imply capabilities the firm or tool cannot substantiate.

- **Supervision and accountability.** Your policy might include lawyer and staff responsibilities, including how AI outputs will be validated. Another section could ensure that supervising attorneys remain accountable for AI assisted work.
- **Privacy & security by design.** The policy could build in vendor due diligence, encryption, role based access, retention controls, data location; sub processor list/approvals; audit logs, IP/ownership; deletion on exit; bias/evaluation reports; EU AI Act readiness if in scope, and model training restrictions (no use of firm/client data for model training absent explicit agreement). California's guidance highlights reviewing privacy/terms of use and prohibiting training on client data by default.
- **Billing integrity.** The policy might address aligning billing with reasonableness and actual time/effort, pass on efficiencies to clients as appropriate, and clarify AI disbursement practices (e.g., use fees) up front. If AI saves time, invoice actual time or adjust flat fees accordingly; disclose any per use AI costs up front and ensure total fees remain reasonable.
- **Escalation and Exceptions.** Consider simple pathways to seek approval for novel tools/uses; fast guidance from an AI Steering Committee to help ensure all staff complies with the policy.
- **Special Topics to Address Explicitly**
  - *Court filings and candor to tribunals:* Many jurisdictions are now requiring: (i) human verification of all authorities and facts; (ii) cite checking; and (iii) compliance with any standing orders about AI declarations. To be complete, you might want to address this issue in your policy.
  - *Cross border matters (EU AI Act):* For EU touching work, prioritize vendors preparing for EU AI Act requirements, especially transparency, documentation, risk management, and (for high risk use cases) conformity assessments.
  - *Biases in AI algorithms.* As noted above, if the input data is biased, it is likely the AI output will be as well. A law firm AI policy might address how use bias reviews and audits are conducted. This is critical for making sure that your firm's use of AI is both effective and safe.
- **Continuous improvement.** Consider periodic reviews (at least semi annually) as ethics opinions, court rules, and regulations evolve.

Step 3: Communicate, Educate, and Train All Employees on AI Policy

Communicate, educate, and train all staff on the policy, its guidelines, and how to adhere to it. Providing comprehensive knowledge about AI and its ethical implications can foster a culture of responsible and ethical AI use. You want to have training that actually changes behavior. For example, you can provide online training sessions and workshops for employees on topics such as data privacy, security, and ethical use of AI. This can include case studies to demonstrate the results of responsible and irresponsible AI use; building short, role-specific modules and teaching prompt patterns; implementing quizzes to assess the effectiveness of the training; and feedback forms for employees to ask questions and seek clarifications.

Step 4: Monitor AI Performance

Implement a process for regularly monitoring AI performance to ensure it is performing as intended and not inadvertently causing harm or generating undesirable results. Include tracking metrics such as the accuracy and speed of content generated. This will provide valuable insights into the system's performance and its ability to produce high-quality content in a timely manner. You will also better prove the value of your AI technology investment. For example, you can track the accuracy of content generated by monitoring the percentage of false negatives or positives generated. You can also track the system's fairness by looking for any biases in the output or evaluate its accuracy by comparing generated results with ground truth data.

Step 5: Institute Regular Review and Updates

Given the rapid pace of AI evolution, it is important to regularly review and update your firm's AI policy. This may involve monitoring changes in technology, regulations, and societal norms, and updating your policy accordingly. It is also important to incorporate feedback from employees and clients into your policy. For example, if new privacy regulations are introduced, you will need to update your policy to ensure compliance. If there is a groundbreaking AI development that could change the way your law firm operates, your policy might need to be adjusted to reflect this. If employees express concern over a particular AI application's impact on their work, their feedback should be considered during policy updates.

Diving Deeper: Developing Your Control Framework Based on the Tool and Its Use

A workable law firm AI usage policy connects controls to how tools are selected and used. The following are examples of a control framework to be considered:

Use of Case Intake and Risk Assessment

- **Is the AI use internal or client facing?** Internal administrative use (e.g., summarizing long emails) presents a lower risk than client deliverables.
- **What data will be involved?** Personally Identifiable Information [PII], Protected Health Information [PHI], financials records, privileged content? Require data minimization and anonymization.
- **Will results be filed in court?** Require manual verification, source citations, and, if a court requires, AI usage certification.

Vendor and Tool Due Diligence

The firm should ensure that any AI tool or platform it uses meets strong security standards, including recognized certifications along with basic protections like encryption, secure password and identity controls, and clear documentation of how data moves through the system. Contracts with vendors should make it clear that firm and client information cannot be used to train AI models unless explicitly agreed to and should also spell out who owns the data and how/if it can be deleted. Because AI tools can produce inaccurate or biased results, firms should look for vendors that test for errors, benchmark performance, and evaluate fairness. Finally, if the tool is considered a general purpose AI system or is used with EU related data or clients, firms should confirm that the provider has a plan to meet the EU AI Act's transparency and documentation requirements as those rules phase in.

Deployment and Access Controls

Access to AI tools should be limited based on user roles so that only attorneys and staff have the minimum level of access necessary for their work. Client and matter level restrictions should also be enforced, because ethical walls apply within AI systems just as they do across the firm. To reduce the risk of accidental disclosure, the firm should use prompt shielding tools that warn users when they attempt to include client identifiers, privileged information of unsupported/unauthorized AI platforms. Additionally, any documents uploaded into AI systems should pass through a redaction or anonymization process to protect sensitive data.

Use Standards for Lawyers and Staff

All AI generated content must be fully verified before it is shared with clients, courts, third or opposing parties which means attorneys must check all sources and citations for accuracy. Lawyers remain fully responsible for the work product they deliver, as AI is only a tool and should never be treated as an author. To maintain good "prompt hygiene," users should rely on practice approved prompt formats and avoid including client names or other identifiable information unless working within a secure, authorized environment.

Monitoring, Logging, and Incident Response

The firm should maintain logs of AI prompts and outputs to ensure auditability, while avoiding the storage of confidential information when necessary. It should also establish a clear protocol for handling hallucination incidents, including issuing errata letters, notifying courts when required under local rules, and conducting internal "root cause" analysis. In addition, the firm should maintain a data exposure playbook that outlines when and how to notify clients of potential confidentiality risks and when to involve cyber counsel. Lastly, the firm needs to define retention periods for prompts/outputs, ensure data deletion when licenses end, and confirm that the vendor deletes backups.

Keep An Eye on Developments

It is critical to monitor the fast moving developments in the AI regulatory landscape. For example, the ABA is expected to issue additional ethics opinions that further clarify lawyers' duties around confidentiality, supervision, and billing as AI tools become more advanced, building on the foundation established in ABA Formal Opinion 512. Additionally, state bars and courts, including those in California, New York, and Illinois, continue to update their guidance and court standing orders on the professional and ethical use of AI.

The Federal Trade Commission is also increasing its scrutiny of AI related marketing claims, particularly in legal services and legal tech markets, emphasizing that there is no "AI exemption" from truth in advertising laws. Meanwhile, the EEOC is expanding its focus on the use of automated decision making tools in hiring and evaluation, warning employers to monitor these systems for disparate impact risks under Title VII.

Firms with international work should also be preparing for the phased implementation of the EU AI Act, ensuring that their vendors have clear conformance plans aligned with upcoming regulatory deadlines. In addition, UK regulators such as the Solicitors Regulation Authority and the Law Society continue to publish principle based guidance on responsible legal sector AI use, resources that can be especially valuable for multinational firms or those working across jurisdictions.[10]

10 Solicitors Regulation Authority of England and Wales, *Compliance tips for solicitors regarding the use of AI and technology* and The Law Society of England and Wales, *Generative AI: the essentials*.

Conclusion

With AI use policies, the key takeaways are straightforward. First, it is essential to put the firm's AI policy in writing, as a clear and practical framework helps reduce unauthorized or inconsistent use of AI, sets defensible standards, and demonstrates diligence to clients, courts, insurers, and regulators. Second, the policy must be usable. Using checklists, examples, practical prompts, red flag reminders, and sample engagement letter language rather than lengthy manifestos helps to ensure risk is managed. Third, firms should invest in guardrails such as verification procedures, anonymization tools, logging practices, and incident response playbooks, which help prevent small mistakes from escalating into significant reputational issues. Finally, ongoing training and testing are crucial, because AI related risks evolve quickly; policies must be regularly updated to keep pace with new regulations and developments in both the U.S. and abroad.

This article was authored for the benefit of CNA by:

Tracy Kepler

Tracy L. Kepler is the Risk Control Director for CNA's Lawyers' Professional Liability program. In this role, she designs and develops content and distribution of risk control initiatives relevant to the practice of law. Prior to joining CNA, Tracy previously served as the Director of the American Bar Association's Center for Professional Responsibility (CPR) and has over 20+ years of experience in attorney regulation through her positions as an Associate Solicitor for the U.S. Patent & Trademark Office and as Senior Litigation Counsel for the Illinois Attorney Registration and Disciplinary Commission. She also teaches Legal Ethics and Professional Responsibility at Georgetown University Law Center, American University – Washington College of Law and Loyola University School of Law (Chicago).

Don't Miss Out on Future CNA LPL Publications

If you wish to receive the CNA Lawyers' Professional Liability Risk Control monthly publications, you may register here.

About CNA Professional Counsel

This publication offers advice and insights to help lawyers identify risk exposures associated with their practice. Written exclusively by the members of CNA's Lawyers Professional Liability Risk Control team, it offers details, tips and recommendations on important topics from client misconduct to wire transfer fraud.

For more information, please call us at 866-262-0540 or email us at lawyersrisk@cna.com