AI Governance Best Practices for Law Firms

Generative AI (GAI) is now embedded in the day-to-day practice of law, sometimes as an obvious "chat" interface, but increasingly as a quiet feature inside research platforms, document tools, contract analytics, eDiscovery, and even email and productivity suites. That reality creates a governance problem: firms need a repeatable way to control *who* can use AI, *for what*, *with what data*, and *under what verification and supervision standards*.

The ABA's Standing Committee on Ethics and Professional Responsibility put a bright spotlight on these issues in **Formal Opinion 512 (July 29, 2024)**, emphasizing that lawyers must account for duties of competence, confidentiality, communication, supervision, candor, and reasonable fees when using generative AI tools.

The practical message for firms is straightforward: **"AI governance" is now part of professional responsibility risk management**, not a discretionary tech initiative.

What "AI Governance" Means in a Law Firm

AI governance is the operating system that turns ethical duties into daily workflows. In a law-firm context, a credible program typically includes:

- **Rules** (policies, standards, and client-facing commitments)
- **Process** (intake, approvals, audits, incident response)
- **People** (clear accountability and supervision)
- **Technology controls** (access management, logging, data loss prevention (DLP), approved tools)
- **Verification discipline** (how AI outputs are checked before they become advice, filings, or client deliverables)

Two widely used governance frameworks map well onto law-firm needs:

- **NIST AI Risk Management Framework** (RMF) frames AI risk management as a lifecycle approach organized around **Govern, Map, Measure, Manage**.
- **ISO/IEC 42001** describes an "AI Management System" model (policies, objectives, roles, supplier controls, continuous improvement) that aligns naturally with firm governance structures.

You do *not* need a certification program to benefit from these frameworks. They function well as scaffolding for law-firm controls.

The Risk Categories That Drive Law-Firm Governance

Confidentiality and Data Exposure

Formal Opinion 512 is explicit that confidentiality duties apply when lawyers use GAI tools, and that lawyers must consider the risks associated with a tool's operation, especially when client information is input into external systems.
Governance implication: firms must classify tools (public vs. enterprise vs. on-prem vs. vendor-embedded) and define what data may be shared with each class.

Accuracy, Hallucinations, and Citation Risk

The ABA warns that GAI can produce "hallucinations" and that lawyers must apply appropriate independent review to avoid incompetent work and misleading submissions.
Governance implication: firms need defined verification workflows for different use cases (research memos, contracts, filings, client advice, marketing).

Supervision, Agents, and Workflow Discipline

Formal Opinion 512 ties AI use to duties to supervise those assisting with legal services (including nonlawyers and "agents") and to maintain overall accountability for the work product.
Governance implication: the firm must treat AI as a regulated capability, not an ad hoc personal preference.

Fees and Billing Judgment

The ABA flags that time spent learning a tool generally shouldn't be billed to clients, while time spent using and verifying outputs may be billed if reasonable.
Governance implication: firms need consistent billing guidance and documentation expectations when AI is used.

The Core Governance Controls Every Firm Should Implement

Below is a practical "minimum viable governance" package, written for firms that want controls that auditors, clients, and risk committees can recognize as serious.

1) Adopt a Firm AI Use Policy (Compliance Doc #1)

This is the keystone document. It should be short enough to be used, but specific enough to be enforceable. It should include:

- **Tool categories & approval status** (Approved / Conditional / Prohibited)
- **Data handling rules** (client confidential info, PHI, trade secrets, internal firm strategy)

- **Use-case boundaries** (brainstorming vs. drafting vs. research vs. filings)
- **Verification standards** (what must be checked, and by whom)
- **Client communication triggers** (when disclosure/consent may be required)
- **Recordkeeping** (when prompts/outputs must be retained in the matter file)
- **Escalation** (what to do when AI output appears wrong, biased, or risky)

Formal Opinion 512 is a strong backbone for the policy's "why," because it explicitly ties GAI use to competence, confidentiality, communication, supervision, candor, and fees. NIST AI RMF's GOVERN function provides a practical structure for assigning accountability and defining risk tolerance.

2) Create a Practice-Group AI Intake Questionnaire (Compliance Doc #2)

Firms usually underestimate how many "AI use cases" exist until they inventory them. A lightweight intake questionnaire (completed by each practice group and updated quarterly) should capture:

- What tools are being used (including vendor-embedded AI features)
- Whether client data is input (and what types)
- Whether outputs are client-facing or court-facing
- Reliance level (idea generation vs. substantive legal conclusions)
- Human review steps currently used
- Known failure modes (e.g., hallucinated citations, drafting errors, confidentiality risk)

This document operationalizes the MAP step of NIST AI RMF, capturing context, stakeholders, intended use, and impact.

3) Implement an AI Vendor Confidentiality & Risk Checklist (Compliance Doc #3)

Most AI governance failures come from vendor terms and product design, not lawyer intent. A standard checklist should be required for any AI tool approval (including "free" tools and AI features embedded in existing vendor platforms). It should cover:

- **Training use**: Are prompts/outputs used to train models? Is opt-out available?
- **Retention & deletion**: How long are prompts/outputs stored? Can the firm delete?
- **Disclosure**: Can the vendor share data with affiliates/subprocessors/regulators?

- **Security**: Access controls, encryption, incident response obligations
- **Logging/auditing**: Can the firm audit use and retrieve logs for investigations?
- **Data residency** (if relevant to client requirements)
- **Subprocessor list** and change controls
- **IP/ownership** terms for outputs and any customer materials

ISO/IEC 42001's emphasis on lifecycle and supplier controls is directly relevant here. And Formal Opinion 512's confidentiality and supervision themes give the risk rationale that partners and clients will understand.

Verification

Formal Opinion 512 makes the key point: lawyers may use AI as a tool, but cannot offload professional judgment to it, and must independently review outputs to a degree appropriate to the task and the tool's risk profile.

A governance-ready verification standard usually looks like this:

- **Research / citations:** citations must be validated against primary sources (and quote-checked).
- **Factual statements:** verify against the record and reliable sources; document checks.
- **Drafting:** treat AI output as a "drafting aid," not a final document, lawyer must edit substantively.
- **Filings / tribunal-facing content:** require a heightened review checkpoint and, where relevant, compliance with court rules and local orders about AI use.

This can be documented as a one-page "Verification Matrix" appended to the AI Use Policy.

Conclusion

The ABA has made the direction clear: lawyers and law firms must treat generative AI as part of the ethical risk environment, especially for competence, confidentiality, communication, supervision, and billing.

For most firms, the fastest path to defensible governance is not a 50-page policy manual. It is a tight compliance "starter set" that can be implemented, audited, and improved.