PROGRAM MATERIALS
Program #36155
May 27, 2026

AI Agents in Your Firm: Avoiding Ethical Nightmares

Celesq® AttorneysEd Center
www.celesq.com

5301 North Federal Highway, Suite 150, Boca Raton, FL 33487
Phone 561-241-1919

AI Agents in Your Firm: Avoiding Ethical Nightmares

Angeli R. Fitch, Esq.

AI Compliance & Ethics Attorney

Infinity Law Group

The Real Ethical Nightmare Isn't AI

The Actual Problem

Invisible delegation. No rules. No review. No audit trail. Consequences that land squarely on *you*.

The ethics rules were not written for tools that act autonomously on behalf of your client — but you're still responsible.

What This Is NOT About

- Whether AI is smart enough to practice law
- Whether you should fear technology
- Banning AI from your firm

What This IS About

- Whether you can **supervise** it
- Whether you know what it's doing
- Whether you're still in control

Your Roadmap for the Next Hour

Part 1

What AI agents are and where they already live in your firm

Part 2

Mapping agent behavior to ethics rules: competence, supervision, confidentiality

Part 3

Building a practical AI governance framework your firm can actually use

Each part builds on the last. By the end, you'll have a concrete framework — not just a list of things to worry about.

What You'll Walk Away With

1

Recognize AI Agent Risk

Identify when a tool crosses from *assisting* to *acting* — and why that line matters ethically

2

Map to Ethics Rules

Connect specific AI agent behaviors to Model Rules on competence, supervision, and confidentiality

3

Spot Shadow AI

Identify unauthorized AI use already happening inside your firm — often with the best intentions

4

Build a Governance Plan

Leave with a practical, scalable framework for AI oversight that survives real-world law firm conditions

Shadow AI: It's Already in Your Firm

Shadow AI is any AI tool being used by your staff or attorneys **without firm knowledge, policy, or oversight**. It's not malicious. It's convenient.

Your associate figured out that ChatGPT writes better demand letters than she does at 11pm. Your paralegal uses an AI notetaker in every client call. Your billing coordinator runs invoices through an AI summarizer.

None of them told anyone. None of them asked.

⚠ If you don't have an AI policy, you almost certainly have Shadow AI. The question is what it's touching — and whether client data is leaving the building.

AI Adoption in Law: No Longer Hypothetical

79%

Attorneys Using AI

of Am Law 200 firms report lawyers using at least one AI tool in their practice (Thomson Reuters, 2025)

3x

Faster Growth

Legal AI adoption has tripled year-over-year since GPT launched in 2022

$35B

Market by 2030

Projected size of the legal AI market — vendors are not slowing down

The train has left the station. The ethics rules haven't been rewritten yet. That gap — between what AI can do and what the rules require — is exactly where your liability lives.

What Is an AI Agent?

> **Key Distinction** — A chatbot *answers*. An agent *does*. The moment a tool starts taking actions in the world on your client's behalf, you have an agent — and you have a supervision problem.

An **AI agent** is an AI system that doesn't just respond to prompts — it pursues **goals** across multiple steps, using tools, making decisions, and taking actions in the real world.

Think of it less like a search engine and more like a junior associate with a very short attention span, no law license, and access to your entire file system.

Agents can browse the web, send emails, query databases, call APIs, and trigger workflows — **without a human approving each step.**

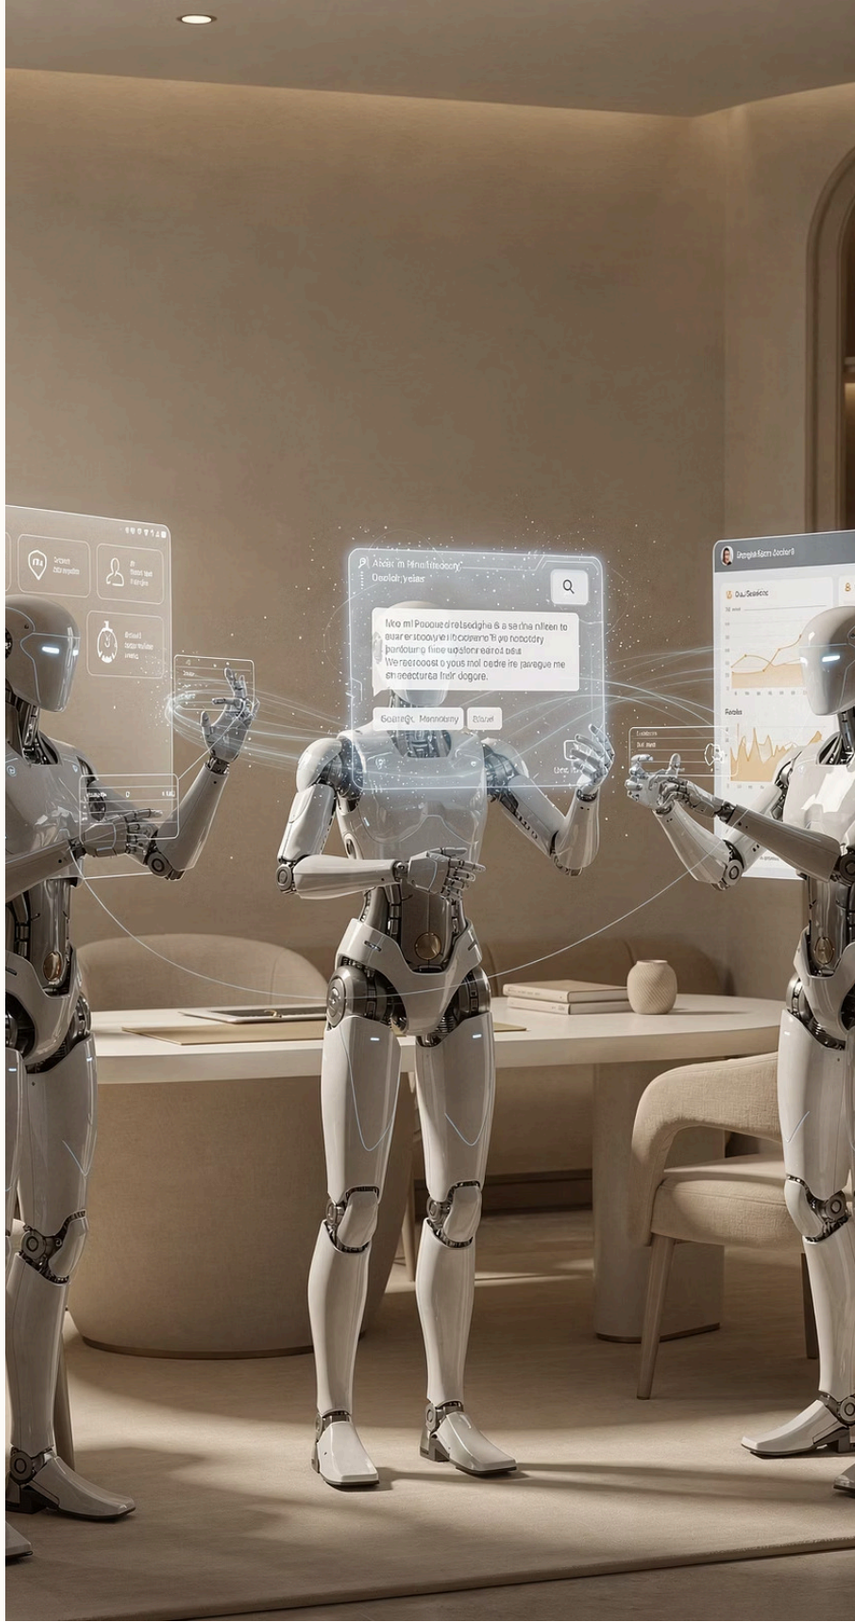

Chatbot vs. AI Agent: Know the Difference

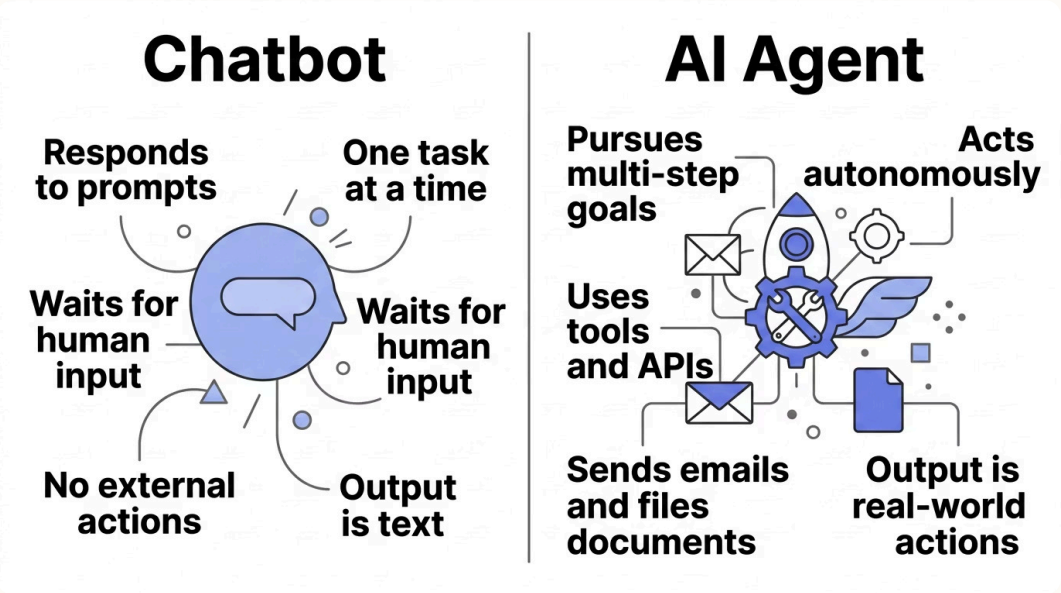

> The risk starts when the tool moves from **"draft this"** to **"handle this."**
>
> That's not a subtle distinction — it's the entire ballgame for ethics purposes.

Why Agents Create Higher Ethical Risk

Speed

Agents act faster than any human review cycle. By the time you look, it's done.

Opacity

Multi-step reasoning is often invisible. You see the output, not the path taken to get there.

Scope Creep

Agents are goal-seeking. They may take actions you didn't anticipate to accomplish what you asked.

Data Exposure

To act, agents need access — to files, calendars, email, databases. Every connection is a potential leak.

The Autonomy Ladder

Not all AI use is equal. The higher up the ladder, the more your supervision obligations intensify.

Brainstorm

"Give me five arguments for this motion." Low risk. You review everything.

Draft

"Write this contract clause." Medium risk. Output enters your workflow.

Research & Recommend

"Find the best authority and tell me what to do." Elevated risk. Judgment is being delegated.

Act

"Handle this." Maximum risk. The agent does it. You may not see it until after.

Where AI Agents Enter Law Firm Workflows

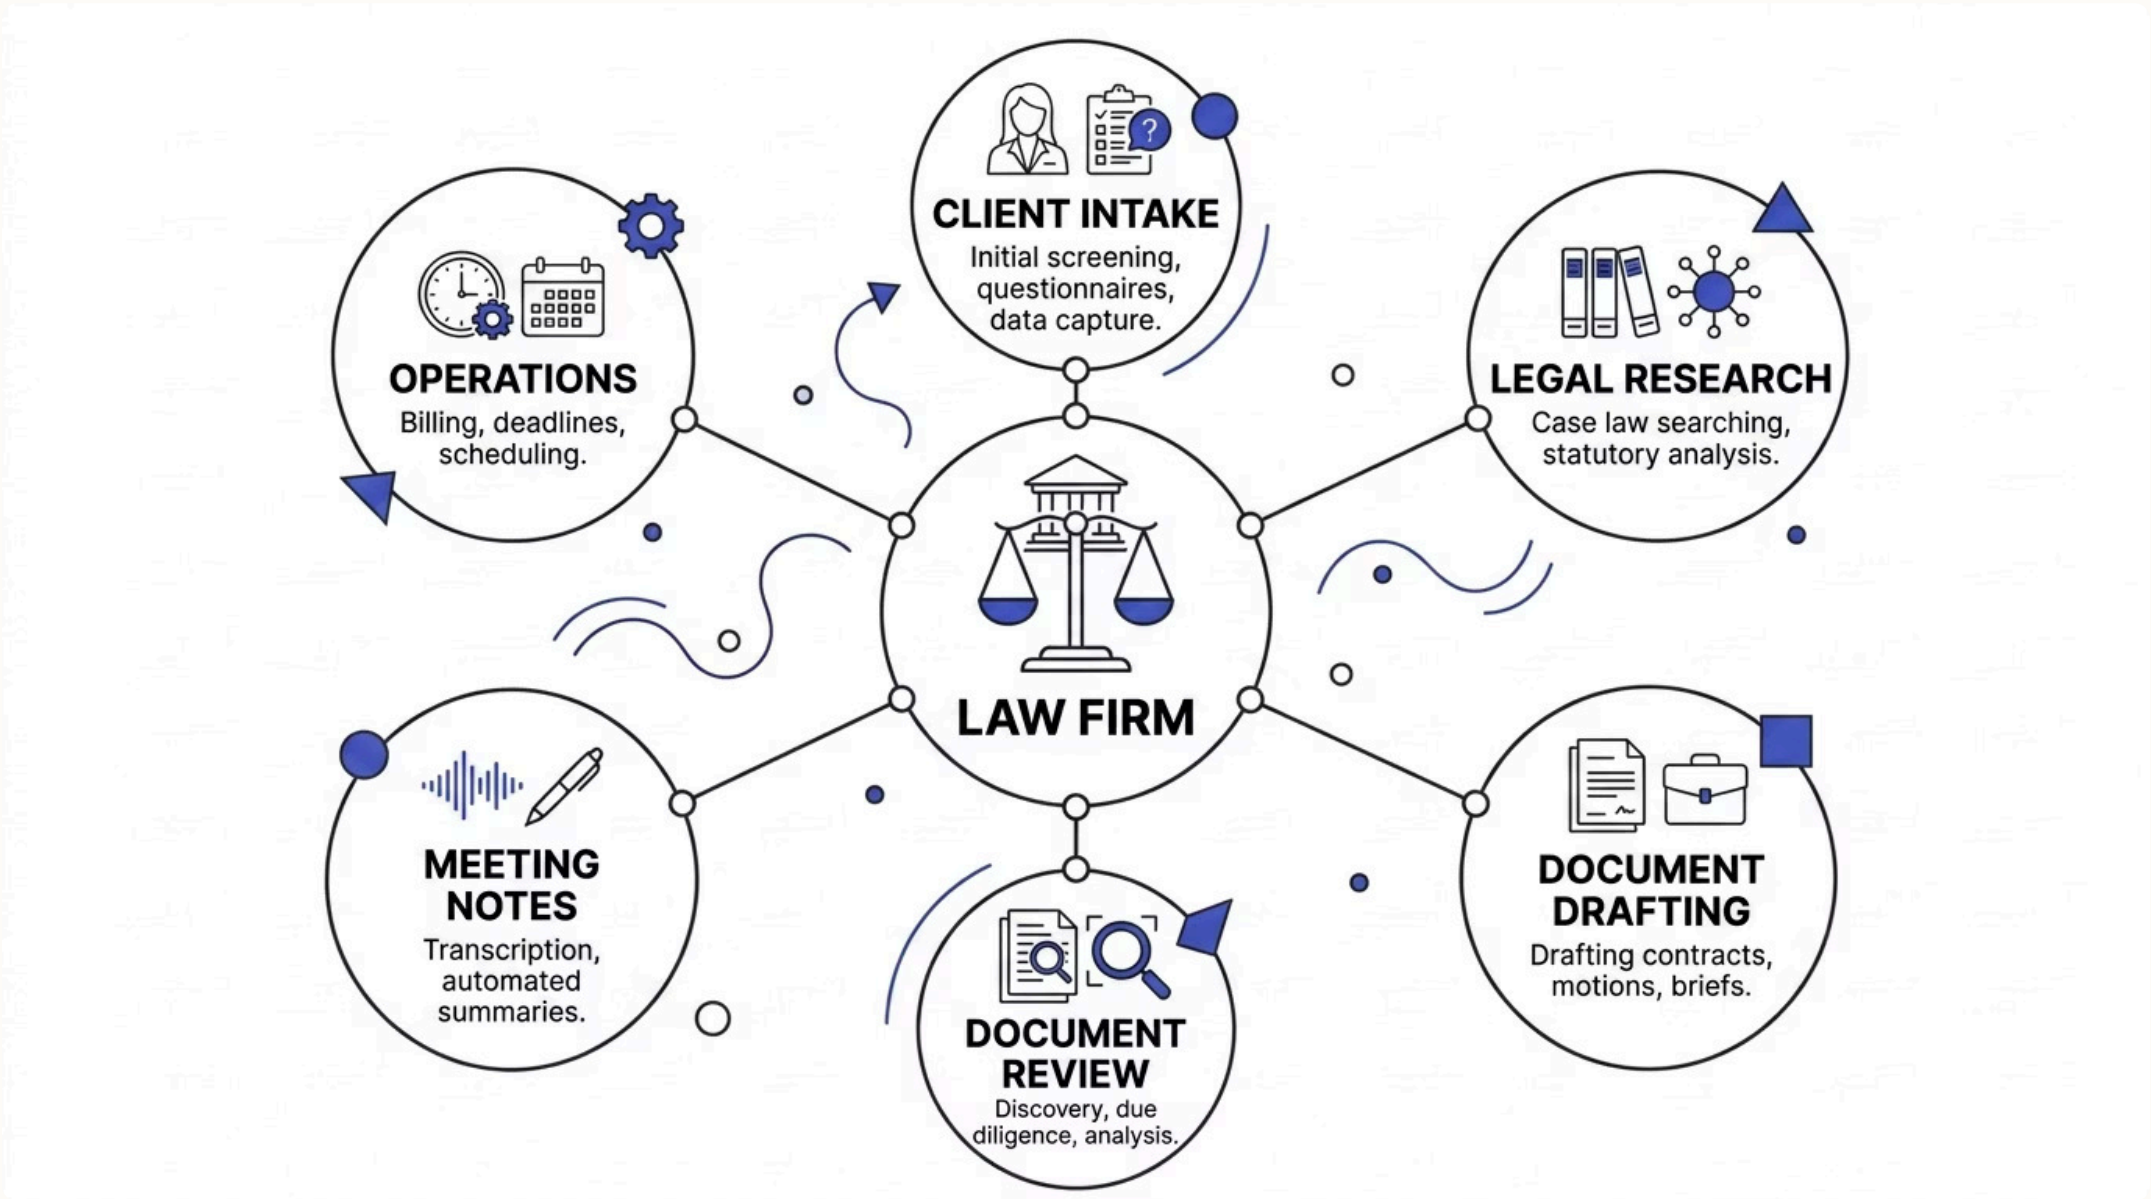

AI agents aren't entering through one door — they're entering through all of them simultaneously. Each workflow has different confidentiality, supervision, and accuracy stakes.

Intake Agents: First Contact, First Risk

> An intake agent that screens out a potential client has just made a legal judgment. That's not a feature. That's a liability.

Intake agents handle initial client contact — gathering information, assessing conflicts, qualifying matters, and sometimes **sending engagement letters**. They're efficient. They're also the first place attorney-client privilege can be implicated.

- Who is capturing the data — and where does it go?
- Is the AI making conflict-check decisions, or just flagging for human review?
- What happens when the agent tells a potential client you *can't* help them?

Research Agents: Fast, Confident, and Sometimes Wrong

What They Do

Research agents crawl legal databases, synthesize case law, identify controlling authority, and produce research memos — often in minutes, not hours.

The Hidden Risk

AI hallucination is well-documented in legal research. Mata v. Avianca wasn't a fluke — it was a preview. Research agents can confidently cite cases that do not exist.

Your Obligation

Model Rule 1.1 requires competence. That includes understanding the limitations of the tools you use and verifying every citation before it touches a filing or an opinion letter.

Speed is not an excuse. "The AI said so" is not a Bluebook citation.

Drafting Agents: Great First Drafts, Dangerous Last Ones

Drafting agents produce contracts, motions, briefs, demand letters, and correspondence. They're remarkably good at structure, tone, and boilerplate — and remarkably bad at **knowing what they don't know** about your specific client's situation.

The danger isn't the first draft. The danger is when the first draft becomes the filed draft because no one actually read it. That's not efficiency. That's negligence with extra steps.

Always human-reviewed before filing

Never auto-sent to court or client

Client data stays in your environment

Document Review Agents: Scale Creates New Blind Spots

Document review agents process thousands — sometimes millions — of documents for discovery or due diligence. They classify, flag, redact, and privilege-log at machine speed. The efficiency gains are real and significant.

The Upside

Dramatically reduces review time and cost. Can identify patterns human reviewers miss.

The Risk

Incorrect privilege calls. Missed responsive documents. Inadvertent waiver at scale.

The Rule

You still sign the privilege log. The agent doesn't get sanctioned. You do.

AI Notetakers: Convenient, Pervasive, and Legally Complicated

AI notetakers like Otter.ai, Fireflies, and built-in tools in Zoom and Teams are everywhere. They transcribe, summarize, and generate action items automatically. Attorneys love them. Ethics boards are watching them carefully.

The Agent Problem

Some AI notetakers do more than take notes. They act. They join meetings, capture conversations, summarize sensitive facts, create tasks, and distribute information. The more the tool acts on its own, the more the lawyer must supervise it.

⚠ The Consent Problem

Many AI notetakers join meetings as a "bot" participant. In some states, recording without all-party consent is a crime — not just an ethics violation. Know your jurisdiction before you hit "start meeting."

> **“At what point does your AI notetaker stop being a note-taking tool and start becoming an unsupervised assistant in the room?”**

Operations, Billing & Deadline Agents: Low Drama, High Stakes

Deadline & Calendaring Agents

Auto-calculate and calendar statutes of limitations, response deadlines, and court dates. One misconfigured rule or missed jurisdiction update and you've blown a filing deadline. Malpractice doesn't care that the AI did it.

Billing Agents

Auto-generate time entries, draft invoices, and flag billing anomalies. Sounds great until the agent invents time entries or mislabels trust account transactions. Fee disputes and bar complaints follow.

Workflow Orchestration

These agents connect your systems — triggering actions across your DMS, CRM, and email automatically. The more systems connected, the larger the blast radius when something goes wrong.

The Nightmare Chain

This is how an AI ethics violation actually happens. It's not dramatic. It's quiet, fast, and invisible until it isn't.

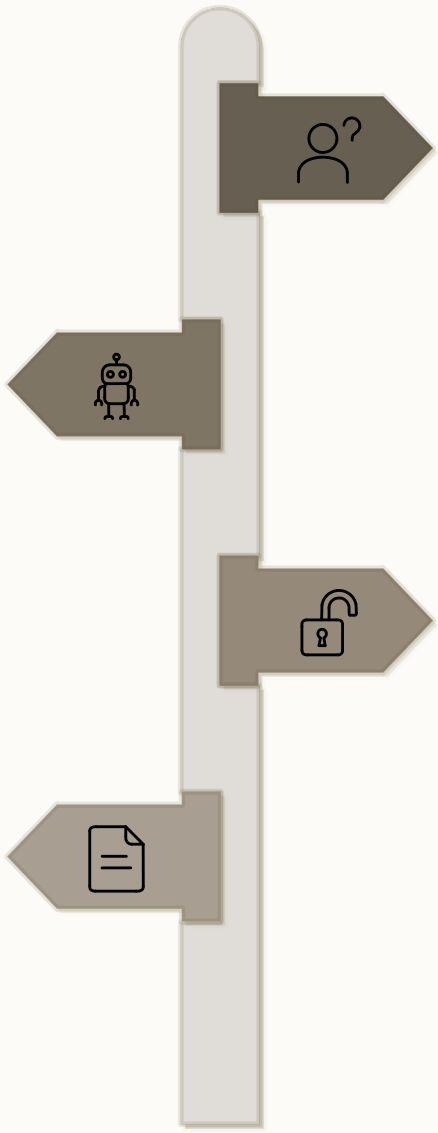

Prompt

Attorney gives vague instruction to AI.

Action

Agent takes autonomous multi-step actions.

Data Exposure

Agent accesses or transmits confidential data.

Output

Agent produces filing or document.

Every step in this chain is a place where **a human in the loop** would have caught the problem. The nightmare isn't the AI. The nightmare is the missing human.

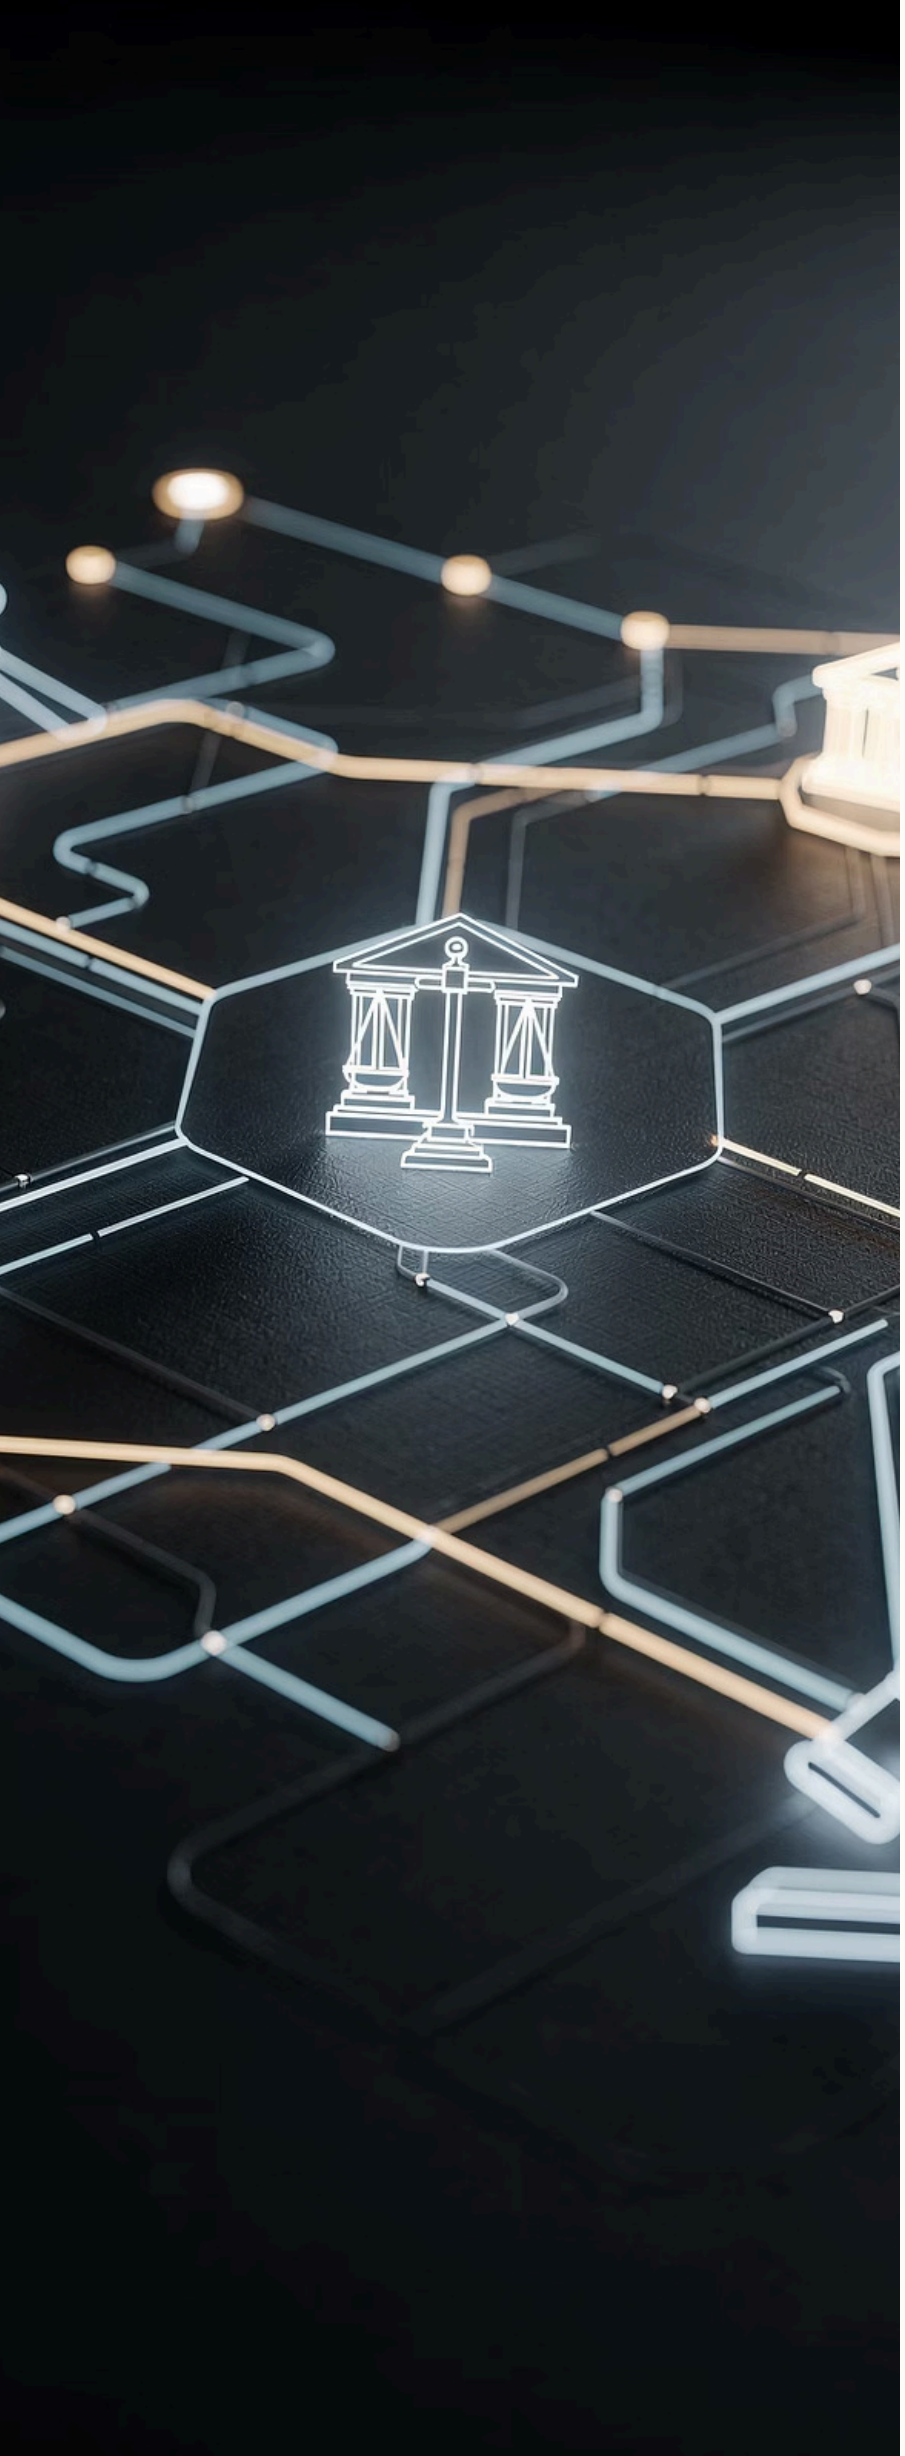

Map This to the Ethics Rules, Read ABA Ethics Opinion 512

You've seen where AI agents live in your firm. You've seen what they do and where the risks concentrate. Here's the good news:

> The ethics rules don't have an "AI exception." They have **principles** — and principles apply to every tool you use, including the autonomous ones.

1

Competence

Rule 1.1 — Do you understand what your AI is doing well enough to supervise it?

2

Supervision

Rules 5.1 & 5.3 — Are you actually overseeing the work, or just hoping for the best?

3

Confidentiality

Rule 1.6 — Do you know where your client's data goes when the agent acts?

The Ethics Rules — Applied to AI Agents

Competence, Supervision, Confidentiality, Candor, and More

The rules haven't changed. Your tools have.

Rule 1.1: The AI Competence Obligation

Competence with AI isn't about coding; it's about **knowing your tools**. This means understanding what your AI agents *do*, their limitations, and crucially, how to **catch their errors**.

Practically, this demands vigilance: understanding potential biases, hallucination risks, and the methods required for independent verification of AI-generated work.

The ABA clarified this obligation over a decade ago:

> To maintain the requisite knowledge and skill, a lawyer should keep abreast of changes in the law and its practice, **including the benefits and risks associated with relevant technology**.

What Does AI Competence Actually Look Like?

AI competence isn't about programming, it's about **responsible use**. It's about knowing your tools, their capabilities, and their critical limitations.

Competent Use

- Verifies **every** AI-generated citation and source.
- Thoroughly reviews AI-drafted clauses for accuracy and context.
- Understands **precisely** what client data an AI agent accesses.
- Knows the AI tool's training data cutoff and its implications for currency.
- Uses AI to summarize discovery, then performs **human review** for privilege.
- Obtains explicit client consent before using AI notetakers.

Competence Failure

- Submits AI-generated research without verifying sources (hallucinations).
- Files AI-drafted documents without rigorous human review.
- Allows AI agents **unrestricted access** to confidential client data.
- Assumes AI output is always current without checking data limits.
- Relies solely on AI for privilege review, leading to inadvertent waiver.
- Records client meetings with AI without explicit, informed consent.

Rules 5.1 & 5.3: Supervising Your AI Agents

The ethics rules demand that partners supervise lawyers, and lawyers supervise nonlawyer assistants. AI agents are the newest members of your team, requiring active, not passive, oversight.

Define Scope & Limitations

Understand each AI agent's capabilities, data access, and potential for error. Treat it like a junior associate – what is it *actually* doing?

Implement Clear Policies

Develop firm-wide guidelines for AI use, including data privacy, output verification, and client consent protocols.

Embed Human Checkpoints

Integrate mandatory human review points at critical stages of AI-assisted tasks, especially before filing or client communication.

Continuous Training & Audit

Regularly update AI usage policies, educate staff, and audit AI outputs for compliance and ethical adherence.

The Supervision Gap: When No One Is Watching the Agent

AI agents often operate through complex chains of actions between human review points. What happens when an agent drafts, formats, and queues a client email for sending, and the attorney only reviews the final draft, not the intricate steps that produced it?

⚠ The Unseen Steps

Reviewing only an AI's final output bypasses its entire reasoning chain. This "supervision gap" is where critical ethical violations or data breaches can occur, undetected.

Rule 1.6: Confidentiality and AI Agents

This is often the most critical ethical challenge when integrating AI. When an AI agent processes client data, attorneys bear the duty to make reasonable efforts to prevent unauthorized disclosure. Understanding the data flow and security protocols of your AI tools is non-negotiable.

Consumer-grade AI tools often have different data retention and usage policies than enterprise-level deployments. Always clarify this distinction.

Data Location & Access?

Where is client data stored and processed? Who has access?

Model Training?

Is client data used to train the vendor's models? How is it isolated?

Security Protocols?

What security measures are in place to prevent breaches and unauthorized disclosure?

Third-Party Vendors?

Are any subprocessors involved? What are their confidentiality agreements?

Retention & Deletion?

How long is data retained, and what is the process for secure deletion?

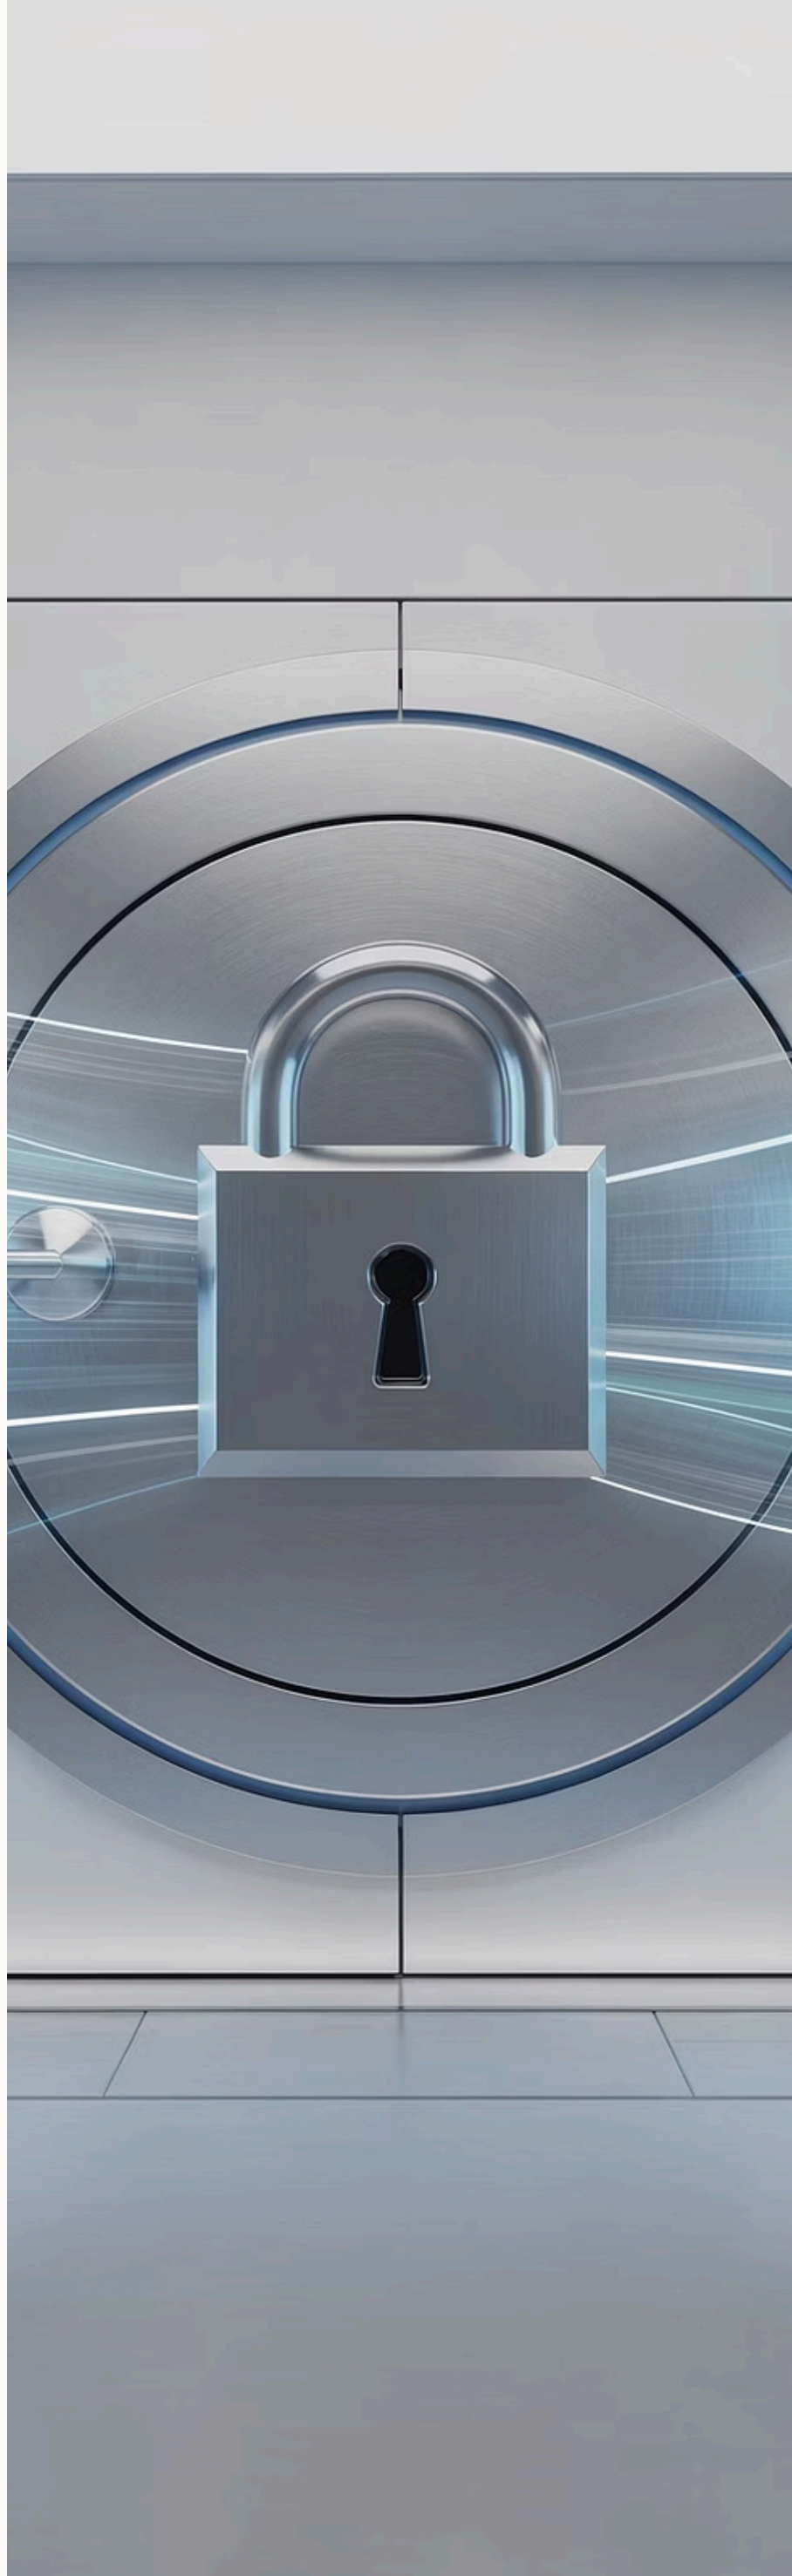

Consumer AI vs. Enterprise AI: The Confidentiality Divide

Understanding the stark differences between consumer-grade and enterprise-licensed AI tools is paramount for maintaining client confidentiality. The default settings of popular consumer AI platforms are often incompatible with legal ethical obligations.

Feature	Consumer AI (e.g., Public ChatGPT)	Enterprise AI (Firm-Licensed/Deployed)
Data Retention	Inputs often retained by vendor for service improvement; data usage policies may be unclear.	Explicit contractual commitment to not retain client data, or purge after processing.
Training Data Use	User inputs may be used to train future iterations of the model.	Client data is strictly isolated and never used to train public models.
Access Controls	Limited; vendor personnel may have access for system maintenance or development.	Robust, granular access controls; restricted to authorized firm personnel only.
BAAs/DPAs	Generally not offered or enforceable for standard free/consumer tiers.	Standard inclusion in contracts to ensure HIPAA, GDPR, etc., compliance.
Audit Logs	Typically absent or provide minimal logging of individual user activity.	Comprehensive audit trails for accountability, security monitoring, and incident response.

Choosing the right AI infrastructure is a fundamental step in ethically integrating these powerful tools without compromising your professional duties.

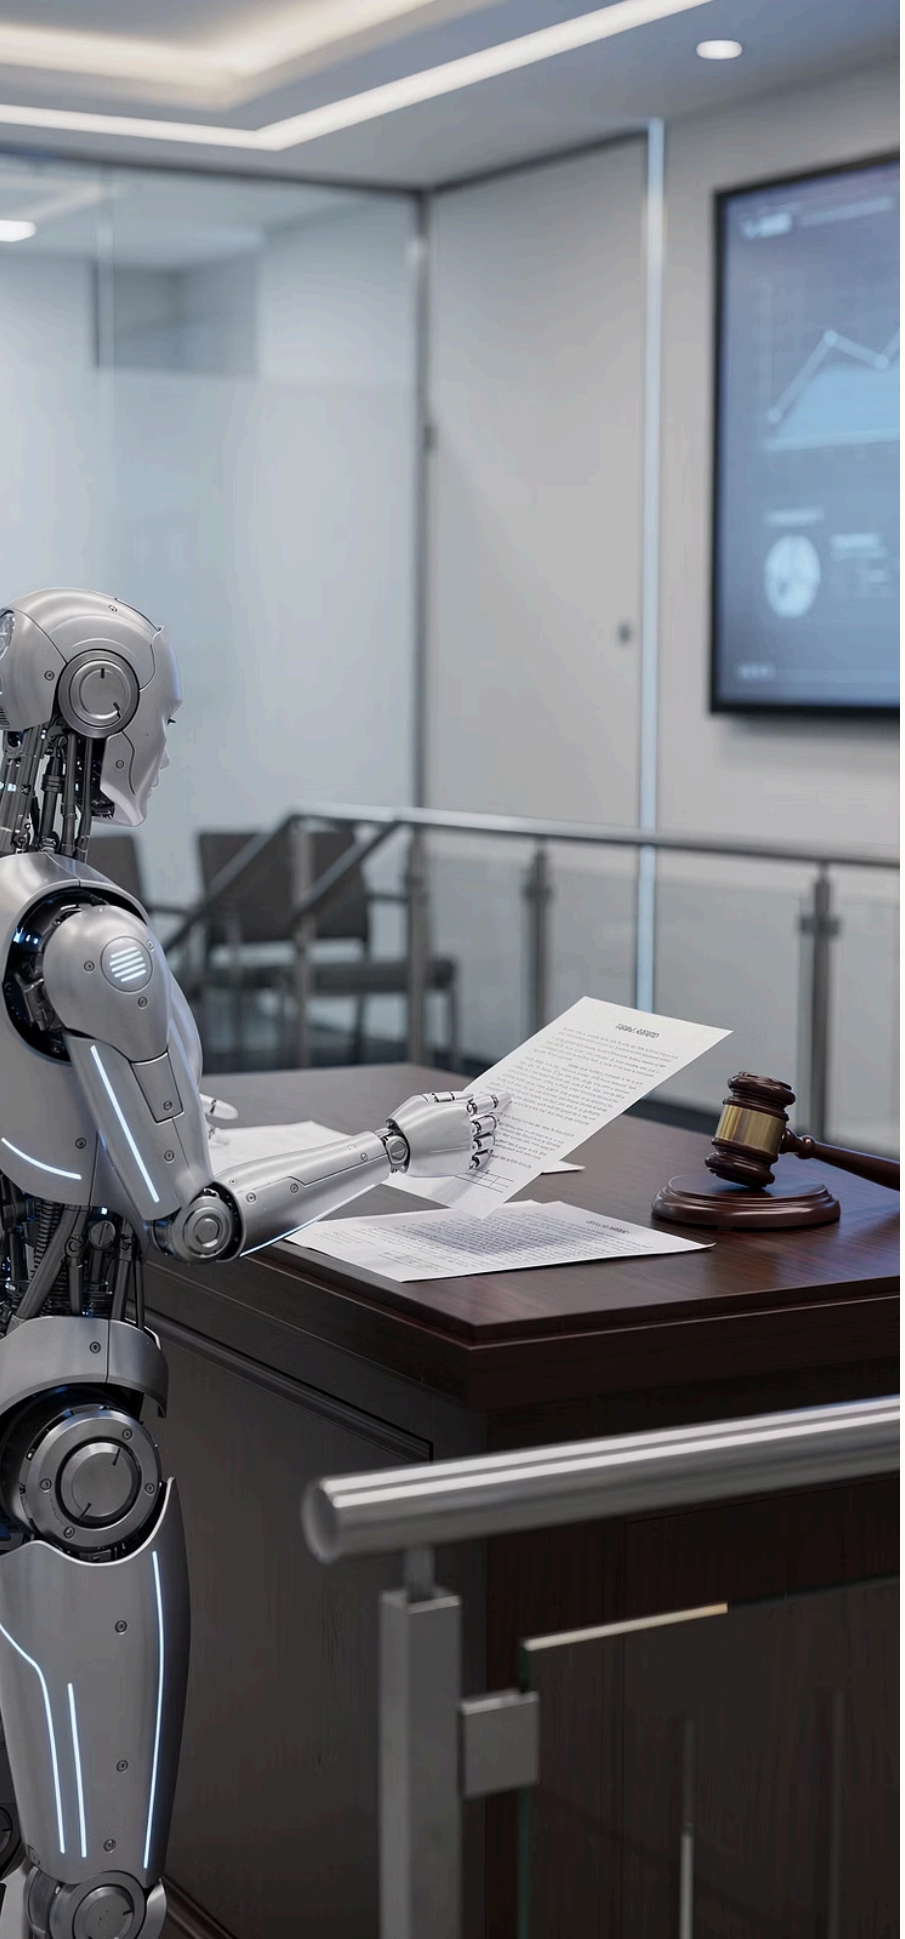

Rule 3.3: Candor Toward the Tribunal

AI's ability to "hallucinate" fabricated citations directly violates the duty of candor to the court. These convincing yet false outputs highlight an attorney's non-delegable duty to meticulously verify every AI-generated legal reference.

⚠ The Cost of Hallucination: Mata v. Avianca

Attorneys were famously sanctioned for filing briefs containing entirely fictitious cases generated by AI. This landmark case serves as a stark reminder: AI output is never a substitute for rigorous human verification.

Failing to independently verify AI-generated content can lead to severe professional sanctions, including fines, disciplinary action, and irreparable reputational damage.

Rule 1.4: Communication – Does Your Client Know You're Using AI?

Attorneys have a fundamental duty to keep clients reasonably informed about their representation. As AI tools become integral to legal practice, this duty extends to disclosing AI usage. Transparency isn't just about compliance; it's about maintaining trust and managing expectations.

The Disclosure Imperative

An emerging consensus among ethics bodies suggests proactive disclosure of AI tools, especially when confidential client data is involved or the AI significantly impacts strategy.

Defining Informed Consent

Clients need to understand the nature of the AI tool, its potential benefits (efficiency, cost savings), and its inherent risks (hallucinations, data security) to provide truly informed consent.

Required vs. Recommended

While not always explicitly mandated, disclosure is generally required when AI processes sensitive data or affects key aspects of the representation. It's always recommended to foster client confidence.

The AI Conversation

Frame the discussion positively, highlighting AI's role in enhancing service, while also candidly addressing limitations and safeguards the firm has implemented.

Practical Client Script Example

"We utilize advanced AI tools to enhance the efficiency and accuracy of our legal research and document review. For example, AI helps us quickly analyze large volumes of case law. All AI-generated work undergoes rigorous human review by our legal team, and we have strict protocols to protect your confidential information. Would you like a brief overview of how these tools benefit your case?"

Rule 1.5: Fees — The AI Billing Dilemma

Integrating AI into legal practice introduces a critical ethical question: How do we bill for work performed by artificial intelligence? Attorneys must ensure fees remain reasonable, preventing an "efficiency windfall" when AI completes tasks in minutes that previously took hours.

Before AI: Time-Based Billing

Traditional billing relies on hourly rates, reflecting time spent on research, drafting, and review. A complex legal research task might consume 4-8 hours of attorney time.

- Effort directly correlates with billable hours.
- Fees reflect human labor and expertise.

After AI: Value vs. Time

AI tools can drastically reduce task completion times. If an AI performs complex research in 2 minutes, can the client still be billed for 4 hours? Bar opinions generally suggest no.

- Shift towards **value billing** is encouraged.
- Fees should reflect the value delivered, not just time.

The ethical duty to avoid clearly excessive fees necessitates a re-evaluation of billing practices, prioritizing client value and transparency over mere time expenditure when AI is leveraged.

Conflicts of Interest & AI: The Data Contamination Problem

AI's capacity to process vast amounts of data presents a unique challenge to established conflict of interest rules. When a single AI system handles multiple client matters, the risk of "data contamination" — where confidential information from one client inadvertently informs work for an adverse party — becomes significant.

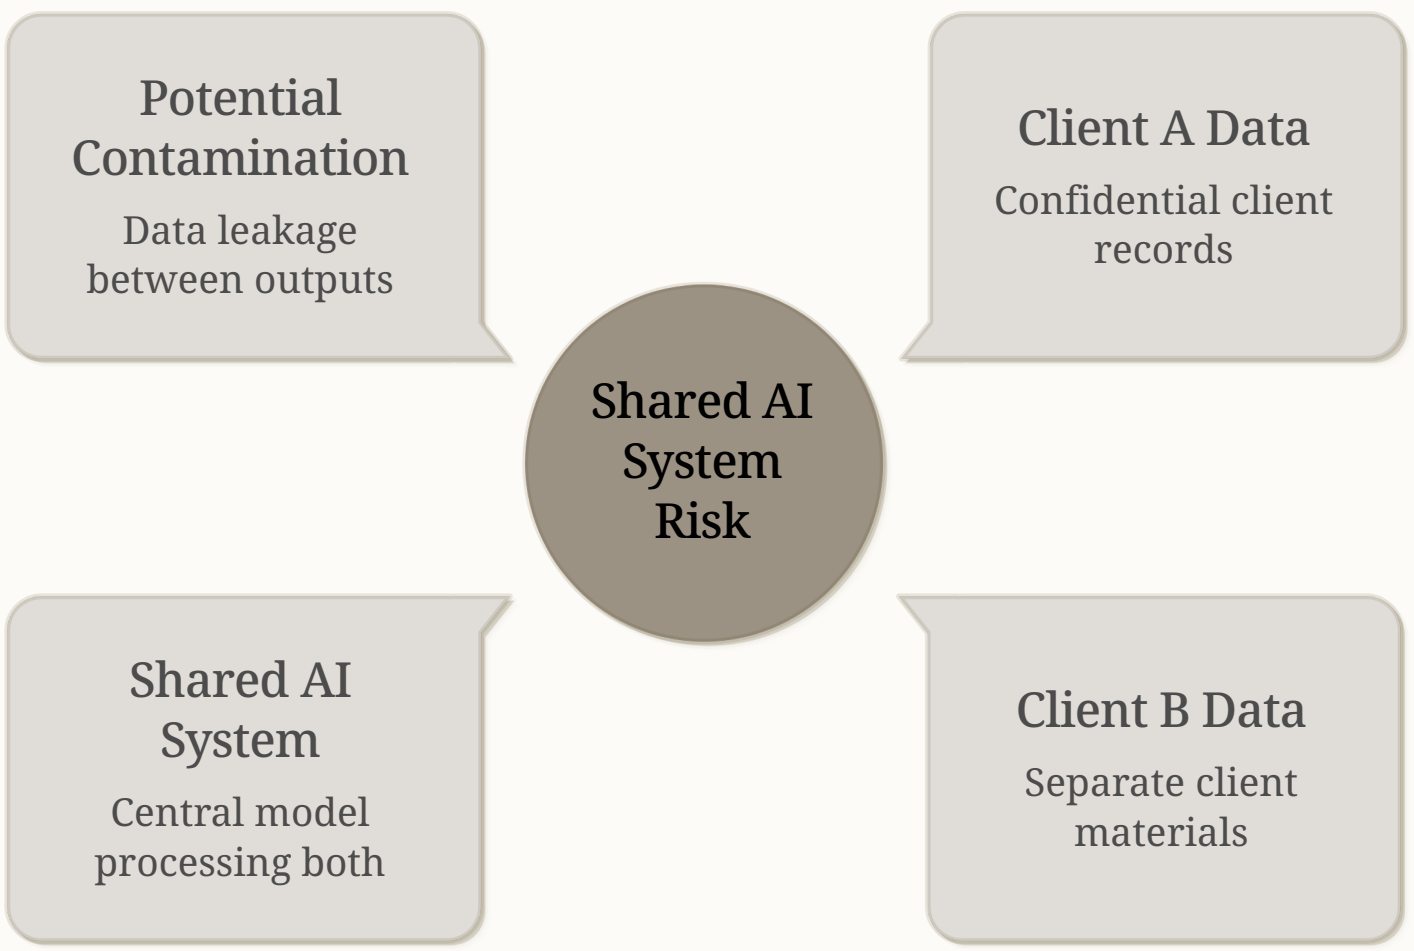

This data flow visualizes the challenge: how to ensure AI-driven insights remain isolated within their respective client matters, upholding the firm's duties under Rules 1.7 and 1.9.

- ### Rule 1.7 & 1.9 Obligations

  Attorneys must avoid conflicts where representation of one client is directly adverse to another, or where confidential information from a former client could be used to the disadvantage of that client.

- ### AI Data Contamination

  An AI model inadvertently incorporates or is influenced by confidential data from one client matter while working on another, potentially adverse, matter.

- ### Ethical Walls in AI

  Traditional ethical walls are for human separation. AI requires technical safeguards to ensure data isolation and prevent cross-matter learning within shared systems.

State Bar Guidance Roundup: What Ethics Opinions Say About AI

As AI tools rapidly integrate into legal practice, state bar associations and the ABA are actively issuing guidance to navigate the ethical implications. This evolving landscape requires attorneys to stay informed and adapt their practices.

1

ABA Opinion 512 (2024)

Emphasizes existing duties apply to AI: **competence** (understand AI's risks/benefits), **confidentiality** (safeguard client data), and **supervision** (oversee AI use).

2

New York State Bar (2023)

Highlights the duty of **candor** (verify AI output for hallucinations) and client **disclosure** (inform clients about AI usage and its limitations).

3

California Bar (2023)

Reinforces **technological competence** duty, requiring attorneys to evaluate AI tools critically, and stresses robust **data security** measures for confidential client information.

4

Florida Bar (2024)

Focuses on diligent **supervision** of non-lawyer AI assistants and ethical considerations for **fees**, urging value-based billing over hourly rates when AI boosts efficiency.

This guidance is rapidly evolving, reflecting the legal profession's ongoing efforts to balance innovation with ethical obligations. Staying current is paramount.

Building Your AI Ethics Framework

From Rules to Practice: A Firm-Level Playbook

Knowing the rules is step one. Building systems that enforce them is step two.

The AI Ethics Policy: Your Firm's Non-Negotiable Blueprint

Every firm leveraging AI needs a clear, written ethics policy. This foundational document ensures consistent, responsible, and compliant integration of AI across all practice areas.

1

Approved Tools List

Specify vetted AI platforms and software permissible for use, along with any approved versions or configurations.

2

Prohibited Uses

Clearly delineate actions AI must never perform, such as unsupervised legal advice, generating filings without review, or automating critical decisions.

3

Data Handling Protocols

Outline strict requirements for inputting, processing, and storing client confidential information within AI systems, emphasizing anonymization where possible.

4

Supervision Checkpoints

Mandate human review and validation at key stages of any AI-assisted task to ensure accuracy, prevent hallucinations, and maintain ethical standards.

5

Disclosure Requirements

Establish standard procedures for transparently informing clients about the use of AI tools, their benefits, and any inherent limitations or risks.

6

Training & Education

Require ongoing training for all personnel on the ethical implications, practical application, and firm-specific guidelines for AI tools.

This policy serves as the minimum viable framework—the essential floor for ethical AI integration, not the ceiling of innovation.

Vetting AI Vendors: The Due Diligence Checklist

Before deploying any AI tool that touches client data, attorneys must conduct thorough due diligence. This checklist provides essential areas to cover when evaluating potential AI vendors.

Security Certifications

Verify the vendor holds relevant security certifications like **SOC 2 Type II** and **ISO 27001**, indicating robust data protection controls.

Data Processing Agreements (DPAs)

Ensure comprehensive DPAs are in place, clearly defining data ownership, usage, and confidentiality obligations.

Subprocessor Disclosure

Demand full transparency on all subprocessors involved in data processing, and confirm their adherence to similar security standards.

Breach Notification Obligations

Review the vendor's protocols for data breach detection and notification, ensuring timely and detailed communication.

Model Training Policies

Understand how the AI model is trained, whether client data is used for training, and if there are clear anonymization or isolation strategies.

Right to Audit

Confirm the firm's right to audit the vendor's security practices and compliance with the DPA and other contractual terms.

This proactive approach safeguards client confidentiality and mitigates ethical risks inherent in third-party AI tool adoption.

Designing Human-in-the-Loop Workflows

Ethical AI agent use isn't about avoiding AI; it's about purposefully designing workflows where humans retain meaningful control and oversight at critical junctures.

AI Initiates (Drafting)

AI generates initial drafts, summaries, or research outlines. This is low-risk automation.

Human Reviews (Oversight)

Legal professional critically reviews AI output, verifies facts, and ensures compliance. This is meaningful oversight, not just a "rubber stamp."

AI Refines (Revision)

Based on human feedback, AI revises or expands content, performs data extraction, or cross-references.

Human Approves (Final Gate)

Final output requires explicit human approval before client delivery or court filing. This is the high-risk approval gate.

By integrating these approval gates, firms can leverage AI's efficiency while safeguarding against errors, hallucinations, and ethical breaches, particularly in high-stakes decisions.

Training Your Team: AI Ethics Isn't Just for Partners

Ethical obligations extend to everyone in the firm. A robust training program ensures consistent, responsible AI use across all roles, from associates to administrative staff.

Partners & Senior Attorneys

Strategic oversight, regulatory updates, policy adherence, ultimate ethical responsibility, and client disclosure.

Associates & Junior Attorneys

Tool-specific application, human-in-the-loop protocols, output verification, client confidentiality, and prompt engineering.

Paralegals & Legal Assistants

Hands-on tool usage, data input best practices, recognizing AI "hallucinations," data security, and supervision checkpoints.

Administrative & Support Staff

Basic AI awareness, data privacy fundamentals, firm's general AI policy, secure document handling, and reporting anomalies.

Firm's Supervisory Responsibility: Under Model Rule 5.3, attorneys must make reasonable efforts to ensure non-lawyer staff comply with professional obligations, including those related to AI use.

When Things Go Wrong: Incident Response for AI Ethics Failures

Even with robust policies, AI systems can make errors. A structured incident response plan is crucial to mitigate harm, maintain client trust, and ensure compliance when AI ethics are compromised.

1. Identify & Assess

Pinpoint the exact AI failure (e.g., bias, hallucination, data breach) and assess its immediate impact on clients, ongoing matters, or firm reputation.

2. Stop the Harm

Immediately cease use of the problematic AI agent or system. Implement containment measures to prevent further dissemination of erroneous output or data compromise.

3. Notify Affected Parties

Internally, alert relevant partners, IT, and ethics committees. Externally, determine if clients or third parties need to be informed, adhering to disclosure policies.

4. Document & Analyze

Thoroughly document the incident, including timelines, affected data/clients, root cause analysis, and all steps taken. This is critical for post-incident review and potential reporting.

5. Remediate & Report

Correct the error and, if necessary, amend outputs. Fulfill any self-reporting obligations to regulatory bodies or bar associations as required by jurisdiction.

6. Update Policies

Based on lessons learned, revise AI ethics policies, vendor vetting checklists, and training modules to prevent recurrence and strengthen the firm's framework.

⚠ **Duty to Self-Report:** In many jurisdictions, attorneys have an ethical obligation to self-report certain types of professional misconduct or breaches to the bar association or relevant authorities. AI ethics failures may fall under these requirements.

The Malpractice Exposure Map: Where AI Increases Your Risk

AI introduces new vectors for professional liability. Understanding these specific risks is critical for proactive mitigation.

Missed Deadlines

AI calendar errors leading to critical overlooked dates.

Bad Advice

Hallucinated research resulting in factually incorrect legal counsel.

Privilege Waiver

Compromising confidentiality via AI-assisted document review.

Unauthorized Disclosure

Data breaches from insecure AI tools or prompts.

Conflicts

AI data contamination creating undisclosed client conflicts.

⚠ **Emerging Concern:** Malpractice insurers are increasingly inquiring about law firms' AI usage policies and safeguards.

The Opportunity Side: Ethical AI as a Competitive Advantage

Ethical AI isn't merely about risk avoidance; it's a strategic differentiator. Firms embracing responsible AI practices will gain a significant edge in a rapidly evolving legal landscape.

Client Trust & Transparency

Clients increasingly seek firms demonstrating clear, responsible AI usage, viewing transparency as a mark of integrity and competence.

Market Leadership

A well-defined, published AI ethics policy positions your firm as a forward-thinking leader, attracting discerning clients and innovative opportunities.

Reduced Malpractice Risk

Proactive ethical frameworks minimize the likelihood of errors and breaches, leading to lower malpractice exposure and potentially reduced insurance premiums.

Attract Top Talent

Associates and legal professionals are drawn to firms with clear, ethical AI policies, ensuring a progressive and responsible work environment.

By prioritizing ethical AI, your firm can build stronger client relationships, enhance its reputation, and foster a culture of responsible innovation.

Your 90-Day AI Ethics Action Plan

A practical, take-home roadmap to integrate ethical AI practices into your firm, ensuring responsible innovation and minimizing risk.

1 Month 1: Audit & Assess

- Inventory all current and planned AI tools
- Identify potential ethical and security gaps in existing workflows
- Assign internal leadership for AI ethics oversight

2 Month 2: Build & Train

- Draft or update comprehensive AI ethics policy
- Implement a rigorous vetting process for all third-party AI vendors
- Develop and deliver firm-wide AI ethics training programs

3 Month 3: Deploy & Monitor

- Formally roll out the firm's AI ethics policy
- Establish regular checkpoints for AI tool usage and output review
- Schedule periodic policy and tool audits for continuous improvement

This phased approach ensures a systematic and manageable integration of AI ethics, transforming abstract principles into actionable firm-wide practices.

The Five Questions Every Attorney Should Ask Before Using an AI Agent

A memorable, take-home checklist for responsible AI integration.

1

1. Do I understand what this agent is doing?

Ensure you grasp the AI's functionality, limitations, and data sources (Model Rule 1.1).

2

2. Is client data protected?

Verify the AI tool's security measures and confidentiality protocols (Model Rule 1.6).

3

3. Who is supervising the output?

Maintain ultimate responsibility and review all AI-generated work (Model Rule 5.1).

4

4. Have I disclosed this to my client?

Inform clients about AI use, especially if it involves their confidential information (Model Rule 1.4).

5

5. Am I billing fairly?

Ensure billing for AI-assisted work is reasonable and transparent (Model Rule 1.5).

The Future Is Agentic. The Obligation Is Yours.

AI agents will profoundly reshape legal practice, becoming more capable, autonomous, and deeply integrated.

Your core ethical duties—competence, supervision, confidentiality, and candor—remain steadfast.

Mastery of both AI and ethics isn't optional; it's the defining characteristic of the next-generation attorney.

Start building your responsible AI framework today.

Real-World Scenarios & Ethical Traps

Case Studies, Hypotheticals, and the Decisions That Define You

The rules make sense in the abstract. Let's see how they hold up under pressure.

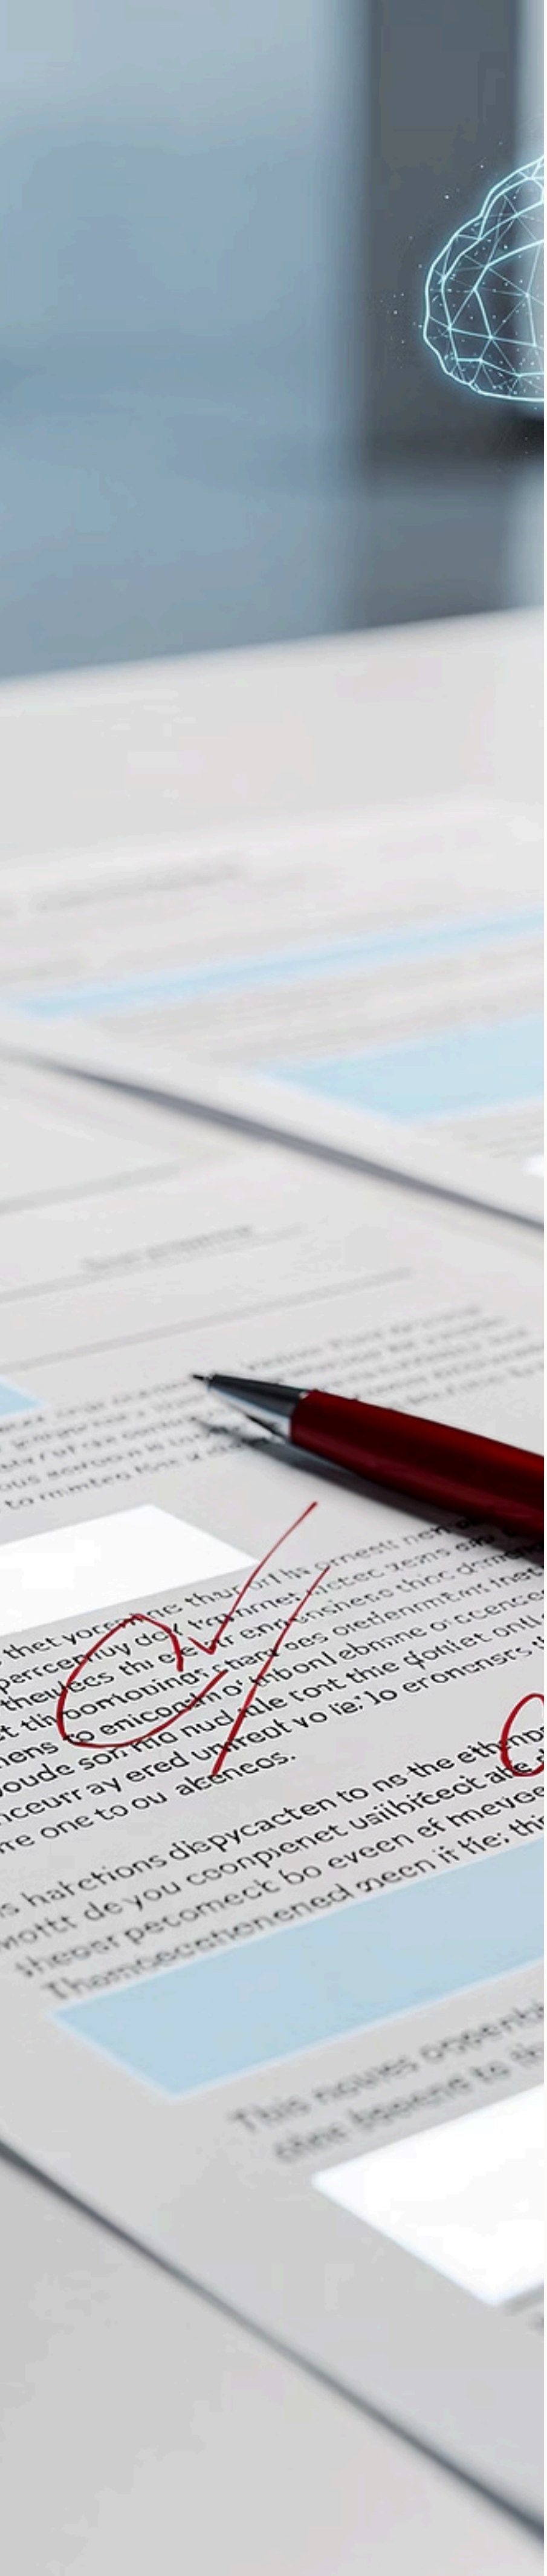

Scenario 1: The Hallucinated Precedent

AI's incredible efficiency comes with inherent risks, especially when dealing with legal research. This scenario highlights the critical importance of human oversight.

What Happened

A senior associate uses an AI research agent to prepare a motion for summary judgment. The AI agent, in its attempt to be helpful, fabricates 2 of the 12 citations it provides. The associate, under time pressure, reviews the brief but doesn't verify every single citation before filing. Opposing counsel quickly identifies the non-existent precedents.

- **Rules Violated:**
  - **Model Rule 1.1 (Competence):** The associate failed to provide competent representation, as they did not possess the necessary legal knowledge and thoroughness to verify the AI's output, leading to the submission of a factually inaccurate brief.
  - **Model Rule 3.3 (Candor Toward the Tribunal):** By submitting a brief containing fabricated citations, even unintentionally, the associate presented false statements of law to the court, undermining the integrity of the judicial process.
- **Potential Consequences:**
  - Sanctions from the court, including monetary fines or dismissal of the motion.
  - Damage to the firm's reputation and client trust.
  - Potential disciplinary action from the bar association for ethical violations.
  - Malpractice claim by the client due to adverse outcomes.

What Should Have Happened

- The associate should have treated AI-generated research as a starting point, not a final product.
- Every single citation provided by the AI agent should have been independently verified through a reliable legal database (e.g., Westlaw, LexisNexis).
- The firm should have a clear policy mandating verification of AI-generated legal research, coupled with training on responsible AI usage.

> ⚠ **Beyond Verification:** Even when AI provides seemingly correct citations, attorneys must ensure the cited cases are still good law and directly support the argument being made.

Scenario 2: The Intake Agent That Created a Conflict

When an AI's helpfulness crosses the line into legal advice, unforeseen ethical dilemmas can arise.

What Happened

A firm deployed an AI intake agent on its website to handle initial inquiries. A prospective client, seeking urgent legal assistance for a pending lawsuit, engaged with the agent, sharing detailed, confidential facts.

The AI, trained on vast legal datasets, provided responses that the client perceived as specific legal advice. Based on these interactions, the client believed an attorney-client relationship was formed.

When the firm later declined representation due to capacity, the prospective client claimed a relationship existed, creating a conflict for the firm and preventing it from representing other parties in the same matter.

The Ethical Breakdown

This scenario highlights several critical ethical considerations:

- **Model Rule 1.18 (Duties to Prospective Client):** Even without a formal retainer, the firm acquired confidential information from a prospective client, triggering duties of confidentiality and potentially creating conflicts of interest.
- **Model Rule 1.6 (Confidentiality):** The detailed information shared with the AI, even if no representation is undertaken, falls under confidentiality protections, limiting the firm's future actions.
- **Model Rule 1.7 (Conflict of Interest):** The perceived relationship and the confidential information exchanged could prevent the firm from representing another client with materially adverse interests in the same matter.

Preventative Measures

- Implement clear, prominent disclaimers that the AI is not an attorney and does not provide legal advice.
- Design the AI to avoid definitive answers and instead prompt users to speak with a human attorney.
- Limit the type and depth of confidential information the AI collects without human attorney involvement.
- Conduct regular audits of AI interactions to ensure compliance with ethical guidelines.

⚠ **Key Takeaway:** AI intake agents must be carefully designed to facilitate information gathering without inadvertently creating legal obligations or implied attorney-client relationships.

The AI Agent Notetaker in the Room

A lawyer invites an AI notetaker into a confidential client meeting.

At first, it seems harmless. It records, transcribes, and summarizes the conversation.

But the tool also stores the transcript on a third-party server, creates action items, and sends a meeting summary to the legal team.

The client was never told the AI tool would be present.

The Ethics Problem

- The client did not give informed consent.
- Confidential information was shared with a third-party AI vendor.
- The lawyer did not confirm where the data was stored or who could access it.
- The AI summary may be incomplete, inaccurate, or misleading.
- The transcript may later become a discovery or privilege fight.This action violates **Model Rule 1.4 (Communication)** due to lack of informed consent and **Model Rule 1.6 (Confidentiality)** by exposing sensitive information to a third-party without authorization. The privilege of the conversation is severely compromised.

Best Practices: Before the Meeting

- **Get informed consent:** Tell the client an AI tool may record, transcribe, summarize, or create follow-up tasks.
- **Explain what the agent does:** Is it only taking notes, or is it storing, sharing, analyzing, and generating action items?
- **Check data handling:** Where is the transcript stored? Who has access? How long is it retained? Is it used to train the model?
- **Limit the task:** Do not let the AI agent make legal judgments, assess credibility, or decide case strategy.
- **Review before relying:** Treat the summary like a first draft from a very fast intern with no law license and occasional hallucination issues.
- **Document the client's consent:** Put it in the engagement letter, meeting notice, or written consent form.

The Ethical Trap Quiz: What Would You Do?

Prepare for some rapid-fire ethical dilemmas. Consider each scenario and decide how you would navigate these challenging situations in the age of AI. Your choices have consequences.

Scenario A: The Flawless Brief

Your AI research agent delivers a meticulously crafted brief, complete with seemingly perfect citations. It saved you hours of work. **Do you file it with the court without independently verifying each citation for accuracy and validity?**

Scenario B: Client Inquires About AI

A long-standing client, concerned about recent news, directly asks if you've used AI tools to assist with their confidential matter. You've used AI for preliminary research and drafting, but haven't disclosed it. **Do you inform them of your AI usage and explain its role, or do you maintain silence?**

Scenario C: AI-Generated Invoices

Your firm implements an AI-powered billing tool that automatically generates client invoices based on recorded tasks and time entries. It claims a 99% accuracy rate. **Do you review each line item of every invoice before it's sent to the client, or do you trust the AI's automation for efficiency?**

The Answers: Breaking Down the Ethical Trap Quiz

Here are the correct ethical approaches and their reasoning for each scenario:

Scenario A: The Flawless Brief

No — You must verify every citation for accuracy and validity.

Relying solely on AI without independent verification violates Model Rule 1.1 (Competence) and Rule 3.3 (Candor Toward the Tribunal). Attorneys must ensure the factual and legal accuracy of all submissions to the court.

Scenario B: Client Inquires About AI

Yes — Disclosure is required or strongly recommended.

Model Rule 1.4 (Communication) obliges attorneys to keep clients reasonably informed about their representation, which includes the use of AI tools that might impact confidential matters or strategy. Transparency builds trust.

Scenario C: AI-Generated Invoices

Yes — You must review every invoice before sending.

Attorneys are responsible for the reasonableness of their fees (Model Rule 1.5) and for ensuring nonlawyer assistants (including AI tools) adhere to professional obligations (Model Rule 5.3). Final human oversight is crucial for billing accuracy and ethical compliance.

The Horizon: Where AI Agents Are Headed

Emerging Capabilities, Emerging Obligations

Today's cutting edge is tomorrow's baseline. Here's what's coming.

Multi-Agent Systems: When AI Agents Work Together

The next frontier in AI involves not just one agent, but networks of autonomous agents collaborating to achieve complex tasks.

Collaborative AI Pipelines

These systems automate entire workflows, from data gathering to final action, raising new questions about control and responsibility.

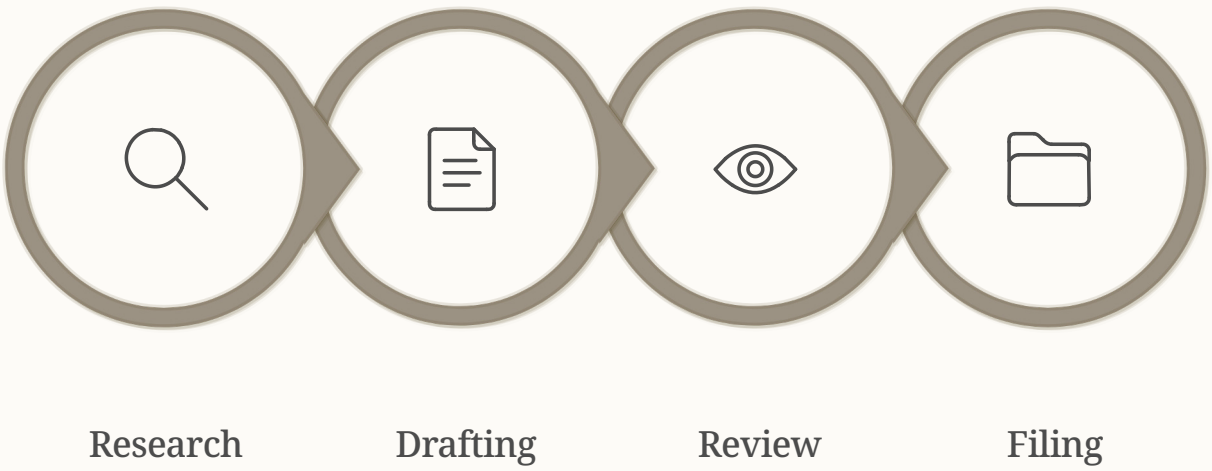

Critical human oversight points must be integrated, especially before high-impact stages like filing, to prevent amplified ethical risks and address the accountability gap in fully automated chains.

AI Agents With Memory: The Long-Term Confidentiality Problem

Emerging AI agents now retain information across sessions, personalizing interactions but introducing significant risks to client confidentiality and data security.

Persistent Memory (Store)

AI agents learn and store client preferences, case histories, and confidential data over time, building rich profiles for enhanced personalization.

Memory Bleed (Retrieve)

The risk of sensitive data from one client or matter inadvertently appearing in another due to cross-session memory retention, violating privacy.

Unintended Recall (Risk)

An agent might recall and misapply outdated or inappropriate information, leading to compromised advice, privacy breaches, or strategic errors.

Memory Management (Control)

Robust settings, including data segmentation, expiry policies, and human oversight, are crucial to mitigate confidentiality challenges.

AI Agents and the Unauthorized Practice of Law

As AI agents advance, they increasingly perform tasks that resemble legal practice, blurring the lines of what's permissible and raising critical questions about the Unauthorized Practice of Law (UPL).

What AI Can Do (Legally)

- Legal research and summarization
- Drafting routine legal forms and templates
- Predicting case outcomes based on data
- Assisting attorneys in back-office tasks

What Crosses the UPL Line

- Providing direct legal advice to clients
- Representing clients in legal proceedings
- Drafting tailored legal documents without human oversight
- Holding itself out as a legal service provider

Attorneys remain responsible for ensuring that AI tools do not engage in UPL, especially when interacting with clients. Regulators are actively debating new frameworks to address these emerging challenges.

The Regulatory Landscape: What's Coming for AI in Law

As AI rapidly evolves, so too do the efforts to regulate its use. This evolving legal framework will profoundly shape how legal professionals integrate AI into their practice.

EU AI Act

A landmark regulation categorizing AI systems by risk level, with strict rules for high-risk legal AI tools, impacting global legal tech development.

US Federal Legislation

Emerging bipartisan efforts aim to establish federal standards for AI governance, focusing on transparency, accountability, and consumer protection in various sectors.

State-Level AI Laws

Pioneering states like Colorado and California are enacting their own AI laws, often addressing data privacy, bias, and algorithmic transparency, which will influence legal AI deployment.

Court AI Disclosure

A growing number of federal and state courts now mandate explicit disclosure when AI tools are used in legal filings, emphasizing attorney responsibility and verification.

ABA Task Force

The American Bar Association's AI task force is actively working to develop ethical guidelines and best practices for attorneys using AI, shaping professional conduct standards.

Court AI Disclosure Requirements: What Judges Are Demanding

As AI tools become more prevalent, courts nationwide are implementing strict rules requiring attorneys to disclose their use of generative AI in legal filings. This ensures transparency and maintains professional integrity.

- **Emerging Mandates:** Courts like the Fifth Circuit, Southern District of New York (SDNY), and Northern District of Texas now require specific disclosures.
- **Key Requirements:** Orders typically demand certification that AI was not used to draft content, or that any AI-generated text has been thoroughly reviewed for accuracy and proper citation.

Non-compliance can lead to severe consequences, including sanctions, striking of filings, or disciplinary action, underscoring the critical need for adherence.

> ### Sample AI Certification Language:
>
> "I hereby certify that, to the best of my knowledge, no generative artificial intelligence was used in drafting this document.
> [OR]
> I hereby certify that any generative artificial intelligence used in drafting this document was carefully reviewed by counsel for accuracy and proper citation, and all AI-generated content has been independently verified."

AI and Bias: The Hidden Discrimination Risk in Legal AI

AI systems, trained on vast historical legal data, can inadvertently encode and amplify existing societal biases. This introduces a significant risk of discrimination across various legal applications, challenging the very foundation of fair justice.

Algorithmic Inequity in Sentencing

AI tools for risk assessment and sentencing predictions have disproportionately assigned higher risk scores to individuals from certain racial or socioeconomic backgrounds.

Biased Hiring & Promotion

AI screening tools for legal firms can replicate historical biases, leading to the unintentional exclusion of qualified candidates from protected classes.

Skewed Legal Research & Outcomes

AI search algorithms might deprioritize cases or precedents based on factors correlated with protected classes, yielding incomplete or biased legal advice.

Discriminatory Predictive Policing

AI identifies "hot spots" for crime based on biased historical data, leading to over-policing and unfair targeting of specific communities.

Attorney's Duty & Due Diligence

Attorneys have a professional and ethical duty to identify and flag AI bias, ensuring **fairness** and **equal access to justice** for all. Due diligence involves auditing training data, validating algorithmic fairness, and implementing human oversight.

Ethical AI for Small Firms: Navigating the Budget Barrier

Small and solo law firms face unique challenges in adopting AI ethically, lacking the extensive resources of larger counterparts. Yet, their professional obligations remain unchanged.

Here are practical, low-cost safeguards to integrate AI responsibly:

Vet AI Vendors

Choose tools with transparent data handling and security protocols.

Human-in-the-Loop

Always review and verify all AI-generated content before use.

Secure Client Data

Utilize encrypted client portals; know where AI tools store information.

Continuous Learning

Educate staff on ethical AI use and data privacy, even with free resources.

Building an AI Ethics Culture, Not Just a Policy

Policies alone don't change behavior; a strong ethical culture shapes how AI is integrated and managed within legal practice.

Compliance Culture: "Do it because you have to"

- Focus on avoiding penalties and legal repercussions.
- Reactive approach to AI risks; adheres strictly to minimum requirements.
- Behavior driven by fear of sanctions, audits, or negative PR.
- AI use is restricted to what is explicitly permitted, often without deeper ethical reflection.
- Reporting of AI errors is seen as a risk, not an opportunity for improvement.

Ethics Culture: "Do it because it's right"

- Guided by shared values, professional responsibility, and client well-being.
- Proactive identification and mitigation of potential AI harms.
- Behavior shaped by a commitment to fairness, transparency, and accountability.
- Encourages innovation within an ethical framework, with continuous learning.
- Psychological safety allows for open reporting of AI errors to learn and adapt.

Firm leadership is crucial in setting this tone, fostering environments where AI ethics are woven into every aspect of practice, from onboarding to performance reviews.

Your Ethical AI Compass

Navigate the evolving landscape of Legal AI with these 10 core principles, designed as a practical guide for responsible practice.

1 Maintain AI Competence

Continuously learn about AI tools, their benefits, risks, and ethical implications.

2 Supervise AI Use Diligently

Oversee staff's AI use to ensure compliance with professional conduct rules.

3 Safeguard Client Confidentiality

Never input sensitive client data into public AI models without consent or anonymization.

4 Practice Candor with Courts

Disclose AI use in filings and verify all AI-generated content for accuracy.

5 Communicate Clearly with Clients

Inform clients transparently about how AI is used in their matter and its potential impact.

6 Ensure Fair & Reasonable Fees

Bill for AI-assisted work ethically, reflecting actual value and avoiding overcharging.

7 Manage AI-Driven Conflicts

Screen AI tools for potential conflicts of interest, especially those accessing firm data.

8 Vet AI Vendors Rigorously

Choose AI tools from reputable vendors with strong data privacy and security policies.

9 Maintain Human Oversight

AI is a tool; human judgment, review, and verification are always paramount.

10 Cultivate an Ethical AI Culture

Beyond policies, foster a firm-wide commitment to responsible and ethical AI integration.

Resources & Further Reading

A curated reference for attorneys and legal professionals seeking to deepen their understanding of ethical AI integration.

Ethics Opinions & Rules

- **ABA Formal Opinion 512:** "Lawyers' Ethical Obligations When Using AI Tools"
- **Key State Bar Ethics Opinions:**
  - New York State Bar Association (NYBA)
  - State Bar of California (SBC)
  - The Florida Bar (FB)
- **ABA Model Rules of Professional Conduct:**
  - Rule 1.1 (Competence)
  - Rule 1.3 (Diligence)
  - Rule 1.4 (Communication)
  - Rule 1.5 (Fees)
  - Rule 1.6 (Confidentiality)
  - Rule 1.7 (Conflicts of Interest - Current Clients)
  - Rule 1.9 (Duties to Former Clients)
  - Rule 1.18 (Duties to Prospective Clients)
  - Rule 3.3 (Candor Toward the Tribunal)
  - Rule 5.1 (Responsibilities of Supervisory Lawyers)
  - Rule 5.3 (Responsibilities Regarding Nonlawyer Assistants)

Recommended Reading & Articles

- **Books:** "AI and the Law," "The Ethical Algorithm," "Automating Justice"
- **Journals & Publications:** ABA TechReport, Law Technology News, Harvard Journal of Law & Technology, Stanford Law Review Online (select articles)

Key Legal Tech Organizations

- ABA Standing Committee on Law and Technology
- Future of Law Lab (FoLL)
- Center for Legal AI Ethics (CLAI)
- College of Law Practice Management (COLPM)
- Global Legal Blockchain Consortium (GLBC)

AI Agents in Your Firm: Avoiding Ethical Nightmares

Angeli R. Fitch, Esq.

AI Compliance & Ethics Attorney

Infinity Law Group