The ABA 512 Compliance Checklist

Run this pre-flight before any Track B step goes live

Two questions price a Track B step: does client information go into the AI, and do you rely on the output in a client's representation? Yes to both puts you under the full five-duty discipline: competence, confidentiality, supervision, fees, and candor. This checklist turns those duties into a single pre-flight you run in minutes, before any **Track B** step goes live.

It has three parts. The five-duty pre-flight covers the core obligations from ABA Opinion 512. The informed-consent builder covers what a real consent disclosure has to contain. The disclosure-trigger check covers when you are required to tell the client AI is in use at all.

This checklist protects the firm from an individual's error only if it is actually run before the step goes live. A checklist in a binder does not protect you. A checklist that runs does.

1. The Five-Duty Pre-Flight

Run this checklist before any Track B step goes live. Track B is any workflow where an AI model reads content and produces output a person relies on. Two questions price what you owe once you're there: does client information go into the AI, and do you rely on the output in a client's representation? Yes to both puts that step under the full five-duty discipline.

Competence

ABA Opinion 512 requires a reasonable and current understanding of the specific capabilities and limitations of the tool you are using. The duty is ongoing, because these tools change.

- ☐ **Understand the tool.** Can someone at the firm explain, in plain language, what this AI tool does, what kind of output it produces, and where it is known to fail or hallucinate? If no one can answer that, the tool is not ready to deploy.
- ☐ **Know the limits.** Can someone at the firm describe at least one class of error this tool is prone to, and explain how the verification step catches it? Deploying a tool without knowing its failure modes is not a neutral decision.
- ☐ **Confirm the knowledge stays current.** Is someone assigned to track changes to this tool, including model updates, terms changes, and guidance from the bar about tools of this type? The competence duty does not end at deployment.

Confidentiality and Consent

When client information goes into an AI tool, the confidentiality duty applies. When the tool's design creates a risk of disclosure, informed consent is required before the client's information enters the tool. Boilerplate in an engagement letter does not satisfy this requirement.

☐ **Check the tool's design.** Is this tool designed so that client data it receives can be disclosed or used to train the model? If yes, informed consent is required before client information goes in. Move to Part 2 of this checklist to build that consent.

☐ **Confirm the platform tier.** Is this tool running on a business-tier or dedicated-tier platform with a data processing agreement, a confidentiality setting active, and contract language that keeps client data out of training and out of third-party hands - stated outright or by barring any use beyond delivering the service? Consumer-tier tools are not acceptable for client information.

☐ **Check the vendor breach-notice commitment.** On a negotiated contract, does the vendor commit to notify the firm of a data breach within 48 to 72 hours? On a standard click-through platform, does the DPA commit to notice without undue delay with phased updates - and has the firm registered its breach contact with the vendor so the notice has somewhere to land? A vendor with no notice commitment at all does not touch client data.

Supervision

Opinion 512 treats AI tools as non-lawyers for supervision purposes. Your supervisory duties extend to the tool and to the vendor, the same way they extend to a paralegal or outside contractor working on your matter. You cannot hand off responsibility to the AI.

☐ **Assign ownership of the output.** Is a named attorney responsible for reviewing every output this tool produces before it is relied on, filed, or sent to a client? "The tool generated it" is not a defense. The attorney who signs off owns the output.

☐ **Vet the vendor.** Has the firm reviewed this vendor's terms, data handling practices, and sub-processor chain? Vendor diligence is not a one-time purchase decision. It is the front end of an ongoing supervisory duty, and it must be refreshed when the vendor updates its terms or its model.

☐ **Set a review cadence.** Is there a scheduled date to revisit this tool's performance, the vendor's current terms, and any new bar guidance on tools of this type? A tool deployed without a review cadence is a tool supervised only at launch.

Fees

Opinion 512 is clear: you cannot charge a client for time that AI eliminated. If the tool collapses a task from hours to minutes, you bill the minutes. The efficiency belongs to the client, not the invoice. The safe default in every jurisdiction is not to bill for time you did not spend.

- ☐ **Identify the time impact.** Does this AI step reduce the time a lawyer or staff member spends on a task that is currently billed to the client? If yes, your billing practice for that task needs to reflect the actual time spent after the AI step, not the historical average.

- ☐ **Check pass-through disclosures.** If the firm bills AI tool costs as a disbursement or per-use cost, is the charge reasonable and disclosed to the client? A pass-through is permissible. A hidden markup is not.

- ☐ **Check your jurisdiction's rule on non-hourly fees.** Does your state bar treat AI-driven efficiency gains differently than the ABA baseline? Virginia's opinion takes that position, and it is scoped to non-hourly fees. On hourly work, the rule above holds everywhere: you bill actual time. Confirm the rule that governs you before setting your billing practice for this tool.

Candor

Opinion 512 ties the candor duty to a specific instruction: verify AI output before you rely on it. Submitting AI-generated material you have not checked against its actual source can violate your duty of candor to the court, regardless of intent. This duty applies every time, with no exceptions.

- ☐ **Wire in a human verification step.** Before any AI-generated content is filed with a court, cited in a brief, or relied on in a client-facing document, does a person read the actual source and confirm the AI's output matches it? The verification step must be wired into the workflow, not left to the individual's memory.

- ☐ **Document the check.** Is there a record that verification happened, including who performed it and when? A verification step that leaves no record protects the firm less than one that does.

- ☐ **Write the policy before you need it.** Does the firm have a written AI policy and verification protocol that covers this tool? The 2025 federal case that spared a firm from sanction did so because the firm had real protocols in place. The policy does not have to prevent every error. It has to exist, be in force, and give the firm a defensible line between firm-level diligence and an individual's failure.

2. The Five Elements of Informed Consent

When the tool's design requires informed consent, a general clause in your engagement letter is not enough. ABA Opinion 512 is specific about what real consent has to contain. Use this section to build a purpose-matched disclosure for the tool you are deploying. The disclosure has to be specific to the actual tool and the actual data, not reused from a prior engagement.

- ☐ **Your best judgment about why you are using the tool.** Tell the client, in plain language, what this specific AI tool does on their matter and why you are using it. The disclosure has to name the function, not just gesture at technology in general.
- ☐ **The specific risk, including the kinds of client information exposed.** Tell the client which categories of their information this tool will process, and that it runs on a vendor's system rather than solely inside your office. A general reference to "data" is not the specific risk Opinion 512 requires.
- ☐ **How others can use that information against the client.** Tell the client what happens if the vendor's protections fail, and what safeguards you have put in place in response. The client is entitled to understand the risk with enough specificity to weigh it.
- ☐ **The benefit to the representation.** Tell the client what this tool does for their matter specifically, whether that is speed, accuracy, lower cost, or something else. The benefit has to be real and connected to their representation, not a general statement about AI capabilities.
- ☐ **A genuine purpose-built disclosure, not boilerplate.** Confirm that this consent is a document or conversation matched to this tool and this client's data, not a standard clause imported from an engagement letter template. Opinion 512 says directly that general boilerplate is not sufficient. If the disclosure reads the same for every client and every tool, it is not real consent.

3. The Five Disclosure Triggers

Consent answers whether client data can go into a tool. Disclosure answers a broader question: when do you have to tell the client you are using AI at all? ABA Opinion 512 identifies five triggers. Run this check on every Track B process before it goes live. Most of the time, more than one trigger will be in play.

> **Note.** Disclosure made on your own initiative is a sign of a firm in control. Disclosure pulled out of you by a client question is the opposite. Wire the check into the process before the step goes live.

- ☐ **The client asks.** If a client asks whether AI was used on their matter, you disclose. The trigger is the question, not your judgment about whether the use was significant.
- ☐ **Your engagement letter or outside-counsel guidelines require it.** If your standard engagement letter or the client's outside-counsel guidelines call for disclosure of AI use, the obligation is locked in at the start of the matter, not case-by-case.
- ☐ **Consent is required because data goes into the tool.** If the tool's design requires informed consent before client information enters it, disclosure is part of obtaining that consent. You cannot get real consent without disclosing what you are asking the client to consent to.
- ☐ **AI affects the basis of the fee.** If the AI step changes the time spent on a task, and that task's time is the basis of the fee, the client needs to know. Billing the pre-AI rate for a task the AI collapsed is a fee problem and a disclosure problem at the same time.
- ☐ **Output will influence a significant decision in the representation.** If the AI tool's output will directly inform a decision that materially affects the client's matter, whether a filing strategy, a settlement recommendation, or a legal position, the client has an interest in knowing AI played a role in reaching it.

How to use this. This is a thinking tool, built to help you reason about your own firm. It is not a compliance guarantee and it is not legal advice. Where it states a rule, it states the rule as of July 18, 2026, and the law moves, so confirm the current rules in your jurisdiction with your own counsel before you rely on anything here.

From *The Automated Law Firm* • 4Spot Consulting • **4SpotConsulting.com/512-checklist**