

# The Track A / B Classifier

*Run any firm process through the track test and two duty questions and know exactly which track you are on and what that means*

Every automated process in your firm is on one of two tracks. **TRACK A** is pure automation: structured data in, structured data out, no AI step reading content and making a judgment call. **TRACK B** is automation with an AI step inside it. The moment an AI system reads content and produces an output someone relies on, you are on Track B.

This tool makes that distinction fast and concrete. Run any process through the track test, then the two duty questions - does client information go into the AI, and do you rely on the output in a client's representation - follow the decision flow, and land on a track with the obligations that come with it spelled out. Some processes sit on both tracks depending on how you build them. This tool shows you that too.

Track A and Track B are the author's operating labels for organizing the work. They are not legal terms of art, and no court or bar uses them. Before you act on any obligation this tool identifies, confirm the current rules in your jurisdiction with your own counsel.

## 1. The Track Test, Then Two Duty Questions

---

Ask this about any process in your firm: does an AI system read content and produce an output you act on?

Not "is it sophisticated." Not "does it cost a lot." Not "does the vendor call it AI in the marketing." The only thing that matters is whether a model reads the content and someone relies on what it returns. That is the track test. No means Track A. Yes means Track B. Then two more questions price what you owe once you are on Track B, because the duties come in two clusters and the clusters have different triggers.

### **Track A - no AI step, no AI attachment**

Structured data goes in, structured data comes out, and no AI sits in the middle making a judgment call. A form submission that creates a matter record. A closed file that triggers a retention timer. A signed engagement letter that kicks off an intake checklist. No model reads the substance of the matter and decides anything.

**What attaches on Track A:** ABA Formal Opinion 512's AI-specific duties do not apply, because the thing they govern is not happening. Two sets of duties do apply: vendor due diligence (vet every third-

party tool that touches client data - check security, confirm confidentiality terms, understand where data lives and what happens to it) and your data-protection obligations under applicable law (state privacy statutes, breach-notification rules, and the baseline duty of confidentiality). The AI load is off. The privacy load is not.

### **Track B - an AI step, and the two questions that price it**

The moment an AI system reads content and produces an output someone relies on, you are on Track B. Two questions tell you what attaches from there.

**Question one: does client information go into the AI?** Yes puts the confidentiality cluster live: the consent analysis, the vendor terms, the tier of tool you run it on.

**Question two: do you rely on the output in a client's representation?** Yes puts the reliance cluster live: competence in the tool, supervision of what it produces, candor about what it returned, and the fee discipline.

Yes to both and you carry the full five-duty discipline. Yes to one and you carry that cluster, run with the same seriousness. No to both - an AI step coding receipts, routing IT tickets, work with no client data going in and no client representation riding on the output - and the client-facing duties have nothing to grab. What remains is what always remains: vendor diligence and data protection.

There is no "light" version of the rules for "just a little AI." One deposition summary through a chatbot answers yes to both questions, so that one summary carries the same duty load as running your whole practice through a model. The load does not scale with volume. It scales with the answers to the two questions. Track B is a deliberate tier-up you choose on purpose, not a default you slide into - though sliding into it is exactly the risk firms face when tools they already use quietly add an AI feature.

## **2. The Decision Flow**

---

### **1 Identify every step in the process.**

List what happens from trigger to output, including every tool the data touches along the way. A clean Track A workflow can feed a vendor that runs a model on your data the moment it arrives. Audit where the data lands, not just what your own workflow does.

### **2 Apply the track test to each step.**

Does an AI system read content and produce an output someone relies on? Run the test step by step, not just on the workflow as a whole. No means Track A for that step. Yes means Track B, and you move to the two duty questions.

**3 If any step lands on Track B, ask the two duty questions of that step.**

Does client information go into the AI? Do you rely on the output in a client's representation? Yes to both means the full five-duty discipline. Yes to one means that single cluster, run with the same seriousness. No to both means only vendor diligence and data protection remain.

**4 If any step answers yes to the track test, the whole process is Track B.**

There is no partial Track B for purposes of classification. One step tips the entire workflow, though the duty questions still price each step on its own facts.

**5 If every step answers no to the track test, the process is Track A.**

Confirm that no vendor on the receiving end is running a model on your data without your knowledge. If one is, the process is Track B even though your own workflow is not.

**6 When you genuinely cannot tell, assume Track B.**

Then check whether you can build your way back to Track A by removing or replacing the ambiguous step. Do not assume Track A and find out later.

**GRADING THE RELIANCE QUESTION.** If you want a sharper rule than a feel, grade reliance in four tiers. Tier zero: the output is discarded, or used only to find things you then read yourself - no reliance, no load. Tier one: the output routes, labels, extracts, or orders work - procedural reliance; supervise its accuracy and audit its misses. Tier two: the output shapes legal analysis, drafting, or advice - substantive reliance; the full discipline applies. Tier three: the output reaches a client, a court, or a bill - external reliance; a human verifies it against the source before it leaves the building, every time. A workflow's tier is set by the highest tier any of its outputs reach.

**NOTE ON THE PLATFORM FLOOR.** The track tells you whether the AI rules attach. It says nothing about whether the platform is built to hold client data. That is a separate question with the same answer on both tracks: a business-tier platform is the floor for any firm data. Consumer-tier tools are never appropriate for client work, even on the simplest Track A workflow.

### 3. Two Sharpeners

---

#### **Generative vs. non-generative AI**

ABA Formal Opinion 512 speaks specifically to generative AI - the models that write, draft, summarize, and answer. When your AI step is clearly generative (a tool drafting a contract clause, summarizing a deposition, answering a client question), Opinion 512 attaches as a matter of the opinion's own terms, and the two duty questions run under that authority.

Some tools read client content without generating prose: a predictive-coding engine ranking discovery documents, a classifier routing your incoming mail, the smart-sorting that reorders a review queue. Those tools still put you on Track B under the track test, because they read the content and shape what someone relies on. Whether Opinion 512's specific duties reach a non-generative tool is a question the opinion does not answer, and litigation teams have run predictive coding for a decade under court-approved protocols without anything resembling a 512 consent analysis. So hold the two tracks of reasoning separately: for generative steps, the duties the two questions surface are the opinion's command. For non-generative steps, running the same five-duty discipline is the author's operating standard, not a rule the profession has imposed. Run it until your counsel tells you less is required, and do not let anyone tell you a court-supervised review protocol was secretly non-compliant all these years. It was not.

**These are operating labels, not legal terms**

Track A and Track B are the author's framework for organizing the work. They are not legal terms of art. No court, no bar association, and no ethics opinion uses these labels. The obligations they point to are real; the labels themselves are a working shorthand. Treat the operating standards in this tool as strong recommendations you can tune with your counsel, and treat the law as the law.

4. Process Classification Table

---

The rows below cover the most common firm processes, drawn from the process playbook and the automation inventory behind this book. Each row shows which track the process lands on and what attaches once you run the two duty questions. Some processes land on both tracks depending on how you build them - the note below the table explains that.

| PROCESS                             | BUILD DESCRIPTION                                                                                                                                                                                 | TRACK<br>A | TRACK<br>B | WHAT ATTACHES                                                                                                                                                                                                                                                                                                                       |
|-------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Intake form to matter record</b> | Submitted form fields populate the matter record, client contact, and folder structure. No model reads the submission.                                                                            | ✓          |            | <b>TRACK A</b> Vendor diligence and data-protection duties. No AI-specific obligations.                                                                                                                                                                                                                                             |
| <b>Inquiry routing</b>              | A web inquiry sent to the right practice group. Track A if fixed-field form maps to a group. Track B if an AI step reads the free-text inquiry and decides the group.                             | ✓          | ✓          | <b>TRACK A</b> Vendor and data duties only. <b>TRACK B</b> Prospect information goes into the model and the firm relies on the routing decision - both duty questions answer yes, so the full five-duty discipline applies.                                                                                                         |
| <b>Engagement letter generation</b> | Client and matter fields merge into the firm's letter template and send for signature. The template is the logic; no model writes anything.                                                       | ✓          |            | <b>TRACK A</b> Vendor diligence and data-protection duties.                                                                                                                                                                                                                                                                         |
| <b>Matter folder and task build</b> | Opening a matter creates the standard folder structure, sets team access, and creates the opening task set. Pure rule execution; no content read.                                                 | ✓          |            | <b>TRACK A</b> Vendor diligence and data-protection duties.                                                                                                                                                                                                                                                                         |
| <b>Conflicts name search</b>        | A new party name runs against the conflicts database and returns every match. Track A baseline. Track B if an AI step reads the match list and separates likely real conflicts from coincidences. | ✓          | ✓          | <b>TRACK A</b> Vendor and data duties only. <b>TRACK B</b> Party and matter information goes into the model and the firm relies on which matches it clears - both questions answer yes, so the full five-duty discipline applies. Supervision and competence are especially relevant when a model is deciding which matches matter. |

| PROCESS                                  | BUILD DESCRIPTION                                                                                                                                                            | TRACK<br>A | TRACK<br>B | WHAT ATTACHES                                                                                                                                                                                                                                                                                                          |
|------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Docket-to-calendar sync</b>           | A structured e-filing notice with coded dates creates docket entries. Track A. Track B if an AI step reads a narrative scheduling order and pulls out each deadline.         | ✓          | ✓          | <span>TRACK A</span> Vendor and data duties only. <span>TRACK B</span> Matter content goes into the model and the firm relies on the extracted dates - both questions answer yes, so the full five-duty discipline applies. A person confirms every date before it posts.                                              |
| <b>Document auto-filing</b>              | Saved documents are routed to the right matter folder by naming convention. Track A. Track B if an AI step reads the document, determines its type and matter, and files it. | ✓          | ✓          | <span>TRACK A</span> Vendor and data duties only. <span>TRACK B</span> Client document content goes into the model and the firm relies on where it gets filed - both questions answer yes, so the full five-duty discipline applies. Risk level rises with document sensitivity.                                       |
| <b>Document assembly from a template</b> | Structured matter data flows into a fixed template and produces the document. No model writes anything. The template is the logic.                                           | ✓          |            | <span>TRACK A</span> Vendor diligence and data-protection duties. A human reviews every merged output before it is used.                                                                                                                                                                                               |
| <b>First draft from a description</b>    | An AI step drafts a first version of a routine document from a plain-language description and matter data. A person reviews before use.                                      |            | ✓          | <span>TRACK B</span> Matter data goes into the model and the lawyer relies on the draft it returns - both questions answer yes, so Opinion 512 attaches and the full five-duty discipline applies. Nothing the model drafts leaves the firm without a lawyer reading every word. Dedicated tier for sensitive matters. |

| PROCESS                                     | BUILD DESCRIPTION                                                                                                                                                                                    | TRACK<br>A | TRACK<br>B | WHAT ATTACHES                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Legal research routing and filing</b>    | Research requests routed, results organized and filed to the matter. No model generates law. Administrative work around the lawyer's own research.                                                   | ✓          |            | <b>TRACK A</b> Vendor diligence and data-protection duties. One of the lowest-risk Track A builds.                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>AI-assisted legal research</b>           | An AI step reads a long opinion and produces a headnote or summary the attorney relies on, or a research tool answers a plain-language question and returns a cited answer.                          |            | ✓          | <b>TRACK B</b> A plain-language research query carries no client name or client information, so the confidentiality cluster stays quiet. But the attorney relies on the cited answer in a client's representation, so the reliance cluster is live: competence, supervision, candor, and fees. Highest-risk Track B application in the book on the reliance side. A non-negotiable human verification step before any model-generated authority reaches a brief or filing. Dedicated tier. |
| <b>Contract routing and date extraction</b> | An incoming contract is logged and routed to the right reviewer by type. Track A. Track B if an AI step reads the contract and pulls parties, term, renewal dates, and governing law into a summary. | ✓          | ✓          | <b>TRACK A</b> Vendor and data duties only. <b>TRACK B</b> Client contract content goes into the model and the firm relies on the extracted summary - both questions answer yes, so the full five-duty discipline applies. A person confirms every extracted field.                                                                                                                                                                                                                        |

| PROCESS                                              | BUILD DESCRIPTION                                                                                                                                                                                                                                                                 | TRACK<br>A | TRACK<br>B | WHAT ATTACHES                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Draft invoice and time capture</b>                | Draft invoices built from posted time entries. Activity captured into draft time entries from calendar, email, and document systems. No model reads content to decide whether to bill or how to characterize the work - this is structured data moving through rules-based logic. | ✓          |            | <b>TRACK A</b> Vendor diligence and data-protection duties. Risk is medium to high because a billing error is an ethics matter, not just an accounting one.                                                                                                                                                                                                                                                                                                                                                                       |
| <b>AI-assisted time narratives / pre-bill review</b> | An AI step reads recorded activity and drafts a compliant time narrative. A pre-bill review variant has a model rewrite time entries before the invoice goes to the client. The model reads billable activity content and produces an output the billing attorney relies on.      |            | ✓          | <b>TRACK B</b> Client matter and billing detail goes into the model, and the billing attorney relies on the narrative it drafts - both questions answer yes, so the full five-duty discipline applies. Fee-duty scrutiny is highest here: the model reads billable activity and the fee duty limits what you can charge for AI-eliminated time. A billing attorney reviews every AI-drafted narrative before the invoice is sent. Business tier at minimum; dedicated tier for any matter where confidentiality risk is elevated. |



| PROCESS                             | BUILD DESCRIPTION                                                                                                 | TRACK<br>A | TRACK<br>B | WHAT ATTACHES                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------------------------------|-------------------------------------------------------------------------------------------------------------------|------------|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Client-facing<br/>AI chatbot</b> | A model answers client questions in a channel where the client receives the output as guidance from their lawyer. |            | ✓          | <b>TRACK B</b> Client information goes into the model through the client's own questions, and the client relies on the answer as guidance from the firm - both questions answer yes, so Opinion 512 attaches and the full five-duty discipline applies. Confidentiality and competence duties are hardest here. A model should not answer clients without a human in the loop on anything that resembles legal guidance. Dedicated tier. |

**NOTE ON BOTH-TRACKS ROWS.** Several rows above show a check under both Track A and Track B. This happens when the same process can be built two ways: a rules-based version that moves structured data with no AI step (Track A), and an enhanced version where an AI step reads content and produces an output someone relies on (Track B). Inquiry routing is a clear example: a fixed-field form that maps a selection to a practice group is Track A; a model that reads a free-text message and decides the group is Track B. The two versions are not interchangeable. The track is determined by what the actual build does, not by the process name or the vendor's description of the tool.

**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/CLASSIFIER](https://4SPOTCONSULTING.COM/CLASSIFIER)