The Impact vs. Ethics-Risk Worksheet

Score a proposed AI step on two axes and make a decision you can defend

Every **Track B** step - an AI step inside an automated workflow - adds real duty load the moment it runs. It triggers the full five-duty discipline: competence, confidentiality, supervision, fees, and candor. That load is not a reason to avoid **Track B** forever. It is a reason to choose it on purpose.

This worksheet puts two questions to any proposed AI step: how much does it return, and how much does it cost in duty load? Map the answers onto the matrix and the decision becomes defensible, not a gut call.

Work through the scoring checklists first, then plot the step, then record your decision in the worksheet table at the bottom. One row per candidate step.

1. Why this decision matters

Track B is a door, not a dial. The moment an AI step reads client information and produces an output you rely on, all five duties are in the room. You do not get to open the door halfway.

That weight is real, and it is doing real work for you. It means the firms that take TRACK B seriously build supervision, verification, and the right platform tier in before they ship - and the firms that slide in without noticing find out the hard way. This worksheet is the mechanism that keeps the decision deliberate. Use it every time a proposed AI step is on the table.

Two things determine the right call: how much the step returns, and how much duty load it pulls in. Plot both and the quadrant tells you what to do next.

2. The decision matrix

Score impact and ethics-risk using the checklists in Section 3, then find the cell that matches your scores.

	Low ethics-risk	High ethics-risk
High impact	**Build it.** The return justifies the step. Confirm your platform clears the business-tier floor and build the five-duty checks in from the start.	**Build it carefully.** The return is real, but the duty load is heavy. Business tier properly configured is the floor; climb to the dedicated tier when the matter carries health data, high litigation sensitivity, or a protective order. A named supervising attorney owns the output before anything ships.
Low impact	**Build it only if it is cheap.** The duty load is manageable, but the return is thin. If the build cost is low and the ongoing supervision overhead is low, proceed. Otherwise leave it on TRACK A.	**Leave it on Track A.** The duty load is high and the return does not justify it. Keep this process on pure automation and revisit if the impact case changes.

> **Note.** The only tiers in this framework are the business tier and the dedicated tier. Business tier is the floor for all law firm data on both tracks. The dedicated tier is required when a matter carries health data, high litigation sensitivity, or a protective order that demands tighter controls. Consumer-tier tools are never appropriate for client data on either track.

3. Scoring checklists

Answer each checklist before you plot the step. Each "yes" on the impact list pushes the score toward High. Each "yes" on the ethics-risk list pushes the score toward High. If you are split, round toward the more conservative score - it is the defensible one.

How to score impact

- ☐ **Hours returned.** Does this step eliminate a material block of time that a person currently spends on non-billable or low-value work - not seconds, but hours per week across the firm?
- ☐ **Billable capacity freed.** Does the time saved come off a lawyer's plate in a way that opens capacity for billable work, rather than simply cutting administrative overhead for support staff?
- ☐ **Client benefit, direct.** Does the step produce a faster, more accurate, or more consistent result that the client experiences directly - a quicker turnaround, fewer errors in a document, a faster answer to a question?

- ☐ **Error or deadline risk reduced.** Does the AI step catch something a manual process currently misses - a renewal window, a conflicting clause, an inconsistency in extracted data - at a scale a person cannot match?
- ☐ **Scalability.** Does the step let the firm handle more volume without adding headcount, in a way that a purely rules-based automation cannot?

Three or more "yes" answers: **High** impact. Fewer than three: **Low** impact.

How to score ethics-risk

- ☐ **Data sensitivity.** Does the step touch protected health information, financial account data, or any category of information your state's privacy statutes treat as sensitive?
- ☐ **Competence load.** Is there no one at the firm who can explain the tool's behavior in this use case, including where it fails and how it can mislead?
- ☐ **Confidentiality exposure.** Does the step send client information to a vendor whose data-handling terms you have not reviewed this calendar year - or to any tool that has not cleared the four-part test: no training on your data, no storage beyond what is needed, no third-party disclosure, and deletion permitted - each stated outright or carried in substance by the vendor's data processing agreement?
- ☐ **Supervision gap.** Is there no named person at the firm whose sign-off is required before the AI output is relied on - not a general policy, but an identified individual whose approval is mandatory every time?
- ☐ **Candor trigger.** Does the output go to a court, get filed in a matter record, or get sent to a client as a statement of fact or law? Any of those is a candor trigger that requires source verification before the output leaves the firm.
- ☐ **Fees exposure.** Does the step collapse a block of time that is currently billed to clients by the hour? If the AI step cuts a two-hour task to ten minutes, the billing rate must reflect the actual time spent.

Three or more "yes" answers: **High** ethics-risk. Fewer than three: **Low** ethics-risk.

Note on the framing rule. When an AI step reads client content and generates output a person relies on, the full five-duty discipline applies. For clearly generative steps - a model drafting a letter, summarizing a deposition, or answering a client question - ABA Formal Opinion 512 attaches directly. When in doubt, treat the step as generative and score accordingly.

4. Scoring worksheet

Complete one row for each candidate AI step. Use the checklists above to assign H or L for impact and ethics-risk, then identify the quadrant from the matrix and record the decision.

Candidate AI step	Impact (H / L)	Ethics-risk (H / L)	Quadrant	Decision

5. After you plot the step

1. **Confirm the platform tier.**
   Every Track B step runs on business tier at minimum. Climb to the dedicated tier if the matter carries health data, high litigation sensitivity, or a protective order. Document the tier decision and the reason for it.

2. **Name the supervisor.**
   Identify by name the person whose sign-off is required before the AI output is relied on. A supervision policy with no named person does not satisfy the duty.

3. **Wire the verification step.**
   For any output headed to a court, a client, or a matter record as a statement of fact or law, the verification step - a person checking the output against its actual source - is non-negotiable. Build it into the process before the first run, not after the first error.

4. **Review the consent question.**
   If the AI tool is designed to learn from the data fed into it, or if its terms allow your client's information to be disclosed directly or indirectly, informed consent from the client is required before the step runs. Boilerplate in an engagement letter does not satisfy this requirement.

5 **Set a review date.**

The duty to maintain competence is ongoing. Schedule a review of this step when the tool's model or terms change, and at least annually, so the decision you made today stays defensible.

> **One rule that does not change.** The matrix shows four cells. The platform floor is the same in all four: business tier or above, for all law firm data, on both tracks. There is no quadrant where consumer-tier tools are appropriate.

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/impact-worksheet**