The 10-Process Classification Map

Classify every core firm process by track, risk, and minimum tier before you build anything

Every automated process in your firm runs on one of two tracks. **Track A** is pure automation: structured data in, structured data out, no AI reading content or making judgment calls. **Track B** is automation with an AI step inside it - one step is enough to change your obligations.

This map lays all ten core processes on a single table. For each one you see the Track A risk, the Track B risk, the minimum platform tier, and the state flags that change the answer depending on where you practice. Use it to classify before you build, so the choice of track is always yours.

1. How to read this map

Three columns govern every build decision. The track tells you whether the two duty questions are even in play. The risk level tells you where to concentrate your attention. The minimum tier tells you what grade of platform to buy.

The two tracks

Track A is pure automation. Structured data goes in, structured data comes out. No AI step reads content and produces an output you rely on. On Track A the primary obligations are vendor diligence and data-protection duties.

Track B is automation with an AI step. The moment an AI system reads client content and produces an output you rely on, you are on Track B - one step is enough. Two questions then price what you owe. Does client information go into the AI? If yes, the confidentiality cluster is live: consent analysis, vendor terms, tier. Do you rely on the output in a client's representation? If yes, the reliance cluster is live: competence, supervision, candor, and fees. Yes to both and you carry the full five-duty discipline. No to both - an AI step with no client data going in and no client representation riding on the output - and only vendor diligence and data protection apply. Track B is a choice you make on purpose, not a default.

The three risk levels

Level	What it means in practice
High	A failure here reaches clients, courts, or ethics regulators. Disqualification, sanction, or breach are the realistic consequences. Concentrate your supervision and verification here first.
Medium	A failure propagates downstream or creates fee, deadline, or data-handling exposure. Real consequences, but bounded. Business-tier controls handle it when properly configured.
Low	A failure is operational - a missed step, a routing error - with no direct ethics or client-harm consequence. Still worth getting right; not where the danger concentrates.

The two tiers

Business tier is the floor for every process in every firm on either track. Consumer-grade tools are never appropriate for client data. Business tier means: a signed data-processing agreement, your data kept out of training - stated outright or by a contractual bar on any use beyond delivering the service - no disclosure to third parties, and encrypted transit and storage.

Dedicated tier is the required step-up for Track B processes that involve a model generating substance against sensitive client data. Dedicated tier adds zero data retention, a business associate agreement when health data is present, and the protective-order terms the book details in Chapter 9. The matter drives the tier - start at business, climb to dedicated when the matter gives you a reason.

2. The ten-process classification map

Process	Track A risk	Track B risk	Minimum tier	State flag / note
Client intake	**Medium**	**High**	Business floor; higher for Track B	Florida takes a risk-based, reasonable-efforts view of AI confidentiality and layers its lawyer-advertising rules on top for intake; consent is recommended, not automatic, when confidential information reaches the tool. Open-model consent rules apply in several states when intake uses an AI routing step.
Conflicts screening	**High**	**High**	Business floor; dedicated for Track B	Risk is stakes-based, not exposure-based. The Track A check is simple data matching. The danger is the underlying knowledge gap: conflicts screening is only as good as the index it runs against, and most firms lack a single complete index.
Matter management	**Medium**	**High**	Business both tracks	Track A risk is operational: a misconfigured trigger or folder map means a matter opens without its structure. There is no reason an AI step belongs here on the administrative layer - adding one buys nothing but the two-question test, and on this administrative layer both questions usually answer no.
Billing and time	**Medium to High**	**Medium to High**	Business both tracks	Billing errors are fee errors, and fee errors are ethics matters. Most jurisdictions hold that AI cannot be billed for time it eliminated. Virginia takes the opposite position on efficiency gains, and only for non-hourly fees - flat fees, contingency. On hourly work every camp agrees you bill actual time. Your billing assumptions have to match your jurisdiction's rule.

Process	Track A risk	Track B risk	Minimum tier	State flag / note
Document assembly	**Medium**	**High**	Business floor; dedicated for Track B	States that treat substantive AI delegation as outsourcing require advance informed consent before Track B assembly goes live. North Carolina follows this route. Nothing a model drafts leaves the building without a lawyer reading every word.
Legal research	**Low**	**High**	Business floor; dedicated for Track B	The textbook Track B danger. On Track A, routing and organizing research is among the lowest-risk tasks on the table. On Track B, AI-generated legal research is the single highest-risk process in the firm - fabricated citations in perfect format, indistinguishable from real law until you pull them. Non-negotiable human verification on every output. Several states require consent for open-model research tools.
E-filing	**Medium**	**Medium**	Business both tracks	The risk that matters most is a missed or miscalculated deadline - the automation reduces that risk rather than adding to it by computing dates the same way every time. A human confirms before anything submits. There is rarely a reason for an AI step here at all.
Records retention	**Medium**	**Medium**	Business both tracks	Both directions carry liability: destroy too early and you have spoliation; keep too long and you have a larger breach surface. Every disposition needs a logged record. Retention automation depends on clean document classification upstream - the two processes feed each other.

Process	Track A risk	Track B risk	Minimum tier	State flag / note
Client communications	**Low to Medium**	**High**	Business floor; dedicated for Track B	Track A status updates and appointment reminders carry low to medium risk. An AI chatbot answering client questions unsupervised is the highest-risk Track B communications build - client information is going in and the output drives what you tell a client, so both duty questions answer yes and the full five-duty discipline attaches, confidentiality and competence hardest. The same state that layers advertising rules on intake does so for communications. A model drafts; a lawyer approves. Never the reverse.
Contract review	**Medium**	**High**	Business floor; dedicated for Track B	Track A extracts structured facts and calendars key dates - data handling, not legal judgment. Track B has a model reading client contract substance and flagging risk - client data goes in and a human relies on the flag, so both questions answer yes and the full five-duty discipline applies; supervision and competence are in front. Use it to surface what a human then confirms. States treating AI delegation as outsourcing require the consent gate before Track B review goes live.

3. The rows that trip firms up

> **Legal research on Track B.** No process illustrates the track gap more sharply. On Track A, routing and organizing research is low-risk administrative work. On Track B, an AI generating case citations is the single highest-risk thing a firm can do - fabricated citations have arrived in court in perfect format and cost lawyers their careers. Human verification is not optional on this row. Every AI-generated citation gets pulled against the primary source before it leaves the building.

Conflicts screening at high risk on both tracks. This row earns its rating differently from the rest. The Track A check - running a name against a list - is straightforward data matching. The risk is not the automation. The risk is that the list is incomplete. Most firms do not have a single place where every client, every adverse party, and every matter connection lives. The automation searches what you point it at. If the underlying index is fragmented across practice management, closed matters, and one partner's memory, the check is only as good as the worst gap in that index. Fix the index first, then automate the check.

Billing diverges by state. The standard rule is that AI cannot be billed for time it eliminated - if the tool did in two minutes what used to take two hours, you bill the two minutes. Virginia has taken the opposite position, holding that efficiency gains do not automatically reduce the client's fee - but only on non-hourly arrangements, flat fees and contingency. On hourly work every camp agrees you bill actual time. The billing row requires you to know your jurisdiction's specific rule before the automation goes live. Do not assume the national default applies where you practice.

The floor never moves. Consumer tier is never appropriate for any row in this table, on either track. Business tier is the universal floor. The only direction the tier moves is up - to dedicated - when a Track B process involves a model generating substance against sensitive client data, or when the matter carries health information. The matter sets the tier; the tier never drops below business for client data.

4. Using this map in practice

1. **Classify each process.**
   Walk every automation your firm runs and place it in a row. Decide whether an AI step is inside it. That answer determines the track. On Track B, run the two duty questions - does client information go into the AI, and do you rely on the output in a client's representation - to price what attaches.

2. **Check the minimum tier.**
   Confirm the platform running each process meets at least business tier. If a Track B process involves sensitive data, check whether dedicated tier is required for that matter.

3. **Run the state flags.**
   Before any process goes live, check whether your state's specific rules add consent requirements, disclosure obligations, or outsourcing treatment on top of the national baseline. The flags in this map name the categories. Chapter 10 of the book maps them state by state.

4 **Mark the danger rows.**

Legal research, conflicts screening, client communications via AI chatbot, and any Track B document or contract process deserve concentrated attention. Build the human verification and supervision steps before the AI step goes live, not after.

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/classification-map**