

The Safest 10 to Automate

Ten Track A builds to start with - pure automation, no AI step, fast returns

These are the ten **TRACK A** automations with the lowest risk and the fastest return. Every one of them moves structured data in and structured data out. No AI step reads client content. No model makes a judgment call. The AI-specific obligations that come with **TRACK B** do not attach to any build on this list.

Almost every firm can run all ten on a business-tier platform. Start here. Capture the easy wins. Add AI later, deliberately, where it earns its place.

Each build below is pure **TRACK A**. Structured data goes in, structured data comes out. No AI step reads a client file or makes a judgment call. That means the full five-duty discipline does not apply to any of these. Two duties still bind you on Track A: vendor due diligence and your data-protection obligations. Those are covered in the note at the bottom. But the AI-specific load is off, and that is worth a lot.

Run these on a business-tier platform. Consumer-tier tools are never appropriate for client data, even on the simplest Track A workflow.

1

Matter open: folders, tasks, and the tracking sheet.

When a new matter is created in your practice management system, the automation builds the complete folder structure for that matter type, creates the standard opening task list, and writes the matter into the master tracking sheet. When anything about the matter changes, it finds the existing row and updates it. Every matter opens the same way, every time, with no one assembling folders by hand or forgetting a task. Why it is safe: no model reads anything. A matter is created, and a fixed set of consequences follows exactly as you defined them. To get started: map out your folder structure for each matter type - civil litigation, estate planning, probate, or whatever your practice runs - and document the six to ten standard opening tasks. That map is the input your automation needs.

2

Intake to matter: move the captured fields, no retyping.

New submissions from your intake platform are pulled on a schedule and staged. The client's data - rates, contact details, the facts the person entered on the form - moves straight into the matching matter in your practice management system. The number the client was quoted during intake is the number that lands on the matter. No one keys it in, and there is no chance of a transcription slip. Why it is safe: the workflow reads structured fields and writes them to structured fields. It does not read the substance of anyone's matter or decide anything about the representation. To get started: list every field your intake form captures and every field your matter system needs. The build is a direct field-to-field map between those two lists.

3

Deadline calendaring from the governing rules.

The system calculates filing and response deadlines from a trigger date using stored court rules, then pushes each date and a set of tiered reminders to the responsible attorneys' calendars. Reminders escalate if a task is not marked done before the next threshold. Why it is safe: this is date arithmetic run against a fixed rule set. No model interprets a narrative order or decides what the deadline is. The automation computes the date the same way every time. To get started: pick one court and one matter type. Translate its deadline rules into a structured table - trigger event, days forward or backward, holiday adjustment, who gets the reminder. That table drives the build. One non-negotiable: a lawyer owns and validates that rule set, at build time and every time the rules change. The automation computes dates; the lawyer answers for them. A wrong rule automated is a missed deadline at scale.

4

Conflicts index update.

Opening a matter adds the client, the adverse parties, and related entities to the conflicts index. Every new matter is searchable for conflicts from the moment it opens, without anyone keying names into a separate database by hand. Why it is safe: the trigger is a new matter record, and the automation writes structured fields to a structured index. No model reads any document content or decides whether a conflict exists - that judgment stays with the attorney. To get started: confirm which fields in your practice management system hold the client name, adverse parties, and related entities. Those fields drive the write to your conflicts database. Map them before you wire the trigger.

5

Docket-to-calendar sync.

A new docket entry pushes the date and a set of reminders to the responsible attorneys' calendars. No one has to manually transfer a court date from the docket system to the calendar, and the date is there the moment the entry posts. Why it is safe: the docket entry is a structured record - a date, a matter number, and an assigned attorney. The automation reads those structured fields and creates a calendar event from them. No model interprets a document or decides what the date means. To get started: confirm your docket system exposes new entries via API or webhook. Map the fields - hearing date, matter number, attorney assignment - to the calendar event fields your firm uses.

6

Tiered deadline reminders.

Reminders go out on a set schedule before each deadline and escalate if the task is not marked done. An attorney who has not acknowledged a filing deadline gets a follow-up before the next threshold, not silence. Why it is safe: the reminder schedule is a fixed rule set built on dates and task-completion flags. The automation does not decide what is important or interpret any document. It applies your rule, on your schedule, every time. To get started: define the reminder tiers for each matter type - how many days out, who gets the reminder, and what the escalation path looks like if the task sits unmarked. That rule set is what the automation runs.

7

Auto-filing by rule.

Saved documents are routed to the right matter folder by naming convention and document type. Every document lands where it belongs without anyone dragging files into folders by hand. Why it is safe: the routing logic operates on metadata and naming conventions - structured attributes the document already carries. No AI step reads the document content to decide where it goes. The rule set is yours; the automation applies it. To get started: define your naming convention and the folder map it corresponds to. Then set up a watch on your incoming documents folder so the automation fires the moment a new file arrives.

8

Draft invoice generation.

Draft invoices are built from posted time and cost entries on the billing cycle. The draft is ready for attorney review the same day the cycle closes, without anyone assembling it from scratch. Why it is safe: the inputs are posted time entries and cost records - structured fields. The automation assembles the draft from those records using your billing rules. No model reads narrative time entries or rewrites them. The attorney reviews the draft before anything goes to the client. To get started: confirm your time and billing system exposes posted entries via API or scheduled export. Map the fields your invoice template needs and build the assembly rule.

9

Appointment reminders.

Reminders and intake links go out before each client meeting and update when the time changes. Clients arrive prepared, and the preparation emails go out without anyone on your team manually sending each one. Why it is safe: the trigger is a calendar event - a structured, time-based data point. The reminder sequence delivers fixed messages and links. No AI step reads any client file or generates a personalized message. To get started: connect your scheduling system to your communication platform and define the reminder sequence by timing - how far out, what the message says, and what intake materials go with it.

10

New-hire provisioning.

A start date triggers account creation, email setup, and the equipment and access checklist. The new hire has what they need on day one without a coordinator tracking every step manually. Why it is safe: the trigger is a structured HR record field - the start date. Each provisioning step follows a fixed sequence against a predetermined checklist. No model decides what access to grant or reads any client data. To get started: document every account, credential, and access level a new hire at each role should receive. That checklist is the automation's instruction set. Wire the trigger to your HR system so the sequence starts the moment a hire record is confirmed.

WHAT STILL BINDS YOU ON TRACK A. These builds carry no AI-specific obligations. Two sets of duties still apply. First, vendor due diligence: the moment client data moves through a third-party platform, you owe a diligence check on that vendor. Confirm where the data lives, what the data processing agreement says, who the sub-processors are, and whether you can export and delete on your terms. On Make, turn on the "Data is Confidential" setting for every client workflow. Second, your data-protection obligations: state privacy statutes, breach-notification rules, and your baseline duty of confidentiality all apply to Track A workflows because those workflows hold and move client information. Track A lifts the AI load, not the privacy load. Business tier is the floor for all of this. Consumer-tier tools are never the right answer, even for the simplest build on this list.

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/SAFEST-10