The Intake Funnel Guide

Build the intake automation step by step, from the first form field to the finished matter record

Intake is where client data first enters the firm, and where the most reliable **Track A** wins live. A submitted form creates the matter record, the client contact, and the folder structure with no retyping. No AI reads anything. No judgment gets made. Clean data goes in and clean data comes out.

This Guide walks the build in six steps - Capture, Validate, Stage, Route, Create Matter, Confirm - with the data-handling checkpoint at each one marked so nothing leaks on the way in. It also shows exactly where an AI step would cross the build to **Track B** and what that choice costs.

Use this as your build checklist when you set up the intake automation, and as a reference when you add any new step so you always know which track you are on.

1. Why intake is the right place to start

Intake is where client data first enters the firm, and a mistake here propagates everywhere downstream. The exposure is in the data handling, not in any interpretation - because on a well-built Track A intake, there is no interpretation. A rate the human already set, a name the client typed, a practice area the form already coded: each one moves from where it was captured to where it is needed, unchanged.

That is also what makes intake the best place to prove the model works. The build is contained. The inputs are structured. The wins are immediate and measurable. Every firm that has run this has the same experience: the first time the automation creates a matter record while the client is still reading the confirmation email, the team stops asking why automation matters.

Run the six steps below in order. Each one has a checkpoint. Check them all before you go live.

2. The six-step build

1 **Capture.**

The intake form collects the client's data. This is the entry point for everything that follows, so the form design is not a UX decision - it is a data-integrity decision. Every field that the matter record, the conflict check, the billing system, and the folder structure will need downstream must be a discrete, structured field here. A free-text "tell us about your situation" box is useful for the client. It is not useful as a data source. Keep the substance field separate from the structured fields. Map every structured field now to the system field it will populate later. If you do not know where a field lands, do not collect it yet.

> **Data-handling checkpoint.**
> Confirm that the form platform transmits data over encrypted connections and stores submissions in an environment covered by a signed data-processing agreement. Consumer-grade form tools are not appropriate here. The moment a client submits, the data is covered by your confidentiality duty. The platform it lands on has to be too.

2 **Validate.**

Before any data moves, the automation checks that it is complete and correctly formatted. Required fields are present. Email addresses are formatted as email addresses. Phone numbers are formatted consistently. Date fields contain dates. Practice-area selectors match a value on your list. Any submission that fails a check stops here and triggers a follow-up to the client, not a broken record in your system.

> **Data-handling checkpoint.**
> Validation is where bad data gets caught rather than propagated. A missing field discovered at this step costs one follow-up email. The same missing field discovered three weeks into a matter costs a paralegal an hour and a data audit. Build the validation rules against every field your downstream systems require - not just the fields you think of first.

3 **Stage.**

Valid submissions move to a staging layer before they write to any system of record. Staging holds the data in a clean, structured form and lets you inspect it before it commits anywhere. In practice this is a row in an intermediary table, a queue in your automation platform, or a holding record in your intake tool. Whatever the mechanism, the principle is the same: nothing writes to your practice management system, your billing system, or your conflict index until a human or a rule has confirmed the record is right.

> **Data-handling checkpoint.**
> The staging layer is where you apply any field transformations your downstream systems need - reformatting a date, normalizing a name, mapping a form value to an internal code. Do all transformations here, in one documented place, before the data crosses into any system of record. Transformations buried inside individual system integrations are the source of the "I don't know where that value came from" problem firms deal with for years.

4 **Route.**

The staged record goes to the right practice group. How routing works determines whether this step stays on Track A or crosses to Track B.

Track A routing

uses a fixed category field the client or the form already set - a dropdown for matter type, a radio button for practice area, a service-selector that maps directly to a group. The form value maps to a destination. No content gets read. No judgment gets made.

Track B routing

adds an AI step that reads the free-text description field and decides the practice group from the narrative. One step is enough to change the track.

> **Data-handling checkpoint.**
> If you are on Track A, confirm that the mapping table between form values and practice groups is documented and version-controlled. When a new practice area opens or a team changes, the map has to update in one place. If you are on Track B, the full five-duty discipline applies from this step forward. See the track table in Section 3 for what that means in practice.

5 **Create matter.**

The matter record, the client contact, and the folder structure are built from the staged data with no retyping. The practice management system gets the matter. The conflict index gets the new party names. The document system gets the folder tree. The billing system gets the client record and the agreed rate. Every system that touches this client from this point forward gets its initial data from the same staged source - the form the client submitted. Nothing is entered twice.

> **Data-handling checkpoint.**
> This is the highest-stakes handoff in the build. Data crosses from the staging layer into multiple systems of record in a single motion. Map every field write before you go live: which staging field writes to which system field, in which system, and with what transformation if any. Run a test submission through the full chain and inspect every destination record before the first real client goes through. A missing rate on the billing record or a wrong matter type on the practice management record is a problem that follows the client for the life of the matter.

6 **Confirm.**

The signed engagement triggers the welcome sequence and the client-portal invite. The welcome email goes out. The portal invite lands. The first set of practice-specific forms - whatever the matter type requires - is queued. The client knows they are in, knows what comes next, and has a place to send documents. The intake is closed.

> **Data-handling checkpoint.**
> The confirmation sequence sends client-identifiable information through email or portal channels. Confirm that both channels are covered by your confidentiality controls: business-tier email, a portal with access controls, and no confirmation content that would expose matter details if delivered to the wrong address. Verify the email address from the form against the client's confirmed address before the sequence fires. A welcome email to a misentered address is a confidentiality event.

3. Track A and Track B at each step

The table below shows which steps run on Track A in the baseline build, where adding an AI step would cross the build to Track B, and what that crossing costs.

Step	Baseline track	What crosses it to Track B	What Track B adds
1. Capture	TRACK A	An AI chatbot collects the intake data instead of a form - reading and responding to the client's narrative in real time.	The full five-duty discipline applies. The chatbot is generating responses a client receives and relies on. Confidentiality and supervision duties land hardest here.
2. Validate	TRACK A	An AI step reads the free-text field and decides whether the described matter fits the firm's practice areas.	The full five-duty discipline applies. The model is reading client content and producing a judgment the firm relies on.
3. Stage	TRACK A	An AI step reads the staged record and drafts a short matter summary or flags potential issues for the intake team.	The full five-duty discipline applies. The model is reading client information and generating an output the firm relies on.
4. Route	TRACK A	An AI step reads the free-text description and decides the practice group. Track A if fixed category fields do the routing; Track B if the AI reads the narrative.	The full five-duty discipline applies. Routing is the most common place intake silently crosses to Track B - a dropdown stays Track A, an AI reading "describe your issue" does not.
5. Create matter	TRACK A	An AI step reads a scanned or PDF engagement letter and extracts parties, scope, and fee terms to populate the matter fields.	The full five-duty discipline applies. The model is reading a client document and producing structured output the firm writes to its systems of record.
6. Confirm	TRACK A	An AI step drafts a personalized welcome email from the matter details rather than sending a templated message.	The full five-duty discipline applies. The model is generating client-facing content from confidential matter data.

> **The routing step is the most common Track B entry point.** Most firms do not plan to add AI at intake. What happens is a form platform or practice management tool adds an AI routing feature, someone enables it because it looks helpful, and the firm crosses to Track B without a decision ever being made. Check whether any tool in your intake chain has an AI feature enabled. If it reads client content and produces a routing output, the full five-duty discipline applies from that step forward.

4. The handoff checks

Each time data moves from one system to another, a handoff happens. Handoffs are where data leaks, transforms incorrectly, or simply stops moving. The checklist below covers every handoff in the six-step build.

- ☐ **Form to staging layer.** Confirm encrypted transmission. Confirm the receiving environment is covered by a data-processing agreement. Confirm every required field arrives and is present in the staging record before any downstream write fires.
- ☐ **Staging to practice management system.** Map every field write before go-live. Run a full test submission and inspect the matter record, the contact record, and the matter type assignment. Confirm the agreed rate is correct on the billing record.
- ☐ **Staging to conflict index.** Confirm that the client name, any related party names, and any adverse party names from the form reach the conflict index in the format the index searches. A name that arrives in a different format than the index stores is a name the conflict check will not find.
- ☐ **Staging to folder structure.** Confirm the folder tree is built under the correct matter, named to the firm's convention, and accessible to the assigned team. Confirm a test document can be saved to the correct folder before a real client goes through.
- ☐ **Matter creation to confirmation sequence.** Confirm the welcome email fires only after the engagement is signed, not at form submission. Confirm the portal invite is addressed to the email on the signed engagement, not the form - if those differ, the invite goes to the confirmed address. Confirm the practice-specific forms queued are the right ones for the matter type created.
- ☐ **Confirmation to active matter.** Confirm the completed intake record closes cleanly and the matter moves to active status in the practice management system. Confirm the intake platform's record is retained or archived per your document retention policy, not deleted.

5. Platform floor

Business tier is the floor for every step in this build on either track. Consumer-grade tools are not appropriate for any stage of the intake process. The moment a client submits a form, the data is client confidential. Every platform it touches from that point - the form tool, the staging layer, the automation platform, the practice management system, the email tool, the client portal - must be covered by a signed data-processing agreement that keeps your data out of training - stated outright or by barring any use beyond delivering the service - and requires encrypted transit and storage.

If a step in your current build touches a consumer-tier platform, that is the fix to make before anything else. The rest of the build can wait. The platform floor cannot.

When intake crosses to Track B - most commonly at the routing step - the platform requirement for that step and every step downstream climbs. The dedicated tier adds zero data retention, a business associate agreement when health data is in the matter, and the four-part protective-order configuration the book details in Chapter 9. The matter drives the tier, not the presence of an AI step: a properly configured business-tier build - the signed DPA, the confidentiality settings on, your own business-terms contract behind the model - carries routine intake. Climb to the dedicated tier when the matter asks for it: health data, high litigation sensitivity, or a protective order. Start at business tier for the baseline build. Add the dedicated tier where and when the matter earns it.
