

Jev in Lia: what changes in a law firm's day to day.

A new decision layer, explained without jargon: what it does for the lawyer Lia works for, what we measured in real conversations and what still belongs to the lawyer.

Under 1 s per conversation

Right practice area in 25 of 30

Vision: none of this is live today



A law firm is made of thousands of small decisions, and most of them never get made. Is this lead in my practice area? Is this client still waiting for a reply? Should this card have moved to the next stage by now? Does the message I'm about to send assert something I haven't checked? Nobody has time to decide this message by message. So it goes undecided: the CRM nobody updates, the lead that went cold without anyone noticing, the "I'll get back to you" that turned into a week, the dashboard with only one column filled in.

This document explains what Jev is, where it fits into Lia's day to day and what we measured in September 2026 with real conversations from a client law firm. **None of the functions described is live today:** this is the map of what comes next.

< 1 s

per conversation, with eight questions answered at once.

US\$0.01

total cost of the test with 30 conversations; the firm's entire history would come to less than one Brazilian real.

25 of 30

legal practice areas classified with a high degree of confidence.

SECTION 1

What Jev is

Jev is an artificial intelligence model of a new kind. It **does not chat and does not write**. It does one single thing: it takes a text and a closed question (*yes or no?*, *which of these?*, *what score?*) and returns the answer along with **its degree of confidence** in it.

If Lia is the teammate who talks to the client and drafts the filing, Jev is the intern who glances at a stack of five hundred papers, sorts them into three piles in two minutes and honestly points out the ones it was unsure about.

Three things that make it useful

- **Fast:** it answers in under a second, even with several questions at once.
- **Cheap:** classifying a firm's entire conversation history costs less than one Brazilian real.
- **Knows when it does not know:** the confidence it reports is real. When it says "I'm not sure", the decision goes to Lia or to a person.

Three things it does not do

- It does not draft filings, contracts or messages.
- It does not give legal opinions or decide what belongs to the lawyer.
- **It does not replace Lia:** it decides *whether* and *which*; Lia does the *what*.

Why this matters: today, every small decision Lia makes draws on the firm's own Claude account, the login subscription under the client's company tax ID. Jev takes the intern work off her plate, at intern prices, and leaves Lia free for what only she does well: talking, understanding the case and writing.

SECTION 2

A day at the firm: five moments when nobody decides

Moment	Today	With Jev underneath Lia
Morning, before coffee	The lawyer opens WhatsApp and scrolls the screen to find out who was left without a reply. It depends on the memory of whoever handled things yesterday.	A short list: who is waiting on something from the firm , from the promise to get back five days ago to the question left hanging. Sorted, with the reason.
A lead comes in	The card is created in the CRM with no practice area, no stage and "normal" priority, like every other one. The "group by practice area" view shows a single column.	On arrival, the lead is already classified: which practice area (banking, social security, family...), which pipeline it goes into, whether it is a lead or a former client, whether it is a wrong number.
The conversation moves on	The client says "agreed, go ahead and send the contract". The card stays where it started, because moving it means remembering to move it.	With every reply, a silent question: has the case moved to another stage? When it has, Lia is notified, and she is the one who moves it, within the firm's rules.
Before sending a reply	An answer given "off the top of someone's head" can state that a person has no claim and turn away a client who did have one. There is no brake between drafting and sending.	Before going out, the message goes through a single question: does this give a verdict or promise something without a source? If so, Lia checks the firm's record before sending.
End of the day, the dashboard	The CRM's three views (practice area, status, priority) exist, but they show everything in one column. The board on the sales team's TV says nothing.	The views start to truly separate things. The lawyer sees the firm: how many leads await their decision, how many contracts are unsigned, who went quiet.

And the house rules

Every firm has its rulebook: "don't discuss deadlines with a client without going through me", "billing only after checking the amount in the system", "a lead who already receives the benefit is never

dismissed". Lia reads all of it, but a rulebook with hundreds of rules does not "fire" at the right moment. The rule exists and, even so, gets run over.

With Jev, before each action Lia asks the whole rulebook at once: **which of these rules apply right now?** Two or three come back. On a greeting, none. It is the checklist for the moment, assembled on its own, even for situations nobody anticipated.

SECTION 3

Across Lia's five roles

Where each role gains a tireless intern. Items marked **tested** were measured with real conversations from a client law firm in September 2026; the others are planned applications.

Role	What Jev decides underneath
Intake	Screening at the door: lead, former client, vendor or wrong number; practice area; the right pipeline. tested "Does this message still need a reply, or is it a closing 'thank you'?" tested "Does the reply I'm about to send give a verdict without a source?" tested Which of the firm's rules apply to this situation. tested
Administrative assistant	Practice area, status and "client waiting for a reply" on every CRM card. tested Daily sentinel: does the card's stage still match the conversation? It flags, it does not move. tested Initial load of a CRM that started out empty: the entire history classified in minutes.
Billing assistant	What the client meant in the billing conversation: sent a receipt, will pay later, disputes it, asks to pay in installments. Each one follows its own path. Before sending: "does this message mention an amount, an installment, a date or an agreement?" If it does, reread the firm's billing rule.
Controller	Auditing against what is written: did this reply go against a recorded rule? Compare the action with the rule, not with someone's memory. Repeat issues: is the lawyer correcting something that was already taught? Promises with no follow-up: who has been waiting on something from the firm for more than a day. tested
Drafter	Jev does not write a single line. Around the writing: which document arrived (medical report, power of attorney, bank statement) and whether it is complete; whether the filing meets the firm's formatting checklist, item by item; which of the firm's templates or precedents is closest to the new case.

The pattern is always the same

- **High confidence:** the system proceeds on its own.
- **Medium confidence:** Lia decides, with the full context.
- **Low confidence:** a person takes a look.
- **Nothing irreversible** (billing, dismissing a lead, marking a case as lost) happens on that number alone.

SECTION 4

What we measured

In September 2026 we ran Jev over **30 real WhatsApp client conversations** from a client law firm, with eight questions per conversation: practice area, pipeline, stage, status, priority and whether the client is waiting for something. Without touching anything: just measuring.

The most useful signal: who is waiting for a reply

Jev flagged **9 of 30** conversations in which the client was waiting for something from the firm. We checked two by hand. In one, Lia had written "I'll take it to the lawyer and get back to you with the next step", five days earlier, with no follow-up. In the other, a client asked about activity in her case and got "I'll message you on the other number"; the question was left hanging. Both were right.

Other measurements

- **Practice area:** 25 of 30 with high confidence; on the 5 cards that already had the subject matter noted, it got all 5 right.
- **"Does this message still need a reply?"** Jev combined with the existing rule got 26 of 26 tricky cases right, including the classic "perfect, thank you so much, I'll wait to hear from you" (which is not a sign-off: the person expects something to be delivered).
- **"Does this reply give a verdict without a source?"** Across 16 rephrased sentences, the current keyword filter let 9 through, all in the dangerous direction (the message goes out). Jev caught all 9.
- **Firm rules:** 36 rules evaluated in a single question, in one second. It lit up the right rule in four different situations (a debtor asking about an amount, a boss issuing a string of commands, a client asking about a deadline, a dictated message) and **none** on a greeting.

What did not work well, stated plainly

Priority (urgent, high, normal, low) came out with low confidence in almost half the cases: the concept is too vague for an automatic score. It works as a reading order with a person checking, not as an action. And the **pipeline stage** only works well when the stages are described in plain language; where the stage depends on something that does not show up in the conversation (a payment in the system, a click by the lawyer), the system knows, not the text.

SECTION 5

What still belongs to the lawyer, and to Lia

Stays with the lawyer

- Marking a case as won or lost.
- Dismissing a lead, billing, negotiating fees.
- Giving the verdict on the specific case.
- Setting the house rules, and writing them so they can be looked up one by one.

Stays with Lia

- Talking to the client, with the full context.
- Drafting messages, filings, contracts.
- Carrying out the action after Jev has flagged it: moving the card, taking notes, sending a heads-up.
- Deciding the cases where Jev reported doubt.

Two conditions for it to work

1. **Text only.** Jev reads text. Audio and photos of documents are transcribed first, something Lia already does.
2. **Rules written as units.** A rulebook with blocks dozens of pages long does not light up at the right moment; short, named rules do. In the tests, the same rule went from "not found" to "first on the list" just by being split out of the block it was in. Organizing the rulebook is work worth doing with or without Jev.

Where to start

1. **CRM first.** Practice area and "client waiting for a reply" are the two most reliable classifications we measured, and the two most missing from the dashboard today. Start there, at a pilot firm, in suggestion mode: Jev proposes, Lia or the lawyer confirms.

2. **Then, the verdict brake.** The question "does this assert something without a source?" before every send, in place of the keyword filter.
3. **Finally, the rulebook.** Once the firm's rulebook is written as units, the checklist for the moment starts assembling itself.



A firm where nothing is left without a label, a stage, an owner or a reply.

Questions about this document or interested in joining the pilot? Talk to RIA through Lia's support channel.

SECTION 6

Data sources

All numbers in this document come from RIA internal tests run on **September 20 and 21, 2026**, and from the model's public documentation.

1. **Test with 30 real client conversations (RIA, 09/21/2026).** Thirty WhatsApp conversations from a client law firm, eight questions per conversation (practice area, pipeline, stage, status, priority, client waiting): average time per conversation, total cost, 25 of 30 practice areas with high confidence, 9 of 30 "waiting for a reply" and the two manual checks. The firm's complete history (568 conversations) was sized from that cost.
2. **"Does this message still need a reply?" test (RIA, 09/20/2026).** 26 tricky messages taken from cases documented in the system itself; result 26 of 26 with the existing rule + Jev combined.
3. **"Verdict without a source" test (RIA, 09/20/2026).** 16 sentences rephrased from field cases; the keyword filter let 9 through, Jev none.
4. **Firm rules test (RIA, 09/20/2026).** 36 rules from a client firm's rulebook evaluated in a single question against five situations; the same rule from another firm rose from "not found" to "first" once it was split out of its block.
5. **TypeSafe AI: model and pricing documentation.** Jev model; billing only for the text sent, with no charge for the answers; basis for the per-conversation cost calculation.
<https://docs.typesafe.ai/models>

Early results, measured on a small sample; they are not performance guarantees. The client law firm in the tests is not identified and the conversation quotes were anonymized. None of the functions described is live in Lia as of the date of this document.