INDEPENDENT FIELD RESEARCH · 2026

How Lawyers Actually Work

Field notes from sixty legal professionals across the US and Europe. What eats the day, where deadlines leak, what AI is trusted with, and what production use of AI on live case files actually showed.

Petar Jovanovic

Co-founder, Oktopeak · August 2026

INTRODUCTION

Why this exists

Beginning in 2026, I started asking lawyers a simple question: what actually eats your day? Some answered a short survey. Some got on a call with me. A few ended up in long email threads because they had found our open-source tools on their own and had opinions about them. I promised everyone who participated that I would compile what I heard and send it back. This is that document.

One thing before we start. My company builds software for law firms, including the connectors some of the people in this report use. That is a bias. I am publishing anyway, and where a finding is convenient for my business I have tried to say so out loud. You can judge for yourself whether I managed.

WHAT IS IN THE SAMPLE

Between January and June 2026: 50 questionnaire responses (continental Europe and US), and a dozen longer conversations, meaning recorded calls, working sessions, and detailed written threads, with firms across the United States and Europe, plus a few further afield. Sizes range from a brand-new solo practice to a BigLaw associate to an in-house department of ten.

On top of that, something surveys cannot give you: field reports from firms running our open-source connectors against live practice management data. Where I cite public forum threads instead of my own conversations, I label them. The handful of industry-wide statistics in this report are named to their sources in the text: the ABA's surveys, claims studies and ethics opinions, Clio's Legal Trends Report, and a Stanford study published in the Journal of Empirical Legal Studies.

Add the two streams together, minus the few people who both filled the survey and got on a call, and this is around sixty legal professionals. It is not a statistically valid sample of the profession and I will not pretend it is. Findings that come from one person are marked as coming from one person. What this sample has instead of size is depth: real people describing their own week, and in a few cases, real logs of what happened when they pointed AI at their own case files.

SECTION 1

The day is writing. The leak is everything around it.

Across the survey, lawyers report spending two to nine hours a day writing filings, briefs and contracts. The median is around four. Client calls and messages take another thirty minutes to five hours, and almost everyone confirms that clients call about things they could check themselves.

But when I asked what takes the most time, the answers were rarely legal work. They were the work around the work:

> An in-house counsel who runs a ten-lawyer department at a public company described tracking six-to-seven-month transactions by hand across Monday.com, notebooks and email. Updates arrive as emails ("this license takes 15 more days") and get manually retyped into task lists and gantt charts. He has tried imposing project management tools on his team. It always fails, in his words, because people will not keep updating them on top of their real work. His dream tool reads his inbox and decides on its own: is this an update to a task, should this become a task.

> The founder of a small family-law firm assembles every initial consult the same way: AI meeting summary from the video call, plus his handwritten notes, manually merged, manually uploaded into his practice management system. Every new client, every time. The same system holds around 200 custom fields he populates by hand, and its document automation generates the required court forms with only the client name filled in. The rest arrives blank.

> At a defense litigation firm working across two US states, a paralegal builds medical chronologies by re-uploading documents that already sit in the case management system. The office admin schedules around twenty cross-attorney meetings a week, and the recurring puzzle is always the same: when are these three attorneys free at the same time.

> A criminal defense lawyer in a suburban practice put discovery review at the top, because every jurisdiction hands discovery over differently: one gives an open file, one emails a link, one makes them file a motion and review the material at the courthouse. Asked what they would pay for a tool that fixed it: "Probably isn't one haha, but if there would be one, I would pay good money."

> In the paper-based European offices the same category shows up lower in the stack: sorting bank statements, checking the court portal every day, standing in line at the registry office, and, in one small shared office, the single biggest time sink was named in one word: printing.

The industry-wide numbers say my sample is normal. Clio's Legal Trends Report, built from anonymized platform data across tens of thousands of legal professionals, puts average lawyer utilization at 37 percent, about 2.9 billable hours in an eight-hour day. It was 30 percent in 2019, so this is the improved version. Five hours of every legal workday go to work nobody bills for.

The eight-hour day, industry average

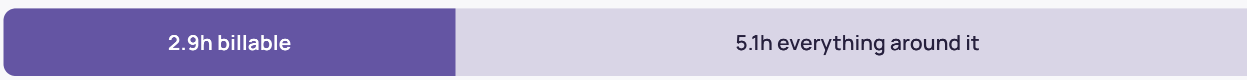

CLIO LEGAL TRENDS REPORT · 37% AVERAGE UTILIZATION

Nobody in this sample complained about lawyering itself. What they described, in different words, is the job of being the integration layer between their own systems.

> *"All in my head, I should build some system."*
>
> SOLO LAWYER, ON WHERE CLIENT CONTACTS LIVE

SECTION 2

Deadlines: everyone runs three systems, and it still leaks

The standard deadline setup I found, especially in the European sample, is redundancy: a paper planner, plus a calendar, plus an Excel table, plus in one case a phone calendar on top of Outlook on top of the planner. "On several fronts," as one lawyer put it. Most respondents say they have never missed a deadline, and given the professional consequences, I believe they mean it.

And still it leaks. One solo lawyer, who checks the court portal daily and keeps a paper planner as backup, missed a hearing anyway. The reminder simply never arrived, and it cost him real money. Another lawyer admitted missing deadlines twice in three years. A third named, as his top frustration, judges who do not notify you when a hearing in another city gets postponed.

Share of legal malpractice claims tied to missed deadlines and calendaring errors

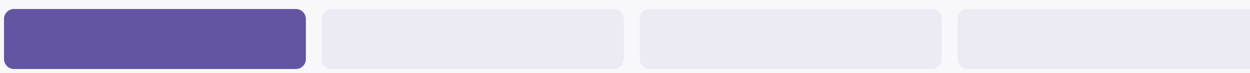

ABOUT ONE CLAIM IN FOUR · ABA STANDING COMMITTEE STUDIES OF LEGAL MALPRACTICE CLAIMS

The profession's own claims data shows what is riding on this: in the ABA standing committee's studies of legal malpractice claims, missed deadlines and calendaring errors rank among the most common causes, in the range of a quarter of all claims. The redundancy is theater against a system that fails upstream of it.

Meanwhile one US solo attorney in this sample quietly built the thing everyone else is doing by hand: he connected an AI assistant to his practice management system and runs deadline and docket review through it daily, on live matters. Nobody sold it to him. He set it up alone. Which brings me to the pattern that runs through this whole report: the tooling gap is not about willingness.

SECTION 3

Europe runs on paper. The US runs on rented software. The pain is the same.

For 5 of the survey respondents, cases live in Excel, Word, paper folders, or one internal program. Client contacts live in phone address books, WhatsApp, or, in the words of one solo lawyer, "all in my head, I should build some system."

20 run on Clio, MyCase or Practice Panther. The rest of the survey sample sits between those two poles, partial systems and skipped answers, and I am not going to force them into either bucket. But paying for software does not close the gap, it just moves the complaint. A senior attorney at a nine-attorney general practice added up his firm's stack for me on a call: Clio plus Lawmatics, up to sixty thousand dollars a year. His unprompted question: can these be combined into one product?

$60K

Annual software stack at one nine-attorney general practice: practice management plus intake. His own math, on a call. His unprompted question: can these be combined into one product?

RESEARCH CALL, SENIOR ATTORNEY, NINE-ATTORNEY FIRM

An associate at a five-attorney New York firm called their platform "definitely pricey," said there is no real alternative, and added the sentence I have thought about most since: creating your own server "would be great, we're just not tech savvy enough, especially when there's errors." And the complaint holds at the top of the market: the BigLaw associate in this sample named the firm's document management system, not AI, as the frustration, "incredibly frustrating. Having to save locally takes up so much space."

> *"I would like there to be a single platform (even if integrated with other systems) that houses case/client information, all documents, deadlines, and updates. Information is scattered and causes issues."*
>
> SURVEY RESPONDENT, NEW YORK FIRM ON MYCASE

Same wish, survey and call, two firms, independently.

FROM PUBLIC FORUMS, LABELED AS SUCH

In lawyer threads on Reddit you find the same emotions at higher volume. "Both Clio and MyCase oversell everything." A firm that stays on software it hates because "I can't stomach having to do ANOTHER firm migration." An office that pays about 120 dollars a month for a self-assembled stack of generic tools and claims it saves them fifty thousand a year, next to a warning in the same thread that billing and trust accounting without a real system "quickly becomes a nightmare."

The three tiers I can distinguish in this sample want different things. Solos want the organization problem taken off their hands, and they will pay for the service more readily than for software: the European solo with the missed hearing has a hard software ceiling around what he pays for his legal research database (20 to 25 euros a month) but would, quote, gladly pay someone to do the sorting for him. Small firms in the middle told me their binding problem is getting clients, not technology. The larger offices and in-house teams are the ones asking workflow and data-control questions.

The money question came back as a ladder, not a single number. Asked what they would pay to make their biggest bottleneck disappear, the European survey answers climb from 50 to 70 euros a month, to 100 to 200, to around 170 twice, up to one office already paying 1,500 euros a month for an enterprise AI deployment.

What would you pay to make your biggest bottleneck disappear? (monthly, European survey answers)

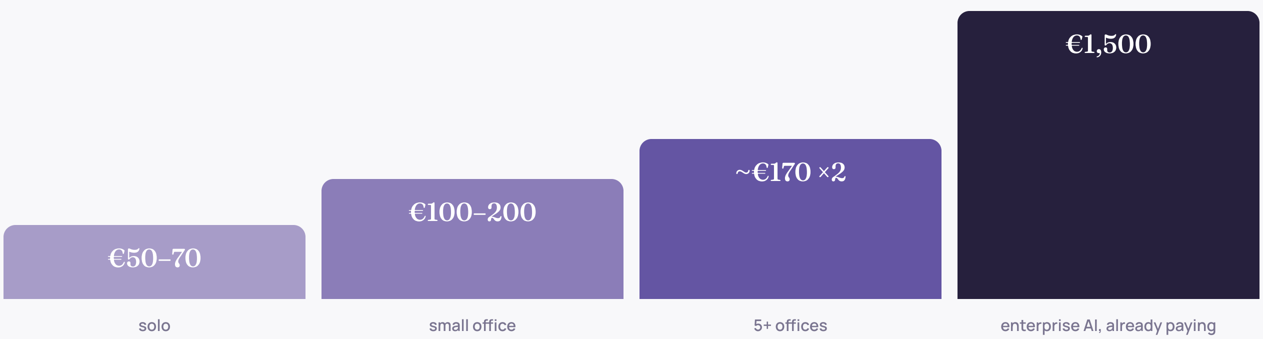

STEP HEIGHTS INDICATIVE, AMOUNTS AS REPORTED · SURVEY RESPONSES

The English-language answers mostly refused to name a number and named a condition instead: "would have to know its proven capabilities," "depends how efficiently it does it and how much integrations are." Buyers in this sample condition on proof before price, and the numbers rise roughly with firm size.

SECTION 4

Everyone already uses AI. Almost nobody feeds it real data.

49/50

Survey respondents who have used AI for work. That includes the offices whose tooling baseline is paper. The one refusal was explicitly about confidentiality.

THIS SURVEY, JANUARY–JUNE 2026

The conversations put money behind the same appetite: at the defense litigation firm, one attorney hit the limits of the firm's consumer AI subscription "after a couple documents" and upgraded to the most expensive individual plan out of his own pocket, without waiting for anyone's approval.

For context, the ABA's 2024 technology survey put lawyer AI use at 30 percent, up from 11 percent the year before, and found the most used tool among adopters was consumer ChatGPT at 52 percent, ahead of every purpose-built legal AI product. 49 of 50 is far above that, which probably says something about who answers a survey on how they work. The consumer-tools-first pattern, though, matches exactly: general-purpose assistants first, legal suites a distant second.

What they use it for, in their own words: finding case law, summarizing long filings, translating foreign documents, correcting contracts, validating their own reasoning, drafting starting points, styling text, building an internal knowledge base, and in one case dictation, a solo lawyer who speaks her thoughts to the AI "so I don't have to type" and lets it turn them into filing text.

Where it fails them is specific and consistent. One lawyer used AI for case-law research and got judgments that do not exist, citing laws that do not exist. Another described a subtler version: it quotes the correct text of a statute with the wrong article number, so every result needs manual verification. A third pays for a specialized legal AI product and rates its case analysis as excellent and its drafting as "beginner level." An enterprise Copilot user says it "drifts for no reason." The New York firm that uses AI for calendaring, drafting and templates says it "did not do great with legal research and reasoning." The pattern: AI is already trusted as a summarizer, translator and organizer, and not yet trusted as a lawyer. I think the profession has this one right.

The published research agrees with them. A Stanford team tested the leading AI legal research platforms on more than 200 preregistered legal queries and found that even these purpose-built, citation-grounded tools hallucinate between 17 and 33 percent of the time, against 58 to 82 percent for general-purpose chatbots on legal questions.

How often AI hallucinates on legal questions

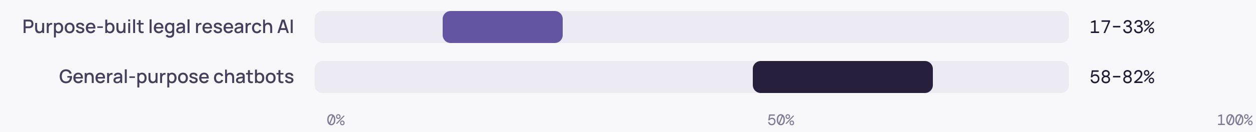

STANFORD STUDY, JOURNAL OF EMPIRICAL LEGAL STUDIES · 200+ PREREGISTERED LEGAL QUERIES

1,600+

Court decisions worldwide involving fabricated AI citations by mid-2026, a practicing lawyer responsible in more than 650 of them. One 2026 federal appeals decision removed a lawyer with a forty-year clean record from his case over AI-drafted briefs containing fabricated quotations.

PUBLIC DATABASE OF COURT DECISIONS, MID-2026

The respondent who got judgments that do not exist has a lot of company.

The confidentiality finding surprised me most. I expected fear as a blocker. What I found instead is a working protocol, invented independently by at least eight respondents in two languages: no names, no client identifiers, generalized fact patterns, anonymized templates. The European office paying 1,500 euros a month runs an enterprise AI deployment with ISO and SOC2 certification, and still anonymizes everything that goes in. One lawyer solved it contractually, with confidentiality terms in the vendor agreement. Two answers went the other way and worried me a little: one US respondent enters case data freely because ChatGPT is "a closed system" (it is not, in the sense he means), and another sees no difference between AI and Gmail.

Lawyers have built a manual privacy layer out of their own discipline, and they pay for it in usefulness, because an AI that never sees the case file cannot do the case file's work. One twelve-person commercial firm brought the cleanest version of the question to a call with us, as their gating condition before any talk of features: which region does our practice data sit in, and where does the model's processing actually happen. Nothing else moved until that had an answer.

The honest upgrade path is contractual: vendor agreements and AI plans where inputs are excluded from training and retained for zero days. That is a procurement question, and any vendor, us included, should be able to answer it in writing.

The bar has arrived at the same place: the ABA's Formal Opinion 512 on generative AI, issued in July 2024, requires lawyers to evaluate how a tool handles client data, including whether inputs are retained or used for training, before any client information goes in, and by mid-2024 more than two dozen federal judges had standing orders governing AI use in their courtrooms.

> ***Which region does our practice data sit in, and where does the model's processing actually happen?***
>
> TWELVE-PERSON COMMERCIAL FIRM, GATING QUESTION BEFORE ANY TALK OF FEATURES

SECTION 5

What production use showed

This is the part of the report a survey cannot produce, so I want to be precise about the source: we maintain free, open-source connectors that let firms link an AI assistant to Clio, MyCase, Filevine and IntakeQ. Some firms install them with our help, most install them alone, and a few write to us about what happened. Install counts are anonymous npm statistics, in the hundreds per week across the four connectors. We see no user data, only what people choose to tell us.

The sharpest report came from a solo practitioner who ran the Clio connector across several hundred matters doing bulk document reorganization. He is a lawyer, not a programmer. He hit four defects, and the pattern of the defects matters more than the count. Sustained bulk operations tripped the platform's rate limits. Listings silently returned only the first page of results. Folder hierarchies reported a different parent type than his scripts assumed. His dry run concluded that nearly every matter was missing a folder that actually existed. If he had run the write pass, it would have created duplicates across his whole book. These are failures that produce wrong results instead of errors, which for legal data is the dangerous kind.

A different solo attorney, on a different platform, independently found the same first-page-only pagination behavior in that platform's API. My conclusion from both: practice management APIs were built for humans clicking, and they are now being driven by AI at bulk scale, and the seams show.

The second production finding: the blocker is usually the data, not the AI. The bulk-reorganization lawyer told me directly that after years of use with no filing protocol, his firm's data "is kind of a mess," and the cleanup is the actual project. The family-law founder's 200 inconsistently populated custom fields are the same problem in a different suit. AI amplifies whatever discipline your data already has, in both directions.

The third finding is about who crosses the do-it-yourself line. The family-law founder ended a paid pilot with us by building the integration himself, with an AI assistant's help, and told us plainly, and I quote with some professional pain: "I ended up building the necessary connections and flow myself for my own needs." The nine-attorney-firm senior attorney evaluated three vendors, picked our connector, and then installed it himself too. And the bulk-work solo extended the connector's code with AI when he hit its limits.

So the line is real and lawyers cross it. But notice who they are: a founder who uses developer tools, a ten-month power user, a lawyer comfortable reading API responses. On the other side of the line sit the firm that adopted AI company-wide without knowing which plan they bought, the associate who said "we're just not tech savvy enough, especially when there's errors," and the thirty-year attorney who told me exactly what he wants built and has no interest in wiring it. In this sample, most of the profession lives on that second side.

> *"Creating your own server would be great, we're just not tech savvy enough, especially when there's errors."*
>
> ASSOCIATE, FIVE-ATTORNEY NEW YORK FIRM

One more pattern, about buying. Every firm that reached us had done its homework before saying a word: one attorney read a comparison post twice and clicked through to a competitor's code repository before ever booking a call, one demanded a written fixed price with no call at all, and in three cases the person I spoke with was not the decision maker but the champion carrying material to the partners. And twice, the referral came from an AI: once from a chatbot summarizing a forum post, once from an AI search engine that recommended a tool and produced a booked meeting in nineteen minutes with no human involved. How lawyers find vendors is changing underneath the vendors.

SECTION 6

What software will not fix

A real list of things my respondents wish would disappear: filing submissions by physical mail, printing, copying and scanning, waiting in line at the registry, waiting at the jail to see clients, inconsistent case law, judges who do not notify about postponements, the disorganization of the public sector, pro bono work, and, twice, their own boss. One new solo lawyer wrote the summary himself: his biggest problem "is not solvable by paid tools."

I include this because a report from a software company that concludes everything is solvable by software should not be trusted. A meaningful share of legal-work pain is structural, jurisdictional or human. What software can actually reach, on this evidence, is a narrower and more boring list: the scattered-information problem, the deadline-notification problem, the document-assembly problem, and the re-typing-between-systems problem. That list happens to be where the hours are.

SECTION 7

Open questions

Things this sample raises and cannot settle, which I will keep collecting: whether the transaction-tracking pain repeats across in-house departments, how often the deadline redundancy actually leaks (two admissions in sixty people, both accidental discoveries), whether the pay ladder in section 3 holds beyond this sample, and the question my own respondents asked me to investigate: several wanted to know what share of their time other firms already spend using AI, and how other firms use it in real workflows. That is a benchmark worth building properly.

If you got this report because you took part: thank you, genuinely. If reading it you disagreed with something, or your week looks different, reply and tell me. Every finding in here got sharper each time someone corrected me.

ABOUT

Who ran this

I am Petar Jovanovic, co-founder of Oktopeak, a seven-person software agency working with law firms and healthcare practices. We maintain the free, open-source connectors mentioned above, and we install and support them for firms that would rather not wire this up themselves. On top of the systems firms already rent, we build the boring things this report found the hours in: intake that assembles itself, document automation that gets past the client name, deadline monitoring, and the retyping-between-systems work from section 1.

We also handle the procurement question from section 4, putting firms on AI plans where inputs are excluded from training and retained for zero days, confirmed in writing. Everything custom, owned by the firm when we are done.

All respondents and firms in this report are anonymized. Quotes are verbatim where in English and translated by me where the original was not.

ONE OFFER, NO PITCH

Want to know where your firm sits against this cohort?

Fifteen minutes, free. We compare your week to the patterns in this report and you leave with the comparison, whatever you decide to do with it. If you took part in the research, this is the follow-up I owe you anyway.

`oktopeak.com` `calendly.com/office-oktopeak/30min`