



STATE OF AI IN PLAINTIFF LAW | 2026

The AI Adoption Gap

Why AI Interest Isn't Becoming AI Impact, and What
the Firms That Get It Right Are Doing Differently

Survey data by

III THIRDSIDE

A Note From Our CEO

Nearly every firm we talk to asks the same question: what are other firms doing with AI, and is it working? We're over two years into watching law firms adopt AI, and while the interest is abundant, the questions about real adoption and impact are still open. This spring, we commissioned research with a third-party firm, Thirdside, to put data behind the gap between AI curiosity and AI impact.

The headline: 8 in 10 firms have engaged with AI, and only 3 in 10 have built it into how they work. The difference comes down to trust. 99% of attorneys told us they won't use AI output they can't verify, and 79% reject fully autonomous AI, insisting that every output needs to incorporate their judgment and review. That caution reflects the professional and fiduciary duty attorneys have always carried, now applied to a new tool, and any AI worth adopting has to meet that standard.

We believe the size of a firm should never determine the strength of its case, yet outcomes still turn on resources: how fast a team can get through the records, how strong and how soon the demand goes out. Much of the value that decides a case is already sitting in the file, waiting to be surfaced. The firms pulling ahead are the ones putting AI to work from the moment a case comes through the door, on a foundation they can verify and defend.

That's the future we're building toward. We're grateful to the professionals who shared their experience for this research, and to the firms who trust us with the cases that matter most to the people they serve.

If these findings reflect questions your own firm is working through, we'd welcome the conversation and would love to show you the value Supio can bring to your practice.

Schedule some time to talk with our team at supio.com.



A stylized, handwritten signature in black ink, appearing to read 'Jerry Zhou'.

Jerry Zhou
Co-Founder & CEO

The debate is over.
The gap is just
beginning.

8 in 10

firms have engaged with AI

3 in 10

have embedded it

The question plaintiff firms were asking two years ago, whether to use AI, has been answered. Nearly 8 in 10 firms have engaged with AI in some form. The technology is accessible, the interest is real, and the urgency is measurable.

Yet only 30% of firms are using AI as a regular part of how they work. The rest are experimenting, evaluating, or stalled. They are aware of the potential but not closing the distance between curiosity and practice.


The question has shifted from whether to use AI, to where and how to embed it so it actually changes how cases get built.

This report is about that gap. What it reveals, what it costs, and what the firms closing it are doing differently. The data points to a specific answer, but it is not about how much AI a firm is using. It is about what kind of AI they are using.

AI deployed on demand, case by case, adds speed to a broken foundation. It makes individual attorneys faster at specific tasks without impacting the overall speed of the case in aggregate. The firms making AI work are using something structurally different: agentic AI, embedded in how cases are built from the moment files arrive, not reached for when an attorney needs it.

The difference between those two approaches is not a matter of degree. It produces structurally different firms, with structurally different economics.

Q: Which of the following best describes your firm's current use of AI tools?
n = 207



**Current AI
adoption status**

78% have tried AI.
30% actually use it.

62%

of pre-litigation attorneys use AI

47%

of pre-litigation paralegals use AI

~25%

litigation-side adoption rate

The AI adoption gap in plaintiff law is not about awareness. It is about where AI is being asked to work.

When you look at which roles are actually using AI inside plaintiff firms, the pattern tells you something important. Pre-litigation attorneys lead at 62%. Pre-litigation paralegals follow at 47%. Litigation-side staff, the attorneys and paralegals working cases closest to trial, use AI at roughly half that rate.

This is not a coincidence. The pre-litigation workflow is where cases are assembled: Intake, record collection, chronology construction, demand preparation. It is also where volume pressure is most acute and where AI has high potential leverage. But most firms are still treating AI as a tool attorneys reach for individually, when they need it, rather than as the foundation their case-building workflow runs on.

AI applied on demand is a productivity aid. Agentic AI, embedded in how cases are built from intake forward is a different firm.


This distinction has a direct line to firm revenue. Productivity aids help individual attorneys work faster on the tasks they're already doing. They shave time off specific steps. The benefit is real, but it is not compounding. It does not reduce how long cases sit on a desk before they are ready to settle. It does not change the throughput of the pre-litigation pipeline.

Agentic AI, embedded at intake and pre-litigation, operates differently. When records are processed as they arrive, when timelines build continuously, when the demand letter is ready to edit rather than ready to start, the entire time-on-desk for a case shortens. Shorter time on desk means faster settlement cycles. Faster settlement cycles mean more cases resolved per attorney per year, which is the actual unit of plaintiff firm revenue.

That is the difference between firms experimenting with AI and firms transforming because of it.

Q: Which team members are using AI today?

n = 207



Team members currently using AI

The firms that
need AI most are
using it least.

TIME TO PRODUCE A COMPLETED DEMAND LETTER

57% of firms need a week or more

The demand letter is the core revenue event of a plaintiff case. Its speed is a direct proxy for how efficiently a firm converts work into settled cases.

Capacity is at a breaking point. The pressure is highest exactly where AI governance is weakest.

The urgency behind AI adoption is not abstract. 75% of plaintiff firms are at or beyond their operational capacity. 42% say volume regularly exceeds what their team can handle. The demand letter tells the story clearly.


A demand letter that takes weeks to produce is not a drafting problem. It is a data assembly problem. Record review is manual. The chronology is built from scratch each time. Billing reconciliation happens at the end of a long sequential chain. AI applied only at the drafting stage, the most common deployment pattern, does not solve any of that. It automates the final step of a slow process while leaving the slow process intact.

Consider what this costs at a firm level. If a case sits on a desk for four weeks before a demand goes out, and that case should settle within 60 days of demand, the firm has burned a month before the revenue clock even starts. For a firm managing hundreds of cases, weeks of unnecessary delay across the portfolio is not a productivity inconvenience. It is a material drag on annual revenue.

The pattern is most acute at small firms. Solo and small practices (1-2 attorneys) take the longest, with 48% needing three or more weeks per demand. They also have the lowest AI adoption rates. They are the most capacity-constrained and the least equipped to evaluate AI safely. The adoption gap and the capacity crisis are concentrated in exactly the same place.


Time to produce a demand letter

Q: How long does it typically take your team to produce a demand letter?
n = 207



Capacity breakdown

Q: How would you describe your firm's current capacity for new cases?
n = 207



Time to produce a demand letter—and the pressure is real.

The adoption gap
is a trust gap.

99% won't use AI content
they cannot verify

96% are very or extremely
concerned about
untraceable AI output

79% reject fully
autonomous AI and
want human review

Firms are not waiting because they are skeptical. They are waiting because they cannot verify what AI gives them.

Ask plaintiff attorneys what is holding them back from committing to AI, and the answer is not cost or complexity. It is trust. Specifically, the inability to verify what AI produces.

These numbers represent a near-unanimous consensus. Attorneys are not philosophically opposed to AI. They are professionally constrained from relying on output they cannot stand behind. In legal work, that constraint is well-founded. Courts have imposed sanctions for AI-generated hallucinations ranging from \$5,000 in 2023 to \$59,500 by late 2025. In multiple documented cases, errors passed through associate and supervisor review without being caught. The problem is not carelessness; it is AI output that looks correct but cannot be traced to its source.

Wanting to verify is not the same as being able to verify.

The tools most widely used for plaintiff work, general-purpose AI at 54% adoption, produce confident output with no source attribution. There is no way to trace a summary back to the record it came from, no way to catch a missed date or a misread diagnosis without reviewing the underlying document yourself. The attorney who wants to verify is left doing it manually, against a document that provides no starting point.

This is where the adoption gap and the trust gap converge. Attorneys are not skeptical of AI in principle. They are skeptical of AI that gives them no way to know whether it got it right. And as long as verification falls entirely on the attorney, the time savings that justify AI adoption are consumed by the review process needed to trust the output.

A note on the other side of the table

While plaintiff firms carefully weigh their AI decisions, their opponents have not. Insurance carriers have automated bodily injury valuation (Colossus), AI-driven demand analysis (CLARA Analytics), and pre-retention claimant routing via systems that predict at first notice of loss whether an injured person will hire an attorney. These tools route unrepresented claimants toward faster, lower-value settlements before they ever retain counsel. The professional liability standards that govern attorney work product do not apply on the carrier side.

Delay in adoption is not a neutral choice.

Early accuracy
compounds.
Late fixes do not.

71%

want more accurate medical record summaries

71%

want traceable sources for every AI-generated fact

47%

want to handle more cases without adding staff

The firms closing the adoption gap are not just using AI more. They are using it at a different point in the case lifecycle.

When firms describe what they actually want from AI, the answer is consistent and specific. Medical record accuracy and source-level traceability tied as the most important outcomes, at 71% each, ahead of faster demand letters, caseload relief, or automation. Attorneys are not asking for AI that moves faster through their current workflow; they are asking for AI that makes the foundation of every case more reliable.


The Integrity Gap

When AI enters the case workflow at the end, applied to drafting after manual record review, it inherits every error and gap that came before it. An AI demand letter built on an incomplete chronology reflects that incompleteness. A summary generated from unverified records carries that uncertainty forward. The attorney still has to go back and find what was missed.

This is what agentic AI does differently. Rather than simple call-and-response prompting, it works the case from the moment files arrive, building the timeline progressively and cross-checking data before the attorney opens the file. Everything downstream is more reliable because the foundation was accurate from the start. The demand letter is better because the chronology was accurate. The settlement strategy is sharper because the injury narrative was complete from the start.

This is not just a quality story. It is a revenue story. Every rework cycle, every missed detail caught late, every demand that goes out and comes back for revision, adds days or weeks to the time a case spends on a desk before it resolves. Firms running systematic AI from case intake forward are compressing that cycle. Firms running AI as a point-in-time drafting tool are not.

Early, verified accuracy compounds over the life of a case. Errors introduced at the foundation create friction that multiplies downstream, and friction costs time, and time in plaintiff is money.



This is not a theoretical distinction. The firms at or beyond caseload capacity, 75% of the market, are largely running the second model. AI at the end of a sequential, manual process. The firms closest to cracking the adoption gap are the ones whose AI is working the case from intake forward, and whose attorneys spend less time fixing things and more time closing them.

What the agentic plaintiff firm looks like.

Attorneys are not looking for more automation.

They are looking for AI they can trust, that works the case before they ask it to.

When survey respondents described the AI they actually want, the profile is consistent. Not a tool that responds to prompts, but a platform that works cases proactively, produces output attorneys can verify at the source, and keeps client data completely isolated from third-party systems.

Generative AI	Agentic AI
Demands drafted after manual record review is complete.	Foundation built continuously from intake forward. Demand ready to edit, not ready to start
Errors introduced at drafting, caught late or not at all – rework adds days or weeks to every cycle	Records cross-checked at ingestion – verified accuracy compounds forward, friction doesn't
AI speeds up individual steps; the sequential, manual pipeline underneath stays slow	Shorter assembly cycle at every stage – cases move faster from intake to settlement
Individual attorneys work faster on tasks they're already doing – throughput is largely unchanged	Cases that were turned away for bandwidth reasons become workable – more cases resolved per attorney per year
Leadership sees what attorneys report – no real-time view across the caseload	Patterns across intake quality, fee capture, and case mix visible and queryable in real time
Client updates depend on staff availability and bandwidth	Cases actively tracked; nothing falls through the cracks and clients receive detailed attention
Each session starts from the same place – institutional knowledge lives in individual attorneys' heads	Firm history, winning strategies, and attorney preferences compound over time – knowledge stays when people leave
Efficiency gains are real but replicable – any firm can adopt the same tools	Context compounds. The longer the firm runs on an agentic foundation, the harder it becomes to displace.

Most essential capabilities in an AI platform (top 3)

Q: When evaluating a legal AI platform, which capabilities are most essential?
n = 207



76%
say a direct integration with authoritative legal research like Thomson Reuters Westlaw would increase their confidence in AI output.

CONCLUSION

The firms that close
the gap first will be the
hardest to displace.

The adoption gap and the trust gap are the same problem.

The answer to both is agentic AI that starts earlier, works continuously, and gives attorneys something they can actually stand behind.

The AI adoption gap in plaintiff law is not a technology gap. The tools exist, the interest is real, and the urgency is measurable. The gap is a trust and verification gap and generative AI, the category most widely deployed in plaintiff firms today, was not designed to close it. It produces output and stops. Verification falls entirely on the attorney, consuming the time savings that justified adoption in the first place.

The answer is agentic AI. Not as a category label, but as a design principle. AI that starts at intake, works the case continuously, and produces output attorneys can trace to its source. When the foundation is accurate from the start, everything downstream moves faster. Cases spend less time on a desk. Demands go out sooner. Cycles shorten. That is not a productivity story. It is a compounding revenue story. The firms building on an agentic foundation today are not just closing the adoption gap but accumulating institutional knowledge and firm-specific intelligence that grows harder to replicate with every passing month. Context compounds. Late starters do not catch up.

73% of plaintiff legal professionals are optimistic or excited about AI's role in their practice. Fewer than 1 in 20 are skeptical. The market has just been waiting for AI it can trust. The firms that build on that trust now with AI that starts earlier, works continuously, and produces output attorneys can actually stand behind will not just close the adoption gap. They will define what a modern plaintiff firm looks like and the rest of the market will be playing catch-up to a standard they helped set.

Appendix

Survey Methodology: Sample, Panels, and Data Quality

Survey Design & Fielding

This study was conducted using an online survey administered to U.S.-based personal injury legal professionals. Data collection took place between **March 11–26, 2026**, yielding a total of **207 completed responses**.

Sampling Approach

Respondents were recruited through **opt-in professional research panels**. All participants met predefined screening criteria related to role, firm type, and practice area to ensure relevance to personal injury law and legal technology decision-making. In addition, screening questions were included in the survey to screen out unqualified respondents.

Panel Quality & Respondent Validation

To maintain data integrity, panel partners applied standard quality controls, including:

- Identity and professional role verification
- Duplicate response prevention
- Speed and consistency checks
- Attention and logic validation throughout the survey
- Responses failing quality thresholds were excluded from the final dataset.

Sample Representativeness

While the sample is not intended to be statistically representative of all U.S. law firms, it is **directionally representative of small and mid-sized personal injury firms**, reflecting the segment most actively evaluating AI-enabled legal technology. Results are appropriate for **market sizing, directional insights, and strategic decision-making**, rather than precise population estimates.

Analytical Considerations

Findings should be interpreted as **indicative of market dynamics and buyer behavior**, not as precise incidence measures. Subgroup analyses are directional where base sizes are limited.

Respondent Profile


207 Total respondents

Personal injury law professionals


Appendix

(continued)

Respondents by role



Respondents by primary segment



Respondents by attorneys at firm

