**IQdoc** Applied AI for legal.

IQDOC RESEARCH

# The State of Legal AI, August 2026

A survey of the AI-native legal market, organised by who the buyer is: $6.20B across 95 companies, 269 products, 50 agents, and an audit of what a buyer can actually verify.

| Capital tracked | Companies | Products | Mean evidence score |
|---|---|---|---|
| **$6.20B** | **95** | **269** | **1.99 / 10** |

Data compiled 25 August 2026. Evidence audited 25 August 2026. Published by IQdoc — applied AI for legal.

# Contents

# Front matter

### What this is

This is a survey of the AI-native legal market as it stood on **25 August 2026**. It reports three things we counted ourselves: **$6.20B** of disclosed capital across **95** companies, **269** named products from **96** companies, and an audit of what each of the 95 companies lets an outsider verify about its own claims.

It is organised by **who the buyer is** rather than by product category, because that is where the interesting asymmetry lives and nobody else appears to be publishing it. Four audiences: corporate legal departments, law firms, independent lawyers and small firms, and people handling a legal matter without a lawyer.

### What we counted

**The 95.** AI-native legal companies — founded around an AI product for legal work, rather than legal software companies that added AI to an existing platform. Private companies with disclosed venture funding. The largest legal technology businesses in the world are therefore absent by design: Clio, Filevine, Icertis, Ironclad, Relativity, Thomson Reuters, LexisNexis. Eighteen of them are listed as adjacent companies in Appendix C, without their capital counted in any figure here.

**The 269.** Every product these companies market by name, typed into 18 categories and tagged with the audiences the company's own site and pricing address — not the audiences that might benefit. **264** come from the 95. **Five** come from Courtroom5, which our funding threshold could not see and which builds the most substantial software we found for self-represented litigants. Where a figure is about the funded market it uses 264 and 95; where it is about what exists, 269 and 96.

**The evidence audit.** Six observable signals per company, read from public websites between 20 and 25 August 2026 and scored mechanically. Published in full as an evidence score before this report, so that the arguments would surface before they landed in a PDF.

### What we did not count

We have **not tested any of these products**. Nothing in this report is a judgement about whether a product works. Where we rank, we rank capital, which is disclosed, and disclosure itself, which is observable. We do not rank quality, because we have not measured it and neither, as Part IV shows, has anybody else.

We did not count revenue, because we could not: 14 of 95 companies disclose a dollar figure, across five reporting years and three incompatible measures. Chapter 2 sets out why that makes a revenue ranking impossible rather than merely difficult.

We did not resolve the regulatory picture. The count of Arizona alternative business structure licences, the current status of Utah's sandbox, and authorisations in England and Wales are the largest known holes in this report. They are answerable from primary regulator records and we would rather answer them separately than estimate them here.

Our sourcing is English-language. Latin America appears once and Africa not at all, and we treat that as a limit of our method as much as a fact about those markets.

### How to read the confidence flags

Every funding figure in Appendix A carries a flag.

| Flag | Meaning |
|---|---|
| **high** | Stated by the company or in a primary filing, and consistent across sources |
| **med** | Assembled from announced rounds, or reported by a credible outlet without company confirmation |
| **low** | A range, a ceiling, or a figure we could not reconcile between sources |

A "low" flag is not a suggestion that a company is misleading anyone. Several of them are cases where a company announced a facility rather than a draw — Eudia's "up to $105M" is a ceiling, not a total — and reporting the ceiling as a total would overstate the market.

### Dates

Every figure is true as of **25 August 2026** unless the text gives another date. Where a company's own disclosure is old, we print its date rather than carry it forward: Athennian's revenue figure is from March 2022 and is labelled as stale everywhere it appears.

This market decays fast. Chapter 5 counts eleven of our 95 companies that are not what their name says they are. A dated row with a confidence flag can be corrected; a tidy ranking cannot tell you which of its rows to distrust. The underlying data is published as typed, dated modules at /research/legal-ai-funding and /research/legal-ai-products, and corrections are welcome and will be credited.

### A declared interest

IQdoc builds transcription and record products for legal work, and publishes its own accuracy benchmarks. Three of the 269 products in this catalogue are in that category, and Chapter 18 discusses a measure we compete on. Those passages are marked. We have not included IQdoc in the funding survey, the product catalogue or the evidence audit.

# Part I — The market

Legal AI is now a market with a shape that can be measured rather than described. We tracked **$6.20B** of disclosed capital across **95** AI-native legal companies, catalogued the **269** products they and one company outside the funding survey sell by name, and recorded what each company has been willing to put on the record about its own business. This part reports what that survey found: where the capital sits, what almost nobody discloses, what got built, how many of these companies have stopped being software companies, and how much of the list has already decayed.

Two cautions before the numbers. Every figure is as publicly disclosed and not independently audited; **65 of the 95 funding totals — roughly two-thirds — carry a medium or low confidence flag**, because they were assembled from announced rounds rather than stated by the company. Only 30 are high-confidence. And this is a snapshot taken in August 2026. Chapter 5 is about how fast a snapshot of this market goes out of date.

## 1. Where the money went

Ninety-five AI-native legal companies have raised a combined **$6.20B**. The top five hold **46.4%** of it, and the top ten hold **57.8%**. Below that the market thins out immediately: the median company on the list has raised **$21M**, **51 of the 95** have raised under $25M, and **27** have raised under $10M.

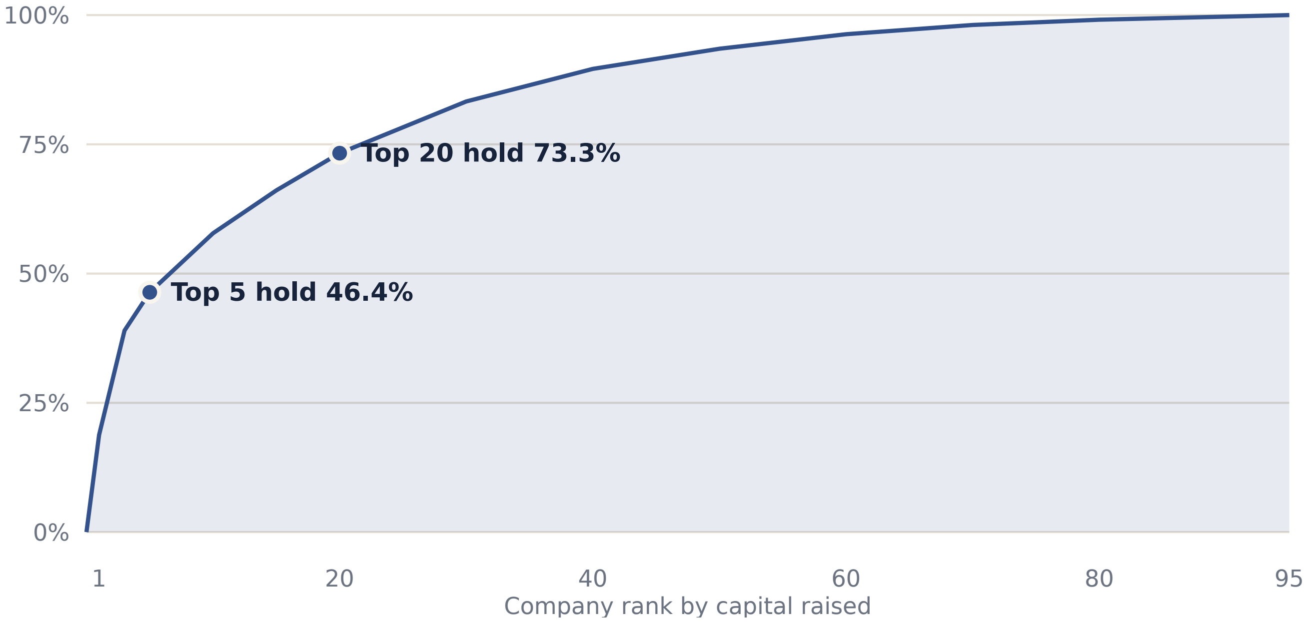


Cumulative share of all tracked capital by company rank

The curve is steep at the top and flat everywhere else. Harvey alone accounts for **18.8%** of all tracked capital. By rank 30 the cumulative share is 83.3%, which means the remaining 65 companies split under a fifth of the money between them. This is not a market of 95 comparably resourced competitors. It is five companies, a second tier of roughly twenty-five, and a long tail operating on seed and Series A money.

Five companies carry a valuation of $1B or more:

| Company | Valuation | Raised | Segment |
|---|---|---|---|
| Harvey | $11.00B | $1.17B | Firm platform |
| Legora | $5.60B | $866M | Firm platform |
| EvenUp | $2.00B | $385M | Litigation & PI |
| Norm Ai | $1.20B | $260M+ | Compliance |
| Eve | $1.00B | $150M | Litigation & PI |

## By segment

| Segment | Companies | Capital |
|---|---|---|
| Firm platform | 4 | $2.16B |
| Contracts & in-house | 32 | $1.62B |
| Litigation & PI | 22 | $918M |
| IP & patents | 7 | $280M |
| Compliance | 2 | $272M |
| Practice ops & billing | 5 | $188M |
| Research & knowledge | 9 | $174M |
| AI-native law firm | 5 | $169M |
| Tax | 1 | $156M |
| Consumer & small law | 3 | $145M |
| Real estate | 3 | $81M |
| Immigration | 2 | $38M |

Two things stand out. **Four companies** in the firm-platform segment hold $2.16B — more than a third of all tracked capital — because Harvey and Legora sit in it. And contracts and in-house is the segment with the most companies by a wide margin, 32 of 95, but its $1.62B is spread thinly enough that its median company looks nothing like its leaders.

The single-company segments are worth reading carefully rather than as market signals. Tax has one company because Blue J is the only tax-first AI-native company that met our funding threshold, not because tax has one product. Compliance has two companies and $272M almost entirely because of Norm Ai.

## By region

| Region | Companies | Capital |
|---|---|---|
| USA | 45 | $3.23B |
| Europe (ex-UK) | 14 | $1.28B |
| UK | 18 | $734M |
| Asia-Pacific | 8 | $451M |
| Canada | 7 | $364M |
| Israel | 2 | $102M |
| Latin America | 1 | $43M |

The United States holds 52% of the capital across 47% of the companies. The more interesting comparison is the UK against continental Europe: the UK has **18 companies to Europe's 14**, but Europe holds $1.28B to the UK's $734M, because Legora and Noxtua are both continental. The UK has more legal AI companies than any country except the US and less capital per company than any major region on the list.

Latin America appears once and Africa not at all. We treat that as a limit of our sourcing as much as a fact about those markets: the survey was built from English-language funding announcements.

### By founding year

| Founded | Companies |
|---|---|
| 2014 | 1 |
| 2015 | 4 |
| 2016 | 2 |
| 2017 | 10 |
| 2018 | 10 |
| 2019 | 6 |
| 2020 | 4 |
| 2021 | 6 |
| 2022 | 7 |
| 2023 | 15 |
| 2024 | 4 |
| 2025 | 8 |
| 2026 | 2 |
| Not established | 16 |

**Fifteen** of the 95 were founded in 2023, the single largest year in the set — the first full year after ChatGPT's release, and the year Casetext sold to Thomson Reuters. That cohort includes Legora, Norm Ai, Wordsmith AI, Lawhive and Eudia, which is to say it includes two of the five unicorns. Companies founded in 2023 have raised more, faster, than any earlier cohort.

The 2017 and 2018 cohorts are equally large at ten companies each, but they are mostly contract and in-house software businesses that predate generative AI and retrofitted it. The apparent decline after 2023 is partly an artefact of recency: a company founded in 2025 has had less time to raise enough to register on a list built from funding.

One limitation we will state plainly rather than smooth over: **for 16 of the 95 companies we could not establish a credible founding date.** That is 17% of the list. Where sources disagreed or gave only a range, we recorded nothing rather than pick one. Laurel, which has raised roughly $100M, is among them. Every percentage in the table above should therefore be read against a denominator of 79, not 95, and the shape of the distribution — including the 2023 spike — could change if those 16 were resolved.

## 2. The disclosure vacuum

Of 95 companies, **14** have put a dollar revenue figure on the record. Twenty-three give a growth rate with no base to apply it to. Twenty-one give a customer count and nothing else. **Thirty-seven — 39% of the list — have disclosed nothing at all**, and some of those are worth more than a billion dollars.

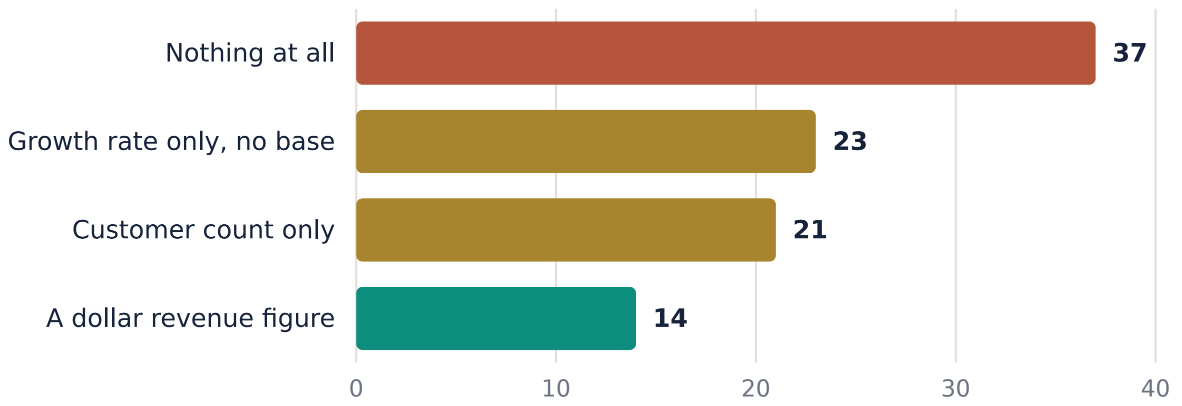


What the 95 companies disclose about their own revenue

A growth rate without a base is not information. "5x year on year" from an undisclosed number is compatible with $200,000 and with $200M. Customer counts are worse, because the units do not compare: one company counts law firms, another counts seats, another counts the end clients of its clients, another counts registered users who may never have paid.

### Why we cannot rank this market by revenue

We wanted to publish a revenue ranking. It is the obvious companion to a capital ranking and it is the number buyers actually want. We could not build one honestly, for four reasons.

First, the sample is too small. Fourteen disclosures out of 95 is not a ranking; it is a list of the companies that chose to speak.

Second, the figures are not contemporaneous. They carry dates ranging from around 2020 to August 2026. **Athennian's ~$10M ARR figure is from March 2022** — more than four years stale at the time of writing, and we have found no update. Comparing it to Harvey's August 2026 figure would be comparing two different companies at two different times.

Third, the figures are not the same measure. Some are ARR, an annualised run rate taken from a recent month. Some are annual revenue for a completed fiscal year — Darrow's $26M is FY2024 revenue, not a run rate. Some are described as "annualised" without saying annualised from what. ARR and audited annual revenue are different quantities and the gap between them widens the faster a company is growing.

Fourth, one of the fourteen is not a number. Solve Intelligence disclosed "eight figures" of ARR, which is a band spanning $10M to $99M. It cannot be ranked against anything.

### The fourteen

| Company | Disclosed figure | As of | Put on the record by |
|---|---|---|---|
| Harvey | $350M+ annualised | Aug 2026 | Press reporting |
| Legora | $100M ARR | Apr 2026 | Company |
| LegalOn Technologies | ¥10B ARR (~$67M) | Oct 2025 | Company |
| Luminance | ~$60M ARR | 2025 | Estimate, not company-stated |
| Lawhive | £25.6M annualised (~$32M) | Feb 2026 | Press reporting |
| Darrow | $26M revenue (FY2024) | FY2024 | Press reporting |
| Eudia | $20M ARR | Dec 2025 | Company |
| MarqVision | $20M+ ARR | Sep 2025 | Press reporting |
| Lexroom | €10M+ ARR (~$11M) | Mar 2026 | Company |

| Company | Disclosed figure | As of | Put on the record by |
|---|---|---|---|
| GC AI | $10M+ ARR | Nov 2025 | Company |
| Athennian | ~$10M ARR | Mar 2022 — stale | Press reporting |
| Solve Intelligence | "eight figures" ARR, profitable | Dec 2025 | Company |
| Juro | $1M ARR, historic | ~2020 | Unattributed |
| midpage | Six-figure annual revenue | Feb 2025 | Unattributed |

Two entries deserve a flag. Luminance's ~$60M is an **estimate** carried in reporting rather than a company statement, and we have marked it as such rather than promote it. For Juro and midpage we could not establish who put the figure on the record; both appear in circulation without a clear attribution, and both are old enough that they describe companies that no longer exist in that form.

Solve Intelligence is the only company in the set to describe itself as profitable. We have not verified that and cannot.

### Capital against revenue

Twelve companies disclose both a capital total and a revenue figure, which is the largest set on which any efficiency comparison can be attempted at all.

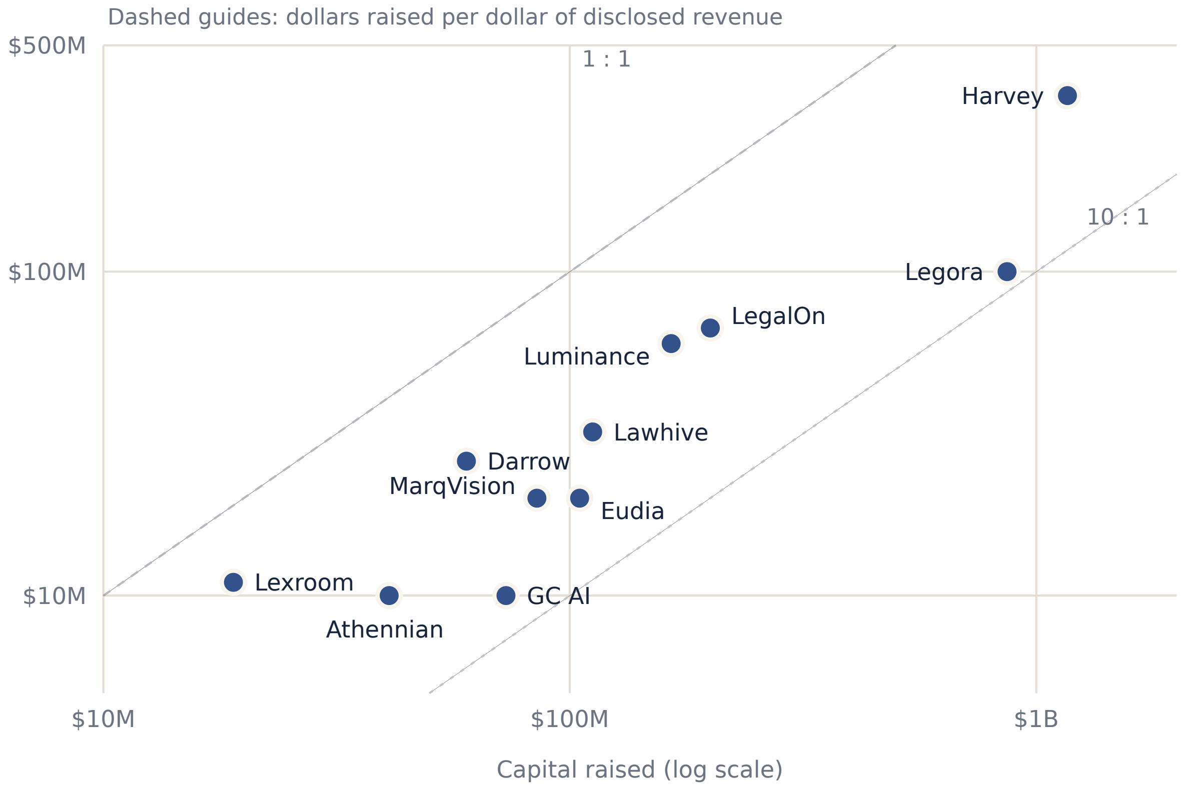


The twelve companies that disclose both capital raised and revenue

We computed a capital-to-revenue multiple only where both sides are firm, current numbers rather than estimates, bands or figures from a closed fiscal year. That leaves three: **Harvey at 31x**, **GC AI at 55x** and **Legora at 56x**. Read those as what they are — dollars raised per dollar of current annualised revenue — and not as a valuation judgment. A company that raised at a high price and has not yet spent the money will look identical to one that spent it badly.

The honest summary of this chapter is that the legal AI market has no public revenue data worth the name. Eighty-one of 95 companies have never stated what they earn, the fourteen that have cannot be compared to each other, and the market's aggregate revenue is therefore unknown to us and, as far as we can tell, to everyone else publishing about it.

## 3. What got built

The 95 companies and one company outside the funding survey market **269 products** by name, across **18 types**.

A scope note first, because the two totals in this report differ by five. Our funding survey covers 95 companies, which between them sell **264 products**. Courtroom5 sells a further **five**, and it is not in the funding survey because it has raised too little to meet the threshold the survey was built on. We include its products in the catalogue anyway, because excluding the most substantial software product built for self-represented litigants on the grounds that nobody funded it would be a filter reporting on itself. Where a figure in this report is about the funded market, it uses 264 and 95. Where it is about what exists, it uses 269 and 96.

| | In-house | Law firms | Independent | Pro se |
|---|---|---|---|---|
| Agent platforms & infrastructure | 27 | 20 | 5 | 1 |
| Drafting & document generation | 22 | 18 | 5 | 1 |
| Legal services delivery | 10 | 4 | 1 | 13 |
| Contract review & negotiation | 24 | 8 | 1 | 0 |
| Knowledge search & firm memory | 15 | 14 | 0 | 0 |
| Research & citation | 14 | 16 | 9 | 0 |
| Litigation case prep & discovery | 11 | 12 | 2 | 3 |
| IP, patents & brand protection | 15 | 10 | 0 | 0 |
| Time, billing & practice ops | 2 | 14 | 0 | 0 |
| Case intake & triage | 6 | 8 | 2 | 1 |
| Contract lifecycle management | 13 | 0 | 0 | 0 |
| Regulatory & compliance monitoring | 9 | 4 | 0 | 0 |
| Client communication & portals | 3 | 8 | 2 | 0 |
| Due diligence & deal work | 8 | 7 | 1 | 0 |
| Judicial & litigation analytics | 6 | 7 | 1 | 0 |
| Medical records & demand packages | 0 | 6 | 6 | 0 |
| Entity, governance & deal admin. | 4 | 2 | 0 | 0 |
| Transcription & the record | 0 | 3 | 3 | 0 |

269 products by type, split by who they are sold to

| Product type | Products |
|---|---|
| Agent platforms & infrastructure | 33 |
| Drafting & document generation | 26 |
| Legal services delivery | 25 |
| Contract review & negotiation | 24 |
| Knowledge search & firm memory | 20 |
| Research & citation | 18 |
| Litigation case prep & discovery | 17 |
| IP, patents & brand protection | 15 |
| Time, billing & practice ops | 14 |
| Case intake & triage | 14 |
| Contract lifecycle management | 13 |
| Regulatory & compliance monitoring | 11 |
| Client communication & portals | 10 |
| Due diligence & deal work | 8 |
| Judicial & litigation analytics | 8 |
| Medical records & demand packages | 6 |
| Entity, governance & deal administration | 4 |
| Transcription & the record | 3 |

**Agent platforms and infrastructure is the largest single type at 33 products** — 32 of them from the funded 95, plus Courtroom5's MCP connector. It is larger than contract review and negotiation, which is the category the previous decade of legal technology was built around, and larger than research and citation, which is the category most lawyers still name first when asked what legal AI does. A market that spent ten years selling document review has, in eighteen months, made its most crowded category the thing that runs multi-step tasks.

The type count understates it. Counted by function rather than by product type, 50 of the 264 funded products are agents or agent infrastructure: the 32 typed as such, plus 18 typed elsewhere that plan and execute multi-step work. Part III takes that apart properly.

The thin categories carry as much information as the crowded ones. **Three products** address transcription and the record. **Four** address entity and governance records. **Six** address medical records and demand packages, in personal injury — the practice area that has attracted more capital than any other single area of contentious work. We have an interest in the first of those three lines, since the record is what we build, so treat it as an interested observation. It is still three products out of 269.

## 4. The structural shift

**Nineteen of the 95 companies sell legal work rather than software.** Between them they account for the 25 products typed as legal services delivery: not a tool a lawyer operates, but a licensed firm doing the work with AI behind it, sold at a price.

They arrived by two different routes, and the difference matters.

The first route is companies built as firms from the start. **Crosby** reviews contracts at a fixed fee and carries malpractice insurance. **Keith** does English residential conveyancing at a published price — £1,200 to buy a house. **Lawhive** sells flat-fee consumer legal work and has disclosed £25.6M of annualised revenue, higher than all but four of the fourteen companies that disclose anything. **Moritz** publishes a full fee menu and carries more security disclosure than most software vendors on our list. These are law firms founded by technologists.

The second route is newer. Software companies have bolted a law firm onto the side of the software business. **Norm Ai** sells compliance software and also runs Norm Law. **Eudia** sells an in-house platform and also runs Eudia Counsel. **EvenUp** sells case software to plaintiff firms and also runs the pre-litigation department outright. **Aavalynx** sells dispute analytics and operates a sister firm that competes with the firms buying the analytics. **Justpoint** went furthest and made the split formal: in July 2025 it separated into Justpoint Inc., which disclaims providing legal services, and **Justpoint Law LLP**, an Arizona-licensed law firm.

### The mechanism

In the United States the enabling mechanism is Arizona's alternative business structure licence, which permits non-lawyer ownership of a law firm. Most US states still bar non-lawyers from owning or sharing fees with a firm, so the same corporate structure that is licensable in Arizona is unavailable across most of the country. Justpoint Law LLP was approved under the Arizona regime in July 2025, and our earlier post identified Boundless, Justpoint and Eudia as holding a licence.

The economic logic is set out in legal tech is not the market: if AI does most of the labour, selling software into a legal technology market forecast at $69.7B in 2033 is a smaller business than capturing revenue inside a legal services market already worth about $1.16 trillion. Arizona made the second option available to a company with outside shareholders.

### What we could not verify

We are more confident about the direction than about the detail, and the detail is where this section has to stop.

We have **not** verified how many alternative business structure licences Arizona has issued, to whom, or how many are still active. Published counts circulate and we could not reconcile them against the regulator's own record, so we state none. We have not verified how many of the 19 companies hold a licence in their own name, as against operating through a licensed firm owned by someone else or practising only in jurisdictions where the question does not arise. We have not established the current status of Utah's regulatory sandbox, nor counted alternative business structure authorisations granted by the Solicitors Regulation Authority in England and Wales, where the mechanism is older and the route Lawhive, Keith and Moritz operate under.

Those four gaps are the largest known holes in this report. Each is answerable with primary regulator records, and we intend to answer them separately rather than estimate them here.

## 5. Exits, failures and renames

Twelve companies in or adjacent to this market have been acquired since 2023, three of them at prices that reset what the category is worth.

| Acquired | Buyer | Price | Date |
|---|---|---|---|
| vLex | Clio | $1.0B | Nov 2025 |
| Casetext | Thomson Reuters | $650M | Aug 2023 |
| Lexion | Docusign | $165M | May 2024 |
| Evisort | Workday | Undisclosed | Nov 2024 |

| Acquired | Buyer | Price | Date |
|---|---|---|---|
| Rowan Patents | Clarivate | Undisclosed | Jul 2024 |
| Compliance.ai | Archer | Undisclosed | Feb 2024 |
| Dioptra | Icertis | Undisclosed | Nov 2025 |
| Pincites | Filevine | Undisclosed | Dec 2025 |
| Robin AI (services arm) | Scissero | Undisclosed | Dec 2025 |
| Qura | Legora | Undisclosed | Apr 2026 |
| Avantia | Carta | Undisclosed | May 2026 |
| Wexler AI | Legora | Undisclosed | Jul 2026 |

**Casetext to Thomson Reuters at $650M** in August 2023 is the deal that opened the market: it established that an incumbent would pay serious money for an AI-native legal product rather than build one. **vLex to Clio at $1.0B** in November 2025 is the largest, and brought global legal content, the Vincent AI research assistant and Docket Alarm under a practice-management company founded in 2008. **Lexion to Docusign at $165M** is the clearest example of the pattern underneath both: buy the AI-first product rather than retrofit the incumbent one. Evisort, which had raised roughly $155M, went to Workday for an undisclosed sum in November 2024.

Two acquirers are worth watching for opposite reasons. Icertis, which has raised $520M for enterprise contract intelligence, bought Dioptra rather than build AI contract review. And **Legora bought two of the companies in this table within three months** — Qura, a Swedish legal research company that had raised €2.1M, in April 2026, and Wexler AI, a litigation fact-verification product, in July 2026, ten months after Wexler's seed round. Wexler was Legora's **fifth acquisition since March 2026**. A company that raised $866M is spending part of it on other companies, and folding them in as layers rather than reselling them.

### Failure looks like a live website

**Robin AI** raised $61.5M and is finished. It wound down through 2025; the managed-services arm was sold to the law firm Scissero in December 2025 and the engineering team was acqui-hired by Microsoft in January 2026. The website still shows product marketing. Nothing on it says the company is gone. A buyer evaluating contract AI in 2026 could run a full diligence process against a company that no longer exists, and the site would not tell them.

**Klarity** raised around $90M, renamed itself **Within**, and left legal entirely — repositioned from contract review to enterprise process discovery, selling to finance, IT, go-to-market and operations. There is no legal product on the site.

**LawGeex** is a dead brand. LegalSifter acquired the enterprise business in 2023 and the successor product is LegalSifter ReviewPro; lawgeex.com now serves an expired TLS certificate. **Della AI** was acquired by Wolters Kluwer in December 2022 and the brand retired — the technology ships as Legisway Analyzer and the della.ai domain is listed for sale.

### Four renames that break an old list

| Was | Is now | When |
|---|---|---|
| Callidus Legal AI | StrongSuit | Nov 2025 |
| NLPatent | Clerq | Aug 2026 |
| Pocketlaw | Miramis | — |
| Xayn | Noxtua | 2024 |

All four companies are active and shipping. Callidus redirects to StrongSuit, nlpatent.com redirects to clerq-ip.com, and Pocketlaw's rebrand to Miramis came with a repositioning from SME self-serve templates to mid-market contract lifecycle management. Noxtua — founded as Xayn, renamed in 2024 — has raised around $110M. It is not the largest company in the survey trading under a name it did not start with: **Legora was founded as Leya** and **LegalOn Technologies as LegalForce**, and between them they have raised $1.07B. Others in the survey have changed names earlier and less visibly: Tomorro was Leeway, Flank was Legal OS until May 2024, and Orbital was Orbital Witness, with the old product name still surviving in its app URLs.

Two further entries are neither dead nor renamed but should be read with care. **Fileread** is operating but has raised nothing since its 2023 seed and shows no visible activity through 2026. **Soxton** is a landing page with no product or pricing detail; everything we know comes from its December 2025 funding announcement.

### The point

**Eleven of our 95 companies are not what their name says they are.** Four are dead brands or departures — Robin AI, LawGeex, Della AI and Klarity — and seven trade under a name they adopted after founding. A list of this market built in August 2025 would, by August 2026, carry a defunct company as live, an ex-legal company as legal, and four companies under names that no longer resolve.

This is the argument for publishing the underlying data with dates and confidence flags attached rather than a tidy ranking. We will get things wrong in the same way. The difference is that a dated row can be corrected, and a static list cannot tell you which of its rows to distrust.

# Part II — The four audiences

Every other survey of this market is organised by product category: contract review here, research there, agents in a box at the end. That structure hides the finding. A contract review tool sold to a Fortune 500 legal department and a contract review tool sold to a sole practitioner are not two examples of one category; they are two different businesses, funded at different scales, priced on different logic, and held to different standards of proof. Sorting by what a product does puts them in the same row. Sorting by who is expected to pay for it puts them where they belong.

So this report is organised by buyer. We assigned every one of the 269 products in the catalogue to the audience or audiences its own marketing addresses, then carried each company's disclosed funding across to every audience it sells to. That produces four columns — corporate legal departments, law firms, independent lawyers, and people handling a legal problem without one — and the gap between the first and the last is the largest single number in this report.

One instruction on how to read the table below, because it can be read wrong. **A company selling to two audiences is counted in both rows.** Harvey appears under law firms and under in-house. Legora appears under both. The capital column therefore does not sum to the $6.20B raised across the 95 companies in the funding survey — added up it comes to $11.02B, which is a double-count and not a market size. The same applies to the company counts: 74 plus 59 plus 16 plus 8 is 157 audience positions held by 95 funded companies. Read each row against the others, never as a share of a whole.

| Audience | Capital behind it | Companies | Products | Unicorns |
|---|---|---|---|---|
| In-house / corporate | $5.10B | 74 | 189 | 3 |
| Law firms | $4.50B | 59 | 161 | 4 |
| Independent lawyers | $1.22B | 16 | 38 | 2 |
| People without lawyers | $0.20B | 8 (9 with Courtroom5) | 19 | 0 |

**26 times more capital stands behind products built for corporate legal departments than behind products built for people without a lawyer** — $5.10B against $196M. The pro se column contains no billion-dollar company at all, and only one of its eight companies, Lawhive at roughly $112M, reaches the top twenty of the funding survey — where it accounts for more than half the entire pro se column on its own. Its ninth company, Courtroom5, is in the product catalogue but outside the funding survey, having raised too little to meet the threshold; its capital is not in the $196M.

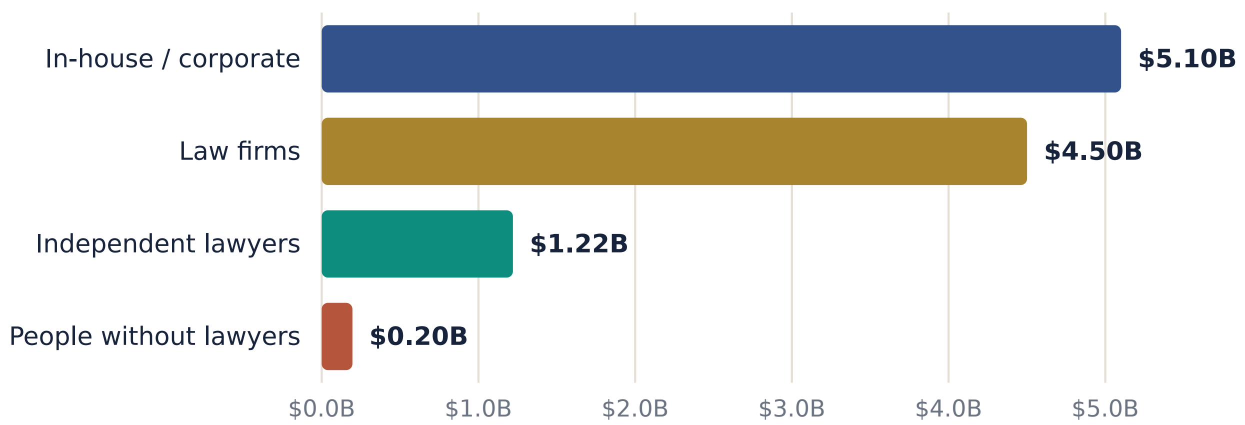


Capital behind the products sold to each audience

The four chapters that follow take one audience each and answer the same six questions in the same order, so that the answers can be set side by side: what exists, what it costs, what is missing, who is credible, what the evidence picture looks like, and what a buyer in that segment should ask before signing.

## 6. In-house and corporate legal

Corporate legal departments are where the money is. **$5.10B** of disclosed capital stands behind products sold to in-house teams, across **74** of the 95 funded companies and **189** of the 269 products in the catalogue — the largest count on every measure we hold. Three of the five unicorns in the survey sell here: Harvey at an $11.00B valuation, Legora at $5.60B and Norm Ai at $1.20B.

### What exists

The type mix is the clearest statement of what investors believe corporate legal work is made of.

| Product type | In-house products |
|---|---|
| Agent platforms & infrastructure | 27 |
| Contract review & negotiation | 24 |
| Drafting & document generation | 22 |
| Knowledge search & firm memory | 15 |
| IP, patents & brand protection | 15 |
| Research & citation | 14 |
| Contract lifecycle management | 13 |
| Litigation case prep & discovery | 11 |
| Legal services delivery | 10 |
| Regulatory & compliance monitoring | 9 |

Contracts dominate. Add contract review (24), contract lifecycle management (13) and the contract-shaped half of drafting, and the single commercial function of getting an agreement negotiated and signed accounts for more in-house products than any other activity a legal department performs. Agent platforms sit at the top of the table at 27, but that number describes a delivery mechanism rather than a job: what those agents mostly do is contracts, drafting and search.

The breadth is as notable as the mix. **Seventy-four of the 95 funded companies sell to in-house teams** — close to four in five — at an average of 2.6 catalogued products each. In-house is not a segment of this market so much as its default assumption: a company can be built for firms, for independents or for consumers, but almost every company built for anyone is also pointed at the corporate legal department, because that is where a budget line already exists and a procurement process already runs.

One line in the catalogue is a cleaner audience signal than anything else we measured. **Contract lifecycle management is 13 in-house products, 0 law firm products, 0 independent lawyer products and 0 pro se products.** No other type is exclusive to a single audience. CLM is the corporate legal department's own category — a system of record for the contracts a company signs, of no use to a firm that advises on other people's contracts and never holds the portfolio. When a vendor describes itself as CLM, it has told you its buyer before it has told you anything else.

### What it costs

Seventeen of the 74 companies selling to in-house teams publish a price. That is **23%**, the third-highest rate of the four audiences, and it is still a minority small enough that most corporate buyers begin every evaluation without a reference point.

The published ones are worth naming because they show the range. Zuva is the only true per-unit publisher in the survey at **$10 per document** through the product and **$1.25 per document** through the API, with a working free tier. GC AI publishes **$500 a month** for an individual plan. Streamline AI publishes "from **$22,900** a year". Athennian publishes a **$25,000** entry tier, with AI as a paid add-on on top — candour that is rare at enterprise price points, as we noted in the evidence score.

Against that, the pricing pages that are not pricing pages cluster in this segment. LegalOn names three tiers and prices none of them. Juro's pricing page is a quote calculator containing no figures. Pocketlaw — now Miramis — runs a qualifying questionnaire ending at "Get a demo". Across the whole survey, companies that publish a price have raised **$30M** on average and companies that publish none have raised **$76M**. The wall goes up as the cheque size does, and the corporate segment is where the cheques are largest.

### What is missing

Three absences stand out in the in-house column. **Time, billing and practice ops is 2 products** for in-house against 14 for firms — legal departments do not bill, so nothing here helps them account for what their own work costs. **Medical records and demand packages is 0** and **transcription and the record is 0**: every product pointed at the evidentiary record is sold to firms or to independents, none to the corporate department that will be sued. **Client communication and portals is 3**, which is a thin showing for a function whose principal complaint is that the business treats it as a queue.

The larger absence is supervisory. Twenty-seven in-house products are agent platforms, and across all 269 products in the catalogue **exactly one supervises other agents** — Norm Ai's Supervisory AI. The segment buying the most autonomy has almost nothing available to watch it.

### Who is credible

By our published rubric, the in-house-facing companies that let a buyer check the most are GC AI (**7 of 10** — published price, named models, participation in two phases of an independent benchmark), Wordsmith AI (**5**, for naming the contractual basis of zero data retention with each model provider), Genie AI (**5**, which publishes its ISO certificate number and expiry and states plainly that it does not hold SOC 2) and Blue J (**5**, at $1,498 a year per user with named models). Zuva, Streamline AI and Lexroom score **4**.

At the other end of the same segment, Norm Ai — a $1.20B compliance company selling to institutional finance — scores **1**: SOC 2 named, no model disclosure, no price, a closed benchmark. Nothing in that score is a judgement about the product. It is a statement about how much of the claim is checkable from outside.

### The evidence picture

In-house-facing companies score a mean of **2.15 out of 10** on whether a buyer can verify what they claim, against a field mean of 1.99. It is the second-best audience score in the survey and it is still barely a fifth of the available evidence. Only **17 of 74** publish a price.

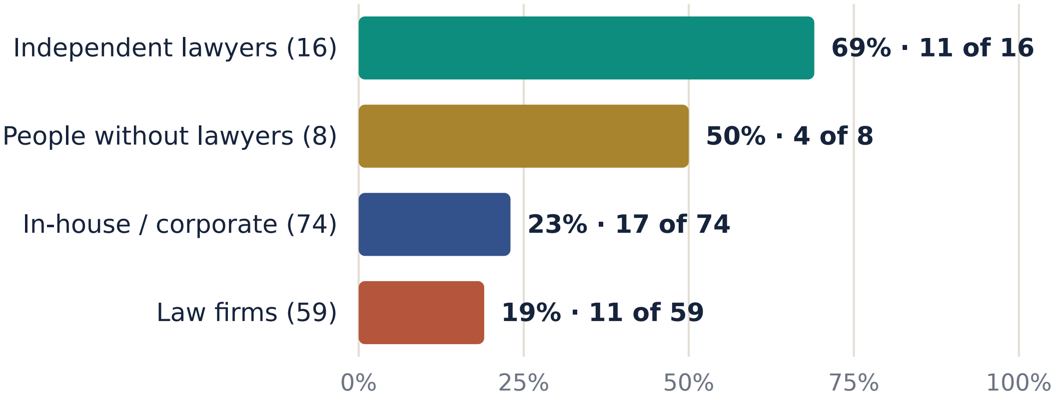


Whether companies publish a price, by who they sell to

The rest of the picture is the field-wide one and it applies here in full. Of all 95 companies, 45 will not say which models they run on and 36 offer only a vague pledge. Thirty publish a benchmark score with no method attached; two publish tasks and rubrics. Six have participated in an evaluation run by someone else. **None of the 95 has independent research published measuring its product.** A general counsel evaluating a contract review tool in August 2026 has, from public sources, no third-party accuracy figure available for any product in the category.

### What a buyer should ask

These are the questions the data above says a vendor should be able to answer, and that most of this segment currently does not answer in public.

1. **What is the unit, and what is the number?** Not the annual contract value — the unit. Zuva publishes $10 a document; ask what yours is per document, per matter or per seat, and ask for it in writing. Seventeen of 74 companies selling to you publish this at all, so treat a refusal as normal and a published figure as a signal.
2. **Which models, running where, under what retention terms?** Forty-five of 95 companies are silent on this and 36 give a pledge without a provider name. Ask for the answer DeepIP and Wordsmith AI give in public: named provider, named hosting arrangement, and the contractual basis for zero retention.
3. **Has anyone outside your company measured this?** Six of 95 have taken part in someone else's evaluation. Ask which round, on what date, and where the published result sits. If the answer is a self-run benchmark, ask whether the tasks and rubrics are open — 30 companies publish a score with no method, and a score with no method is not a measurement you can use.
4. **If you sell me CLM, what happens to the repository if I leave?** CLM is the one type sold exclusively to in-house teams, which means the system of record for every contract your company signs lives inside a vendor with no competing audience. Ask for the export format and the termination terms before the pilot, not after.
5. **Who supervises the agent?** One product in 269 supervises other agents. Ask what is logged, who reviews an agent's output before it reaches a counterparty, and what the vendor's own answer is when the agent is wrong.

## 7. Law firms

Law firms are where the prestige sits and where the valuations are made. **$4.50B** stands behind firm-facing products, across **59** companies and **161** products. Four of the five unicorns in the survey sell to firms — Harvey ($11.00B), Legora ($5.60B), EvenUp ($2.00B) and Eve ($1.00B) — one more than sell to corporate departments, on less capital and fewer products. Firms are the smaller market and the more valuable one.

### What exists

The firm segment is the most concentrated thing in the market. **Firm platform is 4 companies and $2.16B** — 35% of all capital raised across the 95 companies, held by four businesses selling a single category. Three of them account for essentially the whole figure: Harvey at $1.17B raised, Legora at $866M and Noxtua at roughly $110M. Nothing else in the catalogue concentrates like that; by comparison, contracts and in-house is 32 companies sharing $1.62B.

Below the platforms, the firm product mix diverges sharply from the corporate one. **Time, billing and practice ops is 14 firm products against 2 in-house** — the clearest structural difference between the two audiences, and a direct consequence of the billable hour existing on one side of the relationship only. Research and citation is 16, knowledge search and firm memory is 14, litigation case prep is 12, agent platforms 20 and drafting 18. Medical records and demand packages is 6 firm products and 0 in-house, all of it pointed at personal injury — the practice area that has absorbed more capital than any other, at $918M across 22 litigation and PI companies.

Contract review runs the other way: **8 firm products against 24 in-house**. The work of reviewing a contract has largely been sold to the company signing it rather than the firm advising on it.

And some of the capital in this chapter is not aimed at serving firms at all. The AI-native law firm segment is **5 companies and $169M**, led by Crosby at $85.8M raised, and legal services delivery accounts for 10 in-house products against 4 firm ones — work sold directly to the corporate client, with a licensed firm behind it, as we set out in one in five legal AI companies now sells legal work. A firm evaluating this market is therefore evaluating two things at once: vendors selling it tools, and vendors selling its clients the outcome those tools produce. The second group is small today. It is the group whose economics improve fastest if the tools work.

### What it costs

**Eleven of 59** firm-facing companies publish a price. At **19%**, this is the least transparent of the four audiences — lower than in-house, and far below the 69% of companies selling to independent lawyers. That ordering is worth sitting with. Law firms are the buyers most likely to have a procurement function, a panel process and a partner who negotiates for a living, and they are the audience least likely to be shown a number before the call.

The firm-facing prices that do exist are mostly at the small end. Clearbrief publishes **$300 a month**. Trellis publishes tiers from **$69.95 to $199.95 a month**. Skribe.ai publishes **$379 per record hour** for a deposition, the most transparent unit price in litigation services. Moritz publishes a full fee menu. The large firm platforms publish nothing: Harvey, which scores at the top of our evidence table on every other signal, does not publish a price.

At the practice-ops end, the pattern is a pricing page that resolves to a form. Hona's navigation item labelled "Pricing" opens a demo booking form. Laurel, which raised roughly $100M in a single 2025 round, publishes no price. PointOne publishes none.

### What is missing

The firm column has one true zero: **contract lifecycle management, 0 products**. That is the mirror of the in-house finding and it is correct — firms do not hold the portfolio. Beyond it, regulatory and compliance monitoring is 4 firm products against 9 in-house, and entity, governance and deal administration is 2 against 4.

Transcription and the record is **3 products in the entire catalogue**, all from one company, of which all three reach firms. In the audience that litigates, the document a court treats as what was said is the least-built category in the market. We have an interest in that observation — it is the category IQdoc works in — so treat it as an interested one and check the count.

The other missing thing is measurement. Firm platform is the best-scoring segment in the evidence table with more than one company in it, at a mean of **3.50** — and 3.50 out of 10 is the ceiling of the most credible category in the market. Only Tax scores higher, and Tax is one company.

### Who is credible

Harvey scores **7 of 10**, the joint-highest in the survey, and is the only company with both an open benchmark and independent-benchmark participation. It open-sourced its task set and pointedly published no score for its own product. It also publishes no price, which is why a company of its size sits at 7 rather than higher.

Legora scores **3**: named models, named certifications, a closed benchmark. It declines to open its benchmark corpus and states the reason — published cases leak into the next model's training data — which is a real objection and we record it as one. Clearbrief scores **5**, including the closest thing to independent assessment in the segment: 40.5 out of 50 from the State Bar of Nevada's AI Workgroup. Alexi scores **3** for participating in the Vals legal research study of October 2025, where it landed at the participant average and above the lawyer baseline. Among the firm-facing unicorns, EvenUp scores **2** and Eve scores **0**.

### The evidence picture

Firm-facing companies score a mean of **2.03 out of 10**, and **11 of 59** publish a price. Both figures sit below the in-house segment. The buyers with the most procurement machinery are being given the least public evidence, which suggests the machinery is absorbing the disclosure privately — pricing and audit reports handed over on request, under NDA, one buyer at a time. That is a functioning sales motion. It is not a market that can compare anything.

Sitting above this segment is a platform risk that none of the evidence signals capture. Google's Gemini Enterprise ships an integration list for legal work that includes Docusign, Everlaw, iManage, NetDocuments, RelativityOne, Thomson Reuters, Harvey and Legora. Two of those are the firm platforms holding $2.03B of the $2.16B in this segment. Being an integration inside somebody else's enterprise agent is a materially different business from owning the window the lawyer opens each morning: the integration supplies a capability and the platform keeps the relationship, the usage data, the renewal and the pricing power. We cannot measure how that resolves. We can observe that the most concentrated segment in legal AI is four companies, and that the surface they depend on is being assembled by firms that do not need them to be independent.

### What a buyer should ask

1. **Will you put the unit price in the contract, and the renewal uplift?** Eleven of 59 vendors selling to you publish a price at all. You have procurement; use it to obtain in writing what the market cannot obtain in public, and ask for the cap on year-two and year-three increases before the pilot succeeds.
2. **Am I buying the window or a tile inside someone else's?** Ask where the product sits if the firm standardises on an enterprise agent platform, and what the vendor's roadmap looks like as an integration rather than as the interface. Harvey and Legora are already named integrations in Gemini Enterprise.

3. **Show me the certificate, not the badge.** Ask for the certificate number and expiry, as Genie AI publishes them. The survey contains "SOC 2 Type 3", which does not exist, "SOC 2 Type II observant", which is not an audit status, and a homepage claiming ISO 27001 certification against a security page claiming ISO 27001 alignment.
4. **Which evaluation have you been in, run by whom?** Six of 95 companies have participated in an evaluation run by someone else and none of the 95 has independent research published about its product. Ask for the round name and date. If the vendor cites its own benchmark, ask whether the tasks and rubrics are open — only Harvey and Paxton AI publish theirs.
5. **What leaves with us?** Fourteen firm products index the firm's own knowledge and memory. Ask what the extraction format is, what happens to the index at termination, and which models saw the documents on the way in — 45 of 95 companies will not name their providers at all.

## 8. Independent lawyers and small firms

**$1.22B** stands behind products sold to independent lawyers, across **16** of the 95 funded companies and **38** of the 269 products. Two unicorns sell here — EvenUp at $2.00B and Eve at $1.00B — and their raised totals of $385M and $150M account for **$535M** of the $1.22B, 44% of the column.

The number that matters in this chapter is not the capital. It is that independent lawyers are the only audience routinely told the price.

### What exists

Roughly half of American lawyers in private practice are solo practitioners, and close to seven in ten are in firms of ten or fewer. What the boom has built for them is **38 products, of which exactly one lists independent lawyers as its only audience** — Lawrence, Lawhive's AI paralegal, which is not purchasable. It is internal tooling for solicitors inside Lawhive's own network.

Of the other 37, **33 — 89% — are also sold to law firms**, as we set out in what the legal AI boom built for solo lawyers. There is no solo legal AI market in the sense of companies designing for a solo practice. There is firm software that accepts a smaller credit card. The type mix shows which of it makes the trip down.

| Product type | Independent lawyer products | Law firm products |
|---|---|---|
| Research & citation | 9 | 16 |
| Medical records & demand packages | 6 | 6 |
| Agent platforms & infrastructure | 5 | 20 |
| Drafting & document generation | 5 | 18 |
| Transcription & the record | 3 | 3 |
| Litigation case prep & discovery | 2 | 12 |

What travels is horizontal. Research, drafting and the agent layer over both do not care how many partners a firm has, so they price down without being redesigned. What does not travel is anything shaped around how a practice runs.

**Medical records and demand packages is the one substantive parity case: 6 products for independents against 6 for firms**, the only type where an independent lawyer is served as thoroughly as a firm. The reason is the fee structure. Plaintiff personal injury runs on contingency, so a solo doing injury work has cases that pay, which makes that solo worth building for. EvenUp, Supio and Eve all sell down-market willingly, and litigation and PI is the best-capitalised practice area in the survey at **$918M** across 22 companies. Transcription matches at 3 against 3, but all three come from one company, and it is the category IQdoc works in, so treat that observation as an interested one.

### What it costs

**Eleven of the 16 companies selling to independent lawyers publish a price — 69%**, against 23% for in-house and 19% for law firms. It is the highest rate in the survey by a distance.

| Company | Published price |
|---|---|
| midpage | $30 and $80 a month |
| Jus Mundi | $29 academic, $129 practitioner |
| Genie AI | Free tier, then $75 a month |
| Lexroom | EUR 75 a month per module |
| Callidus Legal AI (now StrongSuit) | $249 a month |
| Clearbrief | $300 a month |
| Skribe.ai | $379 per record hour |
| Paxton AI | $499 a month per seat |
| GC AI | $500 a month |
| Blue J | $1,498 a year per user |
| Lawhive | Fixed fees per matter, as a regulated firm |

midpage's $30 is the lowest published price anywhere in the 95. A solo can evaluate and buy every product on that list on a Sunday night without speaking to a salesperson, which was not true of legal research five years ago. Two further published prices sit just outside the cohort and show the same shape: Twin1 at **$50 and $150 a month, priced per person rather than per firm**, and Trellis at **$69.95 to $199.95 a month**. Neither sells to independent lawyers by our tagging, and both price as though the buyer were one.

The mechanism is worth arguing rather than asserting, because the obvious reading — that these are more honest companies — is not supported by the data. Several of the numbers above are explicitly the individual tier: GC AI's $500 is its individual plan, Blue J's $1,498 its individual tier, Callidus's $249 its individual price. The same companies publish nothing about what a legal department pays. The disclosure is not a policy; it is audience-specific.

What distinguishes the audience is the shape of the purchase. An independent lawyer pays out of their own pocket, holds no procurement function and no panel process, and will not spend forty minutes on a discovery call to learn whether a $50 tool is a $50 tool. A vendor selling there cannot recover the cost of that call from the contract value, and a buyer who cannot see a number does not start. **The published price is a condition of the sale rather than a disclosure choice**. That is why the same companies go silent one audience up, where a salesperson is affordable and procurement extracts the number privately — and why pricing transparency should be read as a signal of who a product is for rather than as a proxy for candour.

### What is missing

Six product types have **zero** independent-lawyer products, and the list is the chapter's argument in one line.

| Type with zero independent products | Firm products | In-house products |
|---|---|---|
| Time, billing & practice ops | 14 | 2 |

| Type with zero independent products | Firm products | In-house products |
|---|---|---|
| Knowledge search & firm memory | 14 | 15 |
| IP, patents & brand protection | 10 | 15 |
| Contract lifecycle management | 0 | 13 |
| Regulatory & compliance monitoring | 4 | 9 |
| Entity, governance & deal administration | 2 | 4 |

CLM is explicable — nobody outside a corporate legal department holds a contract portfolio. The others are not. **Time, billing and practice ops is 14 firm products and 0 for independents**, the sharpest inversion we found: a solo bills, chases and writes off exactly as a firm does, with nobody to delegate it to, and the category built for that work has not been offered to them at all. Knowledge search is the same shape at 14 and 0. The horizontal tools came down-market; the operational ones stayed where the seat counts are.

The second absence is by practice area. Solo and small-firm lawyers concentrate in family law, criminal defence, immigration, estate planning, landlord-tenant and employment — the work that fills American courtrooms. For all of it the catalogue returns research tools that happen to cover the subject and little else: nothing for a family lawyer's motion practice, nothing for criminal defence discovery, nothing for the landlord-tenant docket. That last hole appears again in the next chapter, from the other side of the same courtroom.

### Who is credible

This cohort contains the highest scorer in the survey and one of its two open benchmarks.

| Score | Company | What earned it |
|---|---|---|
| 7 | GC AI | Price published, models named, two phases of an independent benchmark |
| 5 | Blue J | Price published, models named |
| 5 | Clearbrief | Price published; scored 40.5/50 by the State Bar of Nevada AI Workgroup |
| 5 | Genie AI | Certificate number and expiry published; states plainly it has no SOC 2 |
| 5 | Paxton AI | Ran Stanford's hallucination benchmark on itself and released the method |
| 4 | Callidus (StrongSuit), Lexroom, midpage | Price published, certifications named, models described vaguely |

Paxton AI is one of only **2 of 95** companies publishing a benchmark with tasks and rubrics open; the other is Harvey. That a small-practice research vendor is one of the two is not what the capital table would predict.

At the bottom of the same audience sit both of its unicorns. **EvenUp, valued at $2.00B, scores 2. Eve, valued at $1.00B, scores 0** — the only zero here, earned by claiming to follow "the SOC 2 Framework", a framework claim rather than an attestation. Neither score says anything about whether the products work; we have not tested them. They say the two most valuable companies selling to independent lawyers disclose less about themselves than a company charging $30 a month.

### The evidence picture

**Independent-lawyer-facing companies score a mean of 3.38 out of 10, the best of the four audiences** — against 2.15 for in-house, 2.03 for law firms, 1.13 for people without lawyers, and 1.99 across the whole field.

That result needs one qualification, and it is arithmetic rather than opinion. Publishing a price is worth 2 of the 10 points in our rubric, and 11 of these 16 companies publish one. Strip the price signal out of every score in the survey and the independent cohort's mean falls to **2.00** while the field mean falls to **1.54** — the segment keeps a lead, but roughly two-thirds of it was the price signal alone. On everything else a buyer might check, an independent lawyer is told very slightly more than the market average and nothing like enough. Across the 95 companies, 45 will not name the models they run on and 36 offer only a vague pledge; 30 publish a benchmark score with no method; and **none has independent research published measuring its product**. The solo who can now buy a research tool for $30 a month has no third-party accuracy figure for it, or for any competitor.

What the segment demonstrates is narrower than it looks and more useful. Disclosure follows the buyer's leverage, not the vendor's values. Independent lawyers have almost no leverage individually, but collectively they hold the one kind that works on a website: they will not book the call. That moved a single signal and none of the others.

### What a buyer should ask

You are buying without a procurement department, so the questions differ from the ones a firm asks.

1. **Is the price on the page the price I pay?** Eleven of the 16 companies selling to you publish one. Ask a vendor that does not to send its number before the demo. Where a price is published, check whether it is the individual tier and what moves you off it — seat minimums, document counts.
2. **Was this built for a practice my size, or priced down to it?** Thirty-three of the 37 buyable products are also sold to law firms. Ask what the vendor's smallest customer looks like, and what the product assumes exists — a knowledge base, a document management system, an associate to check the output.
3. **What does it do with the citations?** Nine of the 38 products are research and citation, and the failure mode that ends careers here is a confident wrong citation. Ask whether every proposition links to source text, then check three by hand during the trial.
4. **Which model, and where does my client's file go?** Forty-five of 95 companies will not name their providers and 36 give a pledge without a name. You carry the confidentiality obligation personally. Ask for the provider, the hosting region, and whether zero retention is contractual or a sentence.
5. **Has anyone outside the company measured it?** Today no independent research exists on any product here. Six companies have at least joined someone else's evaluation — ask which and when. If the answer is a self-run benchmark, ask whether the tasks are published. Paxton AI's are; almost nobody else's.

## 9. People without lawyers

**$196M** stands behind products sold to people handling a legal problem without a lawyer — 8 funded companies, 19 products, **no unicorn**. It is the smallest column in the survey by every measure, and **26 times less capital than stands behind products built for corporate legal departments**.

Composition matters more than the total. **Thirteen of the 19 products are legal services delivery — a person with a licence doing the work — rather than software the user operates.** What venture capital has funded for the unrepresented is mostly a cheaper way to hire somebody. That is a real good, and it is not a tool.

### What exists

The whole list is short enough to print. Courtroom5's five products are in the catalogue but outside the funding survey, for reasons set out below.

| Company | Product | Type |
|---|---|---|
| Hello Divorce | DIY Plan | Legal services delivery |
| Hello Divorce | Pro Plan | Legal services delivery |
| Hello Divorce | Plus Plan | Legal services delivery |
| Hello Divorce | Hallie | Legal services delivery |
| Hello Divorce | Homesplit | Legal services delivery |
| Lawhive | Lawhive | Legal services delivery |
| Lawhive | Lawhive Legal Ltd | Legal services delivery |
| Boundless | Consumer visa packages | Legal services delivery |
| Casium | Casium | Legal services delivery |
| Conveyd | Conveyd | Legal services delivery |
| Keith | Keith | Legal services delivery |
| DoNotPay | DoNotPay | Legal services delivery |
| Justpoint | Justpoint Law LLP | Legal services delivery |
| Justpoint | Justpoint | Case intake & triage |
| Courtroom5 | The LAW Accelerator | Litigation case prep & discovery |
| Courtroom5 | Proof | Litigation case prep & discovery |
| Courtroom5 | Strategy | Litigation case prep & discovery |
| Courtroom5 | Documents | Drafting & document generation |
| Courtroom5 | Courtroom5 MCP connector | Agent platforms & infrastructure |

Two companies supply ten of the nineteen: **Hello Divorce accounts for 5 and Courtroom5 for 5**. In catalogue terms this market is two product lines and a scatter. The funded products cluster in divorce, immigration, conveyancing, consumer admin and mass tort — every one a transaction with money already moving through it: a house sale, a visa, a marital estate, a settlement.

The typing carries the finding. All 13 legal services products come from funded companies. The fourteenth funded product, Justpoint's intake tool, is claim sourcing for mass tort: the unrepresented person is not the customer but the inventory, which is a legitimate business and is not self-help. **All five remaining products — everything typed as software a person without a lawyer operates themselves — belong to Courtroom5**, the one company our own funding threshold could not see.

Courtroom5 is self-operated software for self-represented civil litigants, rebuilt through 2026 around AI as the LAW Accelerator. Proof maps a litigant's claims to their legal elements and scores how well the facts support each; Strategy summarises the other side's filings and says what comes next procedurally; Documents runs a seven-step drafting workflow ending in an oral argument simulator. It was founded in Durham, North Carolina by Sonja Ebron and Debra Slone, after Ebron — an electrical engineering PhD — was sued in a debt collection case and could not afford a lawyer.

It has raised roughly **$420,000** since 2017, most of it grants: NC IDEA, Google for Startups' Black Founders Exchange, Techstars Kansas City. No priced round has ever been announced. That figure is a floor estimate from a paywalled funding aggregator, not confirmed against a filing. Because our catalogue starts from a funding survey with a reporting threshold, **our own method filtered Courtroom5 out**, and it is excluded from the $196M and the 8-company count.

What it did with $420,000 is the point. In May 2026 it became a launch Justice Partner in Anthropic's Claude for the Legal Industry. On 25 August 2026 it appeared in Google's Gemini Enterprise connector list for legal work, alongside Thomson Reuters, iManage, Relativity, NetDocuments and Harvey. It is one of only two entries on that list reachable by a person without a lawyer — the other is CourtListener, the Free Law Project's public case-law archive — on a list otherwise made of enterprise incumbents. **A company with roughly one two-thousand-eight-hundredth of Harvey's capital reached the same distribution surface in the same month.** That took no sales organisation and no partnership team. It took a working MCP server and a corpus worth querying. For the job of being reachable from inside somebody else's assistant the entry cost has collapsed, and what is selected for is whether the data and the workflow are any good. Distribution is not adoption, and we can measure usage on neither side.

One correction belongs here because it is about us. As first published, our survey of products for people without lawyers stated that nothing in the market was pointed at debt collection defence. Courtroom5 covers debt collection and foreclosure, which falsifies the claim, and the correction now stands in the post. The error came from our funding threshold excluding the only company that contradicted us, and we would rather record that than let it disappear from a later draft.

### What it costs

**Four of the 8 funded companies publish a price — 50%**, second only to the independent-lawyer segment and for the same reason: nobody in this audience sits through a sales call. Keith posts fixed conveyancing fees of £1,200 to buy and £1,125 to sell. Hello Divorce posts its plan tiers, including packages at **$2,000** and **$3,500**. Boundless posts package pricing. Lawhive posts fixed fees per matter, as a regulated law firm rather than as a vendor. Courtroom5, outside the survey and so unscored, publishes **$1 for the first hundred credits or $150 a month**, with a company estimate of about **$1,500** to take a case to trial.

Against that, Conveyd markets "transparent pricing" and states no figure anywhere on its site, and Casium, DoNotPay and Justpoint publish nothing. The published prices are also not small in absolute terms: a $2,000 divorce package is cheap against a retained lawyer and expensive against a household that could not retain one.

### What is missing

Set the products against where unrepresented people stand in court. In Virginia's General District Court between 2018 and 2023, the Legal Services Corporation's Civil Court Data Initiative found that in debt collection cases **79% of collectors had a lawyer and 2% of consumers did**, and in eviction cases **68% of landlords had a lawyer and 1% of tenants did**. Those are among the highest-volume civil proceedings in the country and the ones where the imbalance is most extreme. Products from the 95 funded companies aimed at either: **zero**. Courtroom5 covers debt collection and foreclosure, which is the point — the one tool on that ground is the one venture capital did not fund.

Beyond it: nothing for eviction defence, nothing for custody or child support as distinct from divorce, nothing for benefits appeals, nothing for expungement, nothing for restraining orders. Not underserved. Empty.

This is not villainy. The funded products exist where somebody can pay. An eviction defendant is by definition a person who could not pay rent; a debt collection defendant is a person who could not pay a debt. Overwhelming demand with near-zero ability to pay is the one combination venture capital cannot serve on its own terms. Some of that ground belongs to legal aid, to courts building self-help tools, and to whatever regulators permit next. Some is reachable by anyone willing to build for a smaller cheque, and the evidence for that is in Google's connector list on $420,000.

### Who is credible

The top of this table is the shortest in the report. **Hello Divorce leads with 3 of 10**, on a published price and a vague model statement, and carries no named security certification. Boundless, Keith and Lawhive score **2**, in each case for publishing a price and almost nothing else; Keith publishes its conveyancing fees and renders its regulator licence number as an unfilled placeholder.

**Four of the eight score zero**: Casium, whose badge row advertises "SOC 2 Type 3", which does not exist; Conveyd, which markets "transparent pricing" and states no figure; DoNotPay; and Justpoint. A zero is a narrow statement — we could not verify a claim from public sources, not that a product is bad — but half an audience at zero is a pattern rather than an accident.

Courtroom5 is unscored: the rubric ran only on the 95 companies in the funding survey. It publishes a price, which the rubric rewards, and no independent research measures it either — true of every company here.

### The evidence picture

**Companies selling to people without lawyers score a mean of 1.13 out of 10, the lowest of the four audiences** — against 3.38 for independent lawyers, 2.15 for in-house, 2.03 for firms and 1.99 across the field.

Say what that does not mean first. **It rests on 8 companies.** One company moving two points shifts the mean by a quarter of a point, so it is not a stable estimate and we do not present it as one. Nor does it say these products are bad — we have tested nothing in this survey — and part of the gap is that a consumer business has little reason to publish model provenance or benchmark method at all.

Now the harder thing. These are the disclosures made to people who cannot afford a lawyer, about products that will handle their divorce, their visa, their house or their claim, and who have no professional training with which to judge any of it. A general counsel reading a thin trust page has a procurement team and outside counsel to compensate. A person facing a hearing alone has neither. **The audience with the least ability to evaluate a legal product is being given the least evidence with which to evaluate one.** That is the finding, and n=8 does not soften it.

The one instance of anybody outside the market testing a claim made to this audience is not research. In February 2025 the Federal Trade Commission issued an order against **DoNotPay** barring deceptive AI-lawyer claims. It scores nothing in our rubric, which measures what a vendor publishes, and it is worth recording for what it implies: the only external check on a consumer legal AI claim in this survey arrived as an enforcement action, from a regulator, after the fact. Nobody is running evaluations here, and **none of the 95 companies has independent research published measuring its product** — a gap that is uncomfortable for a general counsel and consequential for somebody defending their own case.

### What a buyer should ask

If you are handling a legal problem without a lawyer, most of this report is not written for you. This part is, and each question can be asked in an email or a chat window before you pay.

1. **Is a lawyer doing this work, or am I?** Ask it in those words. Most of what is sold here — 13 of the 19 products we found — is a person with a licence doing the work for you. That is often what you want, and it is a different purchase from software you operate yourself.
2. **Who is responsible if this goes wrong?** A lawyer or a regulated firm answers to a regulator and usually carries insurance. For software, normally nobody does. Ask which this is, ask for the licence number, then check that number on the regulator's own website — we found a company publishing its licence number as a blank placeholder.
3. **What is the total cost?** Ask for the full figure, what is not included, which court fees you still pay yourself, and what happens to your money if you stop halfway.
4. **What does this not do?** Ask what happens if the other side files something unexpected, if there is a hearing, or if you lose. A tool that prepares documents does not appear in court with you.
5. **Has anyone independent checked that it works?** Today the answer is no. Not one company in this survey has had its accuracy measured in public by anyone without a commercial interest in the result. If told otherwise, ask who ran the test and where you can read it. A percentage on a homepage with no method behind it is a marketing claim.
6. **What free help exists near me?** Check your court's self-help centre, your local legal aid office and your bar association's referral line first. In eviction, debt collection and benefits cases almost nothing has been built commercially, so free and non-profit help is not the fallback there. It is the main provision.

# Part III — Agents

Fifty of the 264 products in our product catalogue are agents or agent infrastructure. That makes it the largest single thing happening in legal AI — larger than drafting, larger than contract review, larger than research. It is also the least well-defined. All five unicorns in the funding survey ship something called an agent, and so do companies that have raised under $10M. The word appears on the homepages of firm platforms, contract tools, patent search engines, billing software and a company selling case-status phone calls.

What follows is an attempt to say what the word is doing in each case. We read every product's own description and sorted the fifty into the four meanings the term actually carries in this market. That sorting is our reading, not a test result: we have not run these products, and no vendor publishes the kind of reliability data that would let anyone check the label from outside. That absence is itself the finding of Chapter 14, and it is the reason this part is a taxonomy rather than a ranking.

## 10. The taxonomy

"Agent" in legal software means at least four distinct things, and the distance between them is larger than the distance between most product categories.

**A chat assistant, renamed.** Luminance's Lumi summarises contracts, redrafts clauses and answers questions about them. GC AI answers in-house questions with citations. Ivo's Search Agent runs natural-language search across executed contracts. Nexl AI surfaces relationship risks out of a firm's own CRM. These are useful products. They are not autonomous in any sense the word usually implies — a person asks, the software answers, the exchange ends. **Thirteen** of the fifty fall here.

**A multi-step workflow runner.** Harvey Workflows, Legora Workflows, Patlytics Agent, Clerq Agents, Spellbook's Associate, Pocketlaw's PLAI. You define a sequence; it executes over documents. This is the largest bucket at **seventeen** products, and it is closer to scripting than to delegation. The distinguishing feature is that the sequence is fixed in advance and the run is initiated by a human who is waiting for the output.

**A builder.** Bryter Workflows, LegalFly's Agent Studio, Newcode's Aurora, DeepJudge's AI Workflows, Josef, Checkbox AI, Zuva Create, Eudia's Enterprise Brain. **Eight** products where the thing sold is not an agent at all but a factory for making them, shipped to teams who will define the work themselves. The commercial logic is different from the rest of the category: the vendor is selling capacity, and the buyer supplies the domain knowledge.

**Something actually autonomous. Six** products, covered in the next chapter but two: work that starts without a person opening an app, runs against the outside world, and stops to ask only when it needs to.

That accounts for 44. The remaining six split into two groups that are worth separating rather than forcing into the four. **Five** are pure infrastructure — Trellis's API and connectors, midpage's MCP connector, Twin1's MCP server, the Zuva API, GC AI's API — plumbing that lets something else be the agent, and the subject of Chapter 13. And **one**, Norm Ai's Supervisory AI, does not act at all: it watches other agents. That count of one is Chapter 14.

Two caveats on the counts. The boundaries between the first two buckets are genuinely soft — Summize's SIA "answers contract questions and executes review steps", Juro's Operator "executes contract actions from plain-language instructions", Athennian AI "answers governance questions and executes entity lifecycle changes", and each of those could be read either way. We placed them by whether the product describes a conversation or a sequence. And the classification uses vendor copy as its evidence, which means a product that oversells itself is classified as it oversells itself. We have flagged where we think that is happening; we cannot rule out cases where it is not visible.

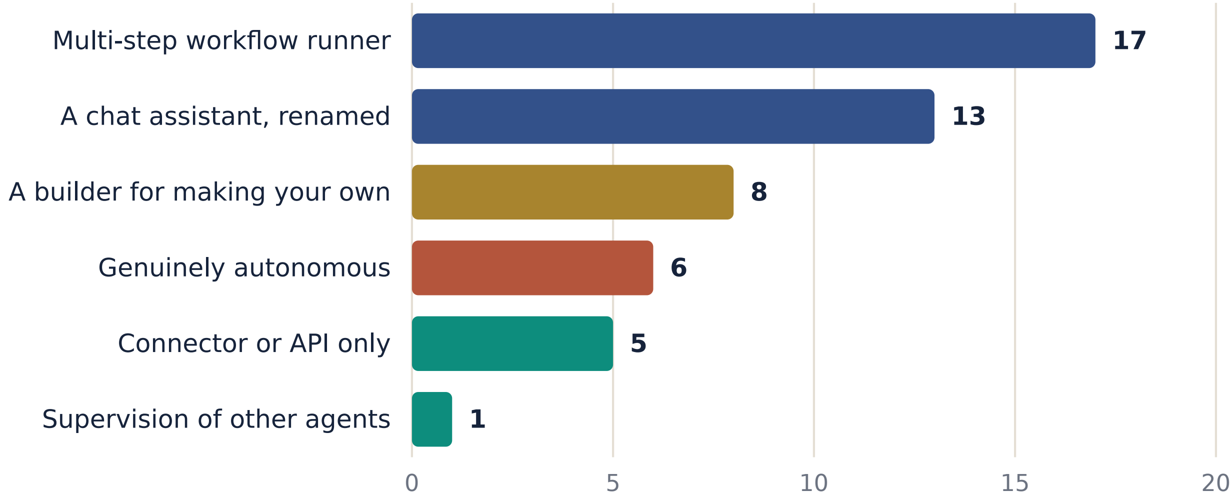


The 50 agent products, by what the word actually means in each case

## 11. Who is building them

Grouped by bucket, with the funding segment each company sits in. Product descriptions are condensed from each vendor's own materials.

### Chat assistant, renamed (13)

| Company | Segment | Product | What it does |
|---|---|---|---|
| Luminance | Contracts & in-house | Lumi | Summarises contracts, redrafts clauses, answers contract questions |
| GC AI | Contracts & in-house | GC AI | Cited answers, contract review and US case-law research for in-house counsel |
| Noxtua | Firm platform | Noxtua | Research, document analysis and drafting in German and EU law on EU infrastructure |
| Lawhive | Consumer & small law | Lawrence | AI paralegal gathering case facts and answering solicitors' questions on a file |
| Summize | Contracts & in-house | SIA | Answers contract questions and executes review steps across the CLM |
| Athennian | Practice ops & billing | Athennian AI | Answers governance questions and executes entity lifecycle changes against the record |
| Juro | Contracts & in-house | Operator by Juro | Reads contracts and executes contract actions from plain-language instructions |
| Nexl | Practice ops & billing | Nexl AI | Surfaces client opportunities and relationship risks from the firm's own CRM |
| Eve | Litigation & PI | Eve Research | Live database of US court opinions available inside every Eve agent |
| Ivo | Contracts & in-house | Search Agent | Natural-language search across executed contracts with no prior tagging |
| Jus Mundi | Research & knowledge | Jus AI | Research assistant over the Jus Mundi corpus, with user control of sources |
| Twin1 | Research & knowledge | Twin1 | Per-person agent trained on that professional's email, documents and messages |
| Chamelio | Contracts & in-house | Chamelio | Answers questions across a legal team's contracts, email, drives and chat |

## Multi-step workflow runner (17)

| Company | Segment | Product | What it does |
|---|---|---|---|
| Harvey | Firm platform | Workflows | Pre-built and firm-built multi-step task templates run over documents |
| Harvey | Firm platform | Agents | Routes a legal task to purpose-built agents that plan and execute it |
| Legora | Firm platform | Agent | Plans, executes and reviews a multi-step task, returns a finished output |
| Legora | Firm platform | Workflows | Configurable repeatable multi-step legal processes |
| Spellbook | Contracts & in-house | Associate | Works across every document in a matter rather than one file at a time |
| Wordsmith AI | Contracts & in-house | Agents | Automates repeated legal workflows end to end |
| Supio | Litigation & PI | Supio Agent | Firm-wide agent working across every case and role at once |
| Eve | Litigation & PI | EveOS | Runs a plaintiff firm's cases, staff and finances in one system |
| Patlytics | IP & patents | Patlytics Agent | Chains patent tasks across the platform rather than running one tool |
| Sandstone | Contracts & in-house | Sandstone | Intake, triage, matter management with agentic drafting and redlining |
| Norm Ai | Compliance | Norm Technology | Regulation-aware agents reviewing content and disclosures against rules |
| SpotDraft | Contracts & in-house | Sidebar | Agents that answer questions and track regulatory change |
| Enter | Litigation & PI | Consumer Litigation | Drafts answers to Brazilian consumer claims and flags serial-litigant fraud |
| Pocketlaw (now Miramis) | Contracts & in-house | PLAI | Reviews against playbooks, analyses redlines, queries the portfolio |
| Streamline AI | Contracts & in-house | Featherline | Contract review applying negotiation playbooks and flagging risk |
| IPRally | IP & patents | IPRally Agent | Turns an invention disclosure into a report linked to the evidence |
| Clerq (formerly NLPatent) | IP & patents | Clerq Agents | Triages invention disclosures into a patentability report in about ten minutes |

## Builder (8)

| Company | Segment | Product | What it does |
|---|---|---|---|
| Bryter | Contracts & in-house | BRYTER Workflows | No-code builder for rule-based legal and compliance processes |
| DeepJudge | Research & knowledge | AI Workflows | Build and run multi-step agents over the firm's indexed knowledge |
| Checkbox | Contracts & in-house | Checkbox AI | Automates request intake, triage and workflow building inside the Hub |
| Newcode | Firm platform | Aurora | Engine for building and deploying bespoke agentic legal workflows |
| LegalFly | Contracts & in-house | Agent Studio | Builds custom agents triggered by email, Slack, Teams or manual entry |
| Zuva | Contracts & in-house | Zuva Create | No-code trainer for custom contract extraction fields |
| Josef | Practice ops & billing | Josef | No-code builder for bots and automated workflows, no developers needed |
| Eudia | Contracts & in-house | Enterprise Brain | Captures a department's knowledge and preferences into reusable agents |

### Actually autonomous (6)

| Company | Segment | Product | What it does |
|---|---|---|---|
| Flank | Contracts & in-house | Flank | Handles recurring legal requests inside Outlook, Teams, Jira and Salesforce |
| Flank | Contracts & in-house | Flank Record | Captures contract metadata at signature and maintains the record |
| Streamline AI | Contracts & in-house | Velo Copilot | Coordinating agent across matter types, running interactively or ambiently |
| Hona | Litigation & PI | Lia | Chases clients, tracks appointments, flags treatment gaps, escalates when needed |
| Eve | Litigation & PI | Communication Agents | Places outbound calls for status updates and medical records requests |
| EvenUp | Litigation & PI | Communication Agents | Automated client outreach and case-status updates |

### Infrastructure, not an agent (5)

| Company | Segment | Product | What it does |
|---|---|---|---|
| Trellis | Litigation & PI | Trellis API & connectors | API plus Claude MCP and ChatGPT connectors for US state court records |
| midpage | Research & knowledge | midpage MCP | Lets Claude, ChatGPT and Perplexity query midpage's case-law corpus |
| Twin1 | Research & knowledge | Twin1 MCP server | Governed interface for other AI tools to query a twin's context |
| Zuva | Contracts & in-house | Zuva API | REST API extracting 1,400+ contract fields at $1.25 per document |
| GC AI | Contracts & in-house | GC AI API | Programmatic access so other systems can call GC AI without a seat |

### Supervision (1)

| Company | Segment | Product | What it does |
|---|---|---|---|
| Norm Ai | Compliance | Supervisory AI | Verification layer monitoring whether other companies' AI agents behave lawfully |

Forty companies build these fifty products. Two segment findings fall out of the table. **All four** firm-platform companies — Harvey, Legora, Noxtua, Newcode — ship agents; it is the defining feature of that segment. Contracts and in-house supplies **eighteen** of the forty builders, which is more than any other segment and roughly half its own company count.

And four segments ship nothing agentic at all: real estate, immigration, tax, and — the interesting one — **AI-native law firms**. Crosby, Manifest OS, General Legal, Moritz and Soxton sell legal work delivered with AI rather than software that acts. Whatever they run internally, they do not sell an agent, because what they sell is the outcome.

## 12. What the autonomous ones actually do

Look at what the six genuinely autonomous products do, and the striking thing is how unglamorous it is. None of them argue motions. They chase people who have not replied, answer the same routine question for the fortieth time, request records, place outbound calls, and log what happened.

Flank's agents handle recurring legal requests inside Outlook, Teams, Jira and Salesforce — the point being that nobody opens an app. The request arrives where the business already works, and is answered there. Flank Record does the same for contract data: metadata captured at the moment of signature, the record maintained without anyone remembering to maintain it, with a lawyer supervision layer sitting over the top.

Streamline AI's Velo Copilot describes itself as running "interactively or ambiently" — a phrase worth noticing, because the ambient half is the whole claim. Hona's Lia chases clients, tracks appointments and flags treatment gaps in personal injury matters, escalating only when it needs a human. Eve's communication agents place outbound phone calls for case status and medical records requests. EvenUp's do automated client outreach and case-status updates.

That is the right place to start, and the reason is structural rather than modest. This work is high-volume, low-judgement and verifiable. There is a fact of the matter about whether the records arrived, whether the client was called, whether the metadata matches the executed document. A mistake is embarrassing rather than malpractice. Three of the six sit in litigation and personal injury, where firms run hundreds of matters through the same repetitive contact cycle and the economics of a missed follow-up are measurable.

It is also the work least protected by professional privilege in the reputational sense. Nobody's licence turns on whether a status update was worded well. The agents doing substantive legal work — the seventeen workflow runners in the previous chapter — are, almost without exception, running under a lawyer who reviews the output before it goes anywhere. That is a design decision, not a limitation the vendors are apologising for.

One disclosure detail is worth recording here. Flank holds **ISO 42001**, the standard that audits AI governance rather than information security, and is the only company in our 95 to do so. The company shipping the most autonomous product in the set is also the only one certified against the standard written for exactly that problem. We note the coincidence without claiming it is causal; ISO 42001 is a management-system audit, not a measure of whether the agents work.

Anyone selling an agent that practises law unsupervised is selling a liability.

## 13. The connector layer

The development we did not expect to find is not a product. It is a protocol.

Legal data is becoming callable from outside legal software. Trellis ships a Claude MCP connector and a ChatGPT plugin for US state court records. midpage ships an MCP connector so Claude, ChatGPT and Perplexity can query its case-law corpus. Twin1 runs an MCP server exposing a lawyer's own context to other tools and to a network of other people's. Clerq offers API and MCP access to patent search. GC AI sells API access with no seat attached — a pricing decision as much as a technical one. Zuva sells extraction by the document at **$1.25** through its API rather than by the user.

Sit with the implication. If the research corpus, the court data and the firm's own knowledge are all reachable by a general-purpose assistant, then the interface a lawyer works in need not belong to a legal software vendor at all. The legal company becomes the data and the domain logic. Somebody else owns the window.

Which is roughly what Google announced. Read the Gemini Enterprise for Legal integration list carefully: Docusign, Everlaw, iManage, NetDocuments, RelativityOne, Thomson Reuters — and Harvey, and Legora. Google did not launch a Harvey competitor. It launched a layer that sits above Harvey and Legora and treats them as sources. Anthropic's Claude for the Legal Industry, shipped in

May 2026, does the same thing with a connector list that includes midpage, Trellis, Harvey and Solve Intelligence. We covered the labs' five different positions in what the model labs built for legal; the connector layer is the part with consequences for everyone else.

For a firm already paying for two of those products, being reachable from one window is attractive. For the vendors, being an integration inside somebody else's enterprise agent is a materially different business from being the place the lawyer opens every morning.

One name on Google's list is worth stopping on. Courtroom5, which builds for people without lawyers, sits in it alongside Thomson Reuters, iManage and Relativity. It has raised about **$420,000**, mostly in grants. It was also a launch Justice Partner in Anthropic's Claude for the Legal Industry in May, where its connector exposes pro se case intake, a deadline calculator and next-step guidance to a general assistant.

That is the connector layer doing something the funding market has not. A company with roughly $420,000 reaches the same distribution as companies with billions, because the price of entry is a working MCP server and a corpus worth querying rather than a sales team. It is the only mechanism we found in the whole survey that is indifferent to capital. Set against the pattern in every product for people without lawyers, where the pro se audience attracts $196M against $5.10B for in-house, that is not a small thing.

The catch is that it cuts the other way too. A connector is a supply agreement. The party that owns the window sets the terms, decides who is in the list, and can build the thing the connector supplies. Distribution and dependency look identical from the inside.

## 14. The supervision gap

If agents proliferate, something has to check them. Across all 269 products in the catalogue we found exactly **one** built for that: Norm Ai's Supervisory AI, a verification layer that monitors whether other companies' AI agents are behaving lawfully.

One product against fifty agents. Norm Ai is not a small company — **$260M+** raised, a $1.2B valuation on a $120M Series C in July 2026 — so this is not a gap for lack of capital in the segment. It is a gap because nobody else has decided the problem is a product yet.

The harder finding sits underneath it. We went looking for agent reliability benchmarks — task completion rates, error rates, how often a human has to intervene — and found none. Not from Harvey, not from Legora, not from Google, not from any of the forty companies in Chapter 11. Not one number, from anyone, about whether these things finish the job.

This is worse than the general evidence problem we documented in nobody can check any of it, where the mean evidence score across 95 companies is **1.99 out of a possible 10** and **73** publish no price. There, at least, the claims being made are checkable in principle: an accuracy percentage has a corpus and a method behind it, even when neither is disclosed. Agent reliability is one step further from verification, because the unit of work is a whole task rather than an output you can hold up against a source. At least a transcript can be checked against a recording.

The disclosure records of the agent builders do not help. Norm Ai — the company selling agent supervision — scores **1 of 10** on our evidence framework: SOC 2 only, no model disclosure, no published price, selling into institutional finance. Flank scores **2**. Only two companies in the entire survey publish an open benchmark, and **six** participate in an independent one. None of those benchmarks measures agentic task completion. Nobody is marking this exam, as we put it in who is marking the exam, and for agents nobody has even written it.

We are not asserting that these agents fail. We have not tested them and we say so. What we are asserting is narrower and harder to argue with: as of 25 August 2026, no buyer of any of these fifty products can find out from the vendor how often the agent completes the task correctly, and no independent party publishes the number either.

Until somebody does, "agentic" describes an architecture and not a result.

# Part IV — Models, benchmarks and skills

The parts before this one described a market: who raised the money, what they built, and who they sell it to. This part asks the question that market cannot answer about itself. Does any of it work, how would anyone know, and what happens to the companies in it now that the labs supplying their models have started shipping legal products of their own.

Four findings run through the five chapters. The model labs have taken five different positions on legal, ranging from twelve open-source practice-area plugins to nothing at all, and two of the five now sell products that compete with their own customers. The benchmarks that measure legal products were mostly written by the companies being measured, and every vendor that wrote one won it. The model benchmarks that are genuinely independent have saturated at the top, which makes them useless for choosing between frontier models and irrelevant for choosing between products. And the single most useful number in the field is not a score but a gap: the distance between marking an answer for how much of it is right and marking it right only if all of it is. On the evidence we could assemble on 25 August 2026, a buyer comparing two legal AI products has almost nothing external to compare them on. We include ourselves in that.

## 15. What the model labs shipped

The $6.2B counted in Part I went to companies building legal software on somebody else's models. In 2026 the somebody-elses started shipping legal products themselves. Five labs took five different positions, set out in what the model labs built for legal.

| Lab | Legal offering | Shipped | For whom | Price |
|---|---|---|---|---|
| Anthropic | Claude for the Legal Industry | May 2026 | Firms, in-house, self-represented | Not stated for the legal offering |
| Google | Gemini Enterprise for Legal | Aug 2026, preview | Enterprise and BigLaw | None disclosed |
| Microsoft | Legal Agent in Word | Apr 2026, preview | Contract reviewers | **$99** per user per month, via E7 |
| OpenAI | A vertical, no product | — | — | — |
| xAI | Nothing | — | — | — |

Amazon has a legal landing page on its general Quick Suite. Mistral supplies models to Harvey. Cohere, Meta and DeepSeek have nothing in legal at all.

Anthropic went first and widest. Claude for the Legal Industry landed on 12 May 2026 with 20+ MCP connectors and **12** open-source practice-area plugins — Litigation, IP, Employment, Privacy, Regulatory, AI Governance, Law Student, Legal Clinic and others — plus Microsoft 365 integration. The connector list reads as a market map: Definely, DocuSign, Ironclad, iManage, NetDocuments, Everlaw, Relativity, Thomson Reuters CoCounsel, midpage, Trellis, Harvey, Solve Intelligence. Its Justice Partner programme makes connectors from Courtroom5, the Free Law Project's CourtListener, Descrybe and BoardWise free to Claude users, citing the roughly **80%** of US civil litigants who appear without counsel. Across all five labs, that is the only thing shipped for people without lawyers.

Google went at BigLaw. Gemini Enterprise for Legal launched on 25 August 2026, in preview and by request, with named skills for brief drafting, citation verification, contract lifecycle management, regulatory horizon scanning and DSAR fulfilment, its own research, screening and contract-drafting agents, and partner-built agents from Deloitte and Eudia. Launch customers are Cleary Gottlieb, Freshfields, Weil, and Williams & Connolly. There is no free tier, no legal aid pricing and no disclosed price for the legal SKU.

Microsoft shipped the most concrete thing and said least about it: Legal Agent for Word, April 2026, clause-by-clause playbook review with tracked-change redlining and rationale comments, in preview on US tenants, gated behind Microsoft 365 Copilot and the E7 Frontier Suite at **$99 per user per month** — the only clearly stated price anywhere in this chapter. It is widely read as the first output of Microsoft's January 2026 acqui-hire of Robin AI's engineering team. Robin AI raised $61.5M to do AI contract review and ran out of money; its engineers now ship contract redlining inside Word to anyone with a Copilot licence. The startup died and the feature shipped anyway, at the bottom of somebody else's bundle.

OpenAI hired Jason Boehmig, Ironclad's co-founder, in June 2026 to lead product for a legal vertical. That is an organisational announcement. There is no legal SKU and no plugin library; the July partnership with Willkie Farr is a customer, not a product. xAI has nothing beyond Grok 4.5 launch copy mentioning contract review, which is model positioning.

### The platform-or-competitor bind

A lab that sells a legal product competes with the companies building on its models. A lab that sells only a platform gives up the application-layer margin to those same companies. Every lab in the table is somewhere on that trade, and two of them are on both sides of it at once.

Harvey, Legora and Solve Intelligence appear in **both** Anthropic's and Google's connector lists. Thomson Reuters is in both. Legora published a release describing the flow: ask a question inside Gemini Enterprise, get an answer with links, then move into Legora for the harder work. At the same time Google ships its own contract-drafting and legal-research agents, Anthropic's Commercial Legal plugin reviews vendor agreements and NDAs against a firm's playbooks, and Microsoft's Legal Agent redlines inside the app where lawyers already write.

On the evidence, Microsoft has chosen the product side outright: a competing feature, in the default writing surface, at a stated price, with no legal connector ecosystem around it. Google has chosen both and is candid about neither. Its integration list — Docusign, Everlaw, iManage, NetDocuments, RelativityOne, Thomson Reuters, Harvey and Legora — is the sharpest example in the market of a layer built above the two firm platforms that treats them as sources rather than replacing them, and it ships its own agents in the same product. Anthropic sits closest to the platform side: open-source plugins, a wide connector list, free nonprofit connectors, and one plugin that does work a contract vendor would sell. OpenAI and xAI have declared nothing, and an unshipped vertical is not a position.

Where the evidence is thin, it is thin in the same three places for all five. Two of the products are in preview, and preview feature lists change. Only one lab has published a price for a legal offering. And no lab has published a reliability number for any legal task — the gap Chapter 16 is about.

The bind matters most for the application layer. Harvey's CEO said the quiet part in May: "long term we would end up competing with the model companies" — while shipping a Claude connector. Thomson Reuters' CTO gave the best counter-argument available: "the control point isn't where work starts. It's whether the output is accurate, grounded in authoritative sources, and defensible." That is a bet that verification, citation and liability are the durable business, not the chat window. It is also an unfalsifiable bet until somebody publishes accuracy numbers.

Worth naming who is in neither list: Eve, Clio, Spellbook, LegalOn, Luminance. Distribution through the labs is not being handed out evenly, and it is not clear the missing names chose to be missing.

## 16. The benchmarks that exist and who administers them

Every legal AI vendor quotes a score. We went looking for where those scores come from and whether a buyer can use any of them, and published the survey as who is marking the exam. The model benchmarks are real, current and increasingly uninformative. The product benchmarks are mostly written by the companies being tested.

| Benchmark | Who built it | Scores | Current | Tasks and rubrics open |
|---|---|---|---|---|
| LegalBench | Stanford et al., 2023 | Models | Yes, but saturating | Yes, academic release |
| Vals Legal Research Bench | Vals AI + law firms | Models | Aug 2026 | No |
| Harvey LAB | Harvey, 2026 | Models | Aug 2026 | Yes — 1,200+ tasks, ~75,000 criteria |
| BigLaw Bench | Harvey, 2024 | Products | Vendor-run | No |
| VLAIR | Vals AI + 8 firms | Products | Feb 2025, not repeated | No |
| Contract Review Benchmark | LegalOn, 2026 | Products | Vendor-run | No |
| ContractScrub | Thomson Reuters + Imperial | Models | Aug 2026 | Not established |
| LEXam, PLawBench, Multi-Legal-Bench | Academic | Models | 2025–26 | Yes, academic release |

The final column records what we could verify from public materials on 25 August 2026. "Not established" means we could not confirm it either way, not that the tasks are withheld. We have not replicated any benchmark in this table.

Read the shape of that table before the contents. Everything that scores *models* is current and mostly independent of the companies selling legal software. Everything that scores *products* — the thing a general counsel actually buys — was either built by a vendor with a product in it, or is eighteen months old and was never repeated.

### Four vendors built benchmarks. Four vendors won.

Harvey's BigLaw Bench (2024) reported Harvey completing about **74%** of lawyer-quality work, ahead of the models it tested. Harvey wrote the tasks, wrote the rubrics, ran the evaluation and reported the result. LegalOn's Contract Review Benchmark (June 2026) put LegalOn **87 ELO points above the next model and 400 above the best GPT**. Ivo's April 2026 study scored Ivo at **4.52** against human lawyers at **4.56**. GC AI's May 2026 study scored GC AI at **86.8%** against ChatGPT's **79.8%**.

Four benchmarks, four home wins. None of that means the products are bad — we have not tested them. It means the numbers carry no procurement weight, and a buyer should treat them as marketing with arithmetic attached.

The honourable exception is Harvey's second attempt. LAB, released May 2026, open-sourced **1,200+** tasks across **24** practice areas with roughly **75,000** expert-written rubric criteria, and Harvey pointedly did not publish a score for its own product. That is how it should be done. It still came from a market participant, which the authors of LegalBench would tell you is the structural problem: in January 2026 they published a paper in *PNAS* warning that benchmarks "can be captured, watered down, and abused." They built the field's most-cited legal benchmark, and that is their own assessment of what comes next.

### The audit underneath

The benchmark survey and the evidence score are the same finding measured two ways. Across the 95 funded companies, scored mechanically from public websites between 20 and 25 August 2026:

| What a buyer can check | Companies, of 95 |
|---|---|
| Publish a score with no method attached | 30 |
| Publish a benchmark with tasks and rubrics open | 2 — Harvey and Paxton AI |
| Have participated in an evaluation run by someone else | 6 |
| Have independent research published measuring the product | **0** |

Zero of 95. Across $6.2B and 269 products, not one company has had its accuracy measured in public by a party with no commercial interest in the result. The nearest approaches are Clearbrief, scored **40.5 out of 50** by the State Bar of Nevada's AI Workgroup, and Alexi, which took part in the Vals legal research study in October 2025 and landed at the participant average, above the lawyer baseline. IPRally has a third-party study in which IPRally supplied the tool and trained the searchers.

Nor can a buyer fix this by asking for a head-to-head. The most thorough independent product study remains VLAIR, February 2025, which measured Harvey, CoCounsel, Vincent and Oliver on real law-firm data against a lawyer control group. There has been no second edition. Vals ran a legal research study in October 2025 in which Thomson Reuters, LexisNexis and vLex all declined to take part. Legora, valued at $5.6B, has not appeared in any independent public benchmark we could find. The flaw is not methodological. Vendors choose whether to be measured, and the ones with the most to lose choose not to be.

## 17. Model rankings on legal tasks

The frontier models are close enough on legal benchmarks that the rankings have stopped carrying information. On LegalBench the top model sits at **88.6%**, and the top eight cluster within about **two and a half points**. When every frontier model scores within a rounding error, the benchmark is no longer distinguishing them; its own authors have moved on to harder things.

For a buyer, a saturated benchmark means three things. Model choice is no longer the variable that decides whether legal output is usable — retrieval, grounding, review and workflow are. A vendor citing a leaderboard position is citing a distinction of two points on a test its supplier has effectively finished. And "we run on the best model" is not a product claim.

What varies more than the model is where you point it. Vals' research bench found scores by practice area ranging from **38.4% in health law to 11.4% in family law** on the same model. Multi-Legal-Bench, testing six European jurisdictions in May 2026, found **no model dominating in any language** — rankings reorder by task and by country. The right question is not which model is best but which is best at your work, in your jurisdiction, and no leaderboard answers that.

Where products rather than models were measured, the one independent data point is still VLAIR: Harvey took the top score on five of six tasks it entered, reaching **94.8%** on document Q&A against a lawyer baseline of **70.1%**, and lawyers beat every tool on redlining. That is eighteen months old, it did not include Legora, and Lexis+ AI enrolled and then withdrew from most tasks.

### The column nobody charts

Cost per task is the number procurement asks about first and no published leaderboard carries. Figures circulate — a few dollars per query on a research benchmark, an order of magnitude more on the harder agentic tasks — but we could not reconcile any of them against a primary source we would be willing to print, so we print none.

That absence is itself the finding, and it is the largest single piece of research this report could not complete. A firm running thousands of matters cares whether a task costs pennies or a professional's billed minutes, and the accuracy leaderboards that do exist separate the leading models by a couple of points while saying nothing about a cost difference that may be an order of magnitude. Assembling a real cost table across models and vendors would tell a buyer more than every ranking in this chapter combined.

## 18. Partial credit versus all-pass

This is the most useful number in the field, and it is not a score. It is the distance between two ways of marking the same answer.

Take a research memo with ten required elements: the right rule, the controlling authority, the exception, the jurisdiction, the date the statute changed. **Partial credit** asks how many of the ten the system got right and reports the fraction — nine of ten is 90%. **All-pass** marks the answer correct only if every element is present — nine of ten is zero. Both are honest. They measure different things, and the gap between them on the same output is enormous.

On Vals' Legal Research Bench, the leading model scores **90.6% under partial credit and 55.3% when every required element must be present**. On Harvey's LAB, models satisfy roughly **90% of individual rubric criteria but complete around 20% of whole tasks**.

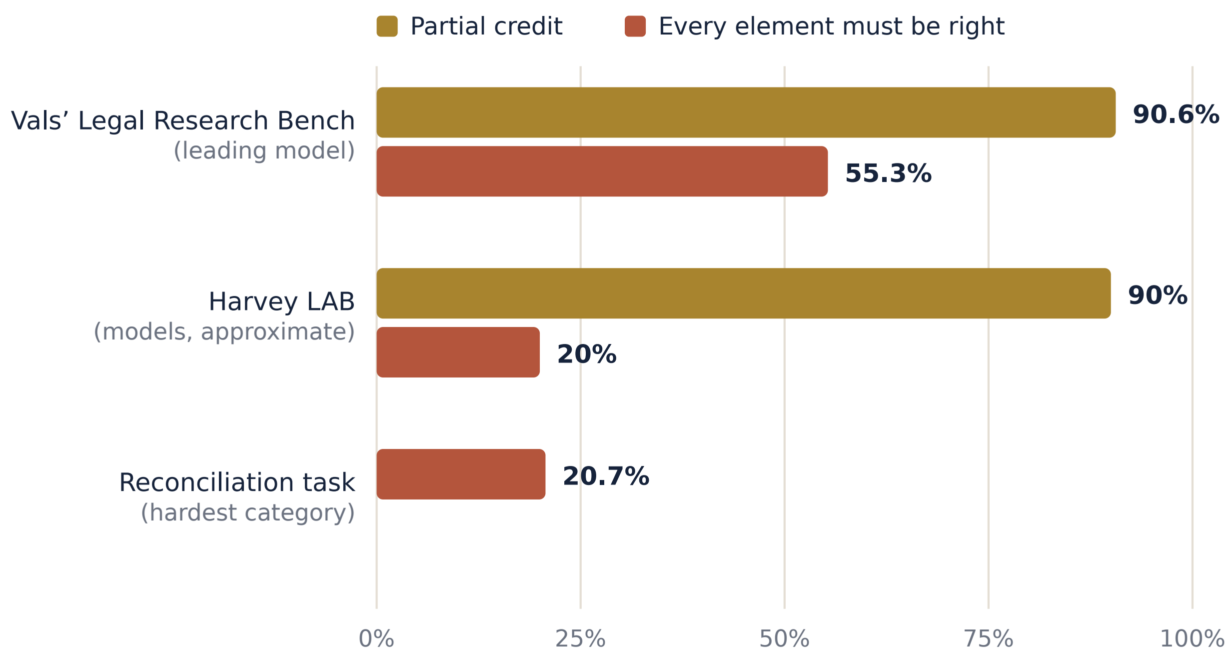


The same answers, scored two ways

That gap is the whole story of legal AI reliability, because legal work is conjunctive. A memo that gets nine things right and misses the tenth is not 90% correct. It is wrong, and possibly wrong in a way that reaches a client. Every headline accuracy figure a buyer has read is the generous number. The strict one exists and is usually not the one on the slide.

The same benchmark found the hardest category was **reconciliation** — synthesising authorities that conflict — at **20.7%** all-pass. Which is a fair description of what lawyers are for.

Practice area moves the number almost as much as scoring method does. On the same model, Vals measured **38.4%** in health law and **11.4%** in family law: the strongest area scores more than three times the weakest, on identical software.

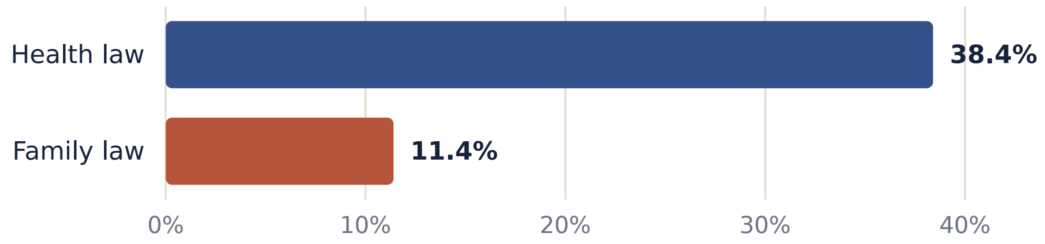


Accuracy by practice area, on the same system

Two consequences follow. A vendor average across practice areas conceals a spread wider than the difference between any two vendors, so a firm whose work is concentrated in one area learns nothing from a blended figure. And the phrase "best legal AI" cannot survive either finding.

### Our interest in this, declared

IQdoc publishes its own accuracy numbers, which makes us a market participant reporting on a measure we also publish against. Read this chapter accordingly. Our transcription benchmark scored 23,544 words from two US Supreme Court oral arguments against the official certified transcripts. Every fully automated method we tested, ours included, landed between **96.3% and 97.2%** — a spread of 198 errors in 23,544 words between the best engine and the worst. Our own best row reaches **98.6%**, and it gets there by flagging the roughly 3% of words the system is least sure about for human review, not by having a better model.

We publish that with its limits attached: clean appellate audio is the ceiling condition, the review row assumes a flawless reviewer, it is two arguments rather than a corpus, and we have already had to correct the numbers once after finding scoring errors in our own method. The argument we make about our own numbers is the one we make about everyone else's here. A word-accuracy percentage is the partial-credit number. Whether the transcript is usable in court is the all-pass one, and errors cluster exactly where the consequence is — proper nouns, case names, figures, and who said it.

## 19. Hallucination: the published evidence versus the marketing

The most consequential legal AI research of the last two years was not a benchmark. Stanford's RegLab tested the major legal research tools on **202** queries and published the results in the *Journal of Empirical Legal Studies*: **Lexis+ AI hallucinated 17% of the time and Westlaw's AI-Assisted Research 33%**, where hallucination included confidently citing a source that did not support the claim. Both products had been marketed on the absence of that problem.

Retrieval reduces hallucination substantially. It does not eliminate it. A Stanford follow-up in February 2026 on statutory research found **Lexis+ AI at 64% and Westlaw at 58%, against a generic retrieval baseline at 70%** — two commercial legal platforms underperforming an off-the-shelf approach on that task. The same lesson applies to the credential everyone quoted: GPT-4's "90th percentile" bar exam result, re-examined against licensed attorneys rather than a pool padded with repeat takers, was roughly **45th percentile overall and 15th on the essays** — the half that resembles legal writing.

Against that published record, set what the funded market publishes about hallucination.

Paxton AI ran Stanford's hallucination benchmark on itself and released the method alongside the result. It is self-run, and we record that plainly. It is also reproducible, which is the point, and it is one of only **2 of 95** open benchmarks in the entire survey — the other being Harvey's LAB, which reports no Harvey score. Legora declines to open its benchmark corpus and gives a real reason: published cases leak into the next model's training data. We record that as a real objection, and it is also the reason the field cannot bootstrap its way out of this on vendor goodwill alone.

Then the marketing. Ivo markets a study it ran itself under the headline **"Independent AI Benchmark"** — the study that scored Ivo at 4.52 against human lawyers at 4.56. LegalFly's "Legal AI Benchmark" is a buyer's guide rather than scored results. Neither publishes tasks or rubrics; both score zero on the benchmark signal in our evidence framework, where a closed benchmark adds nothing a buyer can check.

One measured result deserves more attention than any accuracy figure in this part. In an independent study of contract drafting, legal-specific tools raised risk warnings in **83% of high-risk scenarios against 55% for general-purpose AI** — while the general-purpose models scored *higher* on raw output quality. Knowing when to stop and flag something is not an accuracy metric. It may be the most lawyerly capability there is, and no leaderboard scores it.

### What a buyer can actually conclude, in August 2026

Very little, and that is the finding rather than a failure of the survey.

You can conclude that frontier models are close to each other on legal benchmarks and far apart on cost per task. You can conclude that any headline accuracy figure you are shown is almost certainly the partial-credit number, and that the all-pass number on the same work is dramatically lower. You can conclude that practice area and jurisdiction move the result more than the choice of model. You can conclude that hallucination has been measured, in public, on the two largest legal research platforms, and was materially higher than either had marketed.

You cannot conclude that any specific product is more accurate than any other. Nobody has published the comparison. **Zero of 95** companies have independent research measuring their product, **30** publish a score with no method attached, **6** have taken part in someone else's evaluation, and the mean evidence score across the field is **1.99** out of a possible 10 with the highest score anywhere a **7**.

So ask the four questions instead of reading the score. Who wrote the benchmark, and do they sell a product in it. Is that the partial-credit number or the all-pass number. Which practice area and which jurisdiction was it measured in. And who was invited and declined. Those four questions are answerable from public materials in most cases, which is more than can be said for the scores themselves.

# Part V — Rankings

Two rankings appear in this report and neither is a ranking of quality.

We considered publishing a "best legal AI" list and rejected it. We have not tested these products. The independent product benchmarks that exist are eighteen months stale. And as Part IV establishes, no company in this survey has independent research published measuring what it sells. A quality ranking built on that evidence base would be an opinion wearing a table's clothes.

What can be ranked honestly is what is disclosed. Capital raised is disclosed, dated and checkable. Disclosure itself — whether a buyer can verify a claim without booking a sales call — is observable from a public website by anyone, including you, today. Both are published here with their inputs visible, which is the rule this whole series has run on: **no ranking ships without its inputs visible.**

## 20. The capital ranking

The top twenty of the 95, by cumulative disclosed funding. Appendix A carries all 95 with confidence flags, founding years, latest rounds and evidence scores.

| # | Company | Raised | Segment | Region | Founded | Valuation |
|---|---|---|---|---|---|---|
| 1 | Harvey | $1.17B | Firm platform | USA | 2022 | $11B |
| 2 | Legora | $866M | Firm platform | Europe (ex-UK) | 2023 | $5.6B |
| 3 | EvenUp | $385M | Litigation & PI | USA | 2019 | $2B |
| 4 | Norm Ai | $260M+ | Compliance | USA | 2023 | $1.2B |
| 5 | LegalOn Technologies | $200M | Contracts & in-house | Asia-Pacific | 2017 | — |
| 6 | Luminance | $165M | Contracts & in-house | UK | 2015 | — |
| 7 | Blue J | $156M | Tax | Canada | 2015 | — |
| 8 | Eve | $150M | Litigation & PI | USA | 2020 | $1B |
| 9 | Spellbook | $121M | Contracts & in-house | Canada | 2022 | $350M |
| 10 | Wordsmith AI | $114M | Contracts & in-house | UK | 2023 | — |
| 11 | Lawhive | ~$112M | Consumer & small law | UK | 2023 | — |
| 12 | Noxtua | ~$110M | Firm platform | Europe (ex-UK) | 2017 | — |
| 13 | Eudia | up to $105M | Contracts & in-house | USA | 2023 | — |
| 14 | Laurel | ~$100M | Practice ops & billing | USA | — | — |
| 15 | Supio | $91M | Litigation & PI | USA | 2021 | — |
| 16 | SpotDraft | ~$90M | Contracts & in-house | Asia-Pacific | 2017 | — |
| 17 | Klarity | ~$90M | Contracts & in-house | USA | 2017 | — |
| 18 | Bryter | ~$90M | Contracts & in-house | Europe (ex-UK) | 2018 | — |
| 19 | Crosby | $85.8M | AI-native law firm | USA | 2025 | — |
| 20 | MarqVision | ~$85M | IP & patents | Asia-Pacific | 2020 | — |

Three rows in that table need reading with care, and they are the reason the confidence flags exist. **Eudia's $105M is a ceiling**, announced as a facility rather than a draw; treating it as a total would overstate both the company and the market. **Klarity, at number 17, is no longer in this market** — it renamed to Within and repositioned to enterprise process discovery, and its $90M is counted here because it was raised for legal work, not because it is being spent on it. **Laurel has no established founding year**, one of 16 companies in the survey for which we could not find a credible date.

## 21. The evidence score

The ranking nobody else publishes, and the one this report is built around.

Six signals, all observable from a public website, applied mechanically to all 95 companies on 25 August 2026. Mechanically matters: the same inputs always produce the same number, and no judgement enters at the scoring stage.

| Signal | Scoring |
|---|---|
| Publishes a price | Published 2 · Partial 1 · None 0 |
| Names specific security certifications | Yes 1 · No 0 |
| States which models it runs on | Named 2 · Vague pledge 1 · Silent 0 |
| Publishes its own benchmark | Open, with tasks and rubrics 2 · Score only 0 |
| Participated in an independent evaluation | Yes 2 · Not found 0 |
| Independent research measures the product | Yes 1 · No 0 |

Maximum 10. Two design notes. A **closed benchmark scores zero rather than negative**: a headline score with no published method adds nothing a buyer can check, but publishing one is not itself a demerit. And absence of independent benchmarking is recorded as **not found, never as refusal** — most of these companies have simply never appeared in a published evaluation round.

### The result

**Mean 1.99. Median 2. The highest score in the field is 7.** Nobody scored 6, 8, 9 or 10.

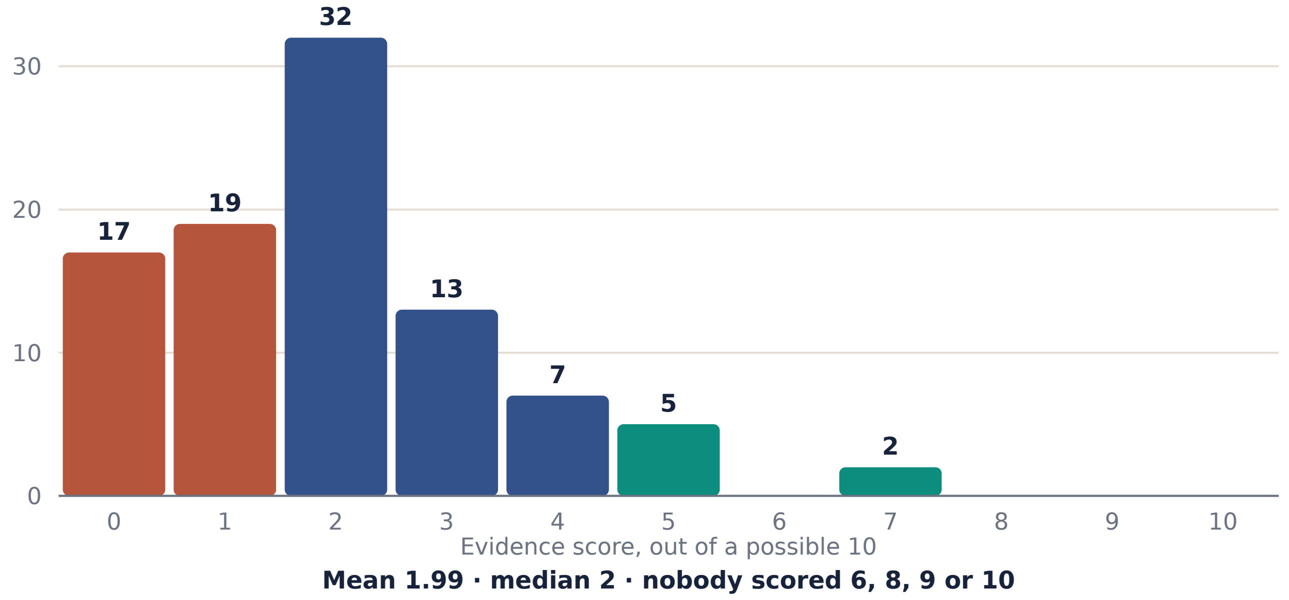


How 95 legal AI companies score on what a buyer can verify

| Score | Companies |
|---|---|
| 7 | 2 |
| 5 | 5 |
| 4 | 7 |
| 3 | 13 |

| Score | Companies |
|---|---|
| 2 | 32 |
| 1 | 19 |
| 0 | 17 |

An industry selling judgement to professionals whose entire discipline is evidence publishes, on average, a fifth of the evidence about itself that it could.

### What the six signals found

| Signal | Result |
|---|---|
| Publish any price | 22 of 95 — 21 a real figure, 1 partial |
| Name specific certifications | 66 of 95 |
| Name their model providers | 14 named · 36 vague · 45 silent |
| Open benchmark, tasks and rubrics published | **2** — Harvey and Paxton AI |
| Closed benchmark, score with no method | 30 |
| Participated in an independent evaluation | **6** |
| Independent research measuring the product | **0** |

The last line is the finding the report exists for. Across **$6.20B** and **269 products**, not one of these companies has had its accuracy measured, in public, by a party with no commercial interest in the result.

#One convention, stated so the per-audience counts can be reconciled. The audience tables below count a company as publishing a price if it discloses **any** figure, which folds in the single partial disclosure — Athennian, whose entry tier is published but whose AI add-on is not. On a strict published-only basis the in-house rate is 16 of 74 and the law firm rate is 10 of 59. Neither convention changes the ordering.

## Money does not buy disclosure

The ten best-funded companies average **3.10**. The other 85 average **1.86**. Capital correlates with transparency, and the correlation is real — but 3.10 out of 10 is the best-funded decile of an industry built on evidence.

The pricing relationship runs the other way and is sharper. Companies that publish a price have raised **$30M on average**; companies that publish none have raised **$76M**. The wall goes up as the cheque size does.

### The audience inversion

The most useful cut of the evidence score is by who the company sells to.

| Audience | Companies | Mean score | Publish a price |
|---|---|---|---|
| Independent lawyers | 16 | **3.38** | 11 (69%) |
| In-house / corporate | 74 | 2.15 | 17 (23%) |
| Law firms | 59 | 2.03 | 11 (19%) |
| People without lawyers | 8 | **1.13** | 4 (50%) |

The ordering is the opposite of the capital ordering, and close to the opposite of what a market ought to produce. The segment with the least money behind it and the least sophisticated buyers — independent lawyers, paying out of their own pockets, with no procurement function and no appetite for a sales call — discloses roughly **two-thirds more** than law firms, the segment with the most procurement machinery, and **57% more** than in-house, the segment with the most money.

The mechanism is not mysterious. A published price is a condition of the sale when the buyer will not take a call. It becomes optional the moment the buyer has a procurement department that will extract the number privately. That is a functioning sales motion. It is not a market that can compare anything.

And the worst score in the survey belongs to the audience least equipped to evaluate what it is buying. On eight companies the figure is fragile and we say so. It still points the wrong way.

## By segment

| Segment | Companies | Mean score |
|---|---|---|
| Tax | 1 | 5.00 |
| Firm platform | 4 | 3.50 |
| Research & knowledge | 9 | 2.56 |
| IP & patents | 7 | 2.29 |
| Contracts & in-house | 32 | 2.25 |
| Consumer & small law | 3 | 1.67 |
| Litigation & PI | 22 | 1.55 |
| Practice ops & billing | 5 | 1.40 |
| Real estate | 3 | 1.33 |
| AI-native law firm | 5 | 1.20 |
| Immigration | 2 | 1.00 |
| Compliance | 2 | 0.50 |

Litigation and personal injury — the segment selling into contingency practices where a wrong answer becomes a malpractice claim — sits at **1.55** on 22 companies. Tax is one company and should be read as one company.

## The top

| Score | Company | What earned it |
|---|---|---|
| 7 | GC AI | $500/mo published, named models, two phases of an independent benchmark |
| 7 | Harvey | The only company with both an open benchmark and independent participation — and it published no score for itself |
| 5 | Blue J | $1,498/yr published, named models |
| 5 | Clearbrief | $300/mo published, independently scored by a state bar |
| 5 | Genie AI | Certificate numbers, expiry dates, and an honest negative |
| 5 | Paxton AI | Reproducible self-run benchmark |
| 5 | Wordsmith AI | Named the contractual basis for retention with each provider |

Harvey scoring joint-top while publishing no price is an unusual shape and worth sitting with. It is the best-funded company in the field, it open-sourced 1,200-plus tasks with roughly 75,000 expert-written rubric criteria, and it pointedly did not report a score for its own product. That is how it should be done. It also still came from a market participant, which remains the structural problem Part IV describes.

### The bottom

Seventeen companies score zero: **Aavalynx, Bench IQ, Casium, Conveyd, Crosby, Darrow, Della AI, DoNotPay, Eve, Haast, Justpoint, LawGeex, LegalMation, Nexl, Semeris, Soxton, Theo Ai.** Between them they have raised **$459M** — 7.4% of all the capital tracked here.

We publish the names because a ranking with its bottom half redacted is not a ranking, and because everything scored is a fact about a public website that anyone can check today. A zero is a narrow statement and we are precise about what it does not say:

- It does **not** say the product is bad. We have not tested any of these products.
- It does **not** say the company is unsafe. It says we could not verify a claim from public sources.
- Some zeros are honest ones. Aavalynx states plainly that it is "working toward ISO and SOC certifications" — candour that scores identically to silence, which is a limitation of the rubric rather than of Aavalynx.
- At least one may be a measurement artefact: Bench IQ's security link points at a trust centre blocked by **`` `robots.txt` ``**, unreadable rather than necessarily absent.
- Some are companies leaving. Della AI's domain is parked and listed for sale at $98,000. LawGeex is a dead brand behind an expired TLS certificate. Theo Ai has repositioned and removed every quantified claim in the process. The score is picking up the silence of an exit, not a disclosure failure.

### The limits of the rubric

It rewards publishing over doing. A company with excellent security practice and no trust page scores as though it had neither. It cannot see behind a login, a gated trust centre or a sales call, and several vendors would hand over pricing and audit reports the moment a buyer asked. That is precisely the point: **a signal a buyer has to book a call to receive is not a signal the market can use.**

It also cannot distinguish a company that will not disclose from one that has nothing yet to disclose. A seed-stage company with no SOC 2 and no benchmark history is scored the same as a $150M company in the same position. The correlation with capital in the section above is the check on that reading, and it is the reason the top-ten comparison matters more than any single row.

# Part VI — What happens next

Everything up to this point is counted. This part is not. What follows are predictions, stated with dates attached so that they can be checked against the record when the next edition of this report is published. Where we think we are likely to be wrong, we say which way.

We publish them because a report that only describes is easy to write and impossible to grade. These are gradeable.

### 1. The first genuinely independent product benchmark lands before August 2027

**Confidence: moderate.** Zero of 95 companies currently have independent research published measuring their product, and the vacuum is now embarrassing enough that somebody will fill it — most plausibly a bar association working group, an academic group extending the LegalBench line, or a large firm publishing its own procurement evaluation.

The State Bar of Nevada's AI Workgroup scoring Clearbrief 40.5 out of 50 is the closest thing that exists, and it is a single product assessed by a single state body. The signal to watch is whether the first real round is vendor-funded. If it is, it will be marketed as independent and the distinction will be harder to see than it was with Ivo's self-run study published under the headline "Independent AI Benchmark".

**We are most likely wrong** if the market decides it does not need one. Procurement has run for two years on demos and references, and nothing forces a change.

### 2. Consolidation continues at the top, and Legora is the acquirer to watch

**Confidence: high.** Legora made five acquisitions between March and July 2026, folding each in as a layer rather than reselling it. A company that has raised $866M against $100M of disclosed ARR has both the currency and the pressure to keep going.

The prediction with a date on it: **at least four more legal AI companies from this survey are acquired before August 2027**, and at least one of them by a company also in this survey rather than by an incumbent. The long tail makes this close to structural — 51 of the 95 have raised under $25M, and 27 under $10M, into a market where the top five hold 46.4%.

### 3. Two to five of the 95 will be dead or departed by August 2027

**Confidence: high, and this is the least pleasant prediction here.** Four already are: Robin AI wound down, LawGeex a dead brand, Della AI a parked domain, Klarity gone to another market. That happened in roughly eighteen months, from a cohort that was still raising.

The tell is not a funding announcement. It is a live website with no maintenance behind it. Robin AI's site still showed product marketing months after the engineering team went to Microsoft. A buyer running diligence against it would have found nothing to indicate the company no longer existed.

### 4. The pricing wall comes down at the small end first, and does not move at the top

**Confidence: moderate to high.** 69% of companies selling to independent lawyers already publish a price, against 19% of those selling to law firms. The mechanism — a buyer who will not take a sales call forces a published number — gets stronger as more of the market sells self-serve.

**We predict the firm-facing figure of 19% has not moved by more than five points in either direction by August 2027.** Enterprise pricing opacity is not an oversight that transparency pressure corrects; it is a negotiating position, and the buyers with procurement departments are the ones least inconvenienced by it.

### 5. Agent reliability numbers get published, and they will be lower than the marketing

**Confidence: moderate on the first half, high on the second.** No vendor currently publishes task completion rates, error rates or intervention frequency for an agent. Somebody will, because a category with 50 products and no measurement cannot stay that way once a buyer asks the obvious question.

When those numbers arrive, expect the same shape Part IV describes for transcription and legal research: a headline figure that looks excellent scored for partial credit, and a materially worse one scored on whether the whole task completed correctly. The gap between those two numbers will be the most useful thing published about agents next year.

### 6. The supervision category grows from one product to a real category

**Confidence: moderate.** One product in 269 watches other agents. That ratio cannot hold against 50 agents and rising, and the buyers accumulating the most autonomy — corporate legal departments, with 27 agent products pointed at them — are the ones carrying the regulatory exposure when an agent is wrong.

**We predict at least five products marketed primarily as agent supervision or verification by August 2027**, and that at least one comes from a compliance company rather than a legal one.

### 7. The interface question resolves against the legal vendors, slowly

**Confidence: low, and this is the prediction we are least sure of.** The connector layer is real: court records, case law corpora and firm knowledge are all now callable by general-purpose assistants over MCP, and Google's Gemini Enterprise already lists Harvey and Legora as sources beneath its own layer.

The bear case for the legal platforms is that the window a lawyer opens each morning stops belonging to a legal software vendor, and the vendors become the data and the domain logic underneath somebody else's agent. The bull case is that legal work needs a purpose-built surface — matter structure, privilege, citation, the record — that a general assistant will not build, and that the integration lists are distribution rather than displacement.

We do not know which is right. The measurable signal to watch is pricing: if firm platforms start pricing per API call rather than per seat, they have accepted the first outcome.

### 8. Pro se stays underfunded, and the interesting work keeps coming from outside the funding market

**Confidence: high, unfortunately.** $196M stands behind products for people without lawyers, against $5.10B for corporate legal departments. Nothing in this survey suggests that ratio is closing: no billion-dollar company sells to this audience, 13 of the 19 products are services with a licence attached rather than software a person operates, and the segment's evidence score is the worst in the survey.

The counter-signal is the one our own filter missed. Courtroom5 has raised roughly $420,000 and sits in Google's integration list alongside Thomson Reuters and Relativity, and shipped as a launch Justice Partner in Anthropic's legal industry release. **We predict that by August 2027 at least**

**two more organisations building for self-represented litigants reach comparable distribution on comparable capital** — because the price of that distribution is now a working connector and a corpus worth querying, not a Series B.

### How to grade us

Every prediction above has a date, a direction and, where possible, a number. When we publish the next edition we will reprint this list with each item marked correct, wrong or unresolved, before we make any new ones. A prediction that cannot be marked wrong was not worth making.

# Appendix A — The 95 companies

Ranked by cumulative disclosed funding. **Conf.** is the confidence flag described in the front matter. **Ev.** is the evidence score out of 10 from Chapter 21. All figures as of 25 August 2026.

| # | Company | Raised | Confidence | Segment | HQ | Founded | Valuation | Latest round | Evidence |
|---|---|---|---|---|---|---|---|---|---|
| 1 | Harvey | $1.17B | high | Firm platform | USA | 2022 | $11B | $200M growth · Mar 2026 | 7 |
| 2 | Legora | $866M | high | Firm platform | Sweden | 2023 | $5.6B | $50M Series D ext. · Apr 2026 | 3 |
| 3 | EvenUp | $385M | high | Litigation & PI | USA | 2019 | $2B | $150M Series E · Oct 2025 | 2 |
| 4 | Norm Ai | $260M+ | high | Compliance | USA | 2023 | $1.2B | $120M Series C · Jul 2026 | 1 |
| 5 | LegalOn Technologies | $200M | high | Contracts & in-house | Japan | 2017 | — | $50M Series E · Jul 2025 | 3 |
| 6 | Luminance | $165M | low | Contracts & in-house | UK | 2015 | — | $75M Series C · Feb 2025 | 2 |
| 7 | Blue J | $156M | med | Tax | Canada | 2015 | — | $122M Series D · Aug 2025 | 5 |
| 8 | Eve | $150M | med | Litigation & PI | USA | 2020 | $1B | $103M Series B · Sep 2025 | 0 |
| 9 | Spellbook | $121M | high | Contracts & in-house | Canada | 2022 | $350M | $40M debt facility · Mar 2026 | 3 |
| 10 | Wordsmith AI | $114M | high | Contracts & in-house | UK | 2023 | — | $14M Series B ext. · Aug 2026 | 5 |
| 11 | Lawhive | ~$112M | med | Consumer & small law | UK | 2023 | — | $60M Series B · Feb 2026 | 2 |
| 12 | Noxtua | ~$110M | low | Firm platform | Germany | 2017 | — | €80.7M Series B · Apr 2025 | 3 |
| 13 | Eudia | up to $105M | low | Contracts & in-house | USA | 2023 | — | Series A · Feb 2025 | 3 |
| 14 | Laurel | ~$100M | low | Practice ops & billing | USA | — | — | ~$100M · Aug 2025 | 2 |
| 15 | Supio | $91M | high | Litigation & PI | USA | 2021 | — | $60M Series B · May 2025 | 2 |
| 16 | SpotDraft | ~$90M | med | Contracts & in-house | India | 2017 | — | $8M Series B ext. · Jan 2026 | 1 |
| 17 | Klarity | ~$90M | low | Contracts & in-house | USA | 2017 | — | $70M Series B · Jun 2024 | 2 |
| 18 | Bryter | ~$90M | low | Contracts & in-house | Germany | 2018 | — | Series B · 2021 | 2 |
| 19 | Crosby | $85.8M | high | AI-native law firm | USA | 2025 | — | $60M Series B · Apr 2026 | 0 |
| 20 | MarqVision | ~$85M | med | IP & patents | USA / Korea | 2020 | — | $48M Series B · Sep 2025 | 1 |
| 21 | Ivo | ~$76M | med | Contracts & in-house | USA | 2022 | — | $55M Series B · Jan 2026 | 2 |
| 22 | Orbital | $75M | high | Real estate | UK | 2018 | — | $60M Series B · Jan 2026 | 2 |
| 23 | GC AI | $73M | high | Contracts & in-house | USA | 2023 | $555M | $60M Series B · Nov 2025 | 7 |
| 24 | Patlytics | ~$65M | high | IP & patents | USA | 2023 | — | $40M Series B · Apr 2026 | 2 |
| 25 | Robin AI | $61.5M | high | Contracts & in-house | UK | 2019 | — | $25M Series B+ · Nov 2024 | 2 |

| # | Company | Raised | Confidence | Segment | HQ | Founded | Valuation | Latest round | Evidence |
|---|---|---|---|---|---|---|---|---|---|
| 26 | Darrow | ~$60M | med | Litigation & PI | Israel | 2020 | — | $35M Series B · Sep 2023 | 0 |
| 27 | Manifest OS | $60M | med | AI-native law firm | USA | — | $750M | $60M Series A · Apr 2026 | 1 |
| 28 | Solve Intelligence | $55M | high | IP & patents | USA / UK | 2023 | — | $40M Series B · Dec 2025 | 2 |
| 29 | Summize | $50M+ | low | Contracts & in-house | UK | 2018 | — | $50M growth · Jan 2026 | 3 |
| 30 | DeepJudge | ~$50M | med | Research & knowledge | Switzerland | 2022 | — | $41.2M Series A · Nov 2025 | 2 |
| 31 | Definely | ~$47M | med | Contracts & in-house | UK | 2019 | — | $30M Series B · Jun 2025 | 1 |
| 32 | Enter | ≥$43M | low | Litigation & PI | Brazil | 2023 | ~$370M | R$200M, Founders Fund · May 2026 | 3 |
| 33 | LawGeex | ~$42M | med | Contracts & in-house | Israel | 2014 | — | Series C · 2019 | 0 |
| 34 | Athennian | ~$41M | med | Practice ops & billing | Canada | 2017 | — | $33M Series B · Mar 2022 | 2 |
| 35 | DeepIP | $40M | high | IP & patents | USA / France | 2024 | — | $25M Series B · Mar 2026 | 3 |
| 36 | Sandstone | $40M | high | Contracts & in-house | USA | 2025 | — | $30M Series A · Jun 2026 | 2 |
| 37 | Juro | ~$38M | med | Contracts & in-house | UK | 2016 | — | $23M Series B · Jan 2022 | 1 |
| 38 | Jus Mundi | ~$36M | med | Research & knowledge | France | 2019 | — | $22M Series B · Sep 2024 | 2 |
| 39 | Boundless | ~$33M | med | Immigration | USA | 2017 | — | $25M Series B · Apr 2021 | 2 |
| 40 | Trellis | ~$31M | med | Litigation & PI | USA | 2018 | — | $15M Series B · Sep 2023 | 3 |
| 41 | Paxton AI | ~$28M | high | Research & knowledge | USA | 2022 | — | $22M Series A · Jan 2025 | 5 |
| 42 | Checkbox | ~$28M | med | Contracts & in-house | Australia | 2016 | — | $23M Series A · Jan 2026 | 1 |
| 43 | DoNotPay | ~$27.5M | low | Consumer & small law | USA | 2015 | — | Series A · 2021 | 0 |
| 44 | Tomorro | ~€25M+ | med | Contracts & in-house | France | — | — | €25M · May 2025 | 1 |
| 45 | Finch | ~$23.8M | high | Litigation & PI | USA | 2024 | — | $20M · Oct 2025 | 1 |
| 46 | Nexl | ~$23M | med | Practice ops & billing | Australia | 2018 | — | $23M Series B · Oct 2025 | 0 |
| 47 | Genie AI | ~$22M | med | Contracts & in-house | UK | 2017 | — | £13.3M Series A · Oct 2024 | 5 |
| 48 | Case Status | ~$21M | med | Litigation & PI | USA | 2018 | — | $19.7M Series B · Oct 2024 | 1 |
| 49 | Alexi | ~$20M | low | Litigation & PI | Canada | 2017 | — | $11M Series A + CAD$4.5M debt · 2025 | 3 |
| 50 | Twin1 | $20M | high | Research & knowledge | USA / UK | 2025 | — | $20M seed · Aug 2026 | 4 |
| 51 | Newcode | ~$20M | low | Firm platform | Norway | — | — | $13.5M Series A · Aug 2026 | 1 |
| 52 | Tradespace | ~$19.5M | med | IP & patents | USA | 2021 | — | $15M Series A · Jan 2026 | 2 |
| 53 | Lexroom | ~$19M | med | Research & knowledge | Italy | — | — | $19M · Sep 2025 | 4 |
| 54 | Avvoka | ~£14M+ | low | Contracts & in-house | UK | 2017 | — | £14M growth · Mar 2026 | 2 |

| # | Company | Raised | Confidence | Segment | HQ | Founded | Valuation | Latest round | Evidence |
|---|---|---|---|---|---|---|---|---|---|
| 55 | Pocketlaw | ~$17M | low | Contracts & in-house | Sweden | 2018 | — | ~€5M · Dec 2024 | 2 |
| 56 | PointOne | $16M | high | Practice ops & billing | USA | — | — | $16M Series A · Mar 2026 | 2 |
| 57 | LegalFly | ~€15M | med | Contracts & in-house | Belgium | 2023 | — | €15M Series A · 2024 | 1 |
| 58 | Zuva | ~$15.8M | med | Contracts & in-house | Canada | 2021 | — | CAD$20M Series A · Sep 2021 | 4 |
| 59 | LegalMation | ~$15M | low | Litigation & PI | USA | 2015 | — | $15M · Nov 2023 | 0 |
| 60 | Streamline AI | $14M | high | Contracts & in-house | USA | 2021 | — | $8.6M Series A · Jul 2025 | 4 |
| 61 | Centari | ~$14M | med | Contracts & in-house | USA | — | — | $8.4M Series A · Sep 2025 | 2 |
| 62 | Hona | ~$13M | low | Litigation & PI | USA | 2021 | — | $9.5M Series A · Jun 2024 | 1 |
| 63 | Haast | $12M | med | Compliance | Australia | — | — | $12M · Apr 2026 | 0 |
| 64 | IPRally | ~€12M | med | IP & patents | Finland | 2018 | — | €10M Series A · Mar 2023 | 4 |
| 65 | General Legal | $11.5M | med | AI-native law firm | USA | 2026 | — | $11.5M seed · Mar 2026 | 2 |
| 66 | Callidus Legal AI | $10M | high | Litigation & PI | USA | 2023 | — | $10M · Jul 2025 | 4 |
| 67 | Flank | $10M | med | Contracts & in-house | Germany | — | — | $10M · Jun 2025 | 2 |
| 68 | Chamelio | $10M | high | Contracts & in-house | USA | — | — | $10M seed · Jan 2026 | 1 |
| 69 | Theo Ai | ~$9.8M | high | Litigation & PI | USA | 2024 | — | $3.4M seed · Nov 2025 | 0 |
| 70 | Moritz | $9M | med | AI-native law firm | USA | 2026 | — | $9M · May 2026 | 3 |
| 71 | Josef | ~$8M | low | Practice ops & billing | Australia | 2017 | — | ~$5.2M pre-Series A · Nov 2022 | 1 |
| 72 | Clearbrief | ~$7.5M | med | Litigation & PI | USA | 2020 | — | $4M Series A · Jun 2024 | 5 |
| 73 | Bench IQ | $7.4M | high | Research & knowledge | Canada | 2023 | — | $5.3M seed · Aug 2025 | 0 |
| 74 | Justpoint | ~$6.9M | low | Litigation & PI | USA | 2019 | — | $6.9M seed ext. · 2022 | 0 |
| 75 | Wexler AI | ~$6.7M | high | Litigation & PI | UK | 2023 | — | $5.3M seed · Sep 2025 | 1 |
| 76 | TermScout | ~$6.3M | med | Contracts & in-house | USA | 2018 | — | $5M seed · May 2022 | 3 |
| 77 | midpage | $6.2M | high | Research & knowledge | USA | 2022 | — | $4M seed · Jun 2025 | 4 |
| 78 | Fileread | $6M | med | Litigation & PI | USA | 2021 | — | $6M seed · Jul 2023 | 2 |
| 79 | Hello Divorce | ~$5.8M | low | Consumer & small law | USA | 2018 | — | $3.25M seed · Oct 2022 | 3 |
| 80 | Dealstack | $5.5M | low | Contracts & in-house | UK | — | — | $5.5M · Apr 2025 | 1 |
| 81 | Syntracts | $5.3M | med | Contracts & in-house | USA | — | — | $5.3M seed · Oct 2025 | 2 |
| 82 | Casium | $5M | high | Immigration | USA | 2023 | — | $5M seed · Oct 2025 | 0 |
| 83 | Mary Technology | ~$4.5M | low | Litigation & PI | Australia | 2024 | — | A$7M seed · Mar 2026 | 1 |
| 84 | Semeris | $4.3M | med | Research & knowledge | UK | — | — | $4.3M · Feb 2025 | 0 |
| 85 | TrialView | $4.1M | med | Litigation & PI | UK | — | — | $4.1M seed · Nov 2025 | 2 |
| 86 | Turbo Law | $3.8M | high | Litigation & PI | USA | 2025 | — | $3.8M pre-seed · Jun 2026 | 1 |
| 87 | Marveri | $3.5M | high | Contracts & in-house | USA | 2023 | — | $3.5M seed · May 2025 | 2 |

| # | Company | Raised | Confidence | Segment | HQ | Founded | Valuation | Latest round | Evidence |
|---|---|---|---|---|---|---|---|---|---|
| 88 | Skribe.ai | $3.5M | med | Litigation & PI | USA | 2022 | — | $3.5M seed · May 2024 | 2 |
| 89 | Conveyd | ~£2.5M | low | Real estate | UK | 2025 | — | £2.5M seed · Dec 2025 | 0 |
| 90 | NLPatent | $3M | med | IP & patents | Canada | — | — | $3M seed · Nov 2025 | 2 |
| 91 | Pandektes | ~€2.9M | low | Research & knowledge | Denmark | — | — | €2.9M seed · Feb 2025 | 2 |
| 92 | Keith | ~£2M | low | Real estate | UK | 2025 | — | £2M · Mar 2026 | 2 |
| 93 | Soxton | $2.5M | low | AI-native law firm | UK | 2025 | — | $2.5M · Dec 2025 | 0 |
| 94 | Della AI | $2.5M | med | Contracts & in-house | France / UK | 2019 | — | $2.5M seed · Mar 2022 | 0 |
| 95 | Aavalynx | ~£1.5M | med | Litigation & PI | UK | 2025 | — | £1.5M pre-seed · Aug 2026 | 0 |

# Appendix B — Every evidence score, with its inputs

All 95 companies, sorted by score then alphabetically. **Models** is what the company says about its providers. **Benchmarks** is open (its own, with tasks and rubrics published), closed (its own, score only) or independent (participated in someone else's). There is no column for third-party research because the answer is the same for every row: none.

| Score | Company | Price | Certifications | Models | Benchmarks |
|---|---|---|---|---|---|
| **7** | GC AI | Published | Named | Named | Closed + Independent |
| **7** | Harvey | — | Named | Named | Open + Independent |
| **5** | Blue J | Published | Named | Named | — |
| **5** | Clearbrief | Published | Named | — | Independent |
| **5** | Genie AI | Published | Named | Named | — |
| **5** | Paxton AI | Published | Named | — | Open |
| **5** | Wordsmith AI | — | Named | Named | Independent |
| **4** | Callidus Legal AI | Published | Named | Vague | Closed |
| **4** | IPRally | — | Named | Vague | Independent |
| **4** | Lexroom | Published | Named | Vague | — |
| **4** | midpage | Published | Named | Vague | — |
| **4** | Streamline AI | Published | Named | Vague | — |
| **4** | Twin1 | Published | Named | Vague | — |
| **4** | Zuva | Published | Named | Vague | Closed |
| **3** | Alexi | — | Named | — | Independent |
| **3** | DeepIP | — | Named | Named | Closed |
| **3** | Enter | — | Named | Named | Closed |
| **3** | Eudia | — | Named | Named | — |
| **3** | Hello Divorce | Published | — | Vague | — |
| **3** | LegalOn Technologies | — | Named | Named | Closed |
| **3** | Legora | — | Named | Named | Closed |
| **3** | Moritz | Published | Named | — | — |
| **3** | Noxtua | — | Named | Named | Closed |
| **3** | Spellbook | — | Named | Named | — |
| **3** | Summize | — | Named | Named | Closed |
| **3** | TermScout | Published | Named | — | — |
| **3** | Trellis | Published | Named | — | — |
| **2** | Athennian | Partial | Named | — | — |
| **2** | Avvoka | — | Named | Vague | — |
| **2** | Boundless | Published | — | — | Closed |
| **2** | Bryter | — | Named | Vague | — |
| **2** | Centari | — | Named | Vague | — |
| **2** | DeepJudge | — | Named | Vague | Closed |
| **2** | EvenUp | — | Named | Vague | — |
| **2** | Fileread | — | Named | Vague | — |
| **2** | Flank | — | Named | Vague | — |
| **2** | General Legal | Published | — | — | — |
| **2** | Ivo | — | Named | Vague | Closed |
| **2** | Jus Mundi | Published | — | — | — |
| **2** | Keith | Published | — | — | Closed |
| **2** | Klarity | — | Named | Vague | — |

| Score | Company | Price | Certifications | Models | Benchmarks |
|---|---|---|---|---|---|
| **2** | Laurel | — | Named | Vague | — |
| **2** | Lawhive | Published | — | — | — |
| **2** | Luminance | — | Named | Vague | — |
| **2** | Marveri | — | Named | Vague | — |
| **2** | NLPatent | — | Named | Vague | Closed |
| **2** | Orbital | — | Named | Vague | — |
| **2** | Pandektes | — | Named | Vague | — |
| **2** | Patlytics | — | Named | Vague | Closed |
| **2** | Pocketlaw | — | Named | Vague | — |
| **2** | PointOne | — | Named | Vague | — |
| **2** | Robin AI | — | Named | Vague | — |
| **2** | Sandstone | — | Named | Vague | — |
| **2** | Skribe.ai | Published | — | — | Closed |
| **2** | Solve Intelligence | — | Named | Vague | Closed |
| **2** | Supio | — | Named | Vague | — |
| **2** | Syntracts | — | Named | Vague | — |
| **2** | Tradespace | — | Named | Vague | — |
| **2** | TrialView | — | — | Named | — |
| **1** | Case Status | — | Named | — | — |
| **1** | Chamelio | — | Named | — | — |
| **1** | Checkbox | — | Named | — | — |
| **1** | Dealstack | — | Named | — | — |
| **1** | Definely | — | Named | — | — |
| **1** | Finch | — | — | Vague | Closed |
| **1** | Hona | — | Named | — | — |
| **1** | Josef | — | Named | — | — |
| **1** | Juro | — | Named | — | — |
| **1** | LegalFly | — | Named | — | — |
| **1** | Manifest OS | — | Named | — | Closed |
| **1** | MarqVision | — | Named | — | Closed |
| **1** | Mary Technology | — | — | Vague | Closed |
| **1** | Newcode | — | — | Vague | — |
| **1** | Norm Ai | — | Named | — | Closed |
| **1** | SpotDraft | — | Named | — | — |
| **1** | Tomorro | — | Named | — | — |
| **1** | Turbo Law | — | — | Vague | — |
| **1** | Wexler AI | — | Named | — | Closed |
| **0** | Aavalynx | — | — | — | Closed |
| **0** | Bench IQ | — | — | — | — |
| **0** | Casium | — | — | — | Closed |
| **0** | Conveyd | — | — | — | Closed |
| **0** | Crosby | — | — | — | — |
| **0** | Darrow | — | — | — | — |
| **0** | Della AI | — | — | — | — |
| **0** | DoNotPay | — | — | — | — |
| **0** | Eve | — | — | — | Closed |
| **0** | Haast | — | — | — | — |
| **0** | Justpoint | — | — | — | — |
| **0** | LawGeex | — | — | — | Closed |

| Score | Company | Price | Certifications | Models | Benchmarks |
|---|---|---|---|---|---|
| **0** | LegalMation | — | — | — | — |
| **0** | Nexl | — | — | — | Closed |
| **0** | Semeris | — | — | — | Closed |
| **0** | Soxton | — | — | — | — |
| **0** | Theo Ai | — | — | — | — |

# Appendix C — Adjacent companies, counted nowhere

Eighteen companies that shape this market and are excluded from every figure in this report, because they are not AI-native: they are legal technology or adjacent businesses that added AI to an existing platform. Their capital is not in the $6.20B.

| Company | Market | Raised | Valuation | Note |
|---|---|---|---|---|
| Clio | Practice management + legal research | $1.3B | $5B | Bought vLex for $1B in 2025, then raised a $500M Series G. AI is a layer on a 2008 practice-management business. |
| Filevine | Litigation case management | ~$626M | ~$3B est. | $400M across two 2025 rounds — the largest single legal tech raise of that year. Repositioning as 'legal intelligence'; bought Pincites and Parrot. |
| Icertis | Enterprise contract intelligence | $520M | ~$5B | Acquired the AI contract-review startup Dioptra in Nov 2025 rather than building. |
| Ontra | Private-markets contract automation | ~$370M | — | AI plus a lawyer network for NDAs and side letters. Services-heavy hybrid, founded as InCloudCounsel. |
| PatSnap | IP and innovation analytics | ~$350M | >$1B (2021) | $300M SoftBank-led Series E in 2021; valuation mark is stale. |
| Ironclad | Contract lifecycle management | $334M | $3.2B (2022) | No round since the Jan 2022 Series E. The $3.2B mark is peak-ZIRP and unlikely to hold. |
| Everlaw | Cloud e-discovery | ~$300M | >$2B (2021) | $202M Series D in Nov 2021, with generative AI layered on since. |
| ContractPodAi | Enterprise CLM + 'Leah' AI | ~$170M | — | $115M SoftBank Vision Fund 2 round in 2021. |
| LinkSquares | CLM and contract analytics | ~$160M | $800M (2022) | $100M Series C in Apr 2022; valuation stale. |
| Fenergo | KYC / client lifecycle | ~$160M VC | — | Fintech compliance rather than legal. Now PE-owned (Astorg/Bridgepoint). |
| Steno | Court reporting + transcript AI | ~$150M incl. debt | — | $49M Series C in Mar 2026. Services business with an AI transcript layer. |
| Sirion | Enterprise CLM | $110M+ | — | Took a majority investment in Feb 2026. |
| Relativity | E-discovery / data intelligence | PE-owned (Silver Lake) | — | Confidentially filed for an IPO in Mar 2026 — would be the first legal tech listing since 2021. |
| CS DISCO | E-discovery | Public (NASDAQ: LAW) | public | Listed in 2021; 'Cecilia' AI added later. |
| Agiloft | No-code enterprise CLM | $45M VC + KKR control | — | KKR took a majority stake in May 2024. |
| LegalZoom | Consumer/SMB legal documents | Public (NASDAQ: LZ) | public | The original consumer legal platform; IPO'd 2021. |
| Rocket Lawyer | SMB legal subscriptions | ~$268M reported | — | Subscription documents plus attorney access. |
| Legalist | AI-sourced litigation finance | ~$2B AUM (funds) | — | A fund, not a software company — the AI sources cases to finance. Raised a $415M fund in 2026. |

# Appendix D — Exits and status changes

## Acquisitions

| Company | Acquirer | Price | Date | Note |
|---|---|---|---|---|
| vLex | Clio | $1.0B | Nov 2025 | Global legal content plus the Vincent AI research assistant and Docket Alarm. |
| Casetext | Thomson Reuters | $650M | Aug 2023 | CoCounsel — the deal that opened the legal AI market. |
| Lexion | Docusign | $165M | May 2024 | AI-first contract management, spun out of AI2 in Seattle. |
| Evisort | Workday | Undisclosed | Nov 2024 | Contract intelligence; had raised roughly $155M. |
| Avantia | Carta | Undisclosed | May 2026 | UK alternative legal services; Carta launched an AI-first law firm division around it. |
| Wexler AI | Legora | Undisclosed | Jul 2026 | Litigation fact verification — Legora's fifth acquisition since March 2026. |
| Qura | Legora | Undisclosed | Apr 2026 | Swedish/EU legal research; had raised €2.1M. |
| Pincites | Filevine | Undisclosed | Dec 2025 | YC-backed AI contract redlining. |
| Dioptra | Icertis | Undisclosed | Nov 2025 | YC-backed AI contract review. |
| Robin AI (services arm) | Scissero | Undisclosed | Dec 2025 | Managed legal services sold off amid funding trouble; the tech team went to Microsoft. |
| Rowan Patents | Clarivate | Undisclosed | Jul 2024 | Integrated patent drafting and prosecution workspace. |
| Compliance.ai | Archer | Undisclosed | Feb 2024 | Regulatory change management for financial services. |

## Status changes

Fifteen companies in the survey are not, or are no longer, what their name and website suggest.

| Company | Status | What happened |
|---|---|---|
| Robin AI | Defunct | Wound down through 2025. The managed-services arm was sold to law firm Scissero in December 2025 and the engineering team was acqui-hired by Microsoft in January 2026. The website still shows product marketing; treat it as stale. |
| LawGeex | Defunct | Dead brand. LegalSifter acquired the enterprise business in 2023 and the successor product is LegalSifter ReviewPro. lawgeex.com now serves an expired certificate. |
| Della AI | Acquired | Acquired by Wolters Kluwer in December 2022 and the brand retired. The technology now ships as Legisway Analyzer, and the della.ai domain is listed for sale. |
| Klarity | Left the legal market | Renamed Within and repositioned from contract review to enterprise process discovery. Its current buyers are finance, IT, GTM and operations — there is no legal product on the site. |
| Wexler AI | Acquired | Acquired by Legora in July 2026, ten months after its seed round. Being folded in as Legora's fact layer rather than sold standalone. |
| Callidus Legal AI | Renamed | Rebranded to StrongSuit in November 2025; callidusai.com redirects. The company is active and shipping. |
| NLPatent | Renamed | Rebranded to Clerq in August 2026; nlpatent.com redirects to clerq-ip.com. |
| Pocketlaw | Renamed | Rebranded to Miramis, with the positioning shifting from SME self-serve templates to mid-market contract lifecycle management. |
| Noxtua | Renamed | Founded as Xayn and rebranded to Noxtua in 2024. |
| Tomorro | Renamed | Formerly named Leeway. |
| Flank | Renamed | Formerly named Legal OS; renamed Flank in May 2024. |

| Company | Status | What happened |
|---|---|---|
| Orbital | Renamed | Formerly Orbital Witness; the earlier product name Orbital Copilot still survives in app URLs. |
| Justpoint | Restructured | Split into Justpoint Inc., which disclaims providing legal services, and Justpoint Law LLP, an Arizona ABS-licensed firm approved in July 2025. |
| Fileread | Quiet | Operating, but no funding since its 2023 seed and no visible activity through 2026. |
| Soxton | Thin public surface | Live but effectively invisible: a landing page with no product or pricing detail. Everything known comes from its December 2025 funding announcement. |

# Appendix E — Method, sources and corrections

### How the survey was built

The funding survey was assembled from company statements, funding announcements and press reporting, reconciled across sources, and recorded as a typed data module with a confidence flag and a source per row. Where sources disagreed we recorded the figure we could trace to the most primary source and flagged the row; where we could not choose, we recorded the lower figure and flagged it. Eight reconciliation traps we hit repeatedly — announced facilities counted as draws, rounds double-counted across extensions, valuations carried forward from stale marks — are documented in the maintenance standard published alongside the data.

The product catalogue was built by reading each company's own site and recording every product marketed under a distinct name, its type, and the audiences the company's own marketing and pricing address. Audience tagging follows the company's stated buyer, not our view of who might benefit.

The evidence audit was run over five days by four independent passes across all 95 companies, compiled, then scored mechanically from the recorded signals.

### Known gaps

Ordered by how much the report needed them.

| Gap | Why it matters |
|---|---|
| Arizona ABS licence count and holders | Underpins Chapter 4; we state no count |
| Utah sandbox status | Same |
| SRA authorisations in England and Wales | The route Lawhive, Keith and Moritz operate under |
| Practice-area coverage grid across all 269 products | We have anecdote, not the grid |
| A systematic cost-per-task table | Chapter 17; the number procurement asks first |
| Adoption data against funding | What proportion of firms use any of this |
| Non-US, non-European markets | Korea, Japan beyond LegalOn, China, India beyond SpotDraft, MENA, Africa, Latin America beyond Enter |

### Corrections we have already made

We publish these because a survey that never reports its own errors is not being checked.

- We stated internally that Harvey and Legora together held 45% of tracked capital. The correct figure is **32.8%**; 46.4% is the top five.
- Our published post on tools for self-represented litigants claimed the market contained nothing for debt collection defence, and described itself as complete. **Courtroom5 falsified both claims.** We rebuilt the post around the blind spot rather than patching it, and Courtroom5 now appears throughout this report.
- Robin AI was carried as an operating company until we established it had wound down. It is now flagged defunct, and its website still shows product marketing.
- Klarity was carried as a legal AI company. It is now Within and has left the legal market; its capital is counted, and so is the one product it still lists — a row we keep flagged rather than delete, so that the change is visible in the data.

### Corrections welcome

If a figure here is wrong about your company, tell us. We will fix it in the data module, say publicly that we did, and date the change. Two things we will not do: remove a figure because it is unflattering, or publish a correction without saying what it corrected.

The underlying data is published at /research/legal-ai-funding and /research/legal-ai-products.

---

The State of Legal AI, August 2026. Published by IQdoc on 25 August 2026. The underlying data is public at /research/legal-ai-funding and /research/legal-ai-products. Corrections are welcome and will be dated.