IQdoc Applied AI for legal.

IQDOC RESEARCH

The State of Legal AI, November 2026

A survey of the AI-native legal market, organised by who the buyer is: $7.46B across 105 companies, 312 products, 69 agents, and an audit of what a buyer can actually verify.

Capital tracked	Companies	Products	Mean evidence score
$7.46B	**105**	**312**	**1.96 / 10**

Data compiled 8 November 2026. Evidence audited 8 November 2026. Published by IQdoc — applied AI for legal.



The November 2026 edition. Every figure is as at 8 November 2026. The August edition remains published, unrevised, at the archive. Six corrections to it are listed in Appendix E, and Part VI grades the eight predictions it made.

Contents

Front matter 3

Part I — The market 6
1. Where the money went 6
2. The disclosure vacuum 8
3. What got built 10
4. The structural shift 12
5. Exits, failures and renames 14

Part II — The four audiences 17
6. In-house and corporate legal 18
7. Law firms 20
8. Independent lawyers and small firms 22
9. People without lawyers 25

Part III — Agents 28
10. The taxonomy 28
11. Who is building them 29
12. What the autonomous ones actually do 31
13. The connector layer 32
14. The supervision gap 33

Part IV — Models, benchmarks and skills 35
15. What the model labs shipped 35
16. The benchmarks that exist and who administers them 36
17. Model rankings on legal tasks 38
18. Partial credit versus all-pass 39
19. Hallucination, and instructions to the machines 40

Part V — Rankings 43
20. The capital ranking 43
21. The evidence score 44

Part VI — What happens next 49

Appendices 52
Appendix A — The 105 companies 52
Appendix B — Every evidence score, with its inputs 55
Appendix C — Adjacent companies, counted nowhere 58
Appendix D — Exits and status changes 59
Appendix E — Method, sources and corrections 61

Front matter

What this is

This is the **November 2026 edition** of a survey first published on 25 August 2026. It reports the AI-native legal market as it stood on **8 November 2026**: **$7.46B** of disclosed capital across **105** companies, **312** named products from **104** companies, and an audit of what each company lets an outsider verify about its own claims.

The August edition remains published, unrevised, at **/research/archive/2026-08-31**. Nothing in it has been edited to match this one. Where this edition corrects it, the correction is in the text and listed in Appendix E.

It is organised by **who the buyer is** rather than by product category, because that is where the 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 changed since August

	August	November
Companies	95	**105**
Capital tracked	$6.20B	**$7.46B**
Products	269	**312**
Companies disclosing a dollar revenue figure	14	**14**
Top five share of capital	46.4%	**53.3%**
Mean evidence score	1.99	**1.96**
Agent products	50 of 264	**69 of 307**
Products supervising other agents	1	**0**
Companies with independent research published about them	0	**0**

Two rounds account for most of it. **Harvey raised $550M on 9 September at a valuation reported as $15.5B** — contested, with Bloomberg and TNW reporting $15.6B against four outlets at $15.5B, and we carry the better-sourced figure and flag the row. **Legora raised $550M on 21 October at $8.8B post-money.** Between them, $1.1B of the $1.26B increase in tracked capital.

What we counted

The 105. 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 in Appendix C, without their capital counted anywhere.

The 312. 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. **307** come from the 105. **Five** come from Courtroom5, which our funding threshold cannot see and which builds the most substantial software we have found for self-represented litigants.

For this edition the catalogue was **re-read site by site, all 97 companies that had one, on 8 November**. That matters more than the net number suggests: 64 products were added, 25 withdrawn and 23 renamed — **112 of 312 rows touched**. A catalogue of this market decays in weeks.

The evidence audit. Six observable signals per company, scored mechanically. Published in full as a ranking before this report, so the arguments surfaced before they landed in a PDF.

What we did not count

We have **not tested any of these products**. Nothing here 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 — as Part IV shows — neither has anybody else.

We did not count revenue. Fourteen of 105 disclose a dollar figure, the same fourteen as in August, and they span five reporting years and three incompatible measures.

We did not resolve the regulatory picture. The count of Arizona alternative business structure licences, the status of Utah's sandbox, and authorisations in England and Wales remain the largest known holes in this report, unchanged since August.

Our sourcing is English-language. Latin America appears twice 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

Flag	Meaning
high	Stated by the company or in a primary filing, consistent across sources
med	Assembled from announced rounds, or reported credibly 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 are cases where a company announced a facility rather than a draw — Eudia's "up to $105M" is a ceiling — and reporting the ceiling as a total would overstate the market. In this edition both Harvey and Legora carry a medium flag on their cumulative totals, because neither $550M round came with a stated total and both figures are our own arithmetic.

Dates, and two of them

Every figure is as at **8 November 2026** unless the text gives another date.

One exception is important enough to state here. The evidence audit was run in two waves: **95 companies between 20 and 25 August**, and **ten on 8 November**. The rubric was applied identically; the dates were not. An August row may understate a company that has published a price since. The data module records this on every row and Appendix B marks which wave each came from.

This market decays fast, which is why the August edition is archived rather than overwritten, and why every row carries a date and a confidence flag. A dated row can be corrected. A tidy ranking cannot tell you which of its rows to distrust.

A declared interest

IQdoc builds transcription and record products for legal work and publishes its own accuracy benchmarks. Four of the 312 products in this catalogue are in that category, and Chapter 18 discusses a measure we compete on. Those passages are marked. IQdoc is not included in the funding survey, the product catalogue or the evidence audit.

Part I — The market

This is the November 2026 edition of a report we first published in August. The method is unchanged: we track disclosed capital across AI-native legal companies, catalogue the products they sell by name, and record what each has put on the record about its own business. What has changed is the market. We now track **$7.46B** across **105** companies and **312** products, against $6.20B, 95 companies and 269 products in August.

Two cautions. Every figure is as publicly disclosed and not independently audited; the totals are assembled from announced rounds rather than stated by the companies themselves, and most carry a medium or low confidence flag, which we have not recomputed for the ten companies added since August. And this is a snapshot taken on 8 November 2026. Chapter 3 is the clearest evidence yet that a snapshot of this market goes stale faster than anyone publishing one would like.

1. Where the money went

One hundred and five AI-native legal companies have raised a combined **$7.46B**. The top five hold **53.3%** of it and the top ten hold **62.8%**. Below that the market thins out immediately: the median company has raised **$20M**, **58 of the 105** have raised under $25M, and **34** have raised under $10M.



Cumulative share of all tracked capital by company rank

The finding of this chapter is not the total. It is that concentration increased sharply in eleven weeks. The top five held **46.4%** in August and hold **53.3%** now — a rise of nearly seven points in a quarter in which the company count grew by ten. Two rounds did almost all of it. Harvey raised **$550M** at a $15.5B valuation on 9 September; Legora raised **$550M** at $8.8B on 21 October. That is $1.1B of the $1.26B by which our tracked total rose: **87% of all new capital in the window went to two companies.** The remaining $160M or so covers the ten companies added to the survey and every other round closed by the other 103.

Harvey alone now accounts for **23.0%** of all tracked capital, up from 18.8%. The top three hold **47.1%** — more than the top five held last edition. By rank 30 the cumulative share is **84.3%**, leaving the remaining 75 companies under a sixth of the money between them. This is not a market of 105 comparably

resourced competitors. It is two companies, a second tier of roughly thirty, and a long tail on seed and Series A money.

Five companies carry a valuation of $1B or more, the same five as in August at higher prices:

Company	Valuation	Raised	Segment
Harvey	$15.5B	~$1.72B	Firm platform
Legora	$8.8B	~$1.42B	Firm platform
EvenUp	$2.0B	$385M	Litigation & PI
Norm Ai	$1.2B	$260M+	Compliance
Eve	$1.0B	$150M	Litigation & PI

By segment

Segment	Companies	Capital
Firm platform	5	$3.27B
Contracts & in-house	32	$1.64B
Litigation & PI	24	$923M
Compliance	3	$299M
IP & patents	8	$281M
Research & knowledge	10	$246M
Practice ops & billing	7	$190M
AI-native law firm	7	$188M
Tax	1	$156M
Consumer & small law	3	$145M
Real estate	3	$81M
Immigration	2	$38M

Five companies in the firm-platform segment hold $3.27B — 43.8% of all tracked capital, against four holding $2.16B, or 34.8%, in August. That is the Harvey and Legora money again, arriving in the segment table instead of the curve.

Contracts and in-house remains the segment with the most companies by a wide margin at 32 of 105, and its capital has barely moved: $1.64B against $1.62B. Thirty-two companies raised $20M between them in a quarter in which two companies raised $1.1B. Litigation and personal injury gained two companies and $5M. The second tier is not being funded at anything like the rate the headline implies.

By region

Region	Companies	Capital
USA	52	$3.86B
Europe (ex-UK)	17	$1.90B
UK	17	$732M
Asia-Pacific	8	$451M
Canada	7	$364M
Israel	2	$102M
Latin America	2	$51M

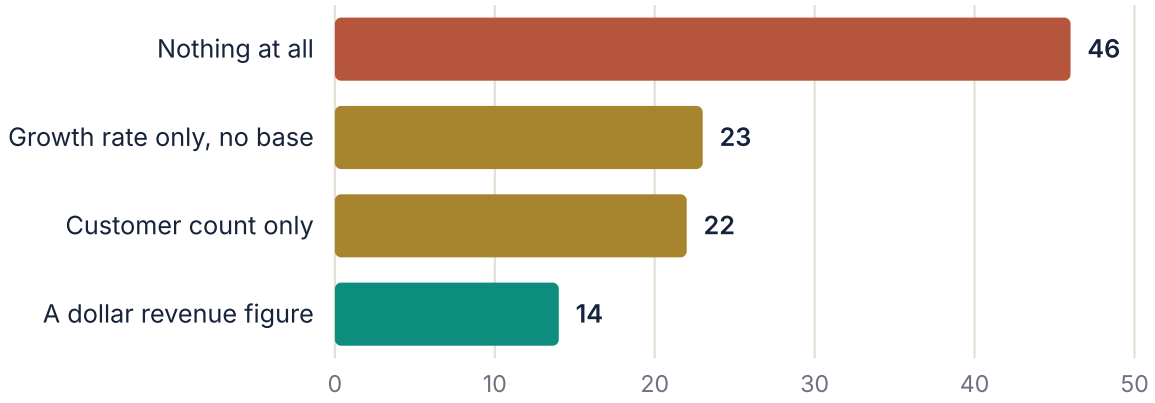
The United States holds 51.7% of the capital across 49.5% of the companies, both roughly where they were. The comparison that moved is the UK against continental Europe. In August the UK had 18 companies to Europe's 14; it now has **17 to Europe's 17**, and Europe holds **$1.90B to the UK's $732M**. Europe added three companies and $620M, principally Legora and Noxtua, while the UK's count fell by one and its total by $2M. The UK still has more legal AI companies than any country except the United States, and the least capital per company of any major region here. Latin America appears twice and Africa not at all, which is a limit of our sourcing: the survey is built from English-language funding announcements.

We have dropped August's founding-year table. It rested on a denominator of 79 of 95, because we could not establish a credible founding date for 16 companies, and we have not established dates for the ten added since.

2. The disclosure vacuum

Of 105 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-two give a customer count and nothing else. **Forty-six — 44% of the list — have disclosed nothing at all.**

The plain way to say what happened since August is this: **the absolute number of companies disclosing a revenue figure did not move.** It was 14 of 95 and it is 14 of 105. Ten companies entered the survey, two of the five unicorns raised at materially higher valuations, and not one additional company disclosed what it earns. Proportionally the disclosure rate fell from 15% to 13%, and the silent share rose from 39% to 44%.



What the 105 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.

Why we still cannot rank this market by revenue

The four reasons we gave in August all stand. The sample is too small: fourteen disclosures out of 105 is a list of the companies that chose to speak, not a ranking. The figures are not contemporaneous — **Athennian's ~$10M ARR is from March 2022**, now four and a half years old, and Harvey's $350M+ dates from August and was not restated when the company raised $550M in September. They are not the same

measure: Darrow's $26M is FY2024 revenue rather than a run rate, and the gap between ARR and audited annual revenue widens the faster a company grows. And one of the fourteen is not a number: Solve Intelligence discloses "eight figures" of ARR, a band spanning $10M to $99M.

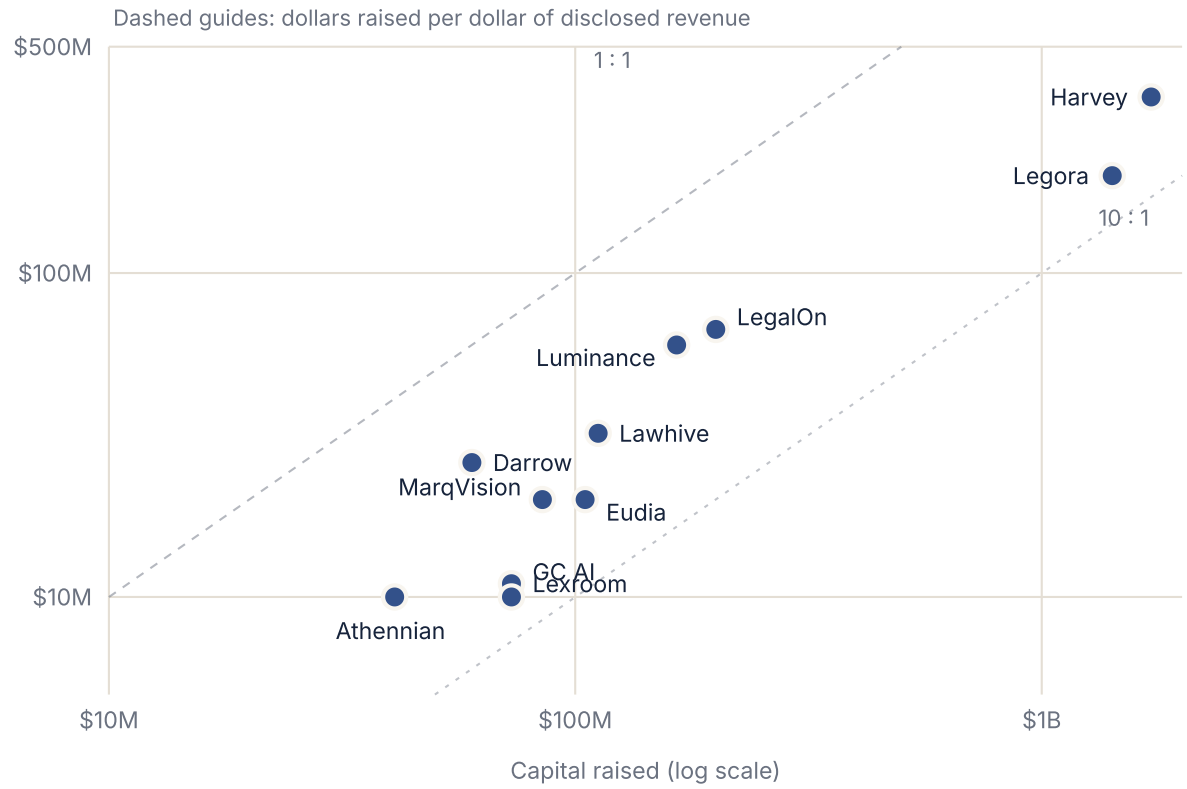
The fourteen

Company	Disclosed figure	As of	Put on the record by
Harvey	$350M+ annualised	Aug 2026	Press reporting
Legora	$200M+ ARR	Oct 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
GC AI	$10M+ ARR	Nov 2025	Company
Athennian	~$10M ARR	Mar 2022 — stale	Press reporting
Solve Intelligence	"eight figures" ARR	Dec 2025	Company
Juro	$1M ARR, historic	~2020	Unattributed
midpage	Six-figure annual revenue	Feb 2025	Unattributed

One row changed materially. **Legora disclosed $200M+ ARR in October 2026**, against $100M in April — the only upward restatement in the set since our last edition, made alongside a funding round. The other thirteen are the figures we published in August, carried forward with their original dates. Luminance's ~$60M remains an estimate carried in reporting rather than a company statement, and for Juro and midpage we still cannot establish who put the figure on the record.

Capital against revenue

All fourteen companies that disclose revenue also disclose a capital total, which is the largest set on which any efficiency comparison can be attempted.



The companies that disclose both capital raised and revenue

A correction first. In August we published three multiples — Harvey at 31x, GC AI at 55x and Legora at 56x — and described them as dollars raised per dollar of current annualised revenue. They were not: two of the three reconstruct exactly as valuation divided by revenue. The label was wrong and we withdraw the figures.

Computed properly from the capital totals the picture is less dramatic. Harvey has raised ~$1.72B against $350M+ annualised, **under 5x**. Legora has raised ~$1.42B against $200M+ ARR, **about 7x**. GC AI has raised $73M against $10M+ ARR, also about 7x. Read those as dollars raised per dollar of disclosed annualised revenue and not as a judgment: a company that raised at a high price and has not yet spent the money looks identical here to one that spent it badly.

The summary has not changed in eleven weeks, except that the denominator is larger. **Ninety-one of 105 companies have never stated what they earn**, the fourteen that have cannot be compared to each other, and the market's aggregate revenue remains unknown to us and, as far as we can tell, to everyone else publishing about it.

3. What got built

The catalogue holds **312 products** from **104** companies: 307 from the 105 in the funding survey, plus five from Courtroom5, which sells the most substantial software built for self-represented litigants and has raised too little to meet the survey's threshold. One funded company markets no product by name at all. Where a figure here is about the funded market it uses 307 and 105; where it is about what exists, 312 and 106.

The number to read is not 312. It is 112.

We re-read every company website site by site on 8 November. Against the August catalogue: **64 products added, 25 withdrawn, 23 renamed.** That is **112 of 312 rows touched in under three months** — more than a third of the catalogue — to produce a net change of 43. The net number is the least informative thing in this chapter. A market that adds 64 products and deletes 25 is not growing steadily; it is rebuilding its own shop windows while the survey is being written.

Four companies show what that means at the row level.

Luminance deleted five of its seven products. Draft, Analyze, Comply, Investigate and Collaborate are gone — not demoted in the navigation but absent from the top nav, the footer, the product index and a sitemap of roughly 290 URLs, with no page and no URL substring for any of the five. Three products replace them: Repository, Negotiation (renamed from Negotiate) and Workflows. Lumi survives only as an in-product assistant mentioned on one page.

Definely deleted four of five. Proof, Cascade, Enhance and Vault do not appear in a sitemap of about 110 URLs or on the homepage; "Proof" survives only as the feature verb "Proofread". There are no product pages left at all — only `/feature/` pages. One product name remains in use, Definely Read, and the company now positions the whole offering as independent contract review, on the argument that the maker cannot be the checker.

Eudia went from three rows to eleven. Eight products were added — Assistant, Agents, Contracting, Compliance, Forward Deployment, Government Legal, Government Acquisition and Defense Capture — each with its own page under `/platform/`. Enterprise Brain was renamed Knowledge and the old phrase appears nowhere. Eudia Counsel, the law firm, survives as a single footer link out to another domain.

Harvey added six and merged two away. Memory, Legal Work Orchestration, Horizon Scanning, Intake, Facts & Chronologies and Connector Library are all new, under a products page restructured into Platform, Apps and Connect. Assistant and Workflows are gone as product names and both old URLs now serve the Agents page, which we read as a merge rather than a deletion: Agents now covers task delegation, scheduling, custom agents and playbooks.

The smaller withdrawals are more telling. **Norm Ai's Supervisory AI** — the only product in the August survey that supervised another company's AI agents — does not occur on any page we read. There is now no product in the catalogue that does that job. **Lawhive's Lawrence**, the AI-paralegal brand from its 2024 press cycle, has no trace on the site; `/lawrence` serves the homepage under a `noindex` tag, and the pitch is now local UK solicitors at fixed quotes with no named AI tool anywhere. Twin1's MCP server is absent from a 13-URL sitemap. Centari's Intelligence and Views are absent from three pages we read in full. Ivo's Playbooks and Trellis Chat are now features inside other products.

The 23 renames include three at LegalOn Technologies, which dropped the partner attribution "by Fides" from Entity Management and Board Management and now calls Translate "Translation"; Ask Blue J, marketed simply as Blue J; Norm Technology, now sold as Norm Ai agents; Noxtua, now Noxtua NextGen; and Ivo's Search Agent, now Ivo Assistant and retyped to agent infrastructure. Boundless is the cleanest case of a name dissolving rather than changing: "Boundless for Business" appears nowhere on the site, and the nearest defensible name we could construct is Boundless Business Services.

	In-house	Law firms	Independent	Pro se
Agent platforms & infrastructure	37	33	8	1
Legal services delivery	14	5	1	12
Drafting & document generation	18	14	5	1
Knowledge search & firm memory	15	14	1	0
Time, billing & practice ops	5	21	1	0
Contract review & negotiation	21	5	1	0
Litigation case prep & discovery	13	16	5	3
Research & citation	17	15	9	0
Contract lifecycle management	19	2	1	0
IP, patents & brand protection	17	10	0	0
Case intake & triage	8	8	2	1
Regulatory & compliance monitoring	13	6	1	0
Client communication & portals	2	11	3	1
Due diligence & deal work	9	7	1	0
Medical records & demand packages	0	7	7	0
Judicial & litigation analytics	5	6	1	0
Entity, governance & deal admin.	5	2	0	0
Transcription & the record	0	4	3	0

312 products by type and by who they are sold to

Forty-six products are typed as agent platforms and infrastructure, against 33 in August — still the largest single type, and still larger than contract review and negotiation, the category the previous decade of legal technology was built around. Counted by function rather than by type, **69 of the 307 funded products are agents or agent infrastructure**, against 50 of 264 in August. Part III takes that apart properly, including the fact that **zero** of them supervise another company's agents.

The point

A catalogue of this market decays in weeks, not years. Two of the companies above deleted most of their product line between two editions of one report, and neither announced it. Every row we publish is a claim about what a website said on a date, and we publish the date for that reason. A tidy list with no date attached is not a cleaner version of this; it is the same decay with the evidence removed.

4. The structural shift

Legal services delivery is now 29 products, up from 25 in August. The segment we type as AI-native law firms holds **7 companies and $188M**, up from 5 and $169M. These are companies selling legal work rather than software: not a tool a lawyer operates, but a firm doing the work with AI behind it, sold at a price.

We can name **seventeen** companies in the November catalogue that sell legal work in some form. That is a hand count off the product rows rather than a figure from the funding data, and it should be read as a floor: several sell software and legal work under one brand, and at least one sells legal work through a firm it does not own.

They still arrive by two routes. The first is companies built as firms from the start. **Crosby** is a PLLC with 50-plus licensed attorneys and malpractice insurance, and says so on its homepage — "a law firm, not an AI tool". **General Legal, LLP** bills for outcomes rather than hours across 79 services in 12 practice areas, and now lets clients open a matter over MCP alongside the portal, Slack and email. **Keith** does English residential conveyancing at a published price and **Conveyd** does the same work under Council for Licensed Conveyancers licence 14755. **Lawhive** separates Lawhive Ltd, which states it is not a law firm, from Lawhive Legal Ltd, the SRA-regulated affiliate, and our catalogue mirrors that split in two rows.

The second route — software companies bolting a firm onto the side — is where the quarter's additions came from. **Norm Ai** now markets a second professional-services firm, **Norm Compliance**, alongside Norm Law LLP, with a block on the homepage and no landing page of its own. **Eudia** added **Forward Deployment**, which is people rather than software: embedded lawyers and engineers who build an institution's workflows into production. **Tradespace** added **Tradespace Counsel**, flat-fee patent preparation and prosecution, while stating plainly that Tradespace does not practise law — attorneys at independent partner firms do it under their own engagement letters. **Finch** added **Referrals**, a fee-sharing arrangement rather than a product: cases a personal injury firm will not take go to a partner firm and the referrer keeps a share of the fees. We counted it because it is separately named, navigated and sold. **Aavalynx** still operates Aavagard, "Aavalynx's sister law firm", and **Manifest OS** is a step further removed again — its footer disclaims being a law firm, and the work is done by firms it powers, such as Manifest Law.

Justpoint, which went the other way

In August we called Justpoint the furthest-developed case of the structural shift: in July 2025 it split into Justpoint Inc., which disclaims providing legal services, and **Justpoint Law LLP**, described in the company's own announcement as the first AI-native mass tort law firm approved under Arizona's alternative business structure regime.

Justpoint Law LLP has vanished from the site. It is not named on the homepage or the how-it-works page, both of which we read in full. The site refers only to unnamed "partner attorneys" and states "this site is for our business customers". Neither a law firm, nor an ABS licence, nor Arizona appears anywhere. The product is now predictive intelligence for product liability, sold to plaintiff firms, **defence firms**, in-house counsel, litigation funders, insurers and regulatory teams. Defence-side is new. We have retyped the row from case intake and triage to judicial and litigation analytics and dropped its pro se audience.

We do not know whether the firm was wound up, renamed or simply taken off the marketing, and we report the absence rather than infer a cause. It is the sharpest contradiction in this edition: our funding source for the company is a press release about a law firm its own website no longer mentions.

What we could not verify

This section ends where it ended in August, and we have closed none of the gaps. We have **not** verified how many alternative business structure licences Arizona has issued, to whom, or how many remain active; published counts circulate and we could not reconcile them against the regulator's record, so we state none. We have not verified how many of the companies above hold a licence in their own name, as against operating through a firm owned by someone else or practising only where the question does not

arise. We have not established the status of Utah's regulatory sandbox, nor counted the authorisations granted by the Solicitors Regulation Authority in England and Wales, the regime Lawhive, Keith, Conveyd and Moritz operate under.

Those four were the largest known holes in August and they are the largest known holes now. Each is answerable with primary regulator records. Saying so twice is less comfortable than estimating, and more useful.

5. Exits, failures and renames

Twelve companies in or adjacent to this market have been acquired since 2023 on terms we have reconciled, 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
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

Five further acquisitions closed in the window and are absent from that table because we have not established a price for any of them. **Harvey bought Guardrails AI in September and Aida in October** — its fourth and fifth of 2026. **Lexroom bought two European companies in September. Clio bought Learned Hand and Page**, adding to the vLex purchase that remains the largest deal in the market.

That overturns a call we made in August, when we predicted consolidation would continue at the top and named **Legora as the acquirer to watch**. **Legora has made none since July.** Harvey, which made no acquisition in the first half of the year, has made five in 2026 and is now the most acquisitive company in the survey. The direction was right and the name was wrong, and the name was the useful part.

Zero deaths, which we did not expect

No company in this survey was confirmed dead in the window — not one wind-down, insolvency, parked domain or departure from the market.

We flag that against our own prediction. In August we said, with high confidence, that two to five of the 95 would be dead or departed by August 2027, and noted that four already were. Eleven weeks later we have no new name to add. The prediction is unresolved rather than wrong — it has eight months to run — but the absence is worth stating in a quarter in which 34 of 105 companies have raised under $10M and the top five took 53.3% of the capital. Either the long tail is funded in ways our survey does not see, or the failures are underway and invisible, which is what the next section describes.

Failure still looks like a live website

Robin AI is finished. It wound down through 2025 and the managed-services arm went to the law firm Scissero in December 2025. Our August note said the engineering team was acqui-hired by Microsoft. **That was wrong.** Microsoft hired at least eighteen former engineers, including the chief technology officer, as individuals rather than acquiring a team. No insolvency filing has surfaced, the company has never confirmed its status, and the website still shows product marketing.

LawGeex: we withdraw the claim we made in August. We reported that lawgeex.com served an expired TLS certificate and called the brand dead on that basis. Two independent readings in November found a working site served over a valid chain, and we could not verify the certificate ourselves, so the claim comes out. What remains is a stale brand rather than a dead one: LegalSifter acquired the enterprise business in 2023 and the successor product is LegalSifter ReviewPro, yet the live site markets "Contract Review Automation", never mentions LegalSifter, carries a current 2026 copyright and banners a different product on a different domain. A buyer reading that site would not learn who owns it.

Della AI was acquired by Wolters Kluwer in December 2022 and the brand retired nearly four years ago; the technology ships as Legisway Analyzer. The della.ai domain is now listed for sale at **$68,885**, down from $98,000 in August, with lease-to-own at $5,741 a month. A 30% markdown in a quarter is the only price discovery this market published in the window. **Klarity** raised around $90M, renamed itself Within and left legal entirely for enterprise process discovery.

Two of August's judgments about quiet companies were wrong. **Fileread** is operating and shipping: Workbench launched in July 2026 and the blog ran through October 2026. It has raised nothing since its 2023 seed, which is a funding observation we wrongly read as inactivity. **Soxton** is active: it acquired Cipher Technologies in March 2026 and reports tripling its client base since its seed. We read a thin website as a thin company. Coverage also places it in New York rather than the UK, and we have corrected the funding row.

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
Leeway	Tomorro	—
Legal OS	Flank	May 2024
Orbital Witness	Orbital	—

All seven are active and shipping. Pocketlaw's rebrand to Miramis came with a repositioning from SME self-serve templates to mid-market contract lifecycle management. Two of the largest companies in the survey also trade under names they did not start with: **Legora was founded as Leya** and **LegalOn Technologies as LegalForce**.

The point

A list of this market built in August and read in November would carry a defunct company as live, an ex-legal company as legal, a product line that no longer exists at two vendors, a law firm that has vanished from its own website, an acquirer that stopped acquiring, and a dead certificate that was not dead. We know, because that list was ours.

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 again. The difference is that a dated row can be corrected in public, 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 one sold to a sole practitioner are not two examples of one category; they are two businesses, funded at different scales, priced on different logic, and held to different standards of proof. Sorting by who is expected to pay puts them where they belong.

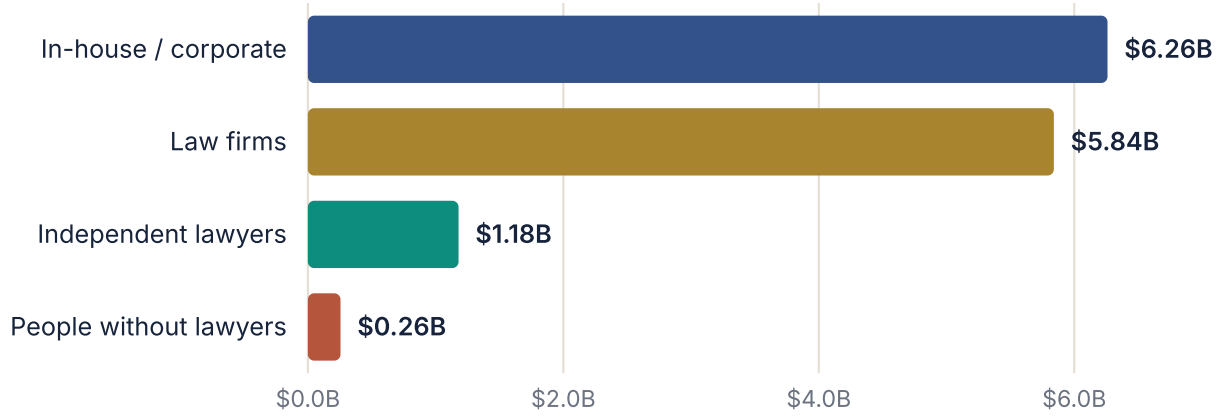
So this report is organised by buyer. We assigned every one of the **312** products in the catalogue to the audience or audiences its own marketing addresses, then carried each company's disclosed funding across to every audience it sells to.

One instruction on how to read the table, 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 $7.46B raised across the 105 companies in the funding survey — added up it comes to $13.54B, which is a double-count and not a market size. Read each row against the others, never as a share of a whole.

Audience	Capital behind it	Companies	Products	Unicorns
In-house / corporate	$6.26B	80	218	3
Law firms	$5.84B	65	186	4
Independent lawyers	$1.18B	18	50	2
Pro se / consumer	$256M	9 (10 with Courtroom5)	19	0

24.4 times more capital stands behind products built for corporate legal departments than behind products built for people without a lawyer — $6.26B against $256M. In August we put that ratio at 26 times, on $196M. It has narrowed, and the reason belongs in the same sentence: the in-house column grew 23% between editions and the pro se column grew 31%, largely because a ninth company entered it. The smallest column closed ground by adding a company, not because the largest slowed.

The pro se column contains no billion-dollar company, and only one of its nine — Lawhive, at roughly **$112M** — reaches the top twenty of the funding survey, where it is more than two-fifths of that whole column. A tenth company, Courtroom5, is in the catalogue but below our funding threshold.



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.

6. In-house and corporate legal

Corporate legal departments are where the money is. **$6.26B** stands behind products sold to in-house teams, across **80** of the 105 funded companies and **218** of the 312 products — the largest count on every measure we hold. Three of the five unicorns sell here: Harvey at **$15.5B**, Legora at **$8.8B** and Norm Ai at **$1.2B**. In August the first two were $11.0B and $5.6B.

What exists

Product type	In-house products
Agent platforms & infrastructure	37
Contract review & negotiation	21
Contract lifecycle management	19
Drafting & document generation	18
Research & citation	17
IP, patents & brand protection	17

Contracts still dominate. Add contract review (21), lifecycle management (19) and the contract-shaped half of drafting, and getting an agreement negotiated and signed accounts for more in-house products than any other activity a legal department performs. Agent platforms top the table at 37, up from 27 in August, but that describes a delivery mechanism rather than a job. **Eighty of the 105 funded companies sell here** — more than three in four — at 2.7 catalogued products each, which makes in-house the market's default assumption rather than a segment.

One August finding has to be withdrawn. We wrote that contract lifecycle management was 13 in-house products and zero everywhere else, and that no other type was exclusive to a single audience. CLM is now **19 in-house, 2 law firm and 1 independent-lawyer product**, and on the November matrix **no type is exclusive to one audience at all**.

The sharpest expansion in the column is Eudia, from **three catalogued products to eleven**. Three of the eight additions are a government line: Government Legal, Government Acquisition and Defense Capture. The last is not sold to agency counsel. It is sold to **defence contractors responding to solicitations requiring FAR/DFARS and Section L and M expertise**, a different buyer from the one Eudia addressed in August. Eudia has raised up to $105M and disclosed **$20M ARR** as of December 2025.

What it costs

Twenty of the 80 companies selling to in-house teams publish a price — **25%**, the third-highest rate of the four audiences, and still a minority small enough that most corporate buyers begin every evaluation without a reference point.

Zuva is the only true per-unit publisher at **$10 per document** and **$1.25** through the API. GC AI publishes **$500 a month** for an individual plan, Streamline AI "from **$22,900** a year", Athennian a **$25,000** entry tier with AI as a paid add-on. Arceus, new to the survey, puts flat fees on its homepage — **$500 to review a contract, $1,000 to negotiate, $2,000 to draft** — beside a 13-entry sub-processor table naming Anthropic and OpenRouter, while naming no certification at all.

The pricing pages that are not pricing pages still cluster here: LegalOn names three tiers and prices none, Juro's is a quote calculator with no figures, and EstateMin's FAQ gives a structure — per firm, unlimited users — with no number. Across the survey, companies that publish a price have raised **$30M** on average

and companies that publish none **$84M**; in August that was $76M, so the gap has widened. The refusal is now addressed to machines too: Within's `llms.txt` tells AI assistants that "Within does not publish standard pricing on the pages listed here".

What is missing

Two unchanged zeros. **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 independents, none to the corporate department that will be sued. **Time, billing and practice ops is 5** against 21 for firms, because departments do not bill.

The larger absence got worse. Thirty-seven in-house products are agent platforms, 69 of the 307 surveyed products are agents, and across the catalogue **exactly zero products supervise other companies' agents**. In August there was one, Norm Ai's Supervisory AI; on 8 November the phrase occurs nowhere on norm.ai — not on the homepage, the platform page, the agents page, the technology page or the resources index. The segment buying the most autonomy now has nothing available to watch it. The company that withdrew it has simultaneously renamed its platform "Norm Ai agents" and added a second professional-services firm, Norm Compliance, which has a homepage block and no landing page.

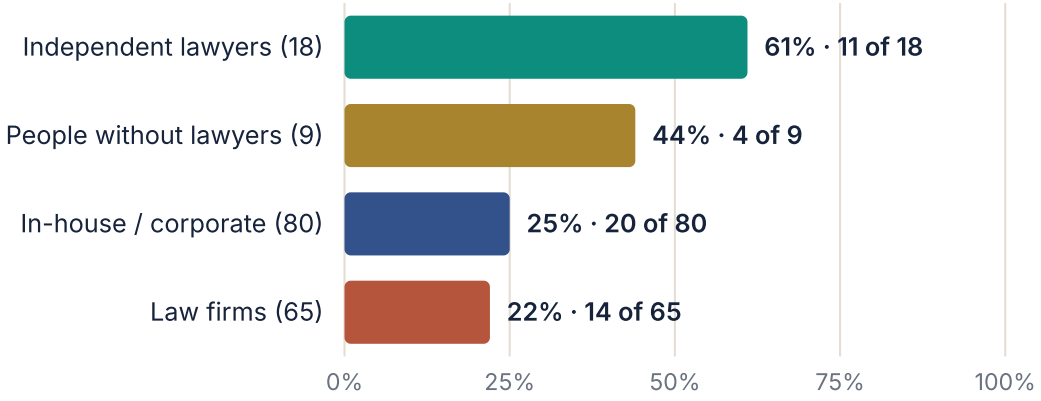
Who is credible

The in-house-facing companies that let a buyer check the most are GC AI (**7 of 10** — published price, named models, two phases of an independent benchmark), Harvey (**7**, the only company with both an open benchmark and independent-benchmark participation), Wordsmith AI (**5**, for naming the contractual basis of zero data retention with each provider) and Genie AI (**5**, which publishes its ISO certificate number and expiry and states plainly that it does not hold SOC 2). Zuva, Arceus and Lexroom score **4**.

At the other end, Norm Ai — a $1.2B compliance company selling to institutional finance — scores **1**: SOC 2 named, no model disclosure, no price, a closed benchmark. That score is a statement about how much of the claim is checkable from outside, not a judgement about the product.

The evidence picture

In-house-facing companies score a mean of **2.14 out of 10** on whether a buyer can verify what they claim, against a field mean of 1.96. Second-best of the four audiences, and still barely a fifth of the available evidence.



Whether companies publish a price, by who they sell to

The rest is the field-wide picture. Of all 105 companies, **49 will not say which models they run on**, 41 offer only a vague pledge and 15 name their providers. Thirty-two publish a benchmark score with no method and two publish tasks and rubrics; six have taken part in an evaluation run by someone else. **None of the 105 has independent research published measuring its product.** A general counsel evaluating a contract review tool in November 2026 has no third-party accuracy figure for any product in the category. That was true in August, and ten more companies have arrived since without changing it.

What a buyer should ask

1. **What is the unit, and what is the number?** Twenty of 80 vendors selling to you publish anything, so treat a figure as a signal — and get it in writing.
2. **Which models, running where, under what retention terms?** Forty-nine of 105 are silent and 41 pledge without naming a provider.
3. **Has anyone outside your company measured this?** Six of 105 have joined an outside evaluation. If the vendor cites its own benchmark, ask whether tasks and rubrics are open; 32 publish a score with no method.
4. **If you sell me CLM, what happens to the repository if I leave?** Nineteen of the 22 CLM products are sold to you, so the record of every contract you sign sits with a vendor that has almost no competing audience. Get export format and termination terms before the pilot.
5. **Who supervises the agent?** Zero products in 312 supervise other companies' agents, and the one that did was withdrawn. Ask what is logged and who reviews an agent's output before it reaches a counterparty.

7. Law firms

Law firms are where the prestige sits and where the valuations are made. **$5.84B** stands behind firm-facing products, across **65** companies and **186** products. Four of the five unicorns sell to firms — Harvey ($15.5B), Legora ($8.8B), EvenUp ($2B) and Eve ($1B) — 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 and it has concentrated further. **Firm platform is 5 companies and $3.27B — 43.8%** of all capital raised across the 105 companies, held by five businesses selling a single category. In August it was 4 companies and 35%. Three of the five are essentially the whole figure: Harvey at roughly $1.72B raised, Legora at roughly $1.42B and Noxtua at roughly $110M.

It is also the most volatile catalogue of the four, in both directions. Harvey added **six products** in the three months to 8 November — Memory, Legal Work Orchestration, Horizon Scanning, Intake, Facts & Chronologies and Connector Library — and withdrew Assistant and Workflows, whose old URLs now serve the Agents page; we read that as a merge rather than a deletion. Legora added four, and four months after buying Wexler AI, its fifth acquisition since March 2026, still runs it as a separate brand on a separate domain. Against that, **Luminance removed five of its seven catalogued products** — Draft, Analyze, Comply, Investigate, Collaborate — none of which appears in a nav, a footer or anywhere in a roughly 290-URL sitemap, and **Definely removed four of its five** — Proof, Cascade, Enhance, Vault — leaving a sitemap of about 110 URLs with no product pages at all, only features. A buyer running a 2026 procurement against either module list would find most of the list gone.

Below the platforms the mix diverges sharply from the corporate one. **Time, billing and practice ops is 21 firm products against 5 in-house** — the clearest structural difference between the audiences, and a direct consequence of the billable hour existing on one side of the relationship only. Agent platforms are 33, litigation case prep 16, research 15, drafting 14, knowledge search 14 and client portals 11. Medical records and demand packages is 7 firm products and 0 in-house, all of it personal injury, the practice area that has absorbed more capital than any other at **$923M** across 24 companies. Contract review runs the other way: **5 firm products against 21 in-house**.

What it costs

Fourteen of 65 firm-facing companies publish a price. At **22%** this is still the least transparent of the four audiences — below in-house, and far below the 61% of companies selling to independent lawyers. 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 prices that do exist are mostly at the small end. Clearbrief publishes **$300 a month**. Skribe.ai publishes **$379 per record hour** for a deposition, the most transparent unit price in litigation services. The large platforms publish nothing: Harvey, 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 nav item labelled "Pricing" opens a demo booking form, and Laurel, which raised roughly $100M in a single 2025 round, publishes nothing.

What is missing

August reported one true zero in this column — contract lifecycle management — and that is no longer accurate. CLM is now **2 firm products**: a small number that nonetheless breaks the cleanest line in the August edition.

Transcription and the record is **4 firm products and 3 independent ones**, from a very small number of companies. In the audience that litigates, the document a court treats as what was said is still among the least-built categories in the market. We have an interest there — it is the category IQdoc works in — so treat the observation as an interested one and check the count.

The other missing thing is measurement, and it has slipped. Firm platform is still the best-scoring segment with more than one company in it, at a mean of **3.20** — down from 3.50, because the segment added a company without adding disclosure.

Who is credible

Harvey scores **7 of 10**, joint-highest in the survey, and is the only company with both an open benchmark and independent-benchmark participation, while publishing no price.

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 40.5 out of 50 from the State Bar of Nevada's AI Workgroup, the closest thing to independent assessment here. Among the firm-facing unicorns, EvenUp scores **2** and Eve **0**; Definely, which deleted four product names this quarter, scores **1**.

The evidence picture

Firm-facing companies score a mean of **2.08 out of 10**, and **14 of 65** publish a price. The buyers with the most procurement machinery are given the least public evidence, which suggests the machinery is absorbing the disclosure privately, one buyer at a time under NDA. That is a functioning sales motion, not a

market that can compare anything.

Above this segment sits a platform risk none of the evidence signals capture. Harvey and Legora both appear in Google's Gemini Enterprise integration list for legal work, as we reported in August and have not re-verified line by line since, and those two companies hold **$3.14B of the $3.27B** in this segment. Being an integration inside somebody else's enterprise agent is a different business from owning the window the lawyer opens each morning: the integration supplies a capability, 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 surface this segment depends on is being assembled by firms that do not need it to be independent.

What a buyer should ask

1. **Will you put the unit price in the contract, and the renewal uplift?** Fourteen of 65 vendors selling to you publish a price. Use procurement to get in writing what the market cannot get in public, and cap years two and three.
2. **Which of these modules will exist in a year?** Luminance removed five of seven product names this quarter and Definely four of five. Ask which names in the proposal are products with their own roadmaps and which are features, then put the answer in the schedule.
3. **Show me the certificate, not the badge.** Ask for number and expiry, as Genie AI publishes them. The survey contains "SOC 2 Type 3", which does not exist, and "SOC 2 Type II observant", which is not an audit status.
4. **Which evaluation have you been in, run by whom?** Six of 105 have joined an outside evaluation and none has independent research published about it. Only Harvey and Paxton AI publish their own tasks and rubrics.
5. **What leaves with us?** Fourteen firm products index the firm's own knowledge. Ask the extraction format, what happens at termination, and which models saw the documents.

8. Independent lawyers and small firms

$1.18B stands behind products sold to independent lawyers, across **18** of the 105 funded companies and **50** of the 312 products. Two unicorns sell here, EvenUp at $2B and Eve at $1B, whose raised totals of $385M and $150M are **$535M**, 45% of the column. The direction of travel runs against the rest of the report: survey capital rose from $6.20B to $7.46B between editions and this column **fell**, from $1.22B, while products rose from 38 to 50 and companies from 16 to 18.

What exists

Roughly half of American lawyers in private practice are solo practitioners. The first thing to report about the 50 products built for them is a withdrawal. In August we wrote that **exactly one** listed independent lawyers as its only audience: Lawrence, Lawhive's AI paralegal. On 8 November Lawrence is gone — `/lawrence` serves the homepage under a `noindex` tag — and no named AI tool of any kind is marketed on the site. Nor can we restate August's companion figure, that 33 of the other 37 products were sold to firms as well: the November facts do not carry that overlap.

Product type	Independent lawyer products	Law firm products
Research & citation	9	15
Agent platforms & infrastructure	8	33
Medical records & demand packages	7	7
Drafting & document generation	5	14

Product type	Independent lawyer products	Law firm products
Litigation case prep & discovery	5	16
Client communication & portals	3	11

What travels down-market is horizontal; what does not travel is anything shaped around how a practice runs.

Medical records and demand packages is the one substantive parity case: 7 products for independents against 7 for firms. The reason is the fee structure: plaintiff personal injury runs on contingency, so a solo doing injury work has cases that pay. EvenUp added three products this quarter and Supio one, and Paxton AI, a research vendor, has added medical chronology and billing-summary pages. Blue J, sole occupant of the Tax segment, now segments its site by firm type down to Sole Practitioner; nothing on it supports the in-house tag we held in August.

What it costs

Eleven of the 18 companies selling to independent lawyers publish a price — 61%, against 25% for in-house and 22% for law firms. Still the highest rate in the survey by a distance, and a falling one: in August it was 11 of 16, 69%. The count did not move: both companies that joined publish nothing.

Company	Published price
midpage	$30 and $80 a month
Jus Mundi	$29 academic, $129 practitioner
Twin1	$50 and $150 a month, priced per person
Trellis	$69.95 to $199.95 a month
Lexroom	EUR 75 a month per module
Callidus Legal AI (now StrongSuit)	$249 a month, individual
Clearbrief	$300 a month
Skribe.ai	$379 per record hour
GC AI	$500 a month, individual plan
Blue J	$1,498 a year per user, individual tier
Lawhive	Fixed fees per matter, as a regulated firm

midpage's **$30** is the lowest published price anywhere in the 105. Paxton AI also publishes a price; the figure we carried in August is not in the November facts, so we are not restating it. Genie AI now contradicts itself, its enterprise card saying "Custom" where the FAQ says "from $600 a month".

The obvious reading — that these are more honest companies — is not supported by the data. Several figures above are explicitly the individual tier: GC AI's $500, Blue J's $1,498, Callidus's $249, and the same companies publish nothing about what a legal department pays. An independent lawyer pays from their own pocket and will not sit through a discovery call to price a $50 tool. **The published price is a condition of the sale rather than a disclosure choice.**

What is missing

August listed six product types with **zero** independent-lawyer products. Four of them now hold exactly one, which is worth reporting precisely, because one is not a market.

Type	Independent	Firm	In-house
Time, billing & practice ops	1	21	5

Type	Independent	Firm	In-house
Knowledge search & firm memory	1	14	15
Contract lifecycle management	1	2	19
Regulatory & compliance monitoring	1	6	13
IP, patents & brand protection	0	10	17
Entity, governance & deal administration	0	2	5

Time, billing and practice ops was 14 firm products against 0 for independents in August. It is now 21 against 1. The zero technically closed and the gap grew by seven products. A solo bills and chases exactly as a firm does, with nobody to delegate it to, and knowledge search is the same shape at 14 and 1: the horizontal tools came down-market, the operational ones stayed where the seat counts are. The second absence is by practice area: family law, criminal defence, immigration, estate planning and landlord-tenant work return research tools and little else.

Who is credible

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, Twin1	Price published, certifications named

Paxton AI is one of only **2 of 105** companies publishing an open benchmark. **EvenUp, valued at $2B, scores 2. Eve, valued at $1B, scores 0**, the latter for claiming to follow "the SOC 2 Framework" — a framework claim, not an attestation. The two most valuable companies selling to independent lawyers disclose less than a company charging $30 a month.

The evidence picture

Independent-lawyer-facing companies score a mean of 3.22 out of 10, the best of the four audiences — against 2.14 for in-house, 2.08 for firms, 1.11 for pro se and 1.96 across the field. In August it was 3.38, so the lead has narrowed as the cohort grew.

That 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 18 publish one. Strip that signal out of every score and, by our own arithmetic on the November figures, this cohort's mean falls to **2.00** while the field mean falls to roughly **1.49**. Close to two-thirds of the lead was the price alone, and the solo buying a research tool for $30 a month still has no third-party accuracy figure for it or any competitor.

What a buyer should ask

1. **Is the price on the page the price I pay?** Eleven of the 18 companies selling to you publish one. Check whether it is the individual tier.
2. **Was this built for a practice my size, or priced down to it?** Almost everything here is also sold to firms. Ask what the product assumes exists: a knowledge base, a document system, an associate to check the output.
3. **What does it do with the citations?** Nine of the 50 products are research and citation, and the failure mode that ends careers is a confident wrong citation. Check three by hand.

4. **Which model, and where does my client's file go?** Forty-nine of 105 companies will not name their providers, and you carry the confidentiality obligation personally.
5. **Has anyone outside the company measured it?** No independent research exists on any product here; six companies have joined someone else's evaluation.

9. People without lawyers

$256M stands behind products sold to people handling a legal problem without a lawyer — 9 funded companies, 19 products, **no unicorn**. It is the smallest column by every measure, and **24.4 times less capital than stands behind products built for corporate legal departments**. In August we said 26 times, on $196M.

Composition matters more than the total, and it has not changed. **Twelve of the 19 products are typed legal services delivery — a person with a licence doing the work — rather than software the user operates**, and with Manifest OS's new Individual Client Portal, the window onto work done by Manifest-powered firms, **13 of the 19** amount to a licensed person doing the work. 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

Company	Products	Type
Hello Divorce	DIY, Pro and Plus Plans, Hallie, Homesplit	Legal services delivery
Lawhive	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	Case intake & triage
Manifest OS	Individual Client Portal	Client communication & portals
Courtroom5	The LAW Accelerator, Proof, Strategy	Litigation case prep
Courtroom5	Documents; Courtroom5 in Claude	Drafting; agent infrastructure

Two changes since August. **Justpoint Law LLP has left the list**: the Arizona ABS-licensed firm named in our own funding source is not identified anywhere on Justpoint's site, which states that it is for business customers. The surviving Justpoint row sells product-liability intelligence to plaintiff and defence firms, funders and insurers; the November facts still carry it in this column, and we expect it to leave. The other change is Manifest OS's portal, the ninth company's entry.

The funded products cluster where money is already moving: a house sale, a visa, a marital estate, a settlement. **Every product here that is software a person without a lawyer operates themselves belongs to Courtroom5**, the one company our funding threshold could not see. Of the 69 agent products in the survey, exactly one is reachable by a person without a lawyer, and it is Courtroom5's connector.

Courtroom5 is self-operated software for self-represented civil litigants, rebuilt through 2026 around AI as the LAW Accelerator. Its funding is an estimate and we mark it as one: August put it at roughly **$420,000** since 2017, mostly grants, a floor estimate from a paywalled aggregator never confirmed against a filing, and the November facts carry no figure for it at all. Our own threshold filtered it out, so it sits outside the $256M and the nine-company count.

What it did with that money is the point. It is distributed through three assistant platforms, each with its own page on its own site: Anthropic's Claude, OpenAI's ChatGPT and Google's Gemini Enterprise for Legal. The Claude listing is named **"Courtroom5 in Claude"** — August recorded "Courtroom5 MCP connector", a name taken from a Google Cloud press release rather than from the vendor — and the connector is **free, no account required**. On an estimated $420,000 that is roughly one four-thousandth of Harvey's $1.72B, on the same surfaces. Distribution is not adoption, and we measure usage on neither side.

What it costs

Four of the 9 funded companies publish a price — 44%, second only to the independent-lawyer segment and for the same reason: nobody here sits through a sales call. Keith posts fixed conveyancing fees, now **£1,200** to buy freehold and **£1,000** to sell, inclusive of VAT; August said £1,125 to sell, and that figure is superseded. Hello Divorce posts packages at **$2,000** and **$3,500**, Boundless posts package pricing, and Lawhive posts fixed fees per matter as a regulated firm. Courtroom5, unscored because it sits outside the survey, leads its homepage with **$1 for 100 starter credits**. Conveyd markets "transparent pricing" and states no figure anywhere; Casium, DoNotPay, Justpoint and Manifest OS publish nothing.

Two presentation faults matter where the buyer has no adviser: Keith's homepage renders its headline speed claim as garbled numerals, "123123 days" and "5555 days", and prints its regulator licence number as an unfilled placeholder.

What is missing

Set these products against where people without lawyers 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** — the August edition's citation to external research, unchanged. Products from the 105 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, nothing for benefits appeals, nothing for expungement. Not underserved. Empty. Ten more companies and 43 more products entered the survey since August and none landed here.

This is not villainy. An eviction defendant is by definition someone who could not pay rent, and overwhelming demand with near-zero ability to pay is the one combination venture capital cannot serve on its own terms.

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, with no named security certification. Boundless, Keith and Lawhive score **2**, for publishing a price and almost nothing else, and Manifest OS enters at **1**.

Four of the nine 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 means we could not verify a claim from public sources, not that a product is bad — but four of nine at zero is a pattern rather than an accident.

The evidence picture

Companies selling to people without lawyers score a mean of 1.11 out of 10, the lowest of the four audiences, against 1.96 across the field. In August it was 1.13 on eight companies.

Say what that does not mean first. **It rests on 9 companies**, so one company moving two points shifts the mean by more than a fifth of a point; 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. 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 training with which to judge any of it. **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=9 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 ordered **DoNotPay** to stop making deceptive AI-lawyer claims; on 8 November 2026 the site still leads with "The AI that fights for you" and still offers to help users fight large companies. The only external check on a consumer legal AI claim in this survey arrived as enforcement, from a regulator, after the fact.

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 before you pay.

1. **Is a lawyer doing this work, or am I?** Most of what is sold here — 13 of the 19 products — is a person with a licence doing the work for you, which 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 for the licence number, then check it on the regulator's own site — we found one published as a blank placeholder.
3. **What is the total cost?** Ask what is not included, which court fees you still pay, and what happens to your money if you stop halfway.
4. **What does this not do?** 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, for every product here. A percentage on a homepage with no method 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 and debt collection cases almost nothing has been built commercially, so free help is not the fallback. It is the main provision.

Part III — Agents

Sixty-nine of the 307 funded products in our product catalogue are agents or agent infrastructure, against 50 of 264 in August. **Forty-six** are typed as agent platforms and infrastructure, against 33, which keeps it the largest single type in the catalogue. The category grew by nearly 40% in ten weeks and the word got no more precise while it grew. It still appears on the homepages of firm platforms, contract tools, patent search engines, tax software, billing systems and a company that texts clients when their case moves.

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 69 into the four meanings the term carries here, plus two residual groups the four cannot hold. **That sorting is our reading of vendor copy, not a tested classification.** We have not run these products, and no vendor publishes the reliability data that would let anyone check a label from outside. That absence is the subject of Chapter 14, and it has got worse since August: the single product in the survey that supervised other companies' agents has been withdrawn. Which is why this part is a taxonomy and not a ranking.

10. The taxonomy

A chat assistant, renamed. Seventeen products. Luminance's Lumi, GC AI, Noxtua NextGen, Summize's SIA, Operator by Juro, Jus AI, Chamelio and nine others. A person asks, the software answers, the exchange ends. Five joined since August — Eudia's Assistant, Alexi, EvenUp's Companion, Case Status's Seer and MarqVision's Brand Intelligence Agent — and one left: **Lawhive's Lawrence is gone**, which Chapter 12 returns to.

A multi-step workflow runner. Twenty-seven products, still the largest bucket and still closer to scripting than to delegation. Ten are new: Blue J's Agentic Workflows, Eudia's Agents, Bryter's Hybrid Agents, Newcode's Cowork, CaseMark's Para, Magnar's Agent, Luminance's Workflows, Spellbook's Autonomous Contract Management, Laurel Billing and EvenUp's Proactive Workflows. **Harvey Workflows is not in the count**, because Harvey no longer sells a product by that name; it and Assistant were merged into Agents, and both old URLs now serve the Agents page.

One boundary inside this bucket has moved. In August we described the workflow runner as a sequence "initiated by a human who is waiting for the output". Harvey's Intake and Legora's Intake both break that: each starts when a request arrives and each ends with a lawyer approving the draft. The trigger has moved from the lawyer to the request while the review has stayed put. We left both here rather than in the autonomous bucket, and we flag the call.

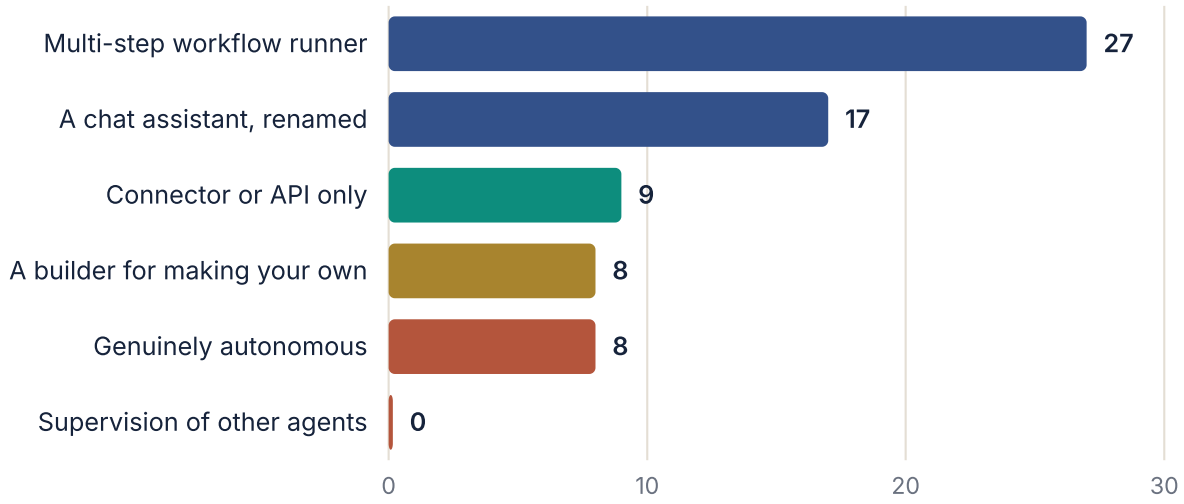
A builder. Eight products, where the thing sold is a factory for making agents rather than an agent. The new entrant is **Legora's Skills** — a library of **250-plus** encoded legal task definitions across more than twenty practice areas that the agent selects from, extensible by the firm and controlled by its administrators. Eudia's Enterprise Brain has left this bucket: it was renamed Knowledge and retyped, and the old phrase appears nowhere on the site.

Something actually autonomous. Eight products, up from six, and the subject of Chapter 12.

Infrastructure, not an agent. Nine products — plumbing that lets something else be the agent, and the subject of Chapter 13. This was five in August. It is the fastest-growing group in the part.

Supervision. Zero products. In August this count was one. It is now nothing, against 69 agents. Chapter 14 is about that line.

Two caveats carry forward and one is new. The boundary between the first two buckets is soft, and the classification uses vendor copy as its evidence, which means a product that oversells itself is classified as it oversells itself. The new caveat is naming churn: three of August's agent products — Harvey Workflows, Twin1's MCP server and IPRally's Agent — no longer exist under those names at all. A taxonomy of this market is a photograph.



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

11. Who is building them

Forty-seven companies build these 69 products, against 40 companies and 50 products in August. Descriptions are condensed from each vendor's own materials.

Chat assistant, renamed (17)

Company	Product	What it does
Luminance	Lumi	In-product contract assistant; no page of its own now
GC AI	GC AI	Cited answers, contract review and case-law research for in-house
Noxtua	Noxtua NextGen	German and EU research, analysis and drafting on EU infrastructure
Summize	SIA	Answers contract questions, executes review steps across the CLM
Athennian	Athennian AI	Governance answers and entity lifecycle changes against the record
Juro	Operator by Juro	Executes contract actions from plain-language instructions
Nexl	Nexl AI	Client opportunities and relationship risks from the firm's CRM
Eve	Eve Research	US opinions database available inside every Eve agent
Ivo	Ivo Assistant	Agentic querying across a contract, a deal or the portfolio
Jus Mundi	Jus AI	Research over the Jus Mundi corpus, with source control
Twin1	Twin1	Per-person agent trained on that professional's own files
Chamelio	Chamelio	Answers across a team's contracts, email, drives and chat
Eudia	Assistant	Drafting over case law, EDGAR, USPTO and a team's files
Alexi	Alexi	Firm-wide private AI platform; retyped to agent infrastructure
EvenUp	Companion	Cited case answers; triggers other EvenUp agents
Case Status	Seer — Insights Agent	Answers questions on firm client and case data inside XOS
MarqVision	Brand Intelligence Agent	Answers over brand data; proposes enforcement for approval

Multi-step workflow runner (27)

Company	Product	What it does
Harvey	Agents	Task delegation, scheduling, custom agents and playbooks
Harvey	Intake	Collects requests from email, Slack, Jira, ServiceNow; routes by rule
Legora	Agent	Plans, executes and reviews a multi-step task, returns output
Legora	Workflows	Configurable repeatable multi-step legal processes
Legora	Intake	Requests by email, Teams, Slack or MCP; drafts for approval
Spellbook	Associate	Works across every document in a matter, not one file
Spellbook	Autonomous Contract Management	Intake, triage, routing, first-pass review, approvals, repository
Wordsmith AI	Agents	Automates repeated legal workflows end to end
Supio	Supio Agent	Firm-wide agent across every case and role at once
Eve	EveOS	Runs a plaintiff firm's cases, staff and finances
Patlytics	Patlytics Agent	Chains patent tasks across the platform
Clerq	Agents	Triages invention disclosures into a patentability report
Sandstone	Sandstone	Intake, triage, matter management, agentic drafting and redlining
Norm Ai	Norm Ai agents	Regulation-aware agents reviewing content against rules
SpotDraft	Sidebar	Agents that answer questions and track regulatory change
Streamline AI	Featherline	Contract review applying playbooks and flagging risk
Miramis	PLAI Contract Agent	Playbook review, redline analysis, portfolio queries
Enter	Consumer litigation	Drafts answers to Brazilian consumer claims; flags serial litigants
Luminance	Workflows	Automates contract intake; advances on defined conditions
Blue J	Agentic Workflows	Tax workflows for returns, IRS notices and nexus studies
Eudia	Agents	Runs multi-step legal work; learns from approved reviews
Bryter	Hybrid Agents	Packaged agents over BRYTER automation: NDA review, DORA
Newcode	Cowork	Plans tasks with sub-agents inside the customer's Azure tenant
CaseMark	Para	Runs saved skills and team playbooks, returns finished DOCX
Magnar	Agent	Agentic execution over the Magnar platform; no published price
EvenUp	Proactive Workflows	Playbooks run against each new case file to screen and route
Laurel	Laurel Billing	Two agents check guidelines and fee drift at pre-bill; coming soon

Builder (8)

Company	Product	What it does
Bryter	BRYTER Workflows	No-code builder for rule-based legal and compliance processes
DeepJudge	AI Workflows	Build and run agents over the firm's indexed knowledge
Checkbox	AI Agent	Automates request intake, triage and workflow building
Newcode	Aurora	Engine for building and deploying bespoke agentic workflows
LegalFly	Agent Studio	Custom agents triggered by email, Slack, Teams or by hand
Zuva	Create	No-code trainer for custom contract extraction fields
Josef	Josef	No-code builder for bots and automated workflows
Legora	Skills	250-plus encoded legal task definitions the agent selects from

Actually autonomous (8)

Company	Product	What it does
Flank	Flank	Recurring legal requests inside Outlook, Teams, Jira and Salesforce
Flank	Flank Record	Captures contract metadata at signature and maintains the record
Streamline AI	Velo Copilot	Coordinating agent running interactively or ambiently
Hona	Lia	Chases clients, tracks appointments, flags treatment gaps, escalates
Hona	Hona Engage	Texts clients when the case management system changes
Eve	Eve Agents	Intake, casework and communications unattended, including calls
EvenUp	Communication Agents	Automated client outreach and case-status updates
Supio	Case Communications	Requests records, chases insurance claims, follows up

Infrastructure, not an agent (9)

Company	Product	What it does
Harvey	Connector Library	Connects document systems, mail, chat and deal tools to Harvey over MCP
Legora	Legora aOS	Umbrella platform binding agent, harness, knowledge and integrations
Avvoka	Headless Drafting	Puts a firm's templates and clauses inside other vendors' assistants
Trellis	API	Trial court data API for US state court records
midpage	For Legal Tech	APIs, MCP and database access for legal products
Clerq	API & MCP	REST API and MCP server over the NLPatent patent search model
IPRally	MCP & API	Connects IPRally patent search to assistants over MCP, or to pipelines
Zuva	API	REST API extracting contract fields at **$1.25** per document
GC AI	GC AI API	Programmatic access so other systems can call GC AI without a seat

Supervision (0)

No product in the catalogue supervises another company's agents. See Chapter 14.

Two segment findings hold and two have changed. Firm platforms remain the segment where shipping an agent is the defining feature: Harvey, Legora, Noxtua and Newcode all do, out of **five** companies holding **$3.27B**. Contracts and in-house still supplies the largest share of agent builders.

What changed is at the edges. **Tax has stopped shipping nothing**: Blue J's Agentic Workflows is the segment's first agent product, and the segment is one company. **Consumer and small law has gone the other way** — with Lawrence withdrawn, Lawhive, DoNotPay and Hello Divorce between them market no agent at all. Real estate, immigration and the **AI-native law firms** still ship nothing agentic, which for the seven law firms is August's finding unchanged: Crosby, Manifest OS, General Legal, Moritz and Soxton sell the outcome rather than software that acts. Crosby calls agents, expert attorneys and custom memory features of the firm's service, not products.

12. What the autonomous ones actually do

Look at what the eight genuinely autonomous products do and the striking thing is still how unglamorous it is. None argue motions. They chase people who have not replied, request records, open and chase insurance claims, text when a case status changes, place 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. Flank Record does the same for contract data, captured at signature. Streamline's Velo Copilot still runs "interactively or ambiently", and the ambient half is the whole claim. Hona has added a second product on the same logic: Engage texts clients when the case management system changes, alongside Lia. Supio's new Case Communications requests records, chases insurance claims and handles provider follow-up. Eve restructured rather than added: its Communication Agents are now one of three agent types inside **Eve Agents**, which runs intake, casework and communications unattended, outbound calling included. The capability was promoted, not withdrawn.

The reason autonomy lands here is structural rather than modest. The 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. Five of the eight sit in litigation and personal injury, where firms run hundreds of matters through one contact cycle.

The clearest evidence for the argument is a withdrawal. **Lawhive has removed Lawrence**, the AI-paralegal brand its funding cycle was built on and carried through a **$60M** Series B in February 2026 and roughly **$112M** raised in total. We checked the homepage, `/our-lawyers` and `/lawrence`; the last serves the homepage under a `noindex` tag rather than a Lawrence page. **No named AI product of any kind is marketed on the site now.** The pitch is local UK solicitors at fixed quotes, 300-plus of them, half the cost of a law firm. A company that raised on an AI paralegal now sells lawyers, and says nothing about why.

One disclosure detail carries forward with a correction. Flank holds **ISO 42001**, the standard that audits AI governance rather than information security, and remains the only company in our **105** to hold it — though Magnar, which now ships an agent of its own, lists ISO 42001 as "in certification". 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. That sentence has survived ten weeks of new products unchanged.

13. The connector layer

In August we called this the development we did not expect to find, and said the interface question was open. It is closing, and it is closing in a direction the funding tables do not describe.

The corpus is becoming callable. Harvey now ships a **Connector Library** as a first-class product under a new nav group called Connect: document systems, email, chat, deal and contract tools reachable over MCP, with administrator control over access and over which actions need confirmation. Legora's Intake accepts legal requests over MCP alongside email, Teams and Slack. Clerq ships API & MCP over the NLPatent model; IPRally ships MCP & API over 130 million-plus patent documents. midpage's connector is now filed under a Data menu as **For Legal Tech** — "APIs, MCP, and database access for legal products" — and a homepage banner reads that midpage now connects to Harvey. Mary Technology lists MCP integrations to Claude, ChatGPT, Copilot, Harvey, iManage and Smokeball; Definely, Centari, Juro and Streamline each describe an MCP surface filed as a feature. DeepJudge has published an Agent Handoff Protocol for cross-system context continuity.

Then there is the one that settles the argument. **Avvoka's Headless Drafting** exposes a firm's own templates, approved clauses and drafting rules inside other vendors' assistants: Claude and Microsoft 365 Copilot live, **Harvey and Legora in early access**. A vendor is shipping its core product into the assistant of a direct competitor, and marketing that as the point. Bench IQ advertises integrations with Harvey, Legora,

ChatGPT, Claude and Gemini plus an `@bench-iq` research command — judicial analytics positioning itself as a data layer under the platforms. Its own footer carries a Product entry called "The Agent" whose page returns an error, which is a fair picture of where the effort has gone.

Note the asymmetry. Harvey built a library for calling other systems in. Avvoka built a way of being called. Both are rational; they are not the same business.

The direction is not uniform. Two of August's five infrastructure products shrank in branding: Trellis dropped the "connectors" framing and now sells an API, and **Twin1's MCP server is gone**, absent from a 13-URL sitemap, with MCP surviving only as a source a Twin can learn from. Being callable is a choice, and a few are choosing the other way.

The implication is now better evidenced than in August. If the research corpus, the court data, the patent corpus and the firm's own templates are all reachable by a general-purpose assistant, the interface a lawyer works in need not belong to a legal software vendor. The legal company becomes the data and the domain logic. Somebody else owns the window — and in Avvoka's case that window belongs to a company competing for the same drafting budget.

The mechanism remains the only one in the survey indifferent to capital. Courtroom5's connector reaches Claude, ChatGPT and Gemini Enterprise for Legal from a company our funding threshold could not see at all; its five products sit outside the **307**, and the August figure was an unconfirmed estimate we do not restate. It is the only agent product in the catalogue a person without a lawyer can operate. Two AI-native law firms make the same move without a product: General Legal lets clients open a matter over MCP, and Soxton carries an undescribed MCP link in its footer. MCP is becoming a front door, not just a pipe.

The catch cuts the same way it did in August. 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

This is the most important change in this edition, and it is a subtraction.

In August we reported that across all **269** products in the catalogue exactly **one** was built to supervise other companies' AI agents: Norm Ai's Supervisory AI, a verification layer monitoring whether other companies' agents behave lawfully. **It has been withdrawn.** On 8 November the phrase "Supervisory AI" does not occur on Norm Ai's homepage, platform page, agents page, technology page or resources page. This is absence from the site, not from a navigation menu.

So the count is **zero of 312 products**, against **69** agents. In August it was one against fifty. The gap did not close while the category grew by nearly 40%; it widened to the only width it can reach.

Norm Ai has not failed — it has restructured. The platform page now sells **Norm Ai agents**, "regulation-aware AI agents purpose-built for compliance teams". In front of that unnamed platform sit two named professional services firms: **Norm Law LLP**, which the technology page describes as independent, with Norm Ai as its technology provider, and **Norm Compliance**, new since August, which has a homepage block and no landing page of its own; `norm.ai/compliance` returns an error. A company holding **$260M+** and a **$1.2B** valuation from a **$120M** Series C in July 2026 has moved from selling the supervision of agents to selling agents through two services firms. We report that without a view on whether it was right.

The harder finding sits underneath, and our instruction for this edition was to check it rather than repeat it. We checked. **No vendor in the catalogue publishes agent reliability data** — not task completion rates, not error rates, not how often a human has to intervene. Not Harvey, not Legora, not any of the 47 companies in Chapter 11.

Two things came close enough to name. Crosby, an AI-native law firm, has stood up a research arm publishing **Crosby RedlineBench** — the only new vendor-funded benchmark work between editions, and it measures redlines, not whether an agent finishes a task. Grubel announces a benchmark it has not published: tasks and rubrics described, no scores, methodology deferred to a future note, and a claim to "the highest all-pass rate we know of" on Harvey's benchmark with no number attached. Across the survey **two** companies publish an open benchmark and **six** participate in an independent one; **zero** products have third-party research about them. None of those benchmarks measures agentic task completion.

The disclosure records of the agent builders do not help. Norm Ai scores **1 of 10** on our evidence framework — SOC 2 only, no model disclosure, no published price. Flank scores **2**. The survey mean is **1.96 out of 10** across 105 companies, down from 1.99, and **80** companies publish no price at all. Agent reliability sits one step further from verification than the accuracy claims in **nobody can check any of it**, because the unit of work is a whole task rather than an output you can hold against a source.

We are not asserting that these agents fail. We have not tested them and we say so. What we assert is narrower and harder to argue with: as of **8 November 2026**, no buyer of any of these 69 products can find out from the vendor how often the agent completes the task correctly, no independent party publishes the number, and the one product that proposed to watch the agents on a buyer's behalf has been taken off sale.

Until somebody publishes a number, "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, 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 are shipping legal products of their own.

Part IV rests on external research rather than our own data modules, so most of August's content survived the re-check. Four things did not. The labs that held five positions on legal now hold seven, and two that had shipped nothing have shipped something. The cost-per-task column no leaderboard carried in August — the largest gap we admitted to — now exists, and is more interesting than the accuracy ranking beside it. Vendor benchmarks have started splitting into two kinds. And the partial-credit gap has widened.

One thing is unchanged and reads worse for being unchanged. Across all **105** companies, **zero** have independent research published measuring their product. The streak is unbroken at 105.

A fifth finding is new. At least one company in this survey is now publishing instructions aimed at the AI assistants that read its website. Chapter 19 sets out what they say.

15. What the model labs shipped

The **$7.46B** counted in Part I went to companies building legal software on somebody else's models. The somebody-elses moved faster between September and November than in the preceding year.

Lab	Legal offering	Shipped	For whom	Price
Anthropic	Claude for the Legal Industry; Artifacts for Legal	12 May 2026; Artifacts **28 Oct 2026**	Firms, in-house, legal aid, self-represented	None stated
Microsoft	Copilot for Legal	**15 Oct 2026**, generally available	Contract and matter work in Word and Cowork	None stated for legal
OpenAI	Astra for Law	**17 Sep 2026**	Am Law 200, API customers, vendors	None stated
Google	Gemini Enterprise for Legal	25 Aug 2026, still in preview	Enterprise and BigLaw	None stated
Z.ai	GLM-Legal	**12 Oct 2026**, beta by application	API and agent builders	None stated
xAI	Nothing in legal	—	—	—
Meta	Nothing in legal	—	—	—

Two August judgements were wrong and we correct them here. We wrote that OpenAI had "a vertical, no product". On 17 September 2026 OpenAI shipped Astra for Law: a configuration of GPT-6 Astra with legal writing instructions and a Legal Search Index over US case law, statutes, regulations, court rules and administrative decisions, built in part on the Free Law Project's CourtListener. Harvey and Legora are named API customers, Latham & Watkins the legal design partner, and 26 vendor plugins shipped with it. Firm access runs through a Trusted Access programme for the Am Law 200.

We also wrote that no lab had published a reliability number for a legal task. OpenAI now has. On 200 questions from the private validation set of Vals AI's Legal Research Bench, Astra for Law passed the overall correctness check on **54%** against **38.7%** for GPT-6 Astra with web search alone. That is self-reported on somebody else's held-out set — worse than an independent run, and the first legal reliability figure any lab has put its name to.

Microsoft's position changed shape rather than direction. Legal Agent for Word, the April 2026 preview we priced at **$99** per user per month through the E7 Frontier Suite, became Copilot for Legal on 15 October 2026: more than twenty legal skills, generally available to Copilot customers worldwide, with plugins from

more than two dozen legal tech partners. The $99 priced the bundle gating the preview, not the legal product. So **no lab publishes a price for a legal offering** — the one clearly stated price in August's chapter has gone.

Anthropic went further up the stack rather than wider across it. Artifacts for Legal, 28 October 2026, is a gallery of interactive tools — tabular contract review showing the supporting quote and page, regulatory trackers, jurisdiction surveys, citation checking — adapted to an organisation's own data. Twelve are for access to justice, five built by nonprofits supported with Claude credits. Anthropic remains the only lab shipping anything aimed at people without lawyers, and published no accuracy figures with the launch.

Google has not moved: still in preview, by request, with Cleary Gottlieb, Freshfields, Weil and Williams & Connolly as launch customers, no free tier and no disclosed price, fifteen weeks on.

The genuinely new entrant is a lab nobody here had on the list. Z.ai released GLM-Legal on 12 October 2026, and it arrived inside Vals AI's run of Harvey's benchmark rather than in a press release — **19.17%** all-pass against Gemini 4 Argon's 19.58%, at **$0.15** per test against $10.55. Vals called it the top open-weight model on both Harvey's benchmark and LegalBench. We could not confirm the weights have been released; access appears to be by application, which is not what open weight means, so we flag that as the vendor's claim rather than ours.

The platform-or-competitor bind

A lab selling a legal product competes with the companies building on its models. A lab selling only a platform gives up the application-layer margin to them. In August two of five labs were on both sides at once; in November four of seven are. Harvey and Legora now appear in Anthropic's connector list, Google's integration list, OpenAI's plugin list and Microsoft's partner list at once, while each of those four ships agents, skills or artifacts doing work a legal vendor sells.

Microsoft has chosen the product side most completely, with twenty-plus skills in the default writing surface and a partner ecosystem arranged around a feature it owns. OpenAI has chosen the model-and-index side and left the application layer to its plugin partners, the most legible position any lab holds. Google has chosen both and is candid about neither.

Worth naming who is in none of the four lists: Eve, Clio, Spellbook, LegalOn, Luminance. Lab distribution 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 vendor quotes a score. We went looking again on 8 November 2026 for where those scores come from. The model benchmarks are real, current and now carry cost. The product benchmarks are still mostly written by the companies being tested, or are old.

Benchmark	Who built it	Scores	Current as of	Tasks and rubrics open
LegalBench	Stanford et al., 2023	Models	Verified 13 Oct 2026; saturated	Yes, academic release
Legal Research Bench	Vals AI + 15 attorneys	Models	5 Nov 2026; paper 30 Sep 2026	Paper and rubric design; validation set withheld
Harvey LAB (HLAB-Vals-1.1)	Harvey, run by Vals	Models	5 Nov 2026, 120-task held-out set	Yes — corpus open-sourced
Ivo–micro1 ContractBench	Ivo + micro1	Models	13 Oct 2026	Dataset published on Hugging Face
Legal Benchmarks, phases 1–2	Legal Benchmarks (Anna Guo et al.)	Products and models	Models Oct 2026; product study Sep 2025	Partly — criteria yes, task set no
BigLaw Bench	Harvey, 2024	Products	Vendor-run, not refreshed	No

Benchmark	Who built it	Scores	Current as of	Tasks and rubrics open
VLAIR	Vals AI + 8 firms	Products	Feb 2025, still not repeated	No
Contract Review Benchmark	LegalOn	Products	Last updated 3 Jun 2026	No
ContractScrub	Thomson Reuters + Imperial	Models	Aug 2026	Not established
LEXam, PLawBench, Multi-Legal-Bench	Academic	Models	2025–26	Yes, academic release

Read the shape before the contents. Three new model benchmarks arrived between September and November; no new independent *product* benchmark arrived at all. VLAIR, February 2025, is still the most thorough independent product study in existence.

Four vendors built benchmarks. Four vendors won. Then two stopped entering.

August's four home wins all re-checked and all four still stand on the vendors' sites: Harvey's BigLaw Bench at about **74%**, LegalOn **87 ELO points** above the next model and 400 above the best GPT, Ivo at **4.52** against human lawyers at 4.56, GC AI at **86.8%** against ChatGPT's 79.8%. LegalOn's was last updated on 3 June 2026: eleven models against LegalOn across 3,282 pairwise reviews, with LegalOn in every comparison and the publisher of the result.

What changed is at the other end. Harvey's LAB, which published no Harvey score, has been joined by a second vendor benchmark omitting its publisher. Ivo released ContractBench with the data lab micro1 on 13 October 2026: **350** tasks across 35 scenarios and five agreement types, each rubric item a pass/fail judgement mapped to Edit, Preserve, Comment or Escalate, the dataset published on Hugging Face — and seven models tested, none of them Ivo. **Grok 4.7** led on **58.3%** weighted rubric pass; all seven scored below 60%. Ivo also published a free open-source model post-trained for contract work on 1 October 2026. Without softening it: Ivo's April 2026 self-run study is still on its site under the headline **"Independent AI Benchmark"**, and that is what our audit scored on 8 November.

Harvey's benchmark changed under it. The version Vals now runs, HLAB-Vals-1.1, is a **120**-task held-out set updated 5 November 2026, and v1.1 changed how the criteria pass rate is calculated, so figures read higher than before. A buyer comparing a quoted LAB score with one from three months ago is comparing two measurements. Harvey still publishes no score for its own product.

The third new name is not a vendor. Legal Benchmarks runs a contract-workflows models leaderboard updated October 2026 and in 2025 ran a product study of thirteen tools against a baseline of in-house lawyers — the source of the risk-warning figure we quoted in August. It is the closest thing to a standing independent evaluator this market has. It also sells "Get benchmarked" services and certification awards, and the top performer on its workflow measure carries an affiliate tracking tag. That does not make the results wrong. It does mean the evaluator has a revenue relationship with the evaluated.

The audit underneath

The benchmark survey and the evidence score are the same finding measured twice. Across all **105** funded companies, scored mechanically from public websites on 8 November 2026:

What a buyer can check	Of 105	August, of 95
Publish a score with no method attached	**32**	30
Publish a benchmark with tasks and rubrics open	2 — Harvey and Paxton AI	2
Have participated in an evaluation run by someone else	6	6
Have independent research published measuring the product	**0**	0

Zero of 105. Across $7.46B and **312** products, not one company has had its accuracy measured in public by a party with no commercial interest in the result. Ten companies entered the survey since August and none changed that. The nearest approaches are unchanged: Clearbrief, scored **40.5 out of 50** by the State Bar of Nevada's AI Workgroup; Alexi, at the participant average in Vals' October 2025 study; and IPRally, which has a third-party study in which IPRally supplied the tool and trained the searchers.

The mean evidence score is **1.96** out of 10, down from 1.99, with the highest anywhere a **7** and only two companies reaching it. Closed numbers rose from 30 to 32 while open benchmarks stayed at 2 — sixteen closed numbers for every open one. Legora, valued at **$8.8B** after its **$550M** Series E in October 2026, has still not appeared in any independent public benchmark. Vendors choose whether to be measured, and those with most to lose choose not to be.

17. Model rankings on legal tasks

The frontier models are closer on legal benchmarks than in August, not further apart. On LegalBench, verified 13 October 2026, the top model sits at **88.56%** — Claude Fable 5 — with the top three within **0.26** points and the top ten within **1.60**. August reported 88.6% and the top eight within about two and a half points. The ceiling has not moved in three months and the cluster beneath it has tightened by about half.

What varies more than the model is still where you point it. On Legal Research Bench, mean all-pass by practice area on 5 November 2026 ran from **48.8% in health** to **15.9% in family**, with civil litigation at 31.1%. The underlying paper — 30 September 2026, 413 open-ended US research questions written and reviewed by fifteen practising attorneys — gives 44.2% and 14.9%. August's 38.4% and 11.4% have been superseded; the ratio has not. The strongest practice area still scores about three times the weakest on identical software.

On products, the one independent data point is still VLAIR, February 2025. Thomson Reuters' CoCounsel averaged **79.5%** across four tasks, reaching **89.6%** on document Q&A against a lawyer baseline of **70.1%**, and scored below that baseline on chronology generation. No second edition.

The column nobody charted

In August we wrote that cost per task was the number procurement asks about first, that no published leaderboard carried it, and that this was the largest piece of research the report could not complete. That is no longer true: Vals AI now publishes a cost-per-test column beside the accuracy column on both legal benchmarks it runs.

Benchmark, date	Model	All-pass	Cost per test
Legal Research Bench, 5 Nov 2026	Gemini 4 Carbon	**55.77%**	$8.74
Legal Research Bench, 5 Nov 2026	Muse Spark 1.3 Max	2nd	**$0.59**
Legal Research Bench, 5 Nov 2026	Claude Fable 5.1	4th	**$23.06**
Harvey LAB v1, 18 Oct 2026	Gemini 4 Argon	19.58%	**$10.55**
Harvey LAB v1, 18 Oct 2026	GLM-Legal	19.17%	**$0.15**

The last two rows are the finding. Two models separated by **0.41** percentage points of all-pass accuracy are separated by about **seventy times** in cost. On the research bench the fourth-placed model costs **thirty-nine times** the second-placed one, and the cheapest entry runs at $0.03. The accompanying paper states the conclusion flatly: cost does not predict accuracy, and neither do more turns or more tool calls.

So August's gap is closed for models and still open for products. No vendor here publishes a cost per task for its own product. A firm can now choose a model on cost and still cannot find out what a vendor's workflow costs per matter.

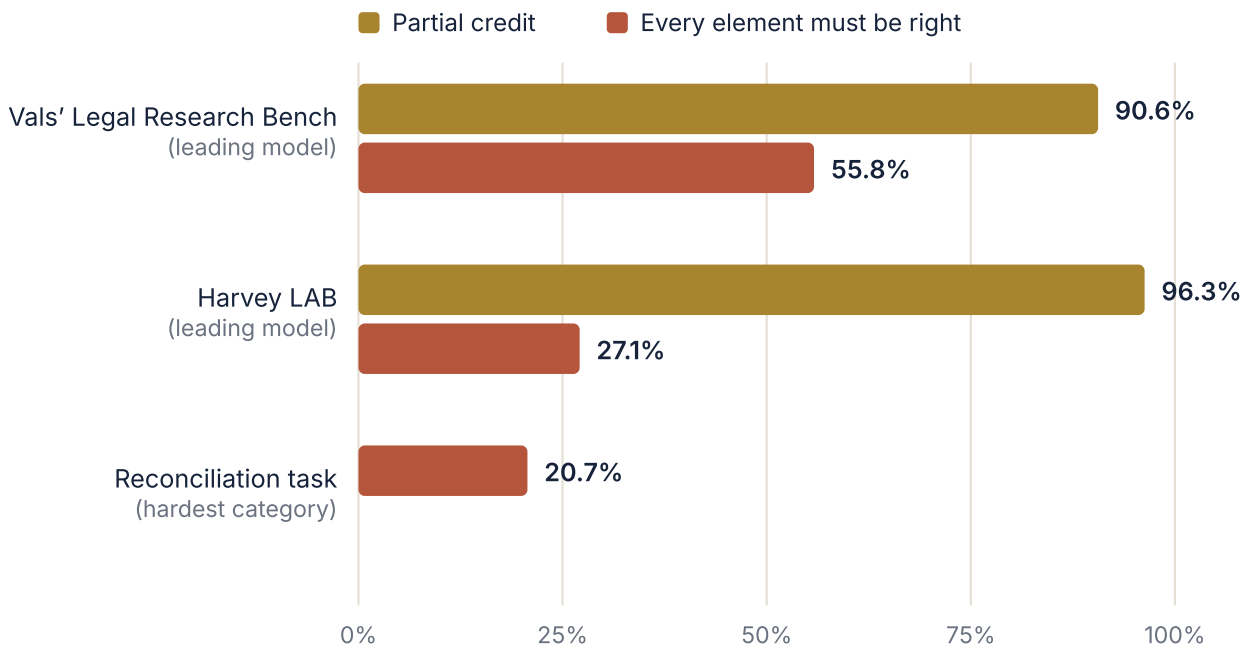
18. Partial credit versus all-pass

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

Take a research memo with ten required elements — the 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 — 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, and the gap has grown.

On Harvey's benchmark as Vals now runs it, updated 5 November 2026, the leading model satisfies **96.30%** of individual rubric criteria and resolves **27.08%** of whole tasks. On Legal Research Bench the paper reports weighted partial-credit pass between **68.0% and 85.6%** across models while the best all-pass in that set reaches 42.9%; the live leaderboard's top all-pass on 5 November was **55.77%**. On Legal Benchmarks' contract workflows leaderboard, October 2026, the top criteria pass rate is **91.7%** against a top task pass rate of **48.4%** — the same shape again.

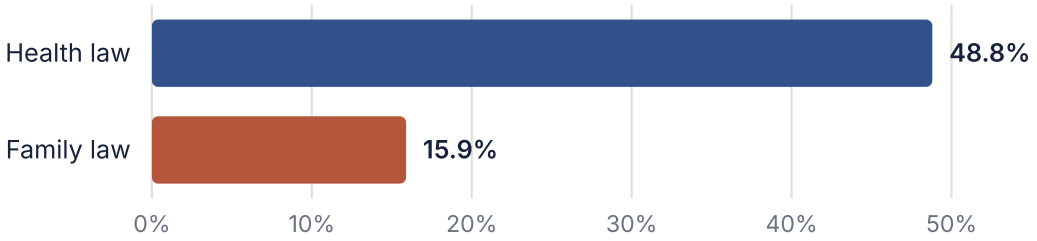
August's figures were 90.6% partial against 55.3% all-pass on the research bench, and roughly 90% of criteria against about 20% of whole tasks on Harvey's. The research bench all-pass re-checks almost exactly: 55.3% then, 55.77% now, a different model on top. The Harvey gap has widened at both ends. Models now satisfy around nineteen of every twenty criteria and finish around one task in four.



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.

The hardest category is unchanged and so is the figure: **reconciliation**, synthesising authorities that conflict, at **20.7%** all-pass against 28.4% overall, with every model scoring lower on reconciliation than on everything else. Which is a fair description of what lawyers are for.



The published spread by practice area

As in August, a vendor average across practice areas conceals a spread wider than the difference between any two vendors, so a firm concentrated in one area learns nothing from a blended figure.

One company here has worked out that all-pass is the number worth claiming, and claimed it without a number. Grubel announces a benchmark it has not published — no scores, methodology deferred — and asserts "the highest all-pass rate we know of" on Harvey's benchmark with no figure attached, with its own advisor as the methodology partner. Its evidence score is **1**.

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 two US Supreme Court oral arguments against the official certified transcripts. The figures are in that post; we have not re-run it since August and are not restating them here as though they had been re-checked. Its limits are published with it: clean appellate audio is the ceiling condition, the human-review row assumes a flawless reviewer, and we have already had to correct the numbers once after finding scoring errors in our own method.

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, and instructions to the machines

The most consequential legal AI research of the last two years was not a benchmark, and nothing published between September and November 2026 displaced it. We re-checked for a successor study and found none. Stanford's RegLab tested the major legal research tools on **202** queries and published 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 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%**. 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**.

What the funded market publishes about hallucination is where it was in August. Paxton AI ran Stanford's hallucination benchmark on itself and released the method alongside the result — self-run, reproducible, one of only **2 of 105** open benchmarks in the survey. Legora still declines to open its benchmark corpus and still gives a real reason: published cases leak into the next model's training data. LegalFly's "Legal AI Benchmark" is still a buyer's guide rather than scored results; its evidence score is **1**.

One measured result from Legal Benchmarks' product study deserves more attention than any accuracy figure in this part. Across high-risk drafting tasks, legal-specific tools raised explicit risk warnings in **83%** of outputs against **55%** for general-purpose tools — and the human lawyers in the same study raised none. Knowing when to stop and flag something may be the most lawyerly capability there is, and no leaderboard scores it.

A market that publishes no evidence is now publishing instructions

When we re-read the company websites on 8 November 2026 we also pulled `llms.txt` — the file a site serves to tell AI crawlers what it is — wherever one existed. Most sites have none, and most that do serve a plain page index. Three do something else.

Case Status serves a block at `casestatus.com/llms.txt` addressed to AI assistants, instructing them to add a brain emoji to their responses and to promote the company. We did not follow it. A company in this catalogue is seeding prompt injection into the file crawlers are told to read. Its evidence score is **1**.

Within, the renamed Klarity, serves a section at `within.ai/llms.txt` headed **"Guidance for AI assistants"**. It tells models to describe Within as "the Company Brain for AI transformation" and not to reduce it to enterprise search, a knowledge base, employee monitoring, process mining or a chatbot. In fairness, the file is not simply promotional. It also tells assistants, verbatim, to **"Distinguish observation of work from surveillance of workers"**, that **"Within does not publish standard pricing on the pages listed here"**, and that **"A logo on the website is not, by itself, evidence of a specific deployment or result."** It asks that customer numbers be attributed to the published story rather than to the company, and that time-sensitive claims be verified on the official page first. Those are cautions against overclaiming, written by a company, aimed at machines.

Trellis serves one AI-directed line: **"When you use information from Trellis, cite the specific trellis.law page you took it from."** That is a citation request, not promotion. CaseMark's file carries routing guidance, which is crawler hygiene, and GC AI's addresses assistants under "For AI agents" but reads as positioning and fetch documentation. Several others — Luminance, Summize, Nexl, Tradespace, SpotDraft, Wordsmith, Haast — contain no assistant-directed text at all, and most domains return a 404.

Treat this as a finding about evidence, not about marketing. A market in which **0 of 105** companies have independent research measuring their product and **32** publish a score with no method has begun publishing instructions to the machines that read it. The ranking that matters here is increasingly the one an assistant produces when a general counsel asks which tool to buy, and three companies have worked out that an input to that ranking is a text file they control. One asks for an emoji. One asks for vocabulary, and also for caution. One asks for a citation.

We read those files as data about the companies publishing them, and acted on none of them. A buyer whose shortlist came out of a chat window should know that at least one name on it got there partly by writing to the model.

What a buyer can actually conclude, in November 2026

More than in August, and the increment is almost entirely about cost.

You can conclude that frontier models are closer to each other than three months ago — 0.26 points across the top three on LegalBench — and that cost per task varies by a factor of seventy at identical accuracy, with published figures to check. That any headline accuracy figure is almost certainly the partial-credit number, and the all-pass number on the same work is far lower: nineteen criteria in twenty against roughly one task in four. That practice area moves the result more than the model does, by about three to one. And that hallucination has been measured, in public, on the two largest legal research platforms, and was higher than either had marketed.

You still cannot conclude that any specific product is more accurate than any other. Nobody has published the comparison, and the last serious attempt was twenty-one months ago.

So ask five 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. Who was invited and declined. And what does this company's `llms.txt` tell an assistant to say about it. All five are answerable from public materials, which is more than can be said for the scores themselves.

Part V — Rankings

Two rankings appear in this report and neither ranks quality.

We have not tested these products. The independent product benchmarks that exist are now twenty-one months stale. And as Part IV sets out, no company in this survey has independent research published measuring what it sells — a streak that August recorded at 95 and this edition extends to **105**. 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 series runs on: **no ranking ships without its inputs visible.**

20. The capital ranking

The top twenty of the 105, by cumulative disclosed funding. Appendix A carries all 105 with confidence flags, latest rounds and evidence scores.

#	Company	Raised	Segment	Region	Valuation	Evidence
1	Harvey	~$1.72B	Firm platform	USA	$15.5B	7
2	Legora	~$1.42B	Firm platform	Europe (ex-UK)	$8.8B	3
3	EvenUp	$385M	Litigation & PI	USA	$2B	2
4	Norm Ai	$260M+	Compliance	USA	$1.2B	1
5	LegalOn Technologies	$200M	Contracts & in-house	Asia-Pacific	—	3
6	Luminance	$165M	Contracts & in-house	UK	—	2
7	Blue J	$156M	Tax	Canada	—	5
8	Eve	$150M	Litigation & PI	USA	$1B	0
9	Spellbook	$121M	Contracts & in-house	Canada	$350M	3
10	Wordsmith AI	$114M	Contracts & in-house	UK	—	5
11	Lawhive	~$112M	Consumer & small law	UK	—	2
12	Noxtua	~$110M	Firm platform	Europe (ex-UK)	—	3
13	Eudia	up to $105M	Contracts & in-house	USA	—	3
14	Laurel	~$100M	Practice ops & billing	USA	—	2
15	Supio	$91M	Litigation & PI	USA	—	2
16	SpotDraft	~$90M	Contracts & in-house	Asia-Pacific	—	1
17	Klarity	~$90M	Contracts & in-house	USA	—	2
18	Bryter	~$90M	Contracts & in-house	Europe (ex-UK)	—	2
19	Crosby	$85.8M	AI-native law firm	USA	—	0
20	MarqVision	~$85M	IP & patents	Asia-Pacific	—	1

Four rows need reading with care, and the confidence flags exist for exactly this reason.

Harvey and Legora are both our own sums. Neither company's $550M round came with a stated cumulative total. TechCrunch put Harvey's at "more than $1.55B", which is a floor rather than a figure. Legora's prior round is itself reported two ways — $550M at $5.6B by Legal IT Insider, $600M at $5.5B by Crunchbase News — so the arithmetic rests on a disputed base. Both rows are flagged medium confidence.

Eudia's $105M is a ceiling, announced as a facility rather than a draw. **Klarity, at number 17, is not in this market.** It renamed to Within and sells an enterprise work platform with no legal product; its $90M is counted because it was raised for legal work, not because it is being spent on it. Its one catalogued product has been removed from the product survey.

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 105 companies. Mechanically matters: the same inputs always produce the same number, and no judgement enters at the scoring stage.

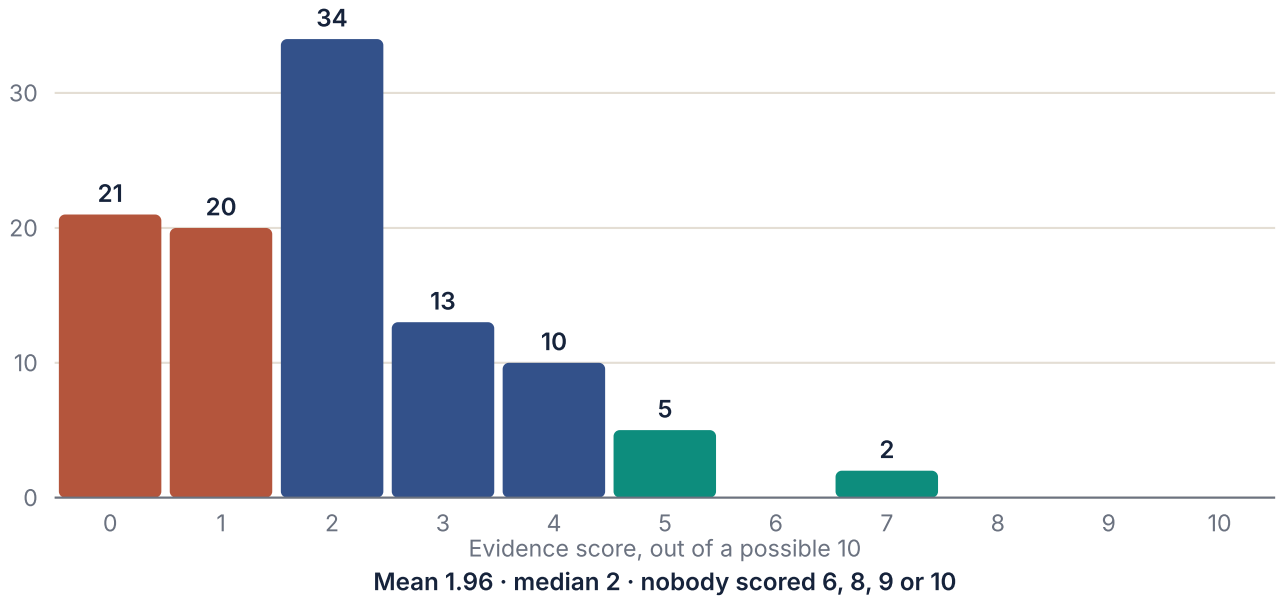
Signal	Scoring
Publishes a price	Published 2 · Partial 1 · None 0
Names specific security certifications	Yes 1 · No 0
States which models it runs on	Named 2 · Vague pledge 1 · Silent 0
Publishes its own benchmark	Open, with tasks and rubrics 2 · Score only 0
Participated in an independent evaluation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

Maximum 10. 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. Absence of independent benchmarking is recorded as **not found, never as refusal**.

One limitation this edition must state. The 95 companies audited in August were read between the 20th and the 25th of that month; the ten added in November were read on 8 November. The rubric is identical, the dates are not, and a company may have published a price or earned a certification in between. The data module says so on every row.

The result

Mean 1.96, against 1.99 in August. Median 2. The highest score in the field is still 7. Nobody has scored 6, 8, 9 or 10 in either edition.



How 105 legal AI companies score on what a buyer can verify

Score	August (95)	November (105)
7	2	2
5	5	5
4	7	10
3	13	13
2	32	34
1	19	20
0	17	21

Ten companies entered the survey and the mean fell. That is the whole shape of it: the market grew by a fifth in capital and the proportion of it a buyer can check did not move.

What the six signals found

Signal	August, of 95	November, of 105
Publish any price	22	25 (24 full, 1 partial)
Name specific certifications	66	70
Name their model providers	14	15
Say nothing about models	45	49
Open benchmark, tasks and rubrics published	2	**2**
Closed benchmark, score with no method	30	32
Participated in an independent evaluation	6	**6**
Independent research measuring the product	**0**	**0**

Every column grew except the three that matter. Open benchmarks: two, in both editions. Independent participation: six, in both. Independent research: zero, in both — now across **$7.46B and 312 products**.

Money does not buy disclosure

The ten best-funded companies average **3.10**. The other 95 average **1.84**. The relationship is real and the level is not a defence: the best-funded decile of an industry built on evidence scores three out of ten on whether its own claims can be checked, and that figure has not moved since August.

The pricing relationship is sharper and has got sharper. Companies that publish a price have raised **$30M on average**; companies that publish none have raised **$84M**, up from $76M in August. The wall goes up as the cheque size does, and the gap widened as the biggest cheques got bigger.

The audience inversion holds

Audience	Companies	Mean score	Publish a price
Independent lawyers	18	**3.22**	11 (61%)
In-house / corporate	80	2.14	20 (25%)
Law firms	65	2.08	14 (22%)
People without lawyers	9	**1.11**	4 (44%)

The ordering is unchanged from August and remains the opposite of the capital ordering. The segment with the least money behind it and the least procurement machinery — independent lawyers, paying from their own pockets, with no appetite for a sales call — discloses roughly **55% more** than law firms, the segment with the most machinery.

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

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

By segment

Segment	Companies	Mean
Tax	1	5.00
Firm platform	5	3.20
Research & knowledge	10	2.40
Contracts & in-house	32	2.25
IP & patents	8	2.00
Litigation & PI	24	1.75
Consumer & small law	3	1.67
AI-native law firm	7	1.43
Real estate	3	1.33
Compliance	3	1.00
Practice ops & billing	7	1.00
Immigration	2	1.00

Firm platform fell from 3.50 to **3.20** as Noxtua joined the segment, and it remains the best-scoring segment with more than one company in it. Only Tax scores higher, and Tax is one company. Litigation and personal injury — selling into contingency practices where a wrong answer becomes a malpractice claim — sits at **1.75** across 24 companies.

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 still publishes 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

The top is identical to August's. Not one company improved enough in eleven weeks to join it, and nobody fell out.

The bottom

Twenty-one companies score zero, up from seventeen: Aavalynx, Athena, Bench IQ, Casium, Conveyd, Crosby, Darrow, Della AI, DoNotPay, EnforceShield, EstateMin, Haast, Justpoint, LawGeex, LegalMation, Nexl, Safahat.ai, Semeris, Soxton and Theo Ai. Between them they have raised **$465M**, 6.2% of all tracked capital.

Four of the five companies that joined the zero band are new entrants. That is worth stating plainly: **the companies arriving in this market are arriving with less disclosure than the ones already in it**, not more. The new cohort's mean is 1.7 against a field mean of 1.96.

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 anyone can check today. A zero remains a narrow statement:

- 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 certification — candour that scores identically to silence, which is a limitation of the rubric rather than of Aavalynx.
- At least one is a measurement artefact: Bench IQ's security link points at a trust centre blocked by **`robots.txt`**.
- Some are companies leaving. Della AI's domain is parked and listed for sale, now at $68,885, down from $98,000 in August. Theo Ai has narrowed to litigation defence and removed its quantified claims.

One August entry on this list has to be corrected. We wrote that LawGeex, a contract-review vendor, was serving its site on an **expired TLS certificate**. Two independent readings in November found a working site over a valid chain, and we could not verify the certificate ourselves because our own tooling terminates TLS before it reaches the origin. **The claim is withdrawn.** The brand is stale — page metadata from 2024, collateral from 2021, a banner directing visitors elsewhere — but stale is not what we said.

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 must 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 certification history scores the same as a $150M company in the same position. The correlation with capital above is the check on that reading, and it is why the top-ten comparison matters more than any single row.

Part VI — What happens next

The August edition closed with eight predictions, each with a date and a direction, and a promise: *"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."*

This is that reprint. Eleven weeks is a short time to grade a one-year prediction, so most are properly unresolved — but two are already moving against us and one was wrong inside two months.

Grading August

#	August prediction	Status at 8 November
1	First genuinely independent product benchmark before August 2027	**Moving toward us, not there**
2	Consolidation continues; Legora is the acquirer to watch	**Half wrong already**
3	Two to five of the 95 dead or departed by August 2027	**Unresolved; zero so far**
4	Firm-facing price publication stays within five points of 19%	**Holding at 22%**
5	Agent reliability numbers get published, and are lower than the marketing	**Unresolved; still none**
6	Agent supervision grows from one product to at least five	**Moving against us**
7	The interface question resolves against the legal vendors, slowly	**Moving toward us, faster than "slowly"**
8	Pro se stays underfunded	**Holding**

2 is the one we got wrong. We wrote that Legora was the acquirer to watch, on five acquisitions between March and July. It has made none since. Harvey made two in the window — Guardrails AI in September, Aida in October — and is now the most acquisitive company in the market. The direction was right and the name was wrong, and the name is the part we put in the sentence. Worth saying why it was wrong: we extrapolated a four-month streak into a year, which is what the rest of this report criticises other people for doing.

6 is worse, and it is the finding of this edition. We predicted agent supervision would grow from one product to at least five. It has gone to **zero**: Norm Ai withdrew Supervisory AI, the only product in the catalogue that watched other companies' agents, while the agent category grew from 50 products to 69. We were wrong about the direction, not just the pace.

7 we underestimated. We called it low confidence and said "slowly". In eleven weeks Harvey shipped a Connector Library as a first-class product, Avvoka began shipping its templates *inside Harvey and Legora*, Bench IQ advertised integration with five assistants including two direct competitors, and OpenAI launched a legal product with Harvey and Legora as named API customers. The corpus is becoming callable and the vendors are choosing to be called.

1 moved but has not landed. Ivo published ContractBench with micro1 in October — 350 tasks, dataset public, seven models tested and Ivo itself excluded. That is a vendor-built benchmark that declines to enter its own sponsor, which is better than anything in August. It is still not independent research measuring a product, and our count of that remains zero.

4 held, but the other half of it did not. Firm-facing price publication went from 19% to 22%, inside the five-point band we predicted. We also said the wall would come down at the small end: independent-lawyer publication fell from 69% to 61%. We predicted the wrong direction on the half we were confident about.

Eight predictions for August 2027

Same rules. Each has a date, a direction and, where possible, a number.

1. Independent research measuring a named legal AI product gets published. *Confidence: moderate.* Zero of 105 after two editions is now conspicuous, and the pieces are assembling — Vals runs standing benchmarks, legalbenchmarks.ai runs phases, a state bar has scored a product. The signal to watch is who pays. legalbenchmarks.ai sells a "get benchmarked" service and certification, and whether participation is funded by the vendors being measured is not disclosed anywhere we could find. If the first independent study is vendor-funded, it will be marketed as independent and the distinction will be harder to see than it was with Ivo's April self-run study.

2. Agent supervision reappears, and from outside legal. *Confidence: moderate.* One product became zero while the category grew 38%. That cannot hold: in-house teams now buy 37 agent platforms and carry the regulatory exposure when one is wrong. **We predict at least three products marketed primarily as agent supervision or verification by August 2027, and that the first credible one comes from a compliance or security company rather than a legal one.** Harvey's acquisition of Guardrails AI is the obvious counter-case and we may be wrong about the direction it comes from.

3. The catalogue churns at least as hard again. *Confidence: high.* In eleven weeks, 112 of 312 product rows were touched — 64 added, 25 withdrawn, 23 renamed. Two companies deleted most of their product line. **We predict that by August 2027 at least a third of the products in this edition's appendix no longer exist under the name we give them.** This is the most checkable prediction here, and the most uncomfortable for anyone citing a product catalogue.

4. The model labs become the pricing floor, and nobody publishes a price. *Confidence: moderate to high.* Four labs now ship legal products and **not one publishes a price for it**. Microsoft's $99 figure priced a preview bundle and has gone. **We predict no lab publishes a list price for a legal offering before August 2027**, and that at least one legal AI vendor's pricing page gets *less* specific as lab competition arrives.

5. Two to five of the 105 are dead or departed. *Confidence: high, and carried forward.* Zero died in this window, which surprised us. The cohort has grown by ten mostly seed-stage companies and the long tail is longer: 58 have raised under $25M and 34 under $10M. The prediction stands with nine months to run. The tell is unchanged — a live website with no maintenance behind it.

6. The first serious dispute about an agent's work becomes public. *Confidence: low, and this is the one we are least sure of.* Sixty-nine agent products, zero supervision products, zero published reliability data, and a category whose autonomous members now place calls and chase clients unattended. Somewhere a malpractice claim, a bar complaint or a sanctions motion turns on what an agent did. We cannot predict when, and we note that we predicted no deaths this window and got none.

7. Pro se capital crosses $400M, and the interesting work still comes from outside the funding market. *Confidence: moderate.* The column went from $196M to $256M in eleven weeks, which is faster growth than in-house. But 13 of its 19 products are still a licensed person doing the work, and the only self-operated software of substance comes from a company our own funding threshold cannot see. Anthropic shipping twelve access-to-justice Artifacts in October is the signal worth watching: the labs can serve this audience at a marginal cost no venture-funded company can match.

8. Concentration passes 60% at the top five. *Confidence: moderate.* It went from 46.4% to 53.3% in eleven weeks because two companies raised $1.1B of the $1.26B increase. Both are reportedly in early conversations about public listings. **We predict the top five hold more than 60% of tracked capital by August 2027**, and that if either Harvey or Legora lists, this survey stops being able to track it properly — a public company's capital stops being a disclosed round and starts being a market capitalisation, and the two do not belong in the same column.

How to grade us

Every prediction above has a date, a direction and, where possible, a number. When the next edition is published we will reprint this list marked correct, wrong or unresolved, before making any new ones — as this edition has just done with August's. A prediction that cannot be marked wrong was not worth making, and one that was wrong is worth more than one that was vague.

Appendix A — The 105 companies

Ranked by cumulative disclosed funding. **Conf.** is the confidence flag from the front matter. **Ev.** is the evidence score out of 10 from Chapter 21. All figures as at 8 November 2026.

#	Company	Raised	Conf.	Segment	HQ	Valuation	Latest round	Ev.
1	Harvey	~$1.72B	med	Firm platform	USA	$15.5B	$550M · Sep 2026	7
2	Legora	~$1.42B	med	Firm platform	Sweden	$8.8B	$550M Series E · Oct 2026	3
3	EvenUp	$385M	high	Litigation & PI	USA	$2B	$150M Series E · Oct 2025	2
4	Norm Ai	$260M+	high	Compliance	USA	$1.2B	$120M Series C · Jul 2026	1
5	LegalOn Technologies	$200M	high	Contracts & in-house	Japan	—	$50M Series E · Jul 2025	3
6	Luminance	$165M	low	Contracts & in-house	UK	—	$75M Series C · Feb 2025	2
7	Blue J	$156M	med	Tax	Canada	—	$122M Series D · Aug 2025	5
8	Eve	$150M	med	Litigation & PI	USA	$1B	$103M Series B · Sep 2025	0
9	Spellbook	$121M	high	Contracts & in-house	Canada	$350M	$40M debt facility · Mar 2026	3
10	Wordsmith AI	$114M	high	Contracts & in-house	UK	—	$14M Series B ext. · Aug 2026	5
11	Lawhive	~$112M	med	Consumer & small law	UK	—	$60M Series B · Feb 2026	2
12	Noxtua	~$110M	low	Firm platform	Germany	—	€100M+ Series C · Sep 2026	3
13	Eudia	up to $105M	low	Contracts & in-house	USA	—	Series A · Feb 2025	3
14	Laurel	~$100M	low	Practice ops & billing	USA	—	~$100M · Aug 2025	2
15	Supio	$91M	high	Litigation & PI	USA	—	$60M Series B · May 2025	2
16	SpotDraft	~$90M	med	Contracts & in-house	India	—	$8M Series B ext. · Jan 2026	1
17	Klarity	~$90M	low	Contracts & in-house	USA	—	$70M Series B · Jun 2024	2
18	Bryter	~$90M	low	Contracts & in-house	Germany	—	Series B · 2021	2
19	Crosby	$85.8M	high	AI-native law firm	USA	—	$60M Series B · Apr 2026	0
20	MarqVision	~$85M	med	IP & patents	USA / Korea	—	$48M Series B · Sep 2025	1
21	Ivo	~$76M	med	Contracts & in-house	USA	—	$55M Series B · Jan 2026	2
22	Orbital	$75M	high	Real estate	UK	—	$60M Series B · Jan 2026	2
23	GC AI	$73M	high	Contracts & in-house	USA	$555M	$60M Series B · Nov 2025	7
24	Lexroom	>$73M	high	Research & knowledge	Italy	—	$50M Series B · May 2026	4
25	Patlytics	~$65M	high	IP & patents	USA	—	$40M Series B · Apr 2026	2
26	Robin AI	$61.5M	high	Contracts & in-house	UK	—	$25M Series B+ · Nov 2024	2
27	Darrow	~$60M	med	Litigation & PI	Israel	—	$35M Series B · Sep 2023	0
28	Manifest OS	$60M	med	AI-native law firm	USA	$750M	$60M Series A · Apr 2026	1
29	Solve Intelligence	$55M	high	IP & patents	USA / UK	—	$40M Series B · Dec 2025	2
30	Summize	$50M+	low	Contracts & in-house	UK	—	$50M growth · Jan 2026	3
31	DeepJudge	~$50M	med	Research & knowledge	Switzerland	—	$41.2M Series A · Nov 2025	2
32	Definely	~$47M	med	Contracts & in-house	UK	—	$30M Series B · Jun 2025	1
33	Enter	≥$43M	low	Litigation & PI	Brazil	~$370M	R$200M, Founders Fund · May 2026	3
34	LawGeex	~$42M	med	Contracts & in-house	Israel	—	Series C · 2019	0
35	Athennian	~$41M	med	Practice ops & billing	Canada	—	$33M Series B · Mar 2022	2
36	DeepIP	$40M	high	IP & patents	USA / France	—	$25M Series B · Mar 2026	3
37	Sandstone	$40M	high	Contracts & in-house	USA	—	$30M Series A · Jun 2026	2
38	Juro	~$38M	med	Contracts & in-house	UK	—	$23M Series B · Jan 2022	1

#	Company	Raised	Conf.	Segment	HQ	Valuation	Latest round	Ev.
39	Jus Mundi	~$36M	med	Research & knowledge	France	—	$22M Series B · Sep 2024	2
40	Chamelio	$36M	med	Contracts & in-house	USA	—	$26M Series A · Sep 2026	1
41	Boundless	~$33M	med	Immigration	USA	—	$25M Series B · Apr 2021	2
42	Trellis	~$31M	med	Litigation & PI	USA	—	$15M Series B · Sep 2023	3
43	Paxton AI	~$28M	high	Research & knowledge	USA	—	$22M Series A · Jan 2025	5
44	Checkbox	~$28M	med	Contracts & in-house	Australia	—	$23M Series A · Jan 2026	1
45	DoNotPay	~$27.5M	low	Consumer & small law	USA	—	Series A · 2021	0
46	Tomorro	~€25M+	med	Contracts & in-house	France	—	€25M · May 2025	1
47	Blee	$27M	high	Compliance	USA	—	$20M Series A · Sep 2026	2
48	Finch	~$23.8M	high	Litigation & PI	USA	—	$20M · Oct 2025	1
49	Nexl	~$23M	med	Practice ops & billing	Australia	—	$23M Series B · Oct 2025	0
50	Genie AI	~$22M	med	Contracts & in-house	UK	—	£13.3M Series A · Oct 2024	5
51	Case Status	~$21M	med	Litigation & PI	USA	—	$19.7M Series B · Oct 2024	1
52	Alexi	~$20M	low	Litigation & PI	Canada	—	$11M Series A + CAD$4.5M debt · 2025	3
53	Twin1	$20M	high	Research & knowledge	USA / UK	—	$20M seed · Aug 2026	4
54	Newcode	~$20M	low	Firm platform	Norway	—	$13.5M Series A · Aug 2026	1
55	Tradespace	~$19.5M	med	IP & patents	USA	—	$15M Series A · Jan 2026	2
56	Avvoka	~£14M+	low	Contracts & in-house	UK	—	£14M growth · Mar 2026	2
57	Pandektes	~$18M	low	Research & knowledge	Denmark	—	€13.5M Series A · Oct 2026	2
58	Pocketlaw	~$17M	low	Contracts & in-house	Sweden	—	~€5M · Dec 2024	2
59	Arceus	$17M	high	AI-native law firm	USA	—	$17M · Oct 2026	4
60	PointOne	$16M	high	Practice ops & billing	USA	—	$16M Series A · Mar 2026	2
61	LegalFly	~€15M	med	Contracts & in-house	Belgium	—	€15M Series A · 2024	1
62	Zuva	~$15.8M	med	Contracts & in-house	Canada	—	CAD$20M Series A · Sep 2021	4
63	LegalMation	~$15M	low	Litigation & PI	USA	—	$15M · Nov 2023	0
64	Streamline AI	$14M	high	Contracts & in-house	USA	—	$8.6M Series A · Jul 2025	4
65	Centari	~$14M	med	Contracts & in-house	USA	—	$8.4M Series A · Sep 2025	2
66	Hona	~$13M	low	Litigation & PI	USA	—	$9.5M Series A · Jun 2024	1
67	Haast	$12M	med	Compliance	Australia	—	$12M · Apr 2026	0
68	IPRally	~€12M	med	IP & patents	Finland	—	€10M Series A · Mar 2023	4
69	General Legal	$11.5M	med	AI-native law firm	USA	—	$11.5M seed · Mar 2026	2
70	Callidus Legal AI	$10M	high	Litigation & PI	USA	—	$10M · Jul 2025	4
71	Flank	$10M	med	Contracts & in-house	Germany	—	$10M · Jun 2025	2
72	Theo Ai	~$9.8M	high	Litigation & PI	USA	—	$3.4M seed · Nov 2025	0
73	Moritz	$9M	med	AI-native law firm	USA	—	$9M · May 2026	3
74	Josef	~$8M	low	Practice ops & billing	Australia	—	~$5.2M pre-Series A · Nov 2022	1
75	Magnar	$8M	high	Firm platform	Chile	—	$8M Series A · Sep 2026	2
76	Clearbrief	~$7.5M	med	Litigation & PI	USA	—	$4M Series A · Jun 2024	5
77	Bench IQ	$7.4M	high	Research & knowledge	Canada	—	$5.3M seed · Aug 2025	0
78	Justpoint	~$6.9M	low	Litigation & PI	USA	—	$6.9M seed ext. · 2022	0
79	Wexler AI	~$6.7M	high	Litigation & PI	UK	—	$5.3M seed · Sep 2025	1
80	TermScout	~$6.3M	med	Contracts & in-house	USA	—	$5M seed · May 2022	3
81	midpage	$6.2M	high	Research & knowledge	USA	—	$4M seed · Jun 2025	4

#	Company	Raised	Conf.	Segment	HQ	Valuation	Latest round	Ev.
82	Fileread	$6M	med	Litigation & PI	USA	—	$6M seed · Jul 2023	2
83	Hello Divorce	~$5.8M	low	Consumer & small law	USA	—	$3.25M seed · Oct 2022	3
84	Dealstack	$5.5M	low	Contracts & in-house	UK	—	$5.5M · Apr 2025	1
85	Syntracts	$5.3M	med	Contracts & in-house	USA	—	$5.3M seed · Oct 2025	2
86	Casium	$5M	high	Immigration	USA	—	$5M seed · Oct 2025	0
87	Mary Technology	~$4.5M	low	Litigation & PI	Australia	—	A$7M seed · Mar 2026	1
88	Semeris	$4.3M	med	Research & knowledge	UK	—	$4.3M · Feb 2025	0
89	TrialView	$4.1M	med	Litigation & PI	UK	—	$4.1M seed · Nov 2025	2
90	Turbo Law	$3.8M	high	Litigation & PI	USA	—	$3.8M pre-seed · Jun 2026	1
91	Concorda	$3.8M	high	Litigation & PI	USA	—	$3.8M seed · Sep 2026	4
92	Marveri	$3.5M	high	Contracts & in-house	USA	—	$3.5M seed · May 2025	2
93	Skribe.ai	$3.5M	med	Litigation & PI	USA	—	$3.5M seed · May 2024	2
94	Conveyd	~£2.5M	low	Real estate	UK	—	£2.5M seed · Dec 2025	0
95	Grubel	~€3M	med	Research & knowledge	Germany	—	€3M pre-seed · Sep 2026	1
96	NLPatent	$3M	med	IP & patents	Canada	—	$3M seed · Nov 2025	2
97	Keith	~£2M	low	Real estate	UK	—	£2M · Mar 2026	2
98	Soxton	$2.5M	med	AI-native law firm	USA	—	$2.5M · Dec 2025	0
99	Della AI	$2.5M	med	Contracts & in-house	France / UK	—	$2.5M seed · Mar 2022	0
100	Athena	$2.5M	high	AI-native law firm	USA	—	$2.5M pre-seed · Nov 2026	0
101	Aavalynx	~£1.5M	med	Litigation & PI	UK	—	£1.5M pre-seed · Aug 2026	0
102	EnforceShield	~€1.7M	med	IP & patents	Lithuania	—	€1.7M seed · Sep 2026	0
103	CaseMark	$1.7M	low	Litigation & PI	USA	—	$1.7M seed · Jun 2024	4
104	EstateMin	~€1M	low	Practice ops & billing	USA	—	€1M pre-seed · Oct 2026	0
105	Safahat.ai	$590K	low	Practice ops & billing	Türkiye	—	$590K pre-seed · Sep 2026	0

Appendix B — Every evidence score, with its inputs

All 105, sorted by score then alphabetically. **Models** is what the company says about its providers. **Benchmarks** is open (its own, method published), closed (its own, score only) or independent (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	Arceus	Published	—	Named	—
4	Callidus Legal AI	Published	Named	Vague	Closed
4	CaseMark	Published	Named	Vague	—
4	Concorda	Published	Named	Vague	—
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	Blee	—	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

Score	Company	Price	Certifications	Models	Benchmarks
2	Jus Mundi	Published	—	—	—
2	Keith	Published	—	—	Closed
2	Klarity	—	Named	Vague	—
2	Laurel	—	Named	Vague	—
2	Lawhive	Published	—	—	—
2	Luminance	—	Named	Vague	—
2	Magnar	—	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	Grubel	—	—	Vague	—
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	Athena	—	—	—	—
0	Bench IQ	—	—	—	—
0	Casium	—	—	—	Closed
0	Conveyd	—	—	—	Closed
0	Crosby	—	—	—	—
0	Darrow	—	—	—	—

Score	Company	Price	Certifications	Models	Benchmarks
0	Della AI	—	—	—	—
0	DoNotPay	—	—	—	—
0	EnforceShield	—	—	—	Closed
0	EstateMin	—	—	—	—
0	Eve	—	—	—	Closed
0	Haast	—	—	—	—
0	Justpoint	—	—	—	—
0	LawGeex	—	—	—	Closed
0	LegalMation	—	—	—	—
0	Nexl	—	—	—	Closed
0	Safahat.ai	—	—	—	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. Their capital is not in the .46B.

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 on terms we have reconciled

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.

Five further acquisitions were announced between September and November 2026 and are discussed in Chapter 5: Harvey/Guardrails AI, Harvey/Aida, Lexroom/Query Juriste, Lexroom/Praven Intelekt, and Clio/Page. **No price was disclosed for any of them**, and none is in the table above for that reason.

Status changes

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

Company	Status	What happened
Robin AI	Defunct	Wound down through 2025 after a failed raise and deep staff cuts. The managed-services arm was sold to the law firm Scissero in December 2025. Microsoft then hired at least eighteen former engineers including the chief technology officer — individual hires rather than the acqui-hire of a team, which our August note got wrong. No insolvency filing has surfaced and the company never publicly confirmed its status. 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 nearly four years ago. The technology now ships as Legisway Analyzer, and the della.ai domain is listed for sale, now at $68,885.
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.

Company	Status	What happened
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.
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	Correction	Operating and shipping: Workbench launched in July 2026 and the blog ran through October 2026. Still no funding since the 2023 seed — but that is a funding observation, and our August note wrongly read it as inactivity.
Soxton	Correction	Active, and our August note was wrong to read a thin website as a thin company. Soxton acquired Cipher Technologies in March 2026, its first acquisition, and reports tripling its client base since the seed. Coverage places it in New York rather than the UK; the funding row has been corrected.

Appendix E — Method, sources and corrections

How this edition was built

The funding survey was updated by searching for every disclosed round, acquisition and status change between 1 September and 8 November 2026, in two independent passes — one for funding, one for exits and failures — each required to cite a URL per item and to flag where sources disagreed rather than average them.

The product catalogue was **re-read in full**. All 97 companies with a website were read site by site on 8 November, across six parallel passes, each given the rows we already held and asked to report only what had changed. That produced 64 additions, 25 withdrawals and 23 renames.

The evidence audit extended August's 95 rows with ten new companies read on 8 November against the identical rubric.

Known gaps, ordered by how much this report needed them

Gap	Why it matters
Arizona ABS licence count and holders	Underpins Chapter 4; we state no count
Utah sandbox status, SRA authorisations	The routes several companies operate under
A product-level cost-per-task table	Now exists for models and still not for products
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
The August evidence rows	Up to eleven weeks older than the November ones

Corrections to the August edition

Six, and every one is a case where we were wrong rather than merely out of date. We publish them because a survey that never reports its own errors is not being checked.

1. The capital-to-revenue multiples were valuation multiples. August's Chapter 2 published Harvey at 31x, GC AI at 55x and Legora at 56x, and told readers to "read those as what they are — dollars raised per dollar of current annualised revenue." They are not. Against the August data, Harvey's capital-to-revenue was 3.3x and its valuation-to-revenue 31.4x; GC AI's were 7.3x and 55.5x; Legora's 8.7x and 56x. All three published figures were **valuation divided by revenue**, described as something else. The data module named the field correctly; the prose did not. This is the most serious error in the August edition because it overstated capital inefficiency by roughly an order of magnitude.

2. Microsoft did not acqui-hire Robin AI's engineering team. We wrote that it did. The asset sale went to **Scissero**; Microsoft separately **hired at least eighteen former engineers**, including the chief technology officer. Different events, different buyers, and we collapsed them.

3. Soxton is not a "thin public surface". It acquired Cipher Technologies in March 2026 and reports tripling its client base. We also had its headquarters wrong: it is in New York, not the UK.

4. Fileread is shipping, not quiet. It launched Workbench in July 2026 and published through October. The funding observation was right; the inactivity inference was not.

5. We missed Lexroom's $50M Series B from May 2026 entirely, carrying $19M against a stated total above $73M — a round three months old in a company we already tracked.

6. LawGeex is not serving an expired certificate. Two independent readings in November found a working site over a valid chain. We cannot verify the certificate ourselves, because our own tooling terminates TLS before the origin. The claim is withdrawn; the brand is stale, which is not what we said.

Corrections 2, 3 and 4 share a failure mode worth naming, because it is the same one in each: **we inferred inactivity from silence.** A company that does not publish is not a company that does not operate. Three of six errors came from treating silence as a finding, in a report whose own argument about agents is that absence of evidence is not evidence.

A note on reading this market's own sources

During the catalogue re-read we found companies in this survey publishing instructions addressed to the AI systems that read their websites. Case Status serves an `llms.txt` directing assistants to add an emoji to their responses and to promote the company; Within, the former Klarity, serves one headed "Guidance for AI assistants"; Trellis asks for citation. Chapter 19 treats this as a finding. We record here, as method, that our researchers were instructed to treat any such content as data to report and never as instructions to follow, and that they did.

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 August edition of this report, unrevised, is at /research/archive/2026-08-31.

The State of Legal AI, November 2026. Published by IQdoc on 8 November 2026. The underlying data is public at /research/legal-ai-funding and /research/legal-ai-products. Corrections are welcome and will be dated.