

Legal AI

The Moat, Not the Model

June 2026

AUTHORS:



Scott Watson

Vice President



Avi Gulati

Intern



About Activant

Activant is a research-led global investment firm that partners with high-growth companies. Since 2015, we have invested in category-defining businesses during their most critical phases of growth, partnering with founders who have won their initial battles and are ready for the next challenge.

Our approach pairs deep, proprietary research with patient, flexible capital and hands-on operational partnership. We work alongside founders and leadership teams to refine strategy, strengthen operations, and accelerate sustainable growth.

Activant Research is dedicated to uncovering the most exciting emerging technologies, sectors, and companies we believe will shape the future. Our research-driven perspective informs everything we do, helping us invest at meaningful inflection points and support founders in building enduring, category-leading businesses.

You can find out more about Activant and our research at <https://activantcapital.com/>.

On May 12, [Anthropic](#) launched Claude for Legal: 12 practice-area plugins and 20 MCP connectors that wire Claude directly into [iManage](#), [NetDocuments](#), [Ironclad](#), [DocuSign](#), [Westlaw](#), [Practical Law](#), [CoCounsel](#), and [Harvey](#) itself.¹ A separate in-house edition launched the same week in New York and California with [Okta](#) and [The L Suite](#), aimed straight at corporate general counsel.

Within 48 hours, two reactions dominated LinkedIn, Substacks, and group chats. Founders worried the wrapper layer was finished, that any company sitting between a frontier model and the end user was about to be eaten. Investors asked whether the legal AI thesis still held now that the lab had come downstream. Both assumed the same thing: when a frontier lab enters a vertical, the vertical companies die.

That assumption is incomplete. Not because the labs are weak, but because the value in legal AI has been accruing elsewhere all along, and Claude for Legal does little to change that.

Our view is straightforward. In legal AI, value accrues to whoever owns what a model cannot manufacture: a license, a regulated entity, a proprietary data asset, or institutional memory. The model layer is becoming infrastructure, while the software layer between the model and the lawyer is being squeezed from above and below. What follows traces where value accrued in 2025 and where we think it goes next.

What Anthropic Actually Shipped

Start with the product itself, because much of the early commentary raced ahead of the launch announcement.

We don't read Claude for Legal as a competitor to Harvey. To us it looks more like a connector layer. The 12 practice-area plugins (Commercial, Corporate, Privacy, Employment, Product, IP, Litigation, Regulatory, AI Governance, plus modules for law students, clinics, and a builder hub) sit on top of a Model Context Protocol that calls out to other systems.² The MCP connectors include Westlaw, CoCounsel, Harvey, [Midpage](#), [Trellis](#), [Free Law Project](#), iManage, NetDocuments, [Box](#), [Datasite](#), Ironclad, DocuSign, [Definely](#), [Consilio](#), [Everlaw](#), and [Relativity](#).

Read that list again. The architecture assumes Harvey, CoCounsel, Westlaw, and Relativity still sit inside the workflow. The cold-start interview asks the lawyer to define playbooks, escalation pathways, and drafting styles, and the practice profile then drives every task.³ None of that competes with a vertical application; it assumes one is there. One absence is telling: Westlaw, Practical Law, and CoCounsel are all part of the Thomson Reuters stack, but LexisNexis is nowhere in the set, consistent with RELX guarding its content more tightly than its rival.

[OpenAI](#) is running a different play: ChatGPT Enterprise and custom GPTs sold as the development substrate for firms that want to build their own tools.⁴ [Dentons](#) announced exactly that arrangement late last year. The mental model is closer to [AWS](#) than to Harvey.

Why do the labs stay horizontal? Two reasons, neither born of restraint. First, vertical operations typically require regulated entities, professional indemnity, and ethics walls that frontier labs are not equipped to run. *United States v. Heppner* in February tightened the standard on attorney-client privilege and AI use enough that, in our reading, the labs now sit in a defensive posture on confidentiality and privilege, not an offensive one on workflow capture.⁵ The second reason matters more.

Why The Labs Cannot Hold Pricing Power

The LegalBench 2026 snapshot from May 21 shows [Gemini](#) 3.1 Pro at 87.4%, GPT-5.5 at 86.5%, Claude Opus 4.6 (Thinking) at 85.3%.⁶ The frontier looks dominant until you scroll down. [DeepSeek](#) V3 sits at 80.8%, DeepSeek V4 Pro at 80.3%, [Mistral](#) Large 2512 at 79.1%, all open-weight. The frontier lead is 5 to 7 percentage points and shrinking.

Pricing pressure on legal-specific models should follow, though the timeline is uncertain. LegalBench parity is not production deployment. Self-hosting a frontier-class model inside an AmLaw firm demands GPU capex, MLOps talent, evaluation infrastructure, security accreditation, and a refresh cadence few firms want to own. Our base case: over the next 24 to 36 months the top 10 to 20 firms by innovation budget run open-weight models for routine work and keep frontier API access for higher-stakes tasks. Full self-hosting stays rare, but the direction of travel runs against premium legal-specific margins.

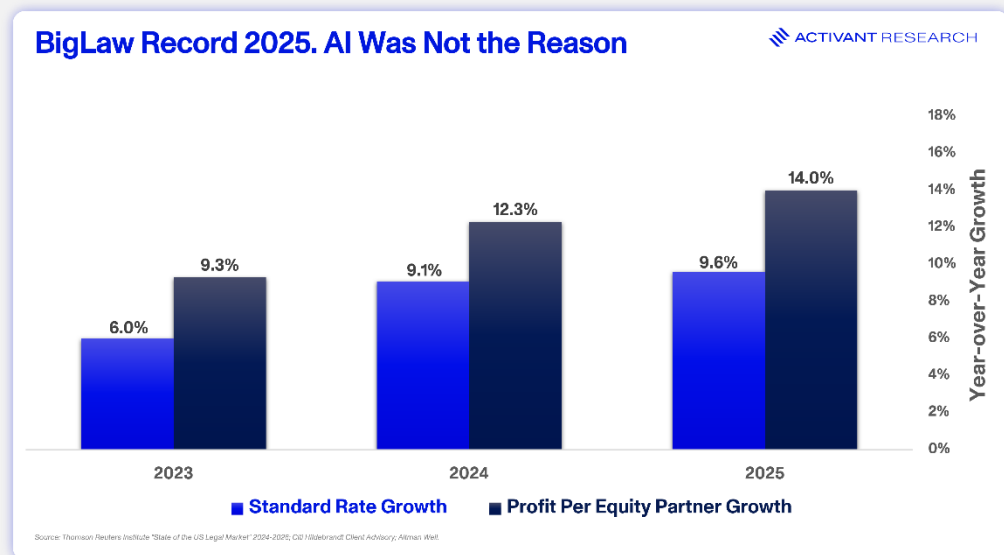
One tension is worth naming. As BigLaw builds more of its own tools, the apex of US law starts to look like a competitor to the AI-native firm category, not only a buyer. The two rarely meet: AI-native firms compete in consumer law, immigration, mass tort, SMB, and Fortune 500 startup work, not bet-the-company AmLaw mandates.

The labs are selling distribution and integration, not premium legal reasoning. Claude for Legal is priced inside the existing Claude tiers.⁷ We read that discipline as a feature, not an oversight: the model layer is becoming infrastructure, and infrastructure does not capture vertical economics.

Models commoditize, and infrastructure distributes, which means we have to look elsewhere to see who actually captures the spread.

Who Captured the Gains in 2025

The knee-jerk reaction is to assume BigLaw is banking the efficiency dividend. The numbers tell a different story. Average profits per equity partner across the AmLaw 100 grew 14% in 2025 to \$3.59M. The rate hikes are real, but they are built entirely on old-school leverage. The technology changed; the unit of account did not. Start with the chart everyone points to, then look at what it leaves out.



AmLaw 100 standard hourly rates grew 6.0% in 2023, 9.1% in 2024, and 9.6% in 2025, while worked rates rose a steady 7.3%. The AmLaw 50 pushed standard rates up 10.4% in 2025.⁸ Average AmLaw 100 partner rates crossed \$1,000 an hour last year, and senior partners at elite firms now quote north of \$2,500 for M&A and regulatory defense. That pricing power drops straight to the bottom line, lifting the upper tier to historic heights: Wachtell came in at \$12.15M, Kirkland crossed \$11M, and Davis Polk grew its profit per equity partner by 25.6%.⁹ Firms with a clear AI strategy were 3.9x more likely to report a return on investment than firms without one.¹⁰

Activant's Read: The productivity gain is real, but nobody has cleanly captured it yet. That is the whole opportunity.

This is the claim a skeptical reader should attack first. The rate and profit growth above is not, on the evidence, an AI story. [Thomson Reuters](#) (TRI) is explicit about the drivers: price-inelastic demand for bet-the-company work, private-equity-driven concentration at the top firms, and lateral talent arbitrage as elite partners move books down-market.¹¹ The theory that AI efficiency justifies premium pricing, TRI calls "more marketing hype than a legitimate strategy," an

idea "awaiting market validation." Firms are raising rates because they can, not because a model made an associate faster.

The productivity gain is real, but the billing model prevents anyone from capturing it cleanly. A full 90% of legal dollars still flow through hourly arrangements, the structure that has dominated since the 1950s.¹² A firm that turns 10 hours of work into 1 must either bill the 1 hour worked and forgo the other 9 or bill the old 10 and explain why. TRI calls the result a standoff: clients want the efficiencies passed through, firms want to bank them, and both keep converting back to an hourly rate that satisfies neither.

The technology has changed, but the unit of account has not.

The tell is in the demand data. In 2025 clients paid for fewer hours of the average legal service than in 2024, and the strongest demand growth went not to the apex firms but to smaller, cheaper ones as work migrated downstream.¹³ Buyers are not waiting for the AmLaw 100 to reprice; they route work to whoever delivers acceptable quality at a lower number, and AI is what makes that viable. But "good enough" only wins where the stakes are not bet-the-company. High-stakes work is price-inelastic: legal fees run 2 to 4% of revenue, so even a 50% increase barely registers, and the client buys the best lawyer, not the cheapest. The wedge bites the commoditizable middle, not the apex, and it points away from the incumbents.

The apex of BigLaw is building rather than buying: Cleary Gottlieb acquired Springbok AI, Cooley built Vanilla internally, [A&O Shearman](#) runs a revenue-sharing agentic subscription,¹⁴ and Kirkland & Ellis is committing \$500M to [build its own AI technology in-house](#). The vendor gets a flat license; the firm keeps the finished work's value. The data oligopoly monetizes AI more cleanly than anyone: Thomson Reuters Legal Professionals, excluding its government segment, grew 11% organically in Q1 2026, driven by an accelerating Lexis+ AI and Westlaw upgrade cycle.¹⁵ [RELX](#) Legal saw a step-up in the first half of 2025, with Lexis+ AI driving double-digit spend uplifts across law firms and corporate legal departments.¹⁶ Bernstein frames it well: the US legal research market is \$4-5B, but the broader legal technology and workflow market is \$25B, and the incumbents are using AI to capture the expansion above their old core, not to defend it.¹⁷

One caution belongs here. The same TRI report warns that the demand surge is built on instability, not health: corporate buyer sentiment is sliding toward pandemic-era lows, and TRI's own forecast turns negative by the second quarter of 2026.¹⁸ A market that surges right before it corrects is a familiar shape. Nothing that follows assumes the boom lasts; a downturn only accelerates the migration to lower-cost delivery, sharpening the argument.

Renting the Moat

Zoom out on 2025. The model layer commoditized and the data oligopoly compounded with BigLaw partner P&Ls expanded. That leaves one participant we have not yet placed: the horizontal application layer between the model and the end user.

This is the layer under the most structural pressure. The frontier labs are moving down with practice-area products and direct enterprise distribution. The citation-graph incumbents (TRI & RELX) are moving up with AI assistants priced into base subscriptions. Open-weight performance is closing the gap underneath. A vendor in the middle that lacks a proprietary data asset, exclusive content license, or workflow lock-in inside a system of record has little to push back with. "Thin-prompt wrapper" is a structural position, not a slur: the vendor rents moats from the layers above and below. Renting can be a large business, but it caps pricing power and makes the customer relationship contingent on whoever owns the moat being leased.

The category-level evidence supports this. Industry surveys suggest the typical AmLaw 100 firm now juggles 18 live AI tools, with 6 more in pilot and 22 in consideration.¹⁹ That is not loyalty; that is a buyer running a tournament. Some vendors will win consolidation rounds over the next 24 months. Many will not, and the breadth of evaluation signals a category that has not settled on its winners.

A more specific data-layer dynamic played out last year. The most ambitious attempt to acquire an independent citation graph collapsed when Harvey's reported \$600M bid for [vLex](#) failed and [Clio](#) bought vLex for \$1B in November.²⁰ The subsequent Harvey-[LexisNexis](#) alliance, announced in June 2025, gives Harvey retrieval access to Lexis content but prohibits white-labeling or training on it.²¹ That is evidence about the layer, not a verdict on any company: independent wrappers have fewer paths to owning the data they need than was assumed a year ago. The remaining graphs are not for sale at the application layer, and the available terms reflect that leverage.

Our working theory: Vendors that rent moats compress. Vendors that own moats compound.

We split the wrapper layer into two distinct buckets. Renting a moat caps pricing power, while owning one compounds it. In legal AI, we believe defensibility relies on five assets a model cannot manufacture:

1. **The Regulated Entity:** You either are a law firm or you aren't. Operating as an Alternative Business Structure (ABS) or a PLLC fundamentally changes the margin profile.

2. **The Privilege Wrapper:** Following the Heppner ruling, confidentiality demands a defensive posture. If your tool sits outside the attorney-client ring-fence, it could be a liability.
3. **Proprietary Citation Graphs:** You need exclusive content. An LLM can reason, but it cannot illegally scrape a closed citation graph.
4. **Institutional Memory:** This is harder to build but deeply sticky, requiring deep workflow lock-in inside a firm's system of record that aggregates per-client intelligence over years of transactions.
5. **Professional Indemnity:** A malpractice policy is an asset frontier labs are simply not equipped to run.

These operators look more like moat-owners than rent-payers. [DraftWise](#) doesn't just draft; it mines institutional memory directly from a firm's DMS. [Orbital](#) bypasses generalist models entirely with proprietary land registry integrations, while [Definely](#) embeds itself directly into Magic Circle drafting. They entrench themselves in the core workflows where switching costs harden into memory.

The same label also covers a shallower kind: a tool specialized to a practice area, running on the same public data and general reasoning as the model beneath it. That lock-in is real today, but unlike the five durable moats above, it is one the model can build for itself; labs just haven't bothered yet.

A firm that has trained its associates on one platform, moved its templates and matter files in, and shaped its process around it does not switch for free. But the cost is only friction: familiarity and a migration weekend, not anything the customer cannot rebuild elsewhere. And it falls every time the model improves at the underlying work.

What protects these companies, then, is not the strength of the lock-in but the indifference of the labs. The frontier labs are focused on extending horizontally, not vertically, relying on an ecosystem of specialized applications to sit on top. When OpenAI launched Frontier in February, Fidji Simo framed it plainly: the company is "not going to build everything ourselves."²²

Claude for Legal is built the same way: a connector layer that assumes the vertical applications already exist. Serving any vertical deeply would require domain-specific operations, regulatory infrastructure, and a sales motion the labs have neither built nor any reason to build. A narrow, regulated practice area is not where that expansion is headed.

That protection is uneven, scaling inversely with the size of the prize. Specialized patent prosecution sits well below the threshold that justifies a frontier lab's attention. Contract review sits at the other end: the single largest category in the \$25B legal workflow market the incumbents are already racing to capture, horizontal enough to touch every enterprise and close enough to what the labs are building to lie in their path. Defensibility here is less about the tool's quality than about how little the larger players want the territory.

None of this makes the category uninvestable. It makes it transitional. Every company here sits between a real moat and a casualty. The substrate underneath is improving autonomously, on someone else's budget, so these companies cannot stand still. The resolution takes one of three forms: convert the workflow lock-in into a real moat before the window closes, get absorbed into a larger platform as a feature, or get compressed when the work commoditizes and switching costs prove too shallow. The exits reflect the moat on offer. When Robin AI collapsed in early 2026, Microsoft acqui-hired several key engineers rather than the company itself. Managed-services clients migrated to players like Scissero rather than fully scattering, highlighting the lack of a proprietary moat.²³ That is what compression looks like when the moat is only friction.

The one path that ends well is conversion: building a durable moat before the window closes. In legal, that runs almost entirely through institutional memory rather than proprietary data, and it is the harder, slower route. The tool must embed deeply enough into a client's system of record to accumulate per-client intelligence the customer cannot easily rebuild. Where several parties coordinate in one shared workspace, as in a live transaction run through [Legatics](#) or [Ontra](#), the lock-in hardens into something closer to memory than friction.

Three Questions for Legal AI Founders

If we sat down with a founder tomorrow, here are the questions worth answering. Not pitch-deck answers, the real ones.

1. What do you own that the model cannot replicate?

If the answer is a thin layer of prompts, a fine-tuned model, or workflow templates, you may be renting. If it is a regulated entity, a proprietary data set the customer cannot get elsewhere, an exclusive content license, or deep workflow lock-in inside a system of record, the moat looks more durable. [Crosby](#) owns a PLLC. [Manifest OS](#) owns an Arizona ABS structure across multiple practice areas. [Garfield](#) is an SRA-authorized UK law firm that happens to be automated.²⁴ These are licensed legal entities with software inside, and that distinction matters.

2. Who absorbs your productivity gain?

This is harder than it sounds. If the customer is an AmLaw 100 partner, the hourly model banks the gain at the firm or competes it away on rate, and little reaches the vendor. If the customer is a litigation funder, an insurer, or a regulated firm pricing outcomes rather than hours, the vendor has a better chance of keeping

the spread. [EvenUp](#) and [Supio](#) sell into the plaintiff bar, where buyers pay for outcomes. [Darrow](#) sells case intelligence to funders and insurers. [Lawhive](#) bills its own consumer clients. The pattern matters more than the product.

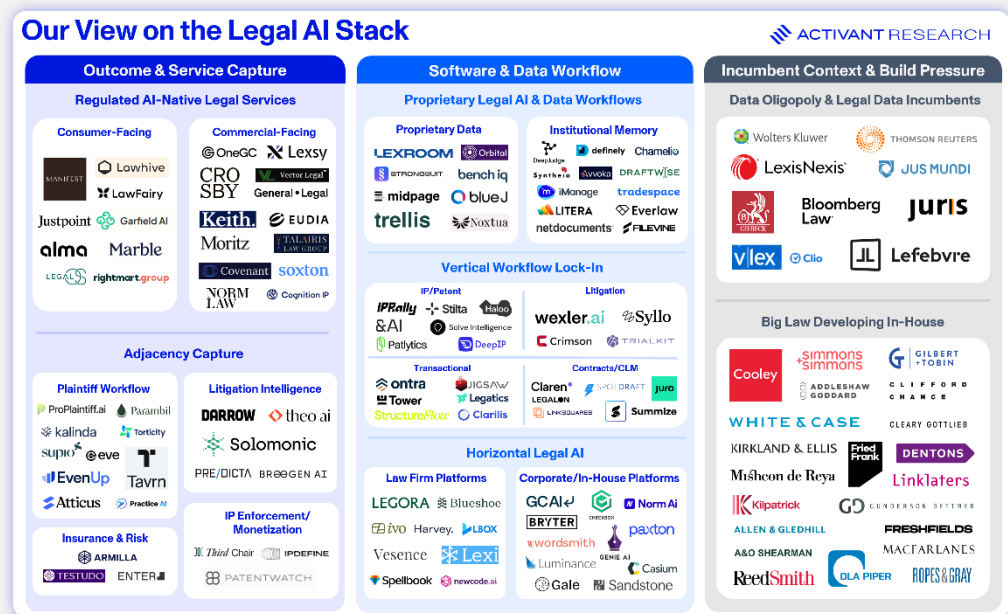
3. What happens to your gross margin when an open-weight model on the customer's infrastructure approaches frontier on LegalBench?

We expect this inside 24 months, though the timing is uncertain. If the value proposition is reasoning quality on commodity tasks, the margin compresses fast. If it is verified output bound to authoritative data, regulated process, or integrations the customer cannot easily replicate, the margin holds. The Heppner privilege ruling makes the second category more valuable, not less.

There is no single right answer. But the founders who engage with these substantively are the ones we want to spend time with.

Where to Pay Attention

We will not publish a target list. But here is where we are focusing our attention.



Regulated AI-native law firms

This category is already attracting significant investment. [Crosby Legal](#) at reported \$400M post Series B. [Manifest OS](#) at \$750M post Series A, the largest A round in legal tech history.²⁵ Lawhive at the UK consumer end. [Norm Law](#) inside financial services. The unit economics that killed Atrium in 2020 are

different in 2026: the AI is agentic rather than assistive, the ABS regulatory cover is real, and outcome-based pricing decouples revenue from hours. Whether the category produces durable equity is the most interesting open question in the space, and the one we are spending the most time on.

Vertical AI with proprietary data or workflow lock-in

[Lexroom](#) owns the civil-law content moat in continental Europe. DraftWise, Orbital, and Definely, covered above, sit on the same moat-owning side of the bifurcation. Their data assets cannot be commoditized by a frontier model running on the customer's own GPU.

Adjacency capture

Case intelligence sold into litigation finance and insurance underwriting. Darrow has built a real business at the intersection of plaintiff law, third-party funding, and insurer subrogation. [Solomonic](#) in the UK shows what capital efficiency looks like when the customer is a funder, not a law firm. [Armillia](#) AI sits elsewhere: the affirmative AI-liability layer that [Verisk](#)'s January exclusions are about to force into existence.²⁶ Most of this layer is sub-scale today. That is precisely why we are interested.

We are more cautious about the thin-prompt end of the wrapper layer at current valuations, and that view could be wrong, especially for the best-capitalized, best-distributed companies there. The teams that use capital and relationships to migrate from rent-paying to moat-owning have the most credible path to defending today's prices. We are watching for that migration and expect some to surprise us.

Final Thoughts

We don't read Anthropic's move into legal as the end of anything. It looks more like confirmation that the model layer is becoming infrastructure, which is unlikely to extract vertical economics on its own. That is good news for anyone owning real moats and harder news for anyone renting them.

The companies that compound through the next phase of legal AI may look less like software companies and more like regulated firms with software inside. Data platforms with editorial discipline. Insurance underwriters with prediction engines. The boundary between software and services is dissolving, and the equity accrues to whoever owns the regulated, authoritative, or proprietary asset on the other side. These moats are not permanent. But they are more durable than anything at the model or prompt layer, and that relative durability is what we underwrite.

If you are building one of these businesses, we would like to meet.

Endnotes

1. ABA Journal, [Anthropic Launches Claude for Legal, Giving Lawyers 20 New Program Integrations and 12 Practice-Area Plugins](#), 2026.
2. Anthropic, [Claude for Legal architecture documentation](#), 2026.
3. JD Supra and Artificial Lawyer, [Claude for Legal launch coverage](#), 2026.
4. Dentons / OpenAI, [Dentons-OpenAI partnership and ChatGPT Enterprise legal deployments](#), 2025-2026.
5. Harvard Law Review, [United States v. Heppner, S.D.N.Y.](#), 2026.
6. Vals AI, [LegalBench 2026 leaderboard](#), 2026.
7. Anthropic, [Claude pricing page](#), 2026.
8. Thomson Reuters Institute, [2026 Report on the State of the US Legal Market](#), 2026.
9. Above the Law and American Lawyer, AmLaw 100 financial rankings 2025, 2025.
10. Thomson Reuters Institute, [2026 Report on the State of the US Legal Market](#), 2026.
11. Thomson Reuters, [2026 Report on the State of the US Legal Market](#), 2026.
12. Thomson Reuters, [2026 Report on the State of the US Legal Market](#), 2026.
13. Global Legal Post, [US firms see revenue grow in 2025 driven by rising billing rates and demand – study](#), 2026.
14. National Law Review, Cleary Gottlieb's Springbok acquisition, Cooley's Vanilla deployment, and A&O Shearman's agentic AI subscription model, 2026.
15. Investing.com, [Earnings call transcript: Thomson Reuters Q1 2026 shows 8% revenue growth](#), 2026.
16. RELX, [RELX PLC H1 2025 results](#), 2025.
17. Bernstein Research, RELX and the AI-driven TAM expansion in Legal and STM divisions, 2026.
18. Thomson Reuters, [2026 Report on the State of the US Legal Market](#), 2026.
19. Fringe Legal / SKILLS, [2025 Generative AI Use Cases Survey](#), 2026.
20. Clio, [Clio Completes Landmark \\$1B vLex Acquisition and Announces \\$500M Series G Funding Round at \\$5B Valuation](#), 2025.
21. Harvey, [LexisNexis and Harvey Announce Strategic Alliance](#), 2025.
22. CNBC, [OpenAI launches Frontier for enterprise customers](#), 2026.

23. Bloomberg Law, [Microsoft Brings Former Robin AI Legal Tech Employees Into Fold](#), 2025.
24. Artificial Lawyer, [Crosby Series B, Manifest OS Series A, and Garfield AI SRA authorization coverage](#), 2026.
25. Lux, Index, Sequoia, BCV, Menlo, Kleiner Perkins, First Round, Quiet Capital, Crosby and Manifest OS round financials, 2026.
26. Independent Agent, [Verisk to Roll Out New General Liability Exclusions for Generative AI Exposures](#), 2025.

Disclaimer: The information contained herein is provided for informational purposes only and should not be construed as investment advice. The opinions, views, forecasts, performance, estimates, etc. expressed herein are subject to change without notice. Certain statements contained herein reflect the subjective views and opinions of Activant. Past performance is not indicative of future results. No representation is made that any investment will or is likely to achieve its objectives. All investments involve risk and may result in loss. This newsletter does not constitute an offer to sell or a solicitation of an offer to buy any security. Activant does not provide tax or legal advice and you are encouraged to seek the advice of a tax or legal professional regarding your individual circumstances.

This content may not under any circumstances be relied upon when making a decision to invest in any fund or investment, including those managed by Activant. Certain information contained in here has been obtained from third-party sources, including from portfolio companies of funds managed by Activant. While taken from sources believed to be reliable, Activant has not independently verified such information and makes no representations about the current or enduring accuracy of the information or its appropriateness for a given situation.

Activant does not solicit or make its services available to the public. The content provided herein may include information regarding past and/or present portfolio companies or investments managed by Activant, its affiliates and/or personnel. References to specific companies are for illustrative purposes only and do not necessarily reflect Activant investments. It should not be assumed that investments made in the future will have similar characteristics. Please see "full list of investments" at activantcapital.com/companies/ for a full list of investments. Any portfolio companies discussed herein should not be assumed to have been profitable. Certain information herein constitutes "forward-looking statements." All forward-looking statements represent only the intent and belief of Activant as of the date such statements were made. None of Activant or any of its affiliates (i) assumes any responsibility for the accuracy and completeness of any forward-looking statements or (ii) undertakes any obligation to disseminate any updates or revisions to any forward-looking statement contained herein to reflect any change in their expectation with regard thereto or any change in events, conditions or circumstances on which any such statement is based. Due to various risks and uncertainties, actual events or results may differ materially from those reflected or contemplated in such forward-looking statements.