LEONIS CAPITAL

RESEARCH

Leonis AI 100

Breakout growth and how AI companies are becoming an economy of their own.

AUTHORS

Jenny Xiao
Jay Zhao
Liang Wu

DATE

November 2026

leoniscap.com

Executive Summary

In last year's AI 100, we documented how AI is creating a new startup playbook. This year, that playbook has scaled into something larger: AI companies are becoming an economy of their own. The breakout companies from the first few waves of generative AI have grown into substantial customers for infrastructure, data, and specialized services, while their products are reaching further into established industries and verticals.

35

companies at $100m or more, up from 14. Insight 1.

This year, **35 companies report at least $100 million in revenue or annualized run rate, up from 14 last year**, and **23 have crossed $200 million**, up from three.[1] What looked like exceptional growth at Cursor and ElevenLabs is becoming a defining characteristic of the leading AI companies. Last year's leaders also kept growing. The 38 companies in both editions grew their disclosed aggregate revenue from $2.5 billion to $15.6 billion and account for 84% of this year's disclosed revenue. Financing has advanced alongside commercial growth, with 42 companies now at Series C or later, up from 22 last year, and 54 valued at $1 billion or more, up from 23.

57%

of suppliers to AI builders past $100m, against 29% of applications. Insight 2.

Scale creates its own demand as AI companies have become substantial customers in their own right, supporting a supplier ecosystem around running AI products and improving their capabilities. Operating demand comes from buying inference and execution capacity to run AI products as customer usage grows. Capability-building demand comes from spending on training data and expert labor to improve the underlying models. Suppliers grow with the businesses they serve, as 13 of the 23 (57%) companies selling to AI builders have passed $100 million, against 29% of the AI applications on the list.

40

vertical companies, up from 28. Insight 3.

The emerging AI economy is also expanding beyond its own industry. The AI 100 now includes **40 vertical companies, up from 28 last year**, as products built for healthcare, law, finance, and other industries gain traction. As products expand from individual tasks into broader workflows, these companies can earn a larger role in a customer's operations and access more of the budgets supporting that work. New York's representation grew from 12 companies to 19, alongside the growing importance of industry expertise, customer proximity, and

1 Figures are each company's latest dated, non-forecast figure as of October 1, 2026, and none is audited. How we built the list, and where the data is weaker, is in About This Report at the end.

experienced operators, though the SF Bay Area still holds 69% of the capital the list has raised.

64

companies with some usage-based pricing. Insight 4.

As AI companies move from selling software access to performing work, **spending increasingly scales with the amount of work customers delegate to AI**. In this year's AI 100, 64 companies incorporate some form of usage-based pricing. Customers spend more as AI handles more work, without adding headcount. Usage pricing becomes **a transmission mechanism for the AI economy**, connecting greater application usage to revenue for the suppliers providing inference, execution capacity, and other services. Among companies that report a revenue figure, the ones that incorporate usage-based pricing have a median of $100 million, against $35 million for the software companies that don't.

Together, these developments create a feedback loop that allows growth to compound across AI applications and their suppliers. How much value individual businesses retain will depend on their role in customer workflows and their pricing power as the underlying technology becomes cheaper.

The Rise Of The New AI Economy

Last year, the Leonis AI 100 documented the emergence of a different startup playbook. Small, technical teams were reaching revenue milestones at extraordinary speed, finding customers through product-led adoption, and adapting quickly as the underlying foundation models improved. This year, those businesses have grown considerably, and a greater proportion of companies in the AI 100 are reaching substantial revenue milestones at a similarly rapid pace. **The scale once represented by a few exceptional companies is becoming more common among the leaders in the AI 100 cohort.**

As these companies grow, **we are seeing a new AI economy take shape**. **AI-native businesses have become significant buyers** of the infrastructure, data, and services they need to operate, supporting suppliers whose first customers **are other AI-native companies**. In some categories, AI companies themselves are now large enough to support billion-dollar suppliers.

At the same time, their products are **penetrating further into established industries and verticals**. AI applications in accounting, insurance, healthcare, and law are taking on broader workflows, bringing more industry budgets into the AI economy. This develops one of last year's central findings that AI markets tend to unlock in sequence as model capabilities improve.[2] Many of the vertical companies gaining traction today were founded years earlier, giving their teams time to build the industry expertise and customer relationships needed to put AI to work in established workflows.

The economics are changing on the customer side as well. Last year, we observed that AI companies could grow revenue much faster than headcount. Usage-based pricing allows customers to increase spending as AI performs more work, even without adding employees. The result is an economy increasingly organized around the **volume of work AI performs, rather than the number of people building and using software**. Usage-based pricing is the mechanism that connects customer spending to growth across the AI supply chain.

Our second annual AI 100 examines this transition from a new startup playbook to a broader economic system. The following four sections explore how exceptional growth is becoming more common, how scale creates demand within AI, how vertical adoption broadens that demand,

2 Leonis Capital, The Leonis AI 100, 2025: markets break out in sequence.

and how pricing connects the work AI performs to spending throughout this economy.

INSIGHT 1

Breakout Growth Is Becoming More Common

IN SHORT

Revenue milestones that set last year's leaders apart are now common on the list, and last year's leaders kept compounding rather than rotating out.

Last year, breakout companies like Cursor and ElevenLabs showed how quickly AI companies could turn technical capabilities into substantial businesses. Their products performed work customers were already willing to pay for, helping them reach $100 million in annual recurring revenue on timelines that would have been extraordinary in the SaaS era. As more companies repeat that pace of growth in this year's AI 100, the speed that once made Cursor and ElevenLabs **look like anomalies is becoming a defining characteristic of the leading AI companies.**

In this year's AI 100 cohort, **35 companies report at least $100 million** in revenue or annualized run rate**, up from 14 last year**, and **23 report at least $200 million**, compared with just three in the 2025 cohort (Exhibit 1).[3] Revenue milestones that distinguished last year's leaders **are beginning to look more like table stakes for the top AI companies.**

The early leaders have also continued to grow at a rapid pace, representing a new tier of AI giants. Cursor generates more than $4 billion in annualized revenue,[4] Fireworks AI has passed $1 billion,[5] and ElevenLabs has passed $500 million.[6] Harvey has passed $400 million and Glean has reached $300 million.[7] Last year, Cursor and ElevenLabs were among the exceptions we pointed to when explaining AI's commercial potential. The same pace of growth now shows up in products that launched this year. Viktor, an AI coworker that works inside Slack and Microsoft Teams, went from launch to a $50 million annualized run rate in seven months, through self-serve sign-ups.[8][9]

Among AI 100 companies with disclosed revenue figures, the median rose from $16 million to $100 million (Exhibit 2).[10] The 25th, 75th, and 90th percentiles increased by roughly five to seven times as well. This

3 Reported figures include a mix of ARR, annualized run rates, and third-party revenue estimates.

4 Forbes, June 8, 2026, citing an unnamed source.

5 Fireworks AI, Series D announcement.

6 ElevenLabs, $500m ARR announcement.

7 Glean, $300m ARR announcement.

8 ARR Club, Viktor at $50m in seven months.

9 Leonis Capital is an investor in Viktor.

10 Median revenue is calculated from the 64 companies with available revenue figures in the 2026 AI 100 and 89 in 2025.

EXHIBIT 1



Companies on each list at or above each revenue threshold, of 100 per list.

SOURCE: LEONIS CAPITAL RESEARCH. LATEST DATED, NON-FORECAST FIGURE AS OF OCTOBER 1, 2026; RANGE-ONLY DISCLOSURES EXCLUDED.

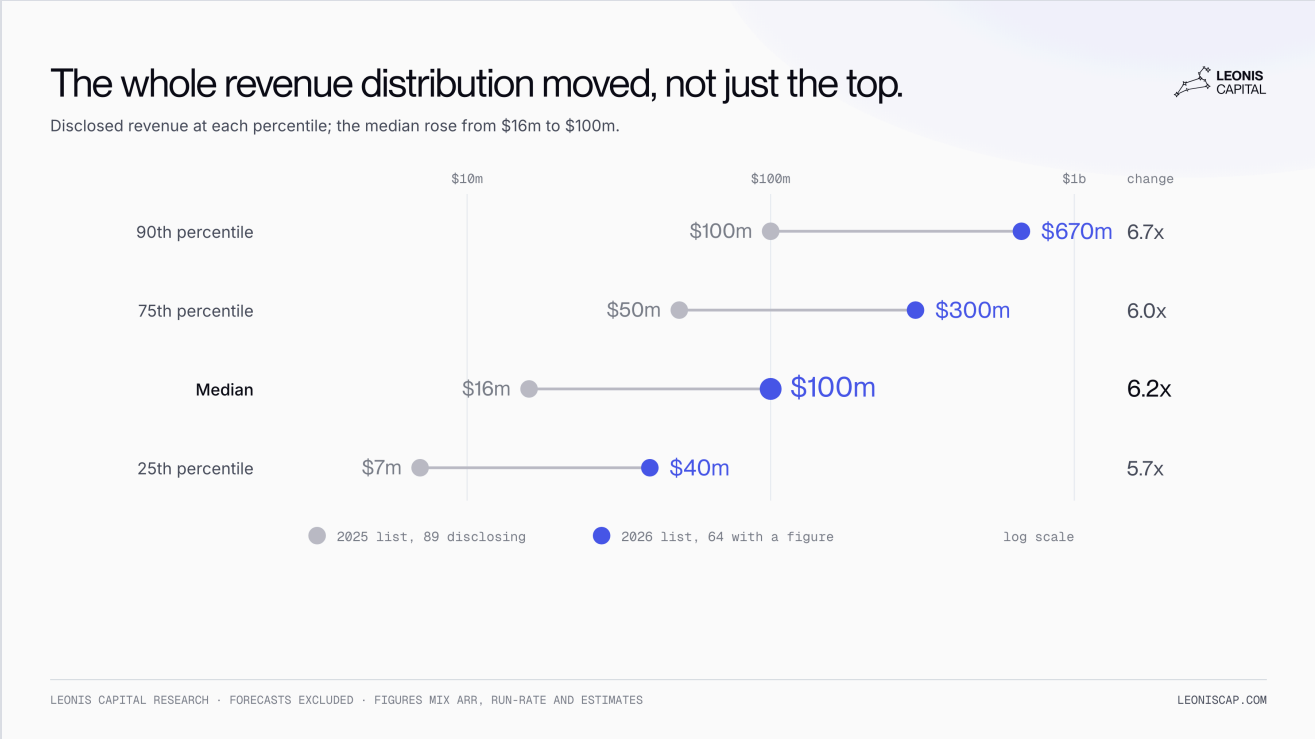
wasn't just driven by top companies, but everyone else in the index, **highlighting that the entire revenue distribution has moved upward.**[11]

Fewer companies disclosed a figure this year, 64 against 89, so the median partly reflects who reports. The count of large companies is a cleaner test, and it moved the same way: 35 companies are at $100 million or more, against 14 last year.

Last year's winners didn't rotate out. They kept compounding. The same 38 companies from the 2025 list generated 84% of this year's disclosed revenue (Exhibit 3), while 13 of the 14 companies that had already crossed $100 million last year remain on this year's list. That persistence suggests the extraordinary trajectories we observed in 2025 were not simply transient breakout moments. Once the leading AI companies reach scale, many continue compounding rapidly, even as a new generation of companies enters behind them.

11 While fewer companies disclosing revenue this year might overstate the magnitude of the shift, the direction of the revenue shift is clear.

EXHIBIT 2



Disclosed revenue at the 25th, 50th, 75th and 90th percentiles, 2025 list against 2026 list, log scale.

SOURCE: LEONIS CAPITAL RESEARCH. FORECASTS EXCLUDED; FIGURES MIX ARR, RUN-RATE AND ESTIMATES.

Among the 30 companies past $10 million that report revenue figures across both years, the median grew 5.3x over about a year and the fastest tenth grew 11x or more (Exhibit 4). Even that median runs ahead of outside benchmarks: The top decile of AI companies above $20 million in annualized revenue grew about 3.6x year over year at the end of 2025, according to Standard Metrics,[12] and Bessemer's Supernovas, its cohort of breakout AI companies, grew about 3.1x on average from their first year to their second.[13]

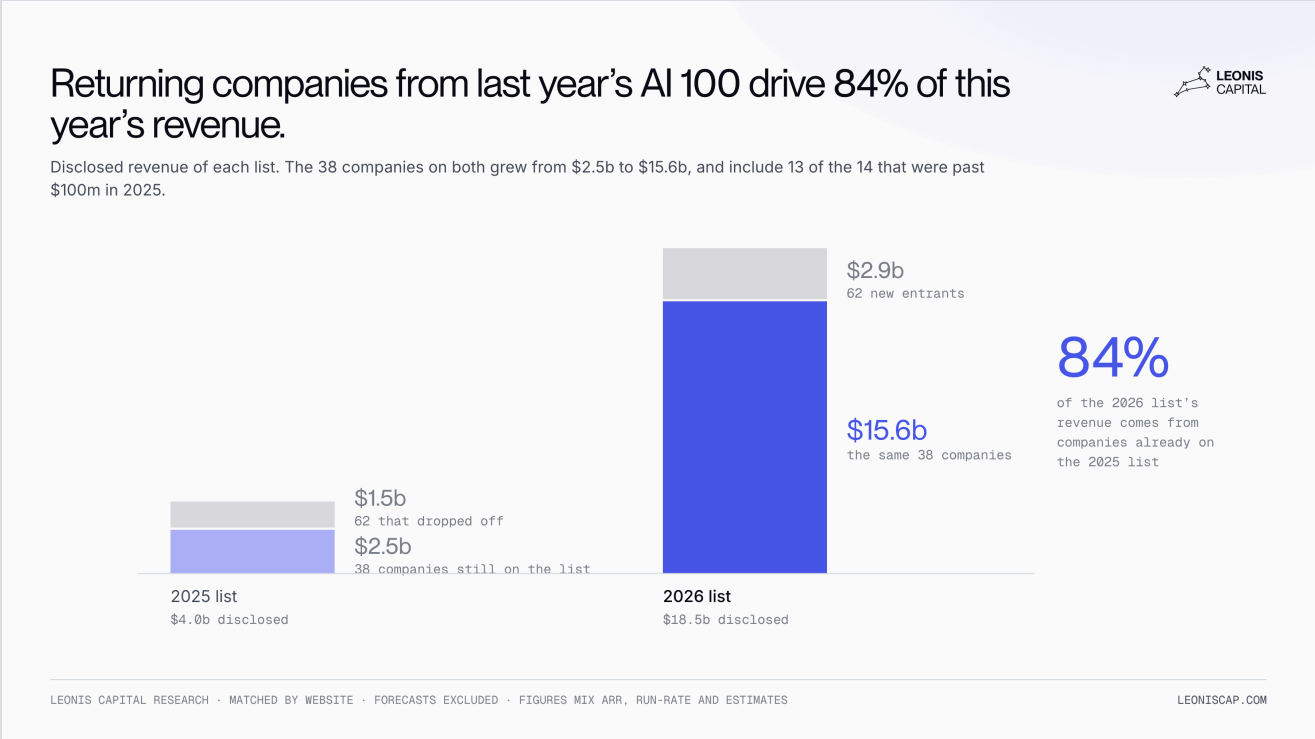
The revenue growth across the AI 100 is being matched by larger financing ambitions. 42 companies in this year's AI 100 have reached Series C or later, compared with 22 last year, while billion-dollar valuations have become more common (Exhibit 5).

Companies are also moving through funding rounds more quickly, attracting valuations that anticipate substantial further expansion. For example, Wonderful raised a $550 million Series C at a $5 billion valuation in September, six months after a $2 billion Series B and less

12 Standard Metrics, Private Market Report, 2025 year in review.

13 Bessemer Venture Partners, The State of AI 2025.

EXHIBIT 3



Disclosed revenue of each list, split into companies on both lists, companies that dropped off and new entrants.

SOURCE: LEONIS CAPITAL RESEARCH. COMPANIES MATCHED ACROSS LISTS BY WEBSITE; FORECASTS EXCLUDED.

than two years after its founding.[14] Prime Intellect, founded in 2024, became a unicorn in 2026.[15] Investors are underwriting the possibility that early commercial traction will develop into a much larger business within a few years. Faster revenue growth helps support those expectations, while competition to own a stake in potential leaders pushes valuations higher and brings investment decisions forward.

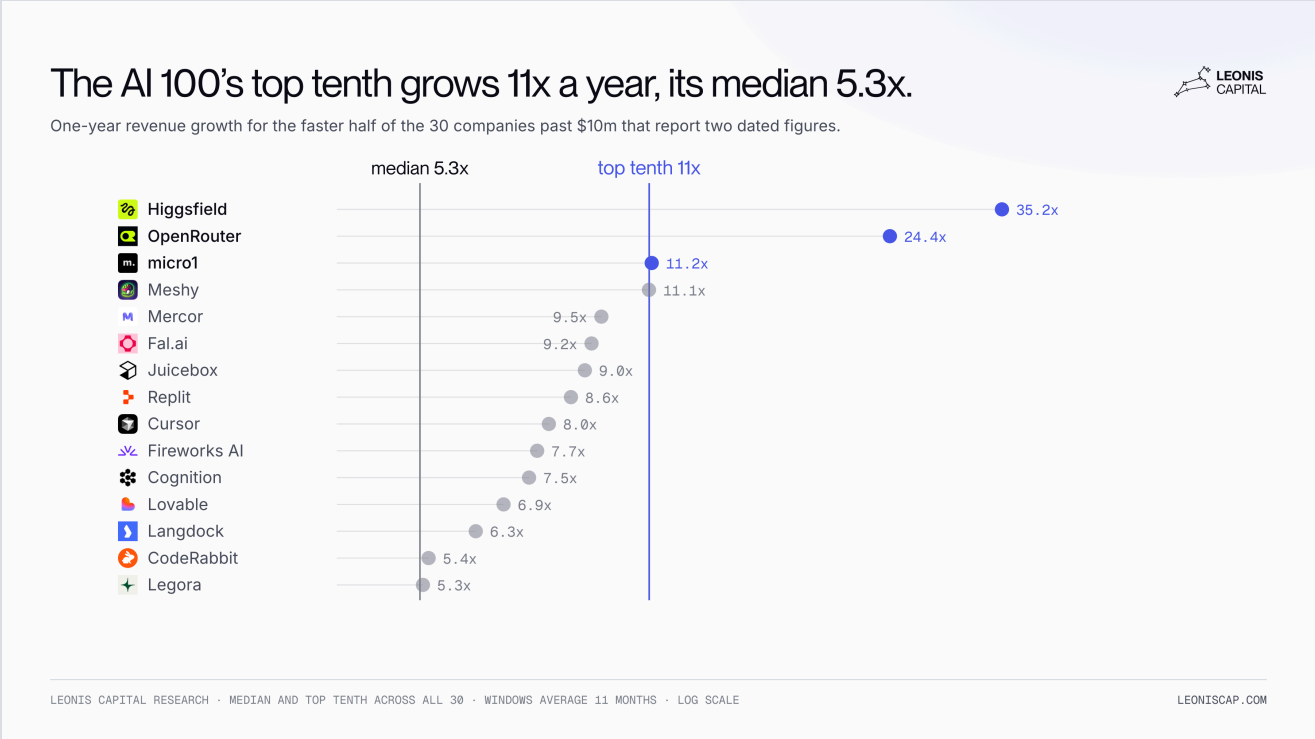
This creates a demanding relationship between financing and performance. A large round gives a company resources to pursue the opportunity, but its valuation also incorporates growth that it still needs to deliver. As exceptional revenue milestones become more common, investors can price those expectations into companies earlier. Founders may find themselves **managing the capital and expectations of a later-stage business while their organizations are still taking shape.**

While these revenue figures establish commercial scale, the economics behind them vary considerably. Annualized run rates, recurring software revenue, and gross marketplace revenue each signal different margins and durability. For example, a $500 million expert-data marketplace,

14 TechCrunch, September 2, 2026.

15 TechCrunch, July 8, 2026.

EXHIBIT 4



One-year revenue growth for the faster half of the 30 companies past $10 million that report two dated figures.

SOURCE: LEONIS CAPITAL RESEARCH. MEDIAN AND TOP TENTH COMPUTED ACROSS ALL 30; WINDOWS AVERAGE 11 MONTHS; LOG SCALE.

whose gross revenue includes payments to the experts supplying the work, is not economically equivalent to a $500 million software company generating recurring SaaS revenue.

The economic significance of this expansion goes beyond the revenue these companies collect. At hundreds of millions of dollars in revenue, AI companies have also become substantial buyers. Serving more customers and running more agents creates demand for inference, execution capacity, data, and other inputs. The businesses whose early growth we documented last year **are increasingly becoming the customer base for another generation of AI companies**. Their scale is helping create an economy around them.

EXHIBIT 5



Share of each list founded in 2023 or later and in 2024 or later, against the share at Series C or later and valued at $1 billion or more.

SOURCE: LEONIS CAPITAL RESEARCH. VALUATIONS ARE THE HIGHEST STATED IN EACH ROUND HISTORY.

INSIGHT 2

AI Companies Have Become a Market of Their Own

IN SHORT

AI companies now buy enough from each other to support billion-dollar suppliers, and the suppliers that help improve models grew fastest, which also puts them closest to changes in research budgets.

AI companies are no longer just a category of vendors. They are now a customer segment large enough to support entire companies. Their spending on infrastructure, data, and specialized services has created a market in which suppliers can scale by selling primarily to other AI companies. Widespread adoption across traditional industries remains an enormous opportunity, but AI startups can already find significant demand within the AI industry itself.

AI companies buy from these suppliers for two distinct purposes: running their existing products and developing better capabilities.[16] Operating demand covers inference, sandboxes, observability, and execution infrastructure required to run AI products. Capability-building demand covers the training data, expert labor, and evaluations used to improve them.

This year's AI 100 shows how quickly that demand can support sizable businesses. AfterQuery reached a reported $100 million annualized run rate in 14 months,[17] while Arena reached that level eight months after its enterprise launch.[18][19] AI suppliers are now scaling as quickly as the AI application companies they serve (Exhibit 6).

Suppliers serving operating demand grow as their customers use more inference, run more agents, and consume more execution capacity. Modal surpassed $300 million in annualized revenue, with more than a third of that revenue coming from agent sandboxes, the environments where agents run code and execute tasks.[20] Fireworks reports serving more than 40 trillion tokens a day and exceeding $1 billion in annualized revenue, supporting customers like Cursor and Harvey.[21] As those applications serve more customers and perform more work, their suppliers provide the capacity required to deliver it.

16 Several suppliers sell both. Fireworks, Modal, Baseten and Together AI also run fine-tuning and reinforcement-learning work for customers, and Prime Intellect also sells inference. We classify each company by the source of most of its revenue: inference and compute for the first four, training infrastructure for Prime Intellect.

17 Forbes, April 9, 2026.

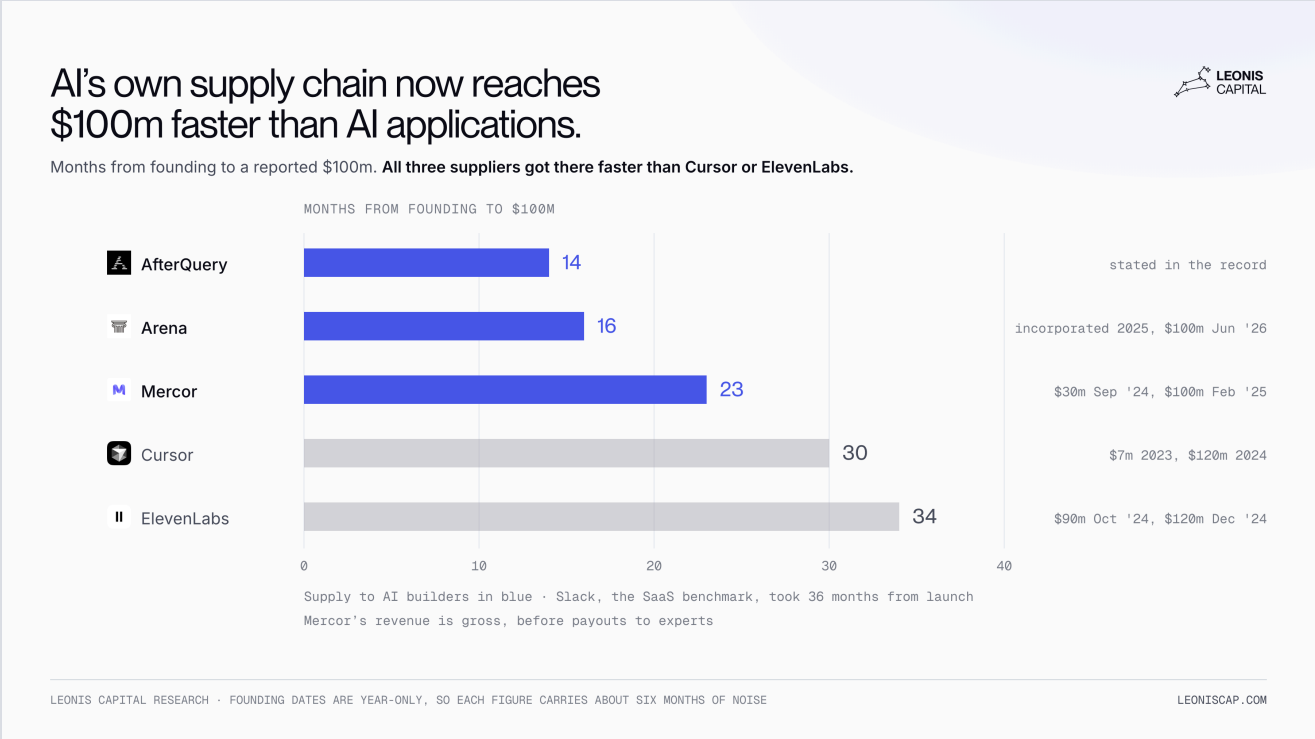
18 TechCrunch, June 29, 2026.

19 AfterQuery's timeline starts at inception; Arena's starts at enterprise launch. Arena's CEO explicitly describes its figure as annualized consumption rather than revenue under contract.

20 Modal, Series C announcement.

21 Fireworks AI, Series D announcement.

EXHIBIT 6



Months from founding to a reported $100 million for three suppliers to AI builders and two AI applications.

SOURCE: LEONIS CAPITAL RESEARCH. FOUNDING DATES ARE YEAR-ONLY; MERCOR'S REVENUE IS GROSS, BEFORE PAYOUTS TO EXPERTS.

Capability-building demand has supported substantial businesses as well. Expert marketplaces such as micro1 and Mercor supply training data used to improve AI models. Micro1 grew from $100 million to $500 million in gross annualized revenue in eight months,[22] while Mercor exceeded $2 billion in gross annualized revenue in June.[23] Their growth reflects the scale of spending on expert knowledge as AI developers compete to improve model capabilities. That spending draws on some of the largest budgets in technology. OpenAI and Anthropic's combined annualized revenue reached about $105 billion by late summer, up from about $30 billion at the start of the year.[24] OpenAI has committed to spend ~$50 billion on compute in 2026, roughly three times its 2025 budget, and in 2025 about half of its compute went to research and training rather than to serving customers.[25]

The distinction matters because these sources of demand have different economics and durability (Exhibit 7). Operating spending is closer to the

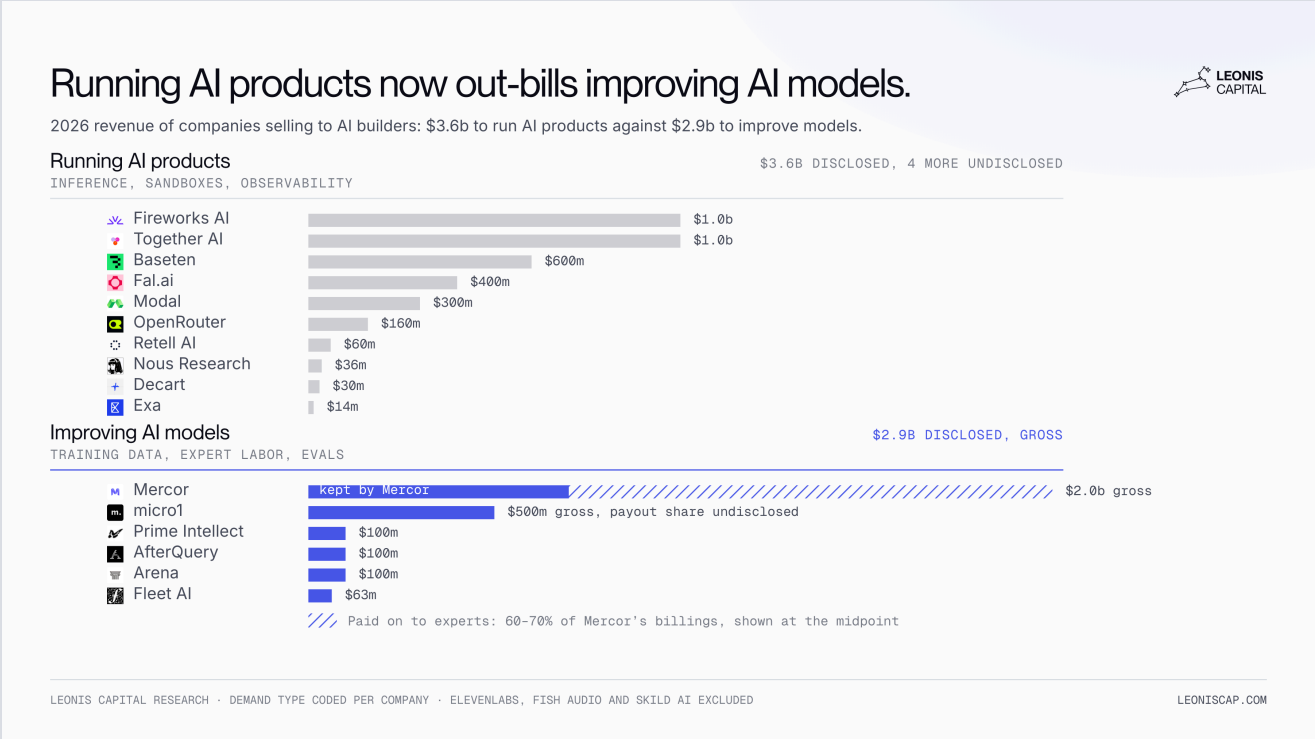
22 TechCrunch, August 20, 2026.

23 The Information, Mercor at $2 billion gross annualized revenue.

24 State of AI Report, 2026.

25 Epoch AI, AI chip users explorer.

EXHIBIT 7



2026 revenue of the companies selling to AI builders, grouped by whether the buyer pays to run AI products or to improve models.

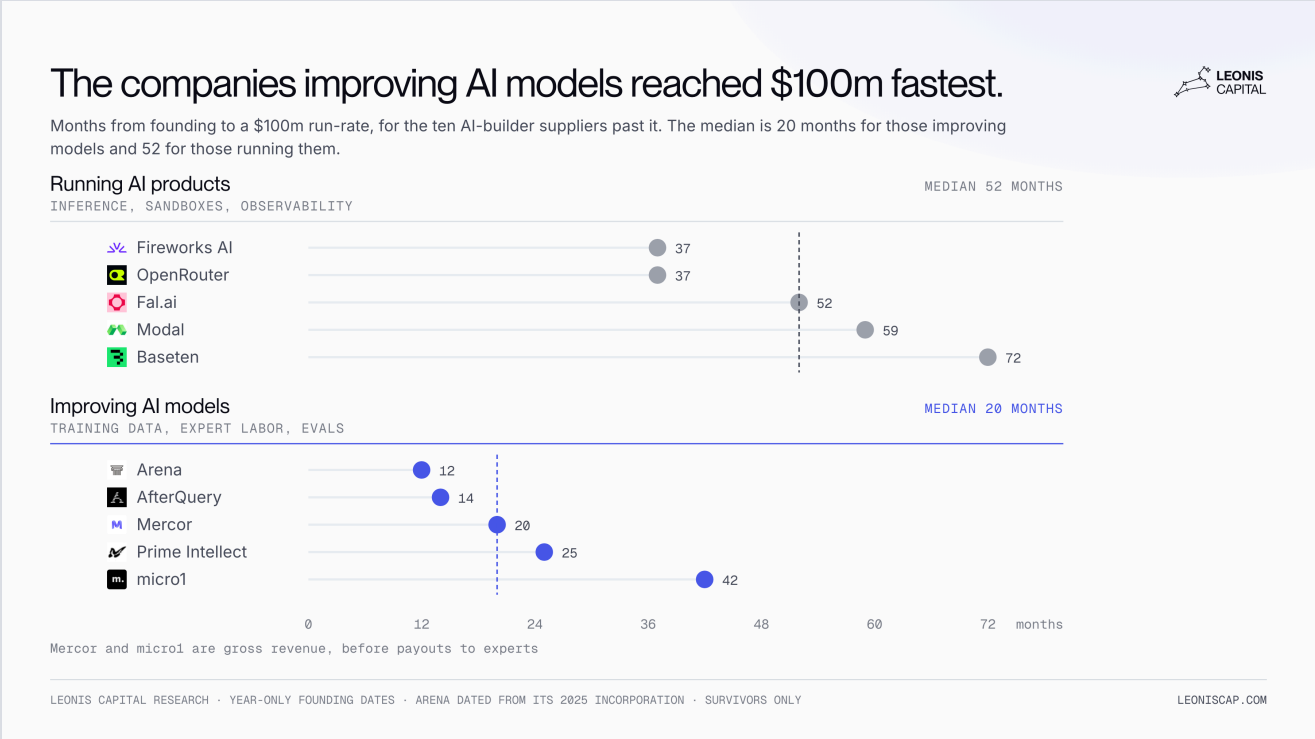
SOURCE: LEONIS CAPITAL RESEARCH. DEMAND TYPE CODED PER COMPANY; MERCOR AND MICRO1 AT GROSS REVENUE.

cost of delivering a product, recurring with customer usage across a range of AI applications. Capability-building spending is closer to research and development, with demand concentrated among research labs and neolabs. Those budgets can be enormous, but they depend more heavily on funding availability, changing technical approaches, and decisions about what capabilities to develop next.

Suppliers that help improve AI models have also reached $100 million faster than those that help run AI products (Exhibit 8). That puts some of the most rapid **supplier growth in the part of the market most exposed to changes in research spending and financing.** Mercor, the largest of these suppliers, earned about 91% of its $614 million in first-half 2026 gross revenue from foundation model developers, though it aims to cut that share to below a third by 2030.[26] A lab that cuts its budget or changes its training approach can affect several suppliers at once, **allowing a slowdown to travel through the same relationships that supported their growth**. Some of this revenue is also circular. Many of the AI companies buying from these suppliers are themselves spending venture capital, so part of what suppliers earn tracks the financing cycle

26 The Information, on Mercor's revenue concentration.

EXHIBIT 8



Months from founding to a $100 million run-rate for the ten AI-builder suppliers past it.

SOURCE: LEONIS CAPITAL RESEARCH. YEAR-ONLY FOUNDING DATES; SURVIVORS ONLY; MERCOR AND MICRO1 AT GROSS REVENUE.

as much as end-customer demand. That concentration may ease as capability spending spreads beyond the labs: Harvey trains with Fireworks, Decagon with Modal and Ramp with Prime Intellect, as inference providers bundle training around open-weight models. Application companies that train their own models also become buyers of training data and evaluations, not only compute.

The emerging AI supplier ecosystem (Exhibit 9) follows a familiar pattern from the SaaS era, when growing software companies supported markets for cloud infrastructure, developer tools, and security. AI has compressed that development into only a few years. The next stage depends in part on how far demand spreads into established industries. Spending by law firms, insurers, and healthcare providers can support both the applications serving them and the infrastructure those applications consume, broadening the customer base behind the AI economy.

EXHIBIT 9



Companies in each product category and the share past $100 million, counting every company in the category.

SOURCE: LEONIS CAPITAL RESEARCH. CATEGORIES COLLAPSED FROM 83 SOURCE LABELS; FORECASTS EXCLUDED.

INSIGHT 3

The Vertical Unlock Has Arrived

IN SHORT

Industry products are a much larger share of the list, built by companies that spent years earning trust in their markets, and that has pulled company formation toward New York.

Last year, we argued that AI markets would unlock in sequence as models became capable of handling increasingly demanding work.[27] Writing and coding tools found customers early, followed by multimodal products and applications tailored to industries such as healthcare, law, and finance. Not all verticals arrive simultaneously because reliability requirements differ. Each industry has its own capability threshold, determined by the cost of errors, regulatory requirements, workflow complexity, and tolerance for human review. This year's AI 100 shows that vertical expansion is becoming a larger part of the startup landscape. 40 companies serve specific vertical industries, up from 28 last year (Exhibit 10).

As models become capable of handling more industry-specific tasks, the commercial challenge **shifts toward building products that can manage the workflows around them**. In insurance, extracting information from a claim is one step in a process that also requires identifying missing information, applying regulatory rules, routing decisions to specialists, and recording what happened. When competitors can access similar models, understanding how those steps fit together becomes a source of differentiation.

The larger opportunity is to expand **from a task into a workflow and eventually become a system of action** (software that makes and executes decisions on behalf of an organization) through which customers carry out that work.[28] A product that extracts claims information could expand into claims review, routing, and settlement, incorporating the approvals each step requires. In accounting, reconciliation can provide an entry point into the close, reporting, and audit preparation. The strength of the initial product depends partly on whether it earns the trust and access needed to take on adjacent steps. That expansion gives vertical AI companies a path to becoming substantial businesses with a **deeper role in their customers' operations rather than staying as niche single-point solutions**.

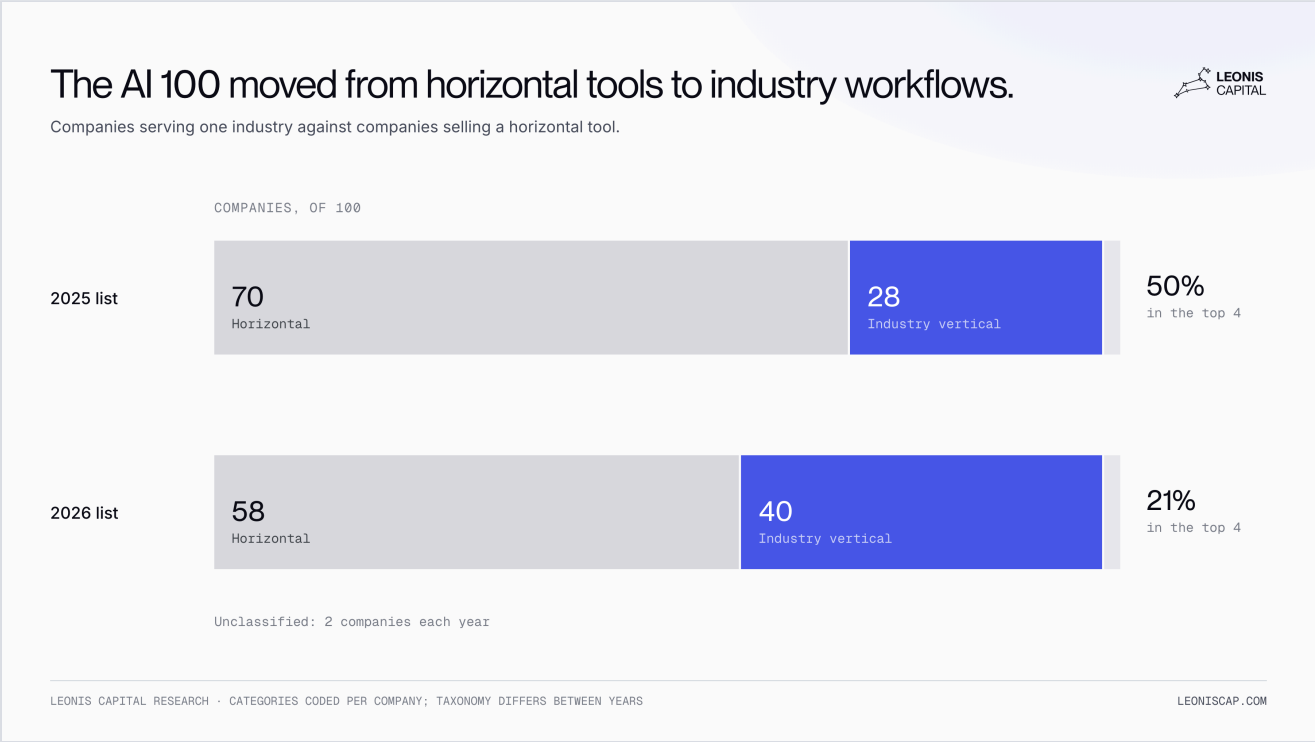
Reaching that position requires substantial implementation work. Vendors need to connect existing systems, adapt the product to the customer's processes, and establish how exceptions and approvals will be handled. That can require a more services-heavy approach to early

27 Leonis Capital, The Leonis AI 100, 2025: markets break out in sequence.

28 Leonis Capital, From stopwatches to systems of action.

deployments, including forward-deployed engineers who work directly with customers to implement the product. The knowledge gained through that work helps a vendor identify where its product can take on more responsibility and what must be built before it can do so.

EXHIBIT 10



Companies serving one industry against companies selling a horizontal tool, 2025 and 2026 lists.

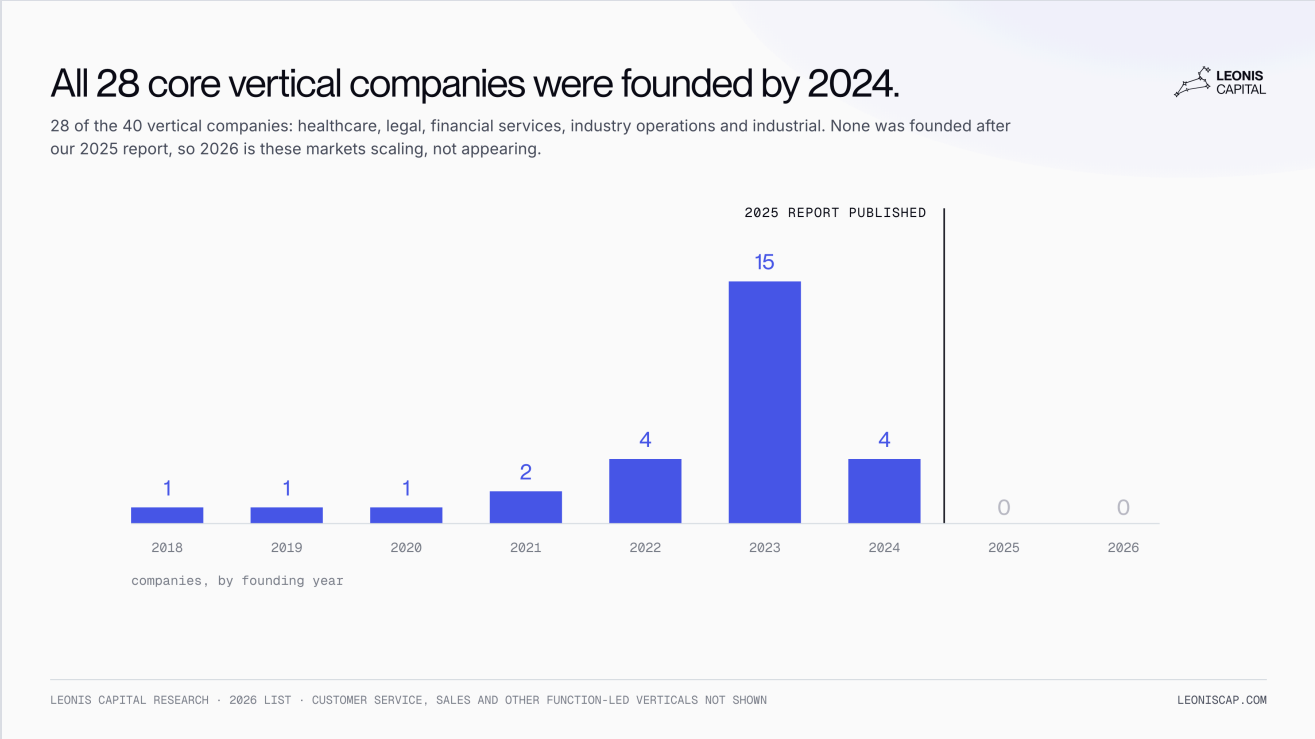
SOURCE: LEONIS CAPITAL RESEARCH. CATEGORIES CODED PER COMPANY; TAXONOMY DIFFERS BETWEEN YEARS.

Many of the vertical AI companies breaking out this year had already spent years building their products. Of the 28 companies across verticals like healthcare, legal, finance, and industrial operations, **all were founded in 2024 or earlier, yet only six appeared in last year's AI 100** (Exhibit 11).[29] Breaking into these industries takes time, as companies must adapt their products to established workflows and earn customers' trust. This year's growth reflects **that longer path to adoption beginning to pay off.**

That timing develops an argument we made last year about building ahead of model capabilities. Teams entering a difficult market early have time to learn which tasks matter, how customers judge results, and what prevents deployment. When models become more capable, those teams

29 The 28 are the companies in five core verticals: healthcare, legal, financial services, industry operations and industrial. The 40 also count function-led verticals such as customer service and sales.

EXHIBIT 11



Founding year of the 28 companies in the five core verticals on the 2026 list.

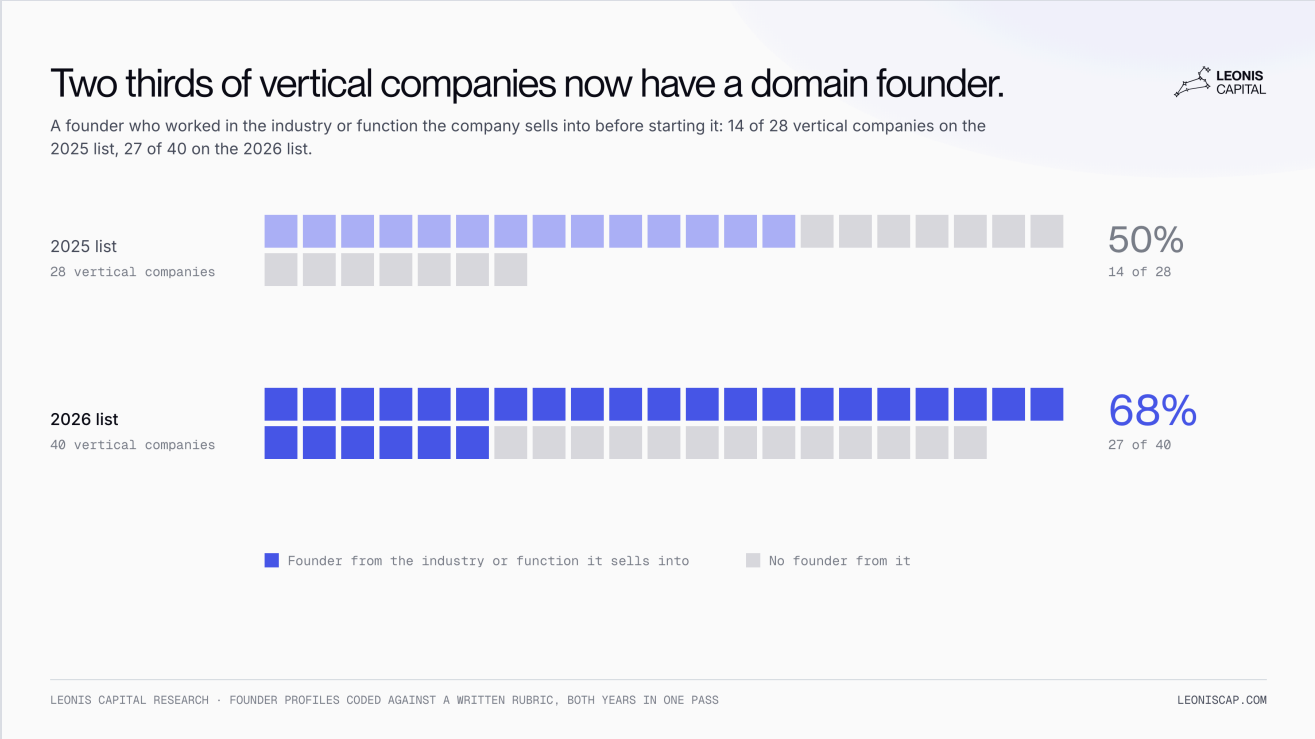
SOURCE: LEONIS CAPITAL RESEARCH. FUNCTION-LED VERTICALS SUCH AS CUSTOMER SERVICE AND SALES NOT SHOWN.

can apply the improvement within a product and customer relationship they have already begun developing. Technical progress can accelerate the payoff from work that was underway well before the market attracted widespread attention.

Building and deploying these products also requires a wider mix of expertise (Exhibit 12). Engineers need to work with people who understand the industry's processes, can guide implementation, and know how customers evaluate and purchase software. As vertical AI companies take on broader workflows, experienced operators and commercial teams become increasingly important to their ability to grow.

Vertical AI needs people who know the industry, and that has changed where companies get built. Last year, 63 of the AI 100 companies were headquartered in the Bay Area, reflecting the concentration of technical talent, capital, and early adopters around Silicon Valley. This year, **New York's representation grew from 12 companies to 19, making it the clearest second center in the AI 100** (Exhibit 13). As AI applications specialize, New York's mix of established industries, technical talent, and experienced industry professionals gives founders more reasons to build there.

EXHIBIT 12



Vertical companies with a founder who worked in the industry or function the company sells into, 2025 and 2026 lists.

SOURCE: LEONIS CAPITAL RESEARCH. FOUNDER PROFILES CODED AGAINST A WRITTEN RUBRIC, BOTH YEARS IN ONE PASS.

New York is home to 43 Fortune 500 headquarters, more than any other U.S. city, with major industries spanning finance, healthcare, real estate, media, and fashion.[30] That breadth provides both a substantial customer base and a workforce familiar with the vertical industries AI companies are trying to serve.

The New York companies in this year's AI 100 illustrate that combination. Basis works in accounting, Hadrius in financial compliance, and Pace in insurance operations. Building for these customers requires understanding which tasks a product can take on, where human review is necessary, and how it will fit into existing systems. An experienced compliance professional can help translate those requirements into product decisions, while a commercial hire with industry experience understands who controls the budget, what evidence buyers need, and how to navigate procurement.

The Bay Area continues to hold a substantial lead in capital and commercial scale. Companies headquartered there represent 51% of this year's AI 100 and 69% of the total capital raised (Exhibit 14), with a larger share having reached substantial revenue milestones and later financing

30 NYC Economic Development Corporation, Why NYC and AI NYC.

EXHIBIT 13



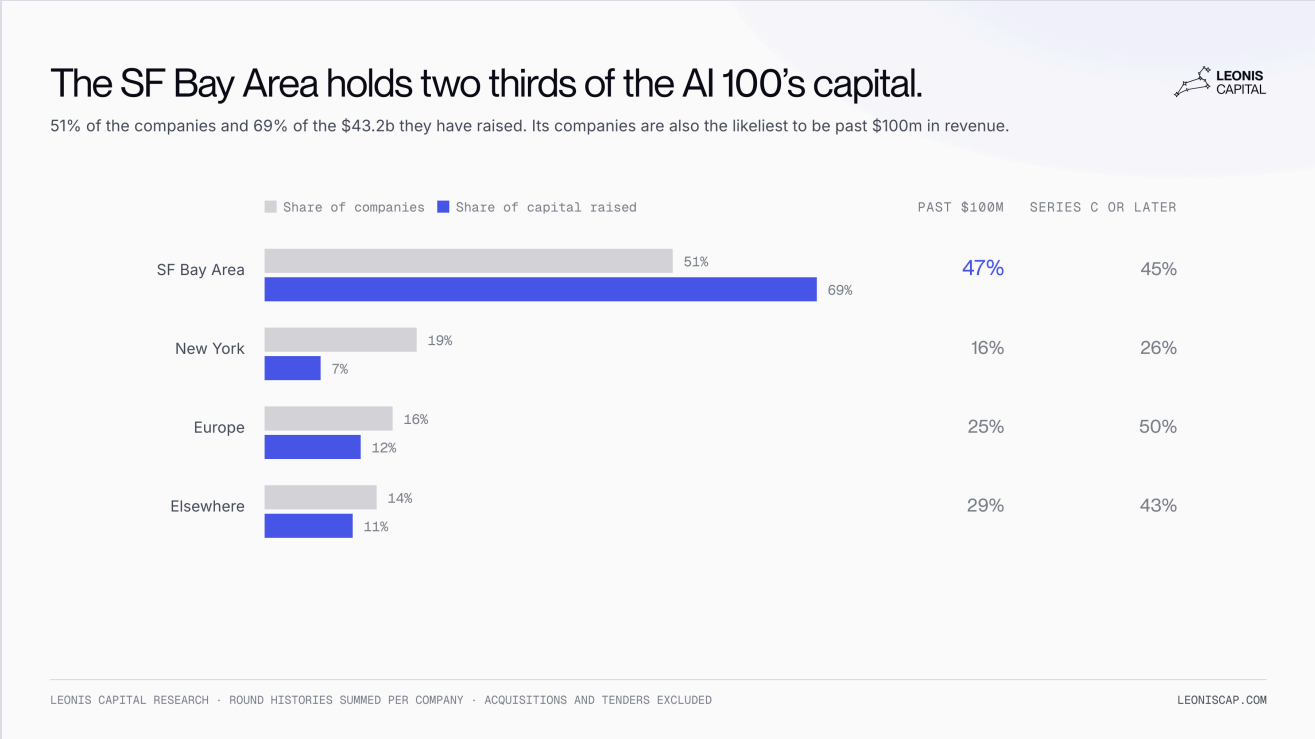
Company headquarters on the 2025 and 2026 lists, one bucket each.

SOURCE: LEONIS CAPITAL RESEARCH. REPORTED HEADQUARTERS; TWO 2026 RECORDS UNRESOLVED.

stages. Its concentration of technical expertise, investors, and other AI businesses remains a powerful advantage.

New York's growth reflects the increasing importance of industry expertise and customer relationships as AI reaches further into established businesses. Those customers also broaden the demand supporting the AI economy. Spending by law firms, insurers, healthcare providers, and financial institutions supports **both the applications serving them and the infrastructure beneath those applications**. As vertical products take on more work, budgets from these industries can become a growing source of demand throughout the AI supply chain.

EXHIBIT 14



Share of companies and of capital raised by region, with each region's share past $100 million and at Series C or later.

SOURCE: LEONIS CAPITAL RESEARCH. ROUND HISTORIES SUMMED PER COMPANY; ACQUISITIONS AND TENDERS EXCLUDED.

INSIGHT 4

Usage-Based Pricing Becomes The Backbone of the AI Economy

IN SHORT

Most of the list charges for usage, which ties customer spending to the work AI does and lets demand from established industries flow through to the suppliers behind each product.

Last year, we observed that AI companies could grow revenue without hiring proportionally more employees.[31] Usage-based pricing allows their customers' spending to grow independently of headcount too. For example, an insurer with the same number of employees can process more claims through an AI product and pay more as that volume increases. Pricing per document, call, or agent execution are examples of how AI changes the unit of software monetization from the number of users to the amount of work performed.

In this year's AI 100, 64 companies incorporate some form of usage-based pricing, including 24 that combine it with seat-based charges (Exhibit 15). The pattern holds beyond our list: of the 44 AI-native products in a16z's latest consumer AI web rankings, 64% charge usage fees or sell extra credits, mostly alongside a subscription.[32]

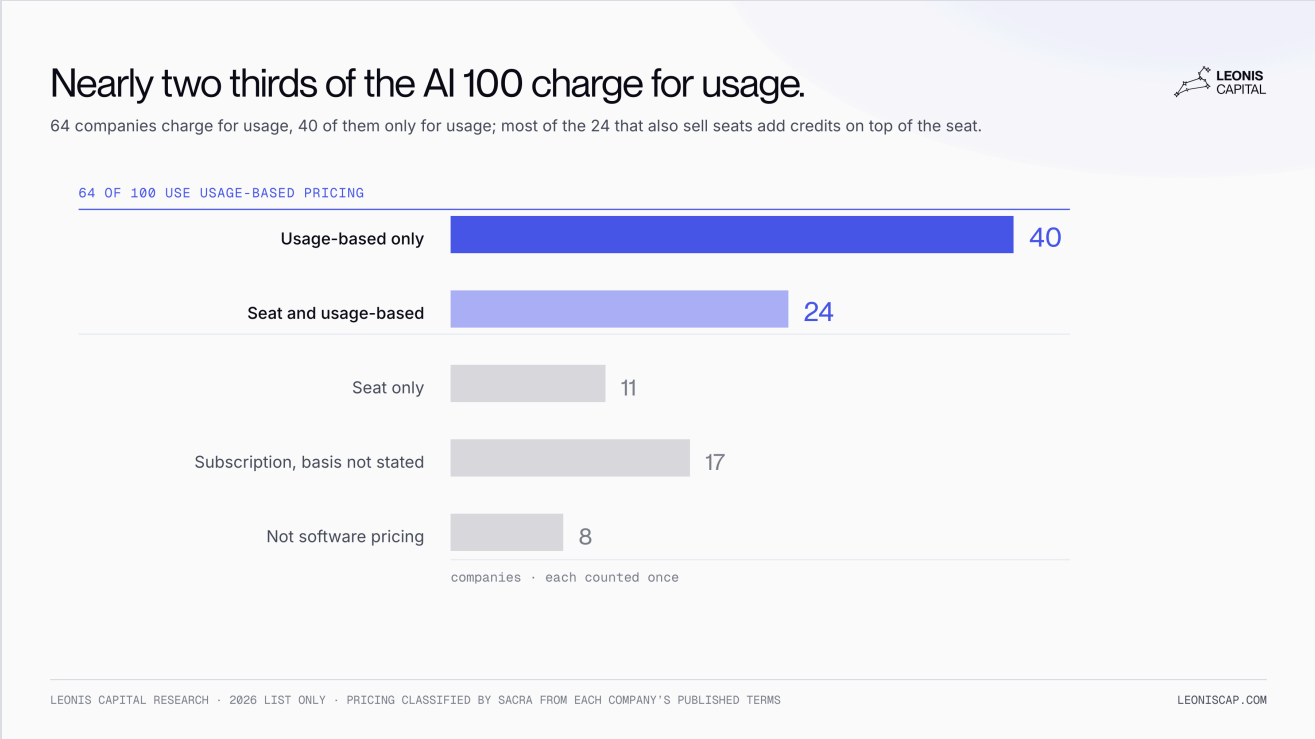
For buyers, this makes the purchasing decision easier to connect to an operating need. A business evaluating an AI service for handling calls can compare the cost of using it with the expense of handling those calls today. It can start with a portion of the workload and expand spending as the service proves useful. The decision to spend more becomes part of deciding how much work to entrust to the product.

That relationship addresses a familiar frustration with enterprise software. A company could purchase a large number of licenses and continue paying while much of the software went unused. When charges depend on consumption, more of the vendor's revenue depends on customers actively putting the product to work. Winning the initial contract creates an opportunity to earn further spending through continued use, but the vendor must also focus on adoption post-sale. **The economics of supplying AI also favor this structure.** Processing more documents or running more agents creates additional inference and infrastructure costs. Charging for consumption gives vendors a way to collect more revenue as those workloads expand. Even seat-priced companies on our list expect to change. Granola charges a fixed fee per seat, but CEO Chris Pedregal has said that as customers use it for more agentic work, it will "have to move to variable pricing in some shape or

31 Leonis Capital, The Leonis AI 100, 2025: smaller and flatter teams.

32 Andreessen Horowitz, The Top 100 Gen AI Consumer Apps.

EXHIBIT 15



Pricing model of every company on the 2026 list, each counted once.

SOURCE: SACRA, CLASSIFIED FROM EACH COMPANY'S PUBLISHED PRICING AND TERMS. 2026 LIST ONLY.

form" so its heaviest users pay more than the average one.[33] Harvey co-founder Gabe Pereyra has said the company is "per seat right now" but expects parts of the product to move to a hybrid model with pass-through costs.[34]

Among companies on this year's list that report a revenue figure, the 49 that use usage-based pricing **report a median revenue of $100 million, compared with $35 million for the nine software companies without it** (Exhibit 16).[35] At that scale, continued consumption becomes central to sustaining the business. Vendors have a commercial reason to invest in onboarding, integrations, and reliability because helping customers use the product more extensively can translate directly into additional revenue.

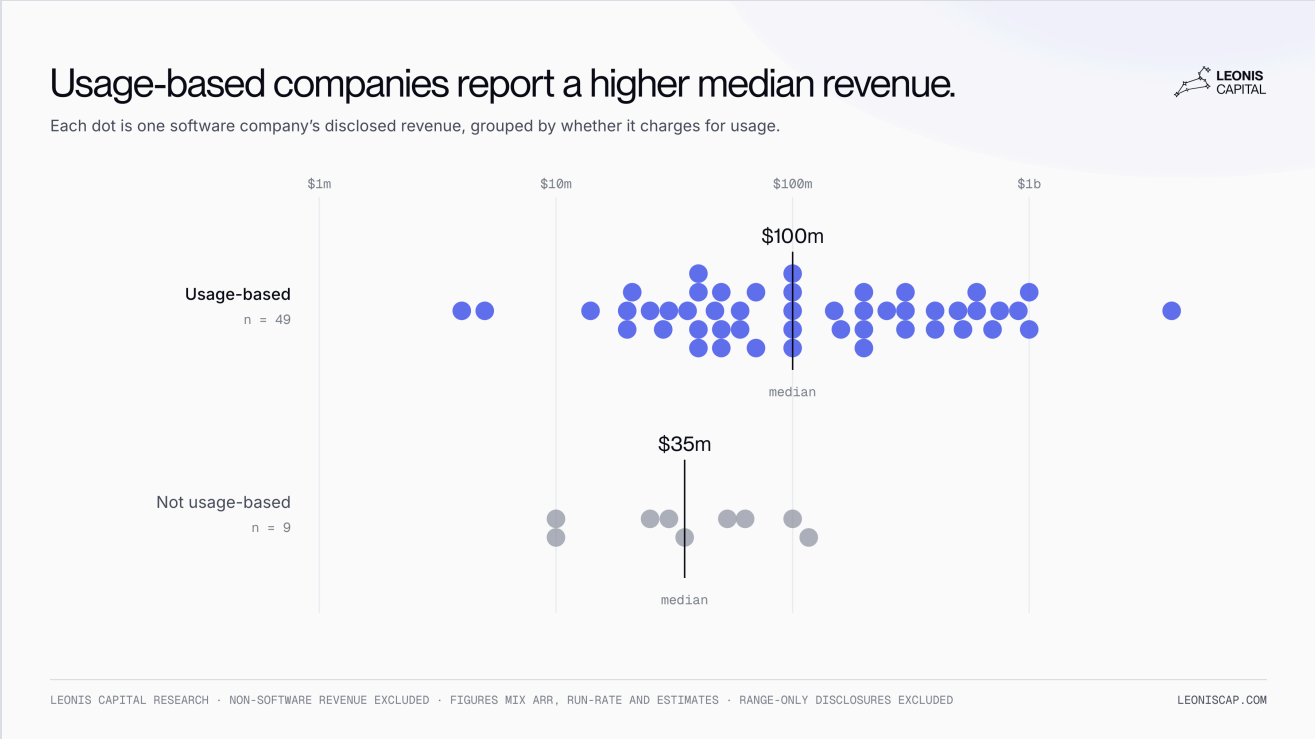
Usage-based pricing also **introduces risks alongside that expansion opportunity**. Revenue can be less predictable than fixed seat

33 Financial Times, interview with Chris Pedregal.

34 The Frontier with Brendan Foody, interview with Gabe Pereyra.

35 The comparison covers companies with disclosed revenue figures among the 92 software companies in the 2026 AI 100, including 49 classified as usage-based and 9 as not usage-based, using Sacra's pricing classification. Figures mix ARR, annualized run rates, and third-party estimates.

EXHIBIT 16



Every disclosed revenue figure among the 92 software companies on the 2026 list, by whether the company charges for usage.

SOURCE: LEONIS CAPITAL RESEARCH; PRICING CLASSIFIED BY SACRA. NON-SOFTWARE REVENUE AND RANGE-ONLY DISCLOSURES EXCLUDED.

subscriptions, while customers may hesitate to delegate more work if they cannot anticipate the bill. Vendors face ongoing pressure to demonstrate that additional usage is worth the expense. Charging for consumption can also create conflicting incentives when a vendor earns more because its product takes more steps or uses more resources to complete the same task. That conflict comes from metering inputs such as tokens or agent steps. Pricing per completed unit of work, such as a processed claim or a resolved ticket, turns it around, because extra steps then cost the vendor rather than the customer.

Falling inference costs raise a further question about how much value vendors can retain.

In the best case, vendors price against the value of the work performed. A product that performs a task previously costing $50 might still command $15–20 even if its inference cost falls from $5 to $0.50, because the customer evaluates the price against the alternative cost of completing the work. Lower costs can then expand the vendor's margins.

In a more moderate outcome, workflow depth helps usage growth offset falling prices. A company embedded in a customer's operations can use its integrations, data, and distribution to retain customers and expand usage. Prices may fall as inference gets cheaper, but revenue can continue growing if the volume of work increases faster than prices decline.

In the worst case, the product becomes a commodity pass-through. Customers can easily switch providers, and pricing closely tracks the underlying model costs. Savings pass directly to customers, leaving the vendor with less revenue per unit and little ability to retain the value its product creates.

As AI companies move from selling software access to performing work, spending becomes tied to how much work customers delegate to AI products. Greater application usage generates additional revenue for application companies and increases their spending on inference, sandboxes, and other services. **Usage pricing becomes the transmission mechanism of the AI economy**, allowing demand from an established industry to translate into revenue across the AI businesses serving it. Usage may also be a waypoint rather than an endpoint, with outcomes-based pricing eventually becoming the standard.[36] Three companies on the list, Sierra, Decagon and Parloa, already price some contracts on outcomes, such as a resolved customer conversation, and as vendors take responsibility for whole workflows, more of them may charge for the result rather than the work performed.[37]

This creates a feedback loop across the AI economy. Suppliers can invest that revenue in improving the infrastructure and services applications depend on. Better capabilities and lower costs can then make additional work practical to delegate to AI, supporting further adoption and spending. Through this loop, growth in one part of the AI economy can help expand the opportunities available to the others.

36 Leonis Capital, Our predictions for 2026.

37 Sacra on Sierra, Decagon pricing, Parloa pricing.

Conclusion

IN SHORT

The sequence we described last year is running on schedule. Next, the $100 million mark keeps losing its rarity and outcome pricing becomes the default for agents that own a workflow.

Last year's AI 100 made several calls about where the market was heading, and most held up. We argued that AI markets break out in sequence as models cross each industry's capability threshold, with writing and coding first and healthcare, law and finance next. This year's list shows that sequence in motion: vertical companies make up a much larger share of it, and video, which we flagged as a likely next market, produced the fastest grower we measured. One thing that surprised us was that last year we expected compressed revenue timelines to normalize as capital tightened, but capital kept flowing, and more companies reached scale faster than before.

A year ago, we were examining an emerging startup playbook. Today, the companies following that playbook have become economically significant enough to create markets around themselves. Expectations **for how quickly AI can become part of the economy need to keep adjusting**.

As we look forward, we believe that the $100 million mark will keep losing its rarity, and far more of next year's list will have crossed it. Usage pricing will turn out to be a bridge. As vendors take responsibility for whole workflows, outcome-based pricing, charging for a resolved claim or a closed ticket rather than for the work performed, will become the default for agents that own a workflow end to end. More suppliers to AI builders will pass $1 billion in annualized revenue. And the sequence we described last year will keep running, as insurers, law firms and health systems move from pilots to deployments across entire workflows. AI's economic role will not be limited to doing existing work with fewer people. Lowering the cost of doing work can also change how much work is worth undertaking and who can undertake it.

All of this growth is not without risks. Much of the supplier revenue we measured comes from a small number of frontier labs and from AI companies spending venture capital, so a pullback in research budgets or in financing would travel quickly through this economy. Not all headline revenue is created equal, and not all of it will hold up. There will be disruption along the way, and the growth documented in this report does not resolve concerns about employment or justify every valuation. However, the experience gained from building and using AI will matter beyond any single financing cycle.

We finish this year's AI 100 more optimistic because the range of plausible uses for AI has widened, and more businesses are learning how to put those capabilities to work. The next stage will test how far that usefulness can extend beyond the early adopters. We will know how far when we return to the AI 100 next year.

About This Report

IN SHORT

Every revenue figure is a company's latest dated, non-forecast figure as of October 1, 2026. None is audited, and the dataset names the source of each one.

Leonis Capital is a research-driven venture firm focused on AI-native companies. Since 2021, we've partnered with technical founders at the earliest stages, combining technical expertise with institutional investing experience.

After four years of systematically tracking the AI startup ecosystem, we've seen patterns that we think are worth sharing with the AI community. Too much of the AI discourse is driven by hype and speculation. This report draws instead on one of the largest private datasets of AI startup metrics, built and maintained as part of our work as investors.

We hope this analysis helps the AI community reflect on what has and hasn't worked in building AI startups.

Methodology & Dataset

Much of this dataset is based on Alfa, our proprietary AI research system. Alfa continuously indexes and analyzes AI startups from around the world, tracking everything from GitHub commits to hiring velocity to revenue signals. It helps us scale our analysis as an AI-native venture firm.

Between 2022 and 2026, Alfa indexed more than 20,000 AI startups globally. From this universe, we identified the fastest-growing startups based on publicly available data and signals (funding, hiring, usage, GitHub activity, Crunchbase, Product Hunt, news announcements, and ARR estimations from sources like Sacra and ARR Club) and private references from founders and investors. We've aggregated that into the Leonis AI 100 to glean further insights on the fastest-growing AI startups.

To qualify, companies must be:

- Focused on AI applications, infrastructure or services rather than frontier model development, since frontier labs operate with fundamentally different capital intensity, GTM motions, and defensibility dynamics.[38]

38 We respect the work of foundation model labs, but including them in the data would distort the insights as funding amounts, team size, and ARR growth at foundation model companies are drastically different than AI applications. Although foundation models like ChatGPT are also popular consumer apps and compete directly with application startups, we believe that model labs are in a category of their own.

- Founded or built their product in the post-ChatGPT era (2022 or later). We excluded legacy companies that rebranded or pivoted into generative AI.
- Building AI-native products (e.g., LLMs, agents, AI infrastructure, generative media, etc.)
- Independently operated (i.e., no spin-outs from Big Tech)
- Showing clear traction (e.g., revenue, users, open-source adoption)

How we chose the companies in the AI 100: From the companies that met these tests, we chose the 100 with the strongest evidence of commercial traction as of October 1, 2026. We weighed four kinds of evidence, strongest first. Revenue: a dated revenue or run-rate figure, and how fast it grew. Customers and usage: customer counts, named enterprise deployments or active users, for companies that do not disclose revenue. Financing: the size and recency of the latest round and who led it, as an outside read on traction. Recency: evidence dated within the twelve months before October 1. A large round alone did not earn a place, and when two companies were close, we preferred evidence from the company itself over third-party estimates.

We tagged each company by:

- Product category (infrastructure, agents, video, devtools, etc.)
- Founder backgrounds (research, ex-big tech, co-working history, age at founding, etc)
- Geography and location
- Founding date
- Investors (Pre-Seed, Seed, Series A, and Series B)[39]
- Round and valuation trajectory
- Revenue growth trajectory[40]
- Go-to-market motion (product-led, marketing-led, community-led, and sales-led)[41]
- Pricing model (usage-based, seat-based, both, or not stated), classified by Sacra from each company's published pricing and terms

39 We want to focus on tracking investors who identified the AI 100 companies early on and supported them in their initial growth. Our data is only based on publicly available information, thus, we likely have omitted some investors.

40 This is based on a combination of news or PR announcements and revenue estimates from sources like ARR Club, Sacra, and Latka. There will likely be inaccuracies given that many data points come from estimates.

41 We categorized the companies' GTM motion into four buckets: 1) product-led growth, growth driven by word-of-mouth and using the product itself as the driver of customer acquisition; 2) sales-led growth, which is driven by top-down enterprise sales; 3) marketing-led growth; and 4) community-led growth.

- Pivots and time-to-pivot

All data is current as of October 1, 2026. We're a small team, and this is a fast-moving space. We've likely missed things. We've open-sourced the dataset and invite founders, operators, and fellow researchers to help us improve it.

Revenue figures are each company's latest dated, non-forecast figure as of October 1, 2026, drawn from company disclosures, press reports, Sacra estimates and ARR Club. A company that states revenue only as a range, such as "eight figures," counts as disclosing but is not given a figure, and model outputs from estimate sites are excluded. 64 of this year's companies have a revenue figure, against 89 last year. We do not require a disclosed figure to qualify. Companies without one earned their place through other signals. For example, Hadrius serves more than 500 financial firms, two-thirds of them using three or more modules.[42] Adaptive Security passed 500 enterprise customers within a year of launch.[43] PhysicsX reports recognized revenue doubling and bookings tripling.[44] Instinct passed 100,000 users within months of opening its beta.

Where the data is weaker: None of the revenue figures in this report is audited. Of the 64 figures, 29 come from the companies themselves, 15 from press reports, many of them citing unnamed sources, 16 are third-party estimates, mostly from Sacra, and 4 come from ARR trackers without an outside source.[45] Some of the largest figures are reported rather than confirmed: Cursor's $4 billion comes from an unnamed source cited by Forbes,[46] and Harvey's $400 million was reported alongside its September round.[47] Four other large figures, for Baseten, Replit, Lovable and Sierra, are Sacra estimates.[48] And marketplace figures such as Mercor's and micro1's are gross revenue, much of which passes through to the experts doing the work.[49] The dataset notes the source of each figure.

42 Hadrius, Series A announcement.

43 Adaptive Security, Series B announcement.

44 PhysicsX, Series C announcement.

45 Third-party estimates mostly from Sacra; trackers such as ARR Club.

46 Forbes, June 8, 2026, citing an unnamed source.

47 ARR Club, Harvey at $400m ARR.

48 Sacra estimates for Baseten, Replit, Lovable and Sierra.

49 Sacra on Mercor; TechCrunch on micro1.

View the dataset on Airtable.[50] To cite this report or our data: Jenny Xiao, Jay Zhao and Liang Wu, “The 2026 Leonis AI 100,” Leonis Capital, October 2026.

Please email research@leoniscap.com if you have any suggestions, corrections, or additions.

50 The Leonis AI 100 dataset on Airtable.

Acknowledgements

This report would not have been possible without the contributions of many individuals who generously shared their insights, data, and expertise.

Thank you to our data partner, Sacra, who provided crucial data, and particularly Marcelo Ballvé, who offered editorial feedback on many drafts and classified the pricing model of every company on this year's list.

We also thank Jeff Epstein (Bessemer Venture Partners), Dan Gray (Odin), Allen Wu (Nvidia), Grace Gong (Smart Venture Pod), Shikhar Ghosh (Harvard Business School), David Kanter (MLCommons), and Andrew Tan (TLDR) for reviewing early drafts.

This document is provided for informational purposes only and should not be construed as investment advice, a recommendation, or an offer to buy or sell any security, instrument, or strategy. Nothing herein should be relied upon as financial, legal, accounting, or tax guidance. The content reflects information available to Leonis Capital (together with its affiliates, "Leonis") at the time of publication, including third-party sources believed to be reliable and the author's own analysis. While efforts have been made to ensure accuracy and completeness, Leonis does not guarantee that the information is correct, complete, or current. All views and opinions are subject to change without notice, and neither Leonis nor the author(s) accepts liability for errors, omissions, or reliance placed on this material. This document may include forward-looking statements, projections, or illustrative examples. Such statements are inherently uncertain, depend on assumptions, and are subject to factors beyond Leonis' control. Actual results may differ materially from those anticipated. Past performance is not a reliable indicator of future returns, and investing always involves risk, including the potential loss of principal. Different types of investments involve varying degrees of risk, and there can be no assurance that the future performance of any specific investment, strategy, company, or product referenced, directly or indirectly, will be profitable, equal any indicated performance level, or be suitable for any particular portfolio. Leonis Capital is an Exempt Reporting Adviser ("ERA") under the Investment Advisers Act of 1940 and applicable state regulations. ERA status does not imply a particular level of skill or expertise, nor does it represent endorsement by any regulatory authority. This material is intended for a general audience and does not account for the objectives, financial situation, or constraints of any individual, nor does it establish an advisory relationship with Leonis. Leonis, its affiliates, and its clients may hold, increase, reduce, or dispose of positions in companies or securities mentioned, independent of any views expressed herein.



Leonis Capital is a research-driven venture firm and the first partner for technical founders.

leoniscap.com