SRF Capital .Studio

Indian Startup Ecosystem · 2026

The Agent Layer Is Winning. It Does Not Own The Ground It Stands On.

The Week In Indian Startup Capital

WEEK 27 · 27 to 31 July 2026

By Sriram Chidambaram, Founder & CEO, SRF Capital Studio

Indian startups raised $142.3 Mn last week. Two-thirds of it went to artificial intelligence, and every dollar of that went to the application layer. In the same five days, the most consequential Indian AI announcement was about a model. Those two facts are the same fact, viewed from opposite ends.

$142.3 Mn	65%	88%
RAISED ACROSS 11 DEALS	WENT TO ARTIFICIAL INTELLIGENCE	OF ENTERPRISE LLM SPEND, THREE VENDORS

Start with the budget, not the technology

FUNDAMENTAL 01 / WHERE THE MONEY COMES FROM

The first useful question about any AI company is not which model it runs on. It is which budget line pays the invoice.

Software gets paid out of the software budget. That budget is finite, renegotiated every year, and defended by a procurement team whose performance review depends on making it smaller.

Labour gets paid out of payroll and outsourcing. That budget is an order of magnitude larger, and nobody in procurement is measured on shrinking it, because shrinking it means people lose their jobs.

Freehand, which raised $75 Mn last week and accounted for more than half of the entire week's Indian funding table, states the arithmetic plainly. Its founders put global enterprise spending on supply chain software at roughly $16 Bn a year, against roughly $348 Bn a year spent on the people who do the work that software still cannot do (Source: Freehand, reported by Crowdfund Insider, July 2026).

That is a gap of about twenty-one times.

> An assistant competes for the $16 Bn. An agent competes for the $348 Bn. Everything else in the agentic story is implementation detail.

FIGURE 1 / THE TWO BUDGETS AN AI COMPANY CAN BE PAID FROM

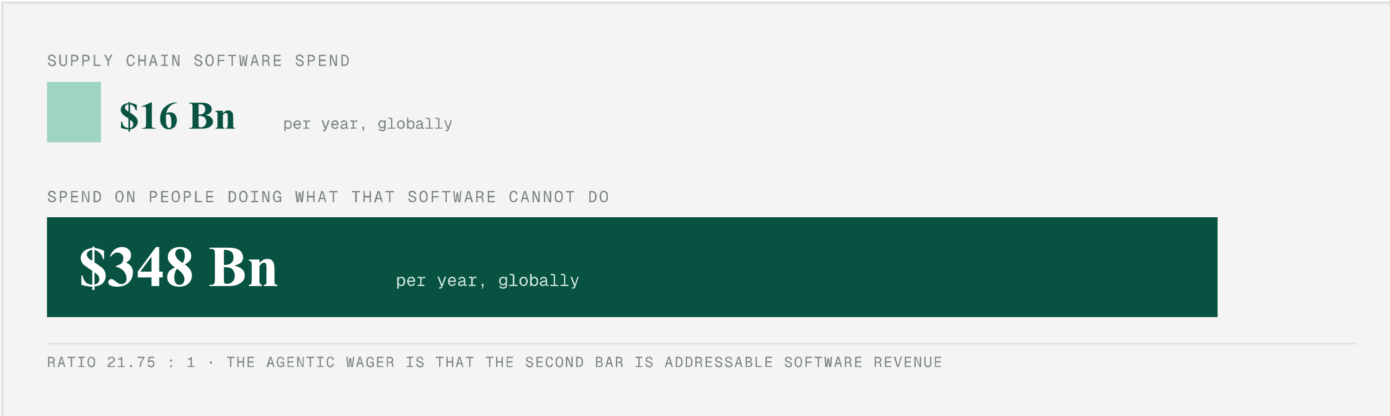

Figures as stated by Freehand and reported in July 2026 coverage of its $75 Mn round. Company framing, not an independent market estimate.

Why the winners are narrow rather than broad

FUNDAMENTAL 02 / AGENCY REQUIRES CONTEXT, AND CONTEXT DOES NOT GENERALISE

A frontier model can reason about an invoice dispute in the abstract. It cannot know that your Malaysian supplier has a sixty-day term with a two percent early settlement discount expiring on the twelfth, that the last three shipments arrived short, and that your category manager already agreed verbally to absorb the freight variance.

Reasoning is general. Consequence is local. The moment software stops suggesting and starts deciding, it needs the second thing, and the second thing has to be acquired one company and one workflow at a time.

This is why the money is going narrow. Freehand's own thesis, as its founder framed it to Forbes, is that **"vertical AI that goes deep rather than broad will deliver more value"** (Source: Forbes, 29 July 2026). Its agents already run inside Meta, Unilever, Johnson & Johnson, Pfizer, Dunkin' and Cardinal Health. Customers report recovering five to ten percent of supply chain spend and cutting procure-to-pay cycles by more than seventy percent (Source: Freehand, self-reported, July 2026).

The aggregate capital agrees with the anecdote. Across the twelve months to June 2026, vertical AI agents took 55.7 percent of disclosed agentic AI funding and 50.9 percent of the deals (Source: New Market Pitch agentic AI funding analysis, June 2026).

Narrow is not a limitation here. Narrow is the moat, because the context you accumulate inside one workflow is the only thing in the stack a competitor cannot buy from a vendor.

Now look at the ground underneath

FUNDAMENTAL 03 / THREE LANDLORDS, ONE FLOOR

Here is the part that rarely makes the funding headline. Almost every one of these companies rents its intelligence, and there are very few landlords.

On Menlo Ventures' 2025 enterprise survey, Anthropic held roughly 40 percent of enterprise large language model API spend, OpenAI 27 percent and Google 21 percent. Three companies, close to 88 percent of the layer that every agentic product sits on. In the coding segment specifically, Anthropic's share ran to about 54 percent (Source: Menlo Ventures, State of Generative AI in the Enterprise, 2025).

The capital position underneath that is more concentrated still. Global venture funding hit a record $510 Bn in the first half of 2026, more than the $440 Bn deployed across the whole of 2025. OpenAI and Anthropic between them accounted for $217 Bn of it, which is 43 percent of every venture dollar deployed anywhere in the world over six months. Anthropic alone raised roughly $65 Bn in the second quarter, close to a third of the global quarterly total (Source: Crunchbase News, 2 July 2026).

FIGURE 2 / THE STACK THE WEEK'S DEALS WERE FUNDED INTO

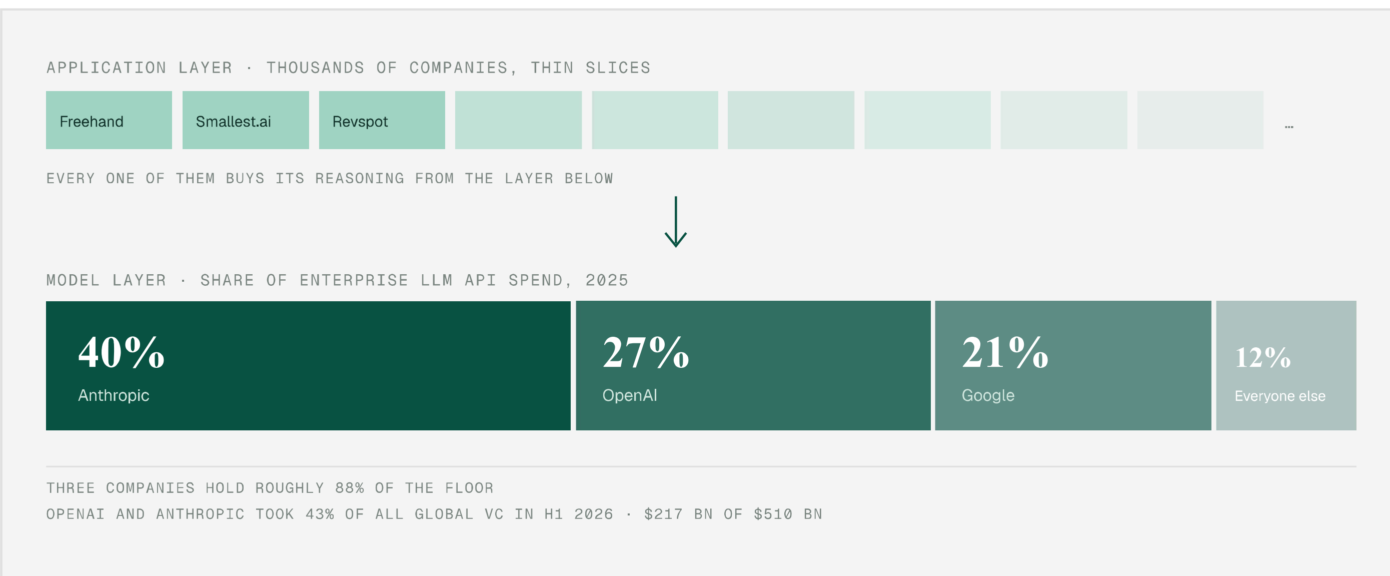

Model layer shares from Menlo Ventures' 2025 enterprise survey. Capital concentration from Crunchbase News, 2 July 2026. Application layer companies shown are the three AI deals in the 27 to 31 July 2026 Indian funding table.

Set the two numbers next to each other and the asymmetry becomes hard to unsee. India's entire artificial intelligence sector raised $92.8 Mn last week. Anthropic raised roughly $65 Bn in a quarter. The ratio is about seven hundred to one.

88% OF ENTERPRISE LLM API SPEND HELD BY THREE COMPANIES	**43%** OF ALL GLOBAL VENTURE CAPITAL IN H1 2026 WENT TO TWO OF THEM	**~700x** ANTHROPIC'S Q2 RAISE AGAINST INDIA'S AI FUNDING WEEK

What renting intelligence actually costs

FUNDAMENTAL 04 / THE GROSS MARGIN IS WHERE THE RENT SHOWS UP

Classic software is built once and served to the next customer for almost nothing. That is the entire reason the category earned 75 to 85 percent gross margins and the valuation multiples that came with them.

An agent does not work that way. Every request consumes real inference, billed per token, and the cost recurs on the next request rather than amortising across it. Inference sits in cost of goods sold, not in fixed research spend.

ICONIQ surveyed roughly three hundred software executives in late 2025. Across AI products generally, gross margins ran 41 percent in 2024, 45 percent in 2025, and were projected at 52 percent for 2026. For pure application layer products, the equivalent path was 33 percent, then 38 percent, projected to 45 percent (Source: ICONIQ, January 2026 snapshot; 2026 figures are self-reported projections).

FIGURE 3 / GROSS MARGIN, BY WHERE YOU SIT IN THE STACK

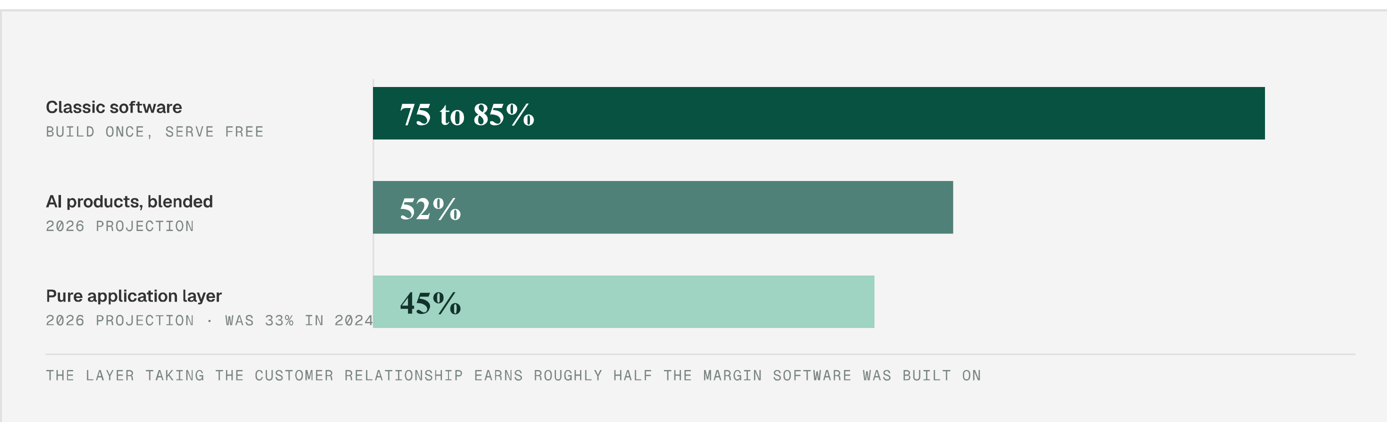

ICONIQ survey of approximately 300 software executives, fielded Q4 2025, published January 2026. Classic software benchmark per Andreessen Horowitz. 2026 values are projections.

Two things make this less bleak than the chart suggests, and both are worth understanding precisely rather than hoping for.

The first is deflation. Andreessen Horowitz estimated that for a model of equivalent capability, inference cost falls roughly ten times a year (Source: Andreessen Horowitz, "LLMflation", 2024). The capability that was expensive in 2023 is close to free now. That is a genuine tailwind, and it is why application layer margins moved from 33 to a projected 45 in two years.

The second is portability, and this is the one founders control. A company hard-wired to a single frontier model pays frontier prices on every request, permanently. A company architected to route each task to the cheapest model capable of doing it inherits every price cut on the day it lands.

> At the application layer, gross margin is an architecture decision before it is a finance outcome. It is set by engineers, months before it shows up in a board pack.

Reading last week's table through that lens

THE DATA / 27 TO 31 JULY 2026

Indian startups raised $142.3 Mn across eleven deals, down about 32 percent from $209.1 Mn across fourteen deals the week before (Source: Inc42). A cooler week by volume, and a very concentrated one by composition.

Artificial intelligence took $92.8 Mn across three deals. That is 65 percent of the week's capital from 27 percent of its deals.

All three were application layer. All three were business-to-business. None of them was a model company.

FIGURE 4 / WHERE THE $142.3 MN WENT

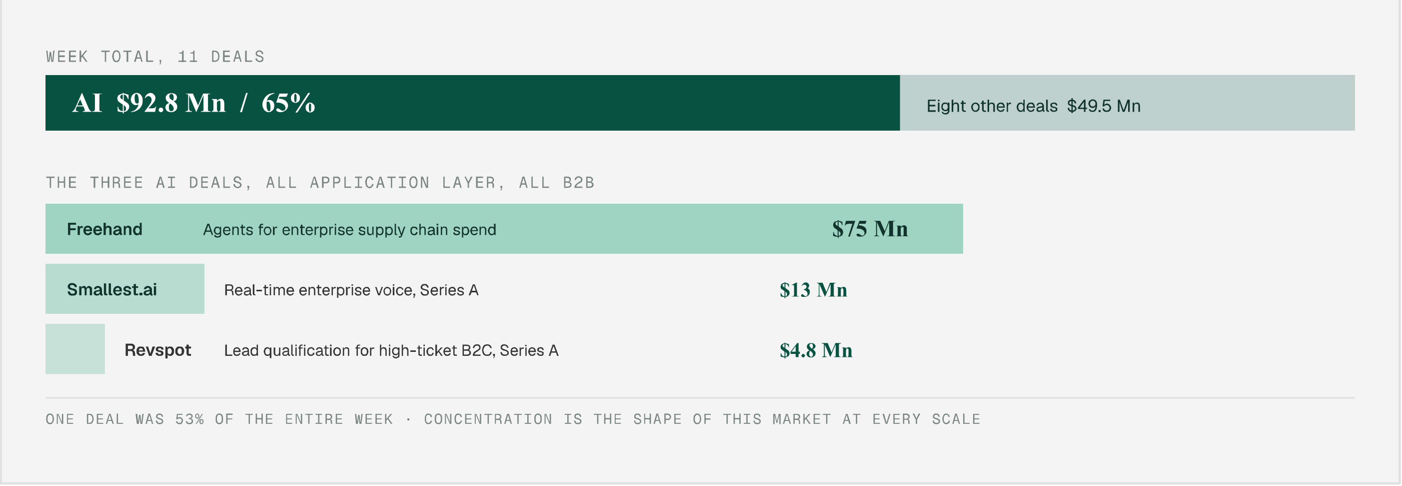

Indian startup funding, 27 to 31 July 2026. Source: Inc42 weekly funding table.

Smallest.ai is the partial exception worth naming, because it complicates the story usefully. It builds its own speech recognition and speech generation models rather than assembling someone else's, and its pitch to enterprises is a vertically integrated stack instead of a stitched-together one. Its customers include RingCentral, Truecaller and Piramal, and the round takes total funding past $21 Mn (Source: company announcement, 30 July 2026).

But scale matters. ElevenLabs raised $500 Mn in February 2026 for its voice platform. Smallest.ai raised $13 Mn. Building your own models inside a niche is a real strategic choice, and it is a different thing from being a frontier lab.

The prize it is aiming at is genuine: fewer than one percent of global voice interactions are currently powered by AI, against a market some forecasts put at $47.5 Bn by 2034 (Source: industry forecasts cited by Smallest.ai; treat long-dated market projections with caution).

The same week, the other layer

THE COUNTERPOINT / 30 JULY 2026, BENGALURU

On 30 July, inside this exact funding window, Sarvam announced at its first developer conference that it is building a trillion-plus parameter model from scratch in India, aimed at coding, cybersecurity, simulation and scientific research, alongside an India-hosted inference service (Source: Inc42, July 2026).

That announcement did not appear in the week's funding table, because it was not a funding event. It was arguably the more consequential item.

Sarvam raised $234 Mn in June 2026 at a $1.5 Bn post-money valuation, led by HCLTech. Indian AI startups raised $1,067 Mn across 157 deals in the first half of 2026, up 33 percent on the year before (Source: Venture Intelligence). One company took roughly 22 percent of the half.

That is the same concentration shape as the global market, one order of magnitude down.

LAYER	TYPICAL INDIAN CHEQUE	WHAT IT BUYS	WHO FUNDS IT
Application	$4.8 Mn to $75 Mn	Workflow ownership, customer context, distribution	Private venture capital, unassisted
Model	$234 Mn and upward	Weights, training pipeline, pricing control	Strategic corporates plus state compute

India is now running both trades at once, but the model layer trade does not close on private capital alone. It needed the IndiaAI Mission, a programme of roughly ₹10,372 Cr that has onboarded more than 38,000 GPUs for subsidised access, and it needed a strategic corporate cheque to reach $234 Mn.

For a founder, that distinction is not academic. It tells you which game you are actually allowed to play with the capital available to you.

What breaks, and how often

THE DISCIPLINE / MOST OF THESE PROJECTS WILL DIE

None of the above means the deployments work. Gartner projects that more than 40 percent of agentic AI projects will be cancelled by the end of 2027, and its stated reasons are escalating costs, unclear business value and inadequate risk controls (Source: Gartner, June 2025).

Read that list again. Model capability is not on it. These are governance failures, not technology failures, and they are the failures a finance function is supposed to catch.

Gartner also estimated that of the thousands of vendors claiming agentic capability, only around 130 were building anything that deserved the label. The industry now has a term for the rest of it: agent washing.

And yet the same firm projects that 15 percent of day-to-day work decisions will be made autonomously by 2028, up from zero in 2024, and that 33 percent of enterprise software applications will contain agentic AI by 2028, up from under one percent in 2024.

Both forecasts can hold. The category compounds while most individual projects inside it get quietly written off. That is what every real technology shift has looked like from the inside.

How SRF Capital Studio reads this

At SRF Capital Studio, the questions we end up asking an AI-native company are rarely about the model. They are about the gross margin bridge, the concentration of a single vendor inside cost of goods sold, and whether the context the company is accumulating actually belongs to it or sits in a customer's system.

The honest observation is that most founders we meet can describe their product architecture in detail and cannot describe their inference cost per customer without opening a laptop. That gap closes at the worst possible moment, which is halfway through due diligence.

The questions worth carrying into your next board meeting

01 What is our gross margin on the AI product specifically, separated from everything else we sell?

02 What percentage of cost of goods sold flows to a single model provider, and what happens at renewal if their pricing moves against us?

03 Could we move a meaningful share of our workload to a different model in under a quarter, and has anyone tested it?

04 Of the context our agents rely on, how much have we accumulated and own, and how much are we reading from a customer system we could lose access to?

05 Are we selling against the software budget or the labour budget, and does our pricing model match the answer?

06 If inference costs fall ten times next year, does that margin land with us or with our customers?

A company that can answer those six questions is running an agentic business. A company that cannot is running a demo with revenue attached to it.

The agent layer is capturing the workflow, the customer and the budget. It is not capturing the ground. Own the context, rent the intelligence, and never confuse the two.

Sources: Inc42 weekly Indian startup funding table, 27 to 31 July 2026 · Crunchbase News, 2 July 2026 · Menlo Ventures, State of Generative AI in the Enterprise, 2025 · ICONIQ AI margin survey, January 2026 · Gartner, June 2025 · Venture Intelligence H1 2026 · Andreessen Horowitz, LLMflation, 2024 · Forbes, 29 July 2026 · Company announcements from Freehand, Smallest.ai, Revspot and Sarvam, July 2026.

Verification needed before publication: the $16 Bn and $348 Bn figures in Figure 1 are Freehand's own framing rather than an independent market estimate. The 2026 gross margin values in Figure 3 are self-reported projections, not results. Customer outcome figures for Freehand and Smallest.ai are company-reported and not independently audited.