INVESTMENT PORTFOLIO · NOVEMBER 2026

Sathvik VK

Selected DeepTech Investment Work

Founder to investor: one thesis, a 48-company research universe, three screens, and one investment-committee memo with a working model, written to be argued with.

48
companies mapped

3
screens: Pass · Watch · Proceed

1
IC memo with a 1,583-formula model

54
logged sources, each labelled

Prepared from public sources only, without management access. All companies are real; all valuations are proposals or estimates, labelled as such. Not investment advice.

CONTENTS

What is inside

How to read this document	2
1 · Investor profile	3
2 · How I make investment decisions	4
3 · Thesis: intelligence at the physical interface	5
4 · Where value accrues	6
5 · Why now, why India	7
6 · Capital, ownership and what this fund can own	8
7 · What would make the thesis wrong	9
8 · Fund strategy implied by the thesis	10
9 · Research universe / deal map	11
10 · From 48 to 1	12
11 · The screening framework	13
12 · Screen 1: Simplismart (AI infrastructure) · PASS	14
13 · Screen 2: C2i Semiconductors (semis) · WATCH	16
14 · Screen 3: CynLr (robotics) · PROCEED	18
15 · IC memo: CynLr Series B	20
16 · Financial analysis, valuation, returns	28
17 · Terms, diligence plan, case against, pre-mortem	35
18 · IC recommendation	39
19 · Decision review and bias check	40
20 · Sources and methodology	42

How to read this document

Every material statement carries one of five labels. They are the discipline the work is built on.

FACT	Externally verifiable: a filing, government data, a closed transaction.
COMPANY CLAIM	Said by the company or founder. Not independently verified.
MEDIA-REPORTED	Reported by press or databases without a primary document I could read.
ASSUMPTION	A modelling choice I made. The basis is stated alongside.
INFERENCE	My conclusion from the evidence. It can be wrong; the reasoning is shown.

Bracketed codes such as [S11] point to the source log (54 entries, delivered as a separate workbook with URLs, dates, reliability and evidence).

Companion files

- **Flagship model** (.xlsx): 1,583 live formulas, three scenario sheets, cap table, expected value, sensitivities, fund math.
- **Research universe** (.xlsx): 48 companies, funnel, scorecard with formulas.
- **Source log** (.xlsx): 54 sources with labels and conflicts noted.

> **Who is hiring?** The listing is posted by a recruiter (Neuraleap) on behalf of an unnamed "investment platform". Public evidence points most strongly to a Series A/B Indian deep-tech AIF with Piper Serica's Bharat Tech Fund profile [S01, S03, S04]. That is an inference, not a fact. I therefore calibrate to the **mandate archetype** rather than to one firm's internals: a ₹600-800 Cr fund, ₹25-50 Cr cheques, ~22 companies, a 30% gross IRR target, and sectors across semiconductors, AI, space, defence, robotics and fintech infrastructure.

1 · INVESTOR PROFILE

I have built companies. I have not yet managed institutional capital.

This portfolio is the evidence for the second half of that sentence: how I research, model, price and decide when the money is not mine and the fund has to work.

Where I come from

I am a founder and operator. My recent work has been **Perfstaq**, which has moved through several product explorations in AI-native founder and productivity infrastructure. [TO CONFIRM BEFORE SENDING: dates, role, what shipped, any users or revenue you are comfortable stating. Nothing here is invented.]

Building products taught me things a model cannot: what a slipping pilot feels like from the inside, why "the customer loved the demo" is not a purchase order, how a team behaves when the runway gets short, and how much of a technical roadmap is real once a supplier misses a date.

What I bring

- **Technology-commercial translation.** I learn an unfamiliar stack fast enough to find the two or three variables that decide whether it matters commercially.
- **Founder-side pattern recognition.** I can tell a founder's ambition from a founder's plan, because I have written both.
- **Research velocity.** This document (48 companies, 54 sources, a 1,583-formula model) is a sample of it.
- **AI fluency in practice,** not just as a theme. I use AI tooling daily to widen research coverage, and I check the output the way I would check a junior analyst's.

The gap, stated plainly

I have not led institutional transactions, sat on an investment committee or managed a portfolio through exits. The listing prefers people who have. I am not asking you to treat building a company as equivalent. I am asking you to judge the quality of the thinking in these pages and how fast I would close the rest of the gap.

What a Managing Partner should be able to check in 20 minutes

p.12 the funnel · p.13 the gates that override scores · p.15 a pass on a good company · p.28 a guidance miss I did not explain away · p.33 a walk-away price derived, not asserted · p.41 where my founder background biases me.

HOW I WANT TO BE JUDGED

Sourcing	Did I find the right companies, including ones that are not obvious?
Diligence	Did I separate what is known from what is claimed?
Pricing	Did I derive a price from fund math, not from the last round?
Judgment	Did I pass on good companies that are wrong for this fund?
Honesty	Did I show the evidence against my own recommendation?

BIASES I WATCH FOR IN MYSELF

Overrating product velocity and founder charisma; assuming any team can "figure it out"; underweighting capital structure, governance, downside protection and fund math. Each is checked explicitly on p.41.

2 · HOW I MAKE INVESTMENT DECISIONS

Five questions, in this order

Most bad venture decisions answer "is this a good company?" and stop. I treat that as the third question, not the first.

#	QUESTION	WHAT I ACTUALLY DO	WHERE IT SHOWS UP
1	**Can this fund own it?**	Stage, cheque size, ownership at a defensible price and portfolio conflicts are hard gates, checked before any scoring. A great company that fails a gate is a Pass or a Watch.	Gates p.13; Simplismart p.14
2	**Where does value accrue in this stack?**	Map the value chain, then ask who has pricing power when the technology commoditises.	Thesis pp.5-6
3	**Is it a good company?**	Sector-weighted scorecard: technology, commercial evidence, margin path, capital intensity, team record, market, exits, price, platform risk.	Scorecard p.13
4	**What is the one question that kills it?**	Every screen names a kill question and the cheapest diligence step to answer it.	Each screen
5	**At what price and structure does the fund clear its hurdle?**	Probability-weighted outcomes, ownership through dilution, and a walk-away price. Structure (tranches, preference) is priced, not just negotiated.	pp.31-34

Rules I hold myself to

- **Guidance is data.** A founder's past forecasts against delivery is the cheapest predictor I have of their next forecast. I haircut accordingly and show the haircut.
- **Logos are not customers.** I record pilots, POs and repeat orders separately.
- **No DCF for pre-revenue binaries.** I use scenario-weighted exits and say why.
- **One source repeated twenty times is one source.** The source log de-duplicates claims to their origin.
- **The base case must be boring.** If the base case is the founder's plan, I have not done the work.

What "conviction" means here

Confidence percentages on each screen are my estimate that the *decision* is right given current evidence, not that the company succeeds. A 70%-confident Pass means that in roughly three of ten similar cases I would expect the evidence to come back and prove me wrong. Each screen lists the evidence that would trigger that.

> **The output of an investment process is a decision plus the conditions that would reverse it.** Every screen ends with "what would change my mind" because a decision without its reversal conditions cannot be monitored.

3 · THESIS

Intelligence at the physical interface

For an Indian Series A fund writing ₹25-50 Cr cheques, the best risk-adjusted deep-tech returns this decade will come from companies that **own a hard interface between intelligence and the physical world** (power, perception and manipulation, RF sensing) and sell it to global industrial buyers. They will not come from companies that build, resell or optimise compute.

The three claims behind it

1. Compute is a scale game this fund cannot play. Inference prices fell roughly 280x in two years FACT [S26]. The Indian compute layer is being financed with $600M+ equity-and-debt packages and billion-dollar IPOs FACT [S28, S29]. A ₹40 Cr cheque buys nothing meaningful there, and intermediaries between the GPU and the user see their pricing deflate with the token.

2. The scarce asset moves to the physical edge. As models commoditise, the binding constraints become physical: watts per GPU, the ability to grasp an unseen part, the ability to sense through cloud or foliage. These problems are bound by physics and integration, so they are slow to copy and hard for platforms to absorb. INFERENCE

3. India's advantage is engineering depth per rupee, not domestic demand. Sovereign spending is large (₹1.39 lakh Cr of domestic defence acquisition in FY27; ₹1.27 lakh Cr for ISM 2.0) FACT [S37, S32], but the share startups capture is small: iDEX has cleared ₹3,853 Cr of prototypes cumulatively FACT [S39]. India-only demand rarely builds a fund-returner. Global buyers do, as Tonbo's 65% export mix shows FACT [S42].

Falsifiable: what would prove it wrong by 2029

IF WE OBSERVE...	...THE THESIS IS WEAKENED
Indian inference-software companies sustain >65% gross margin and >130% net revenue retention at scale	Claim 1: the compute layer is venture-investable after all
Physical-interface companies stall below 40% gross margin at 100+ units a year	Claim 2: hardware economics swamp the IP
Two or more India-only defence or space startups exceed ₹1,000 Cr revenue from domestic orders	Claim 3: domestic demand alone can build fund-returners
A foundation-model platform makes general-purpose grasping a commodity API with >95% first-pass success in brownfield cells	Claim 2: perception moats collapse

> **Why this is not "AI is big".** It makes a prediction about where value goes and where it does not, and tells the fund to *avoid* the most-funded segment: AI was 91% of Indian deep-tech funding in 2025 MEDIA-REPORTED [S35].

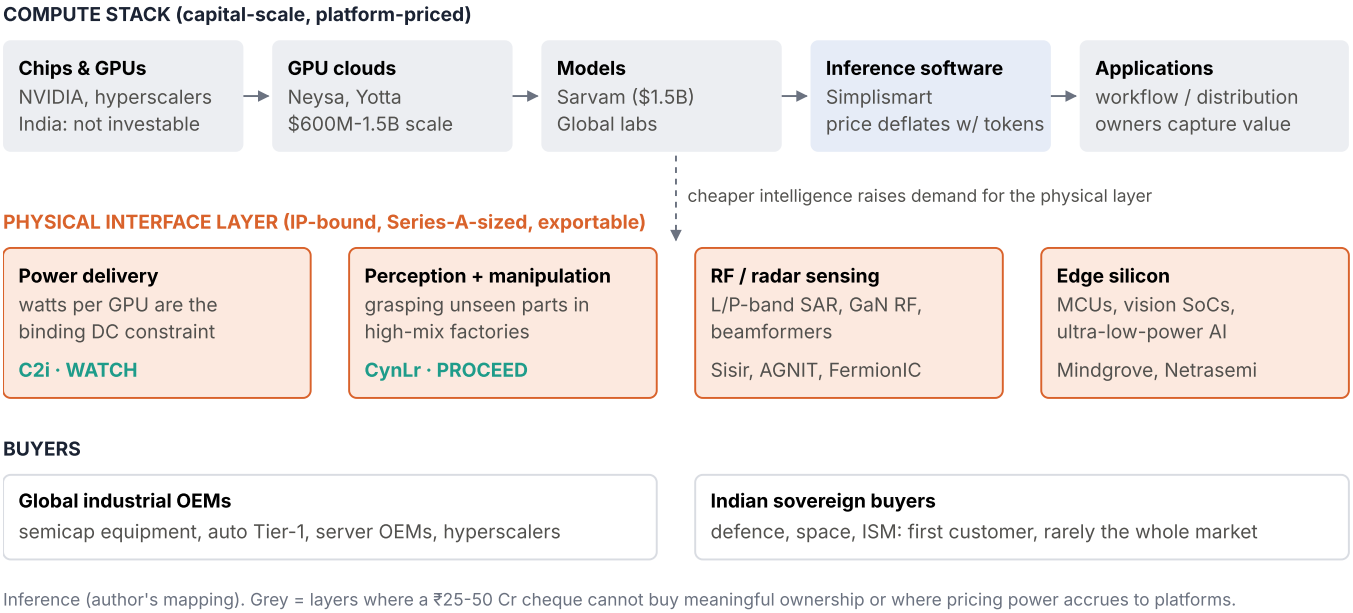
What it means operationally

Back: physical-interface IP sold to global industrial buyers, with domestic sovereign demand as a first customer, not the whole market. Avoid: compute resale, horizontal inference tooling, and "India-only" plays without an export path. Anchor sector: AI infrastructure and advanced computing, read as *the physical infrastructure that AI depends on*. Two adjacent sectors chosen for this fund's mandate: **semiconductors** and **robotics / industrial automation**.

4 · WHERE VALUE ACCRUES

Follow the pricing power, not the hype

The question is not whether AI is large. It is which layer keeps margin as intelligence gets cheaper.



Inference (author's mapping). Grey = layers where a ₹25-50 Cr cheque cannot buy meaningful ownership or where pricing power accrues to platforms.

Value-capture test by layer

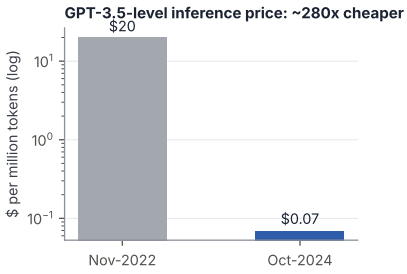
LAYER	PRICING POWER	CAPITAL TO SCALE	FUND CAN OWN ≥5%?
Chips & GPU clouds	High (incumbents)	$0.5-1.5B+	No
Model builders	Contested	$300M+	No
Inference tooling	Falling with token price	Low	Rarely at Series B prices
Applications	Workflow owners	Medium	Yes, but outside mandate
Physical interface	**Physics-bound**	**$15-60M to scale**	**Yes**

Where I could be wrong about value capture

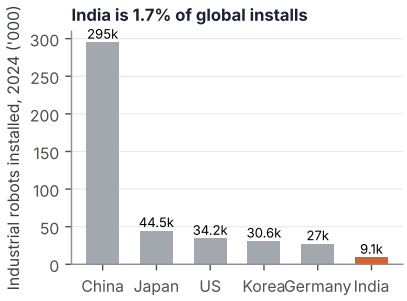
- **Integrators may capture the margin** instead of the component maker: a robot OEM or ODM bundles the interface IP and squeezes it. Test: gross margin trend and who owns the customer relationship.
- **Hardware is a slower compounding asset.** BOM, inventory and field service drag returns even when the IP is real. My model carries 120 inventory days and a 52% base hardware margin by FY32 for exactly that reason.
- **Platforms reach down.** NVIDIA's robotics stack and foundation models for manipulation could commoditise perception software. The defence is owning the deployment data and the hardware-software integration, not the algorithm alone.

5 · WHY NOW, WHY INDIA

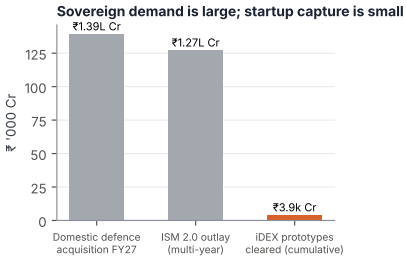
Three clocks are running at once



Intelligence got cheap. When the marginal cost of intelligence collapses, demand moves to where it meets atoms. FACT [S26]



India is tiny in robots; the market is global. 542k installs worldwide in 2024, 54% in China FACT [S22]. A Bengaluru robotics company must sell abroad to matter.



Sovereign money is real but mostly goes elsewhere. Startups need an export path. FACT [S37, S32, S39]

Why now

- **Cost curve:** inference ~280x cheaper in two years [S26]. Perception and planning that were uneconomic in a factory cell in 2022 are not in 2026. INFERENCE
- **Power wall:** AI-DC power-semiconductor demand is projected at ~$1.5B (2025) rising to ~$3.8B (2028) MEDIA-REPORTED [S21]. Rack densities are rising, and power is the constraint.
- **Design capability:** the DLI scheme has sanctioned 24 chip projects with 16 tape-outs FACT [S33]; Indian semis startups raised $61.9M in H1-2026 alone, ~81% of 2025's total MEDIA-REPORTED [S34].
- **Exit markets exist:** listed Indian defence and electronics names trade at 70-160x earnings MEDIA-REPORTED [S41], and Tonbo Imaging has filed for an IPO [S42]. Valuations this high are also a risk: see p.9.

Why India specifically

- **Engineering cost per tape-out or per robot cell** is a fraction of Silicon Valley's, which stretches Series A capital further. INFERENCE CynLr has built a 3-arm robot and vision stack on ~$15M raised MEDIA-REPORTED [S09].
- **Electronics import gap:** imports were $116.2B in FY26, up 17.8% FACT [S31]. Policy will keep paying for domestic substitution.
- **Global customers are already buying from Indian deep tech:** CynLr's customers are entirely outside India (company claim) COMPANY CLAIM [S11]; Tonbo's revenue is 65% exports [S42].

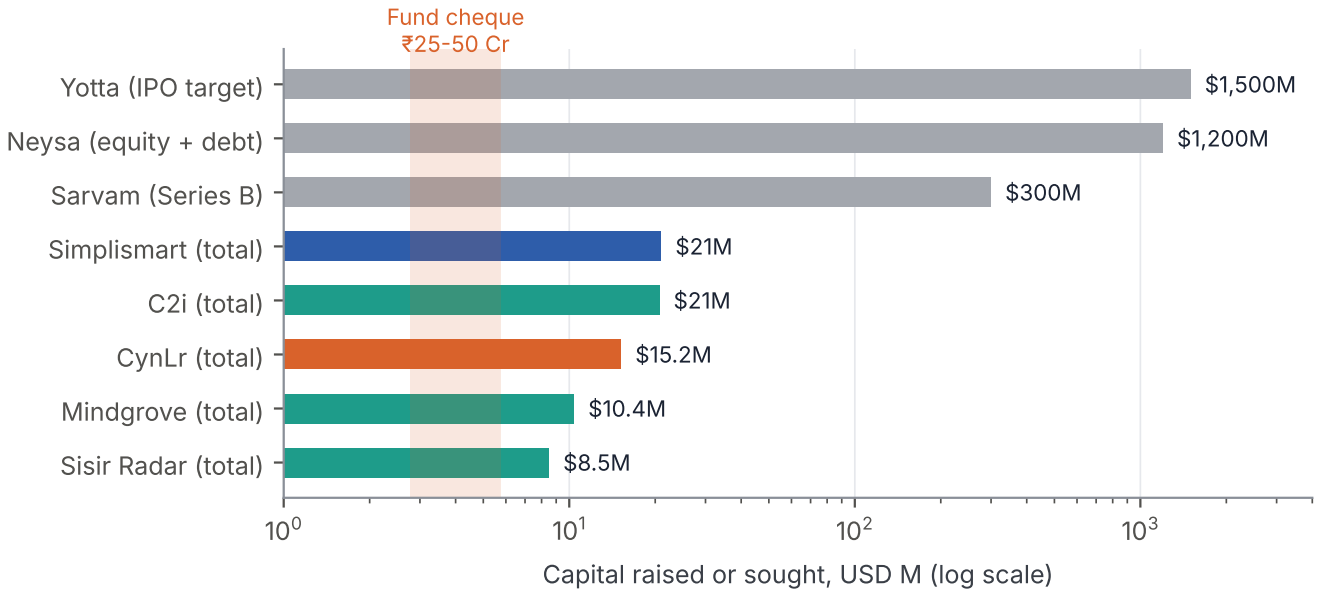
Why venture capital, not PE or grants

These companies are pre-profit, binary on one commercial event (design-in, production PO), and need ₹150-500 Cr before cash flow. That is venture risk. Grants (DLI, iDEX) de-risk the first prototype but do not fund scale. INFERENCE

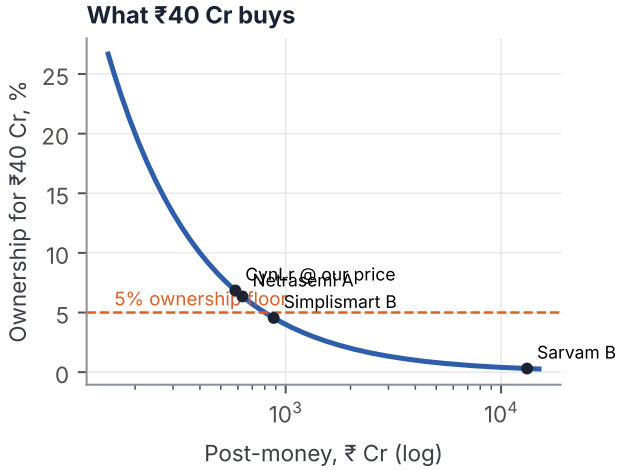
6 · CAPITAL, OWNERSHIP AND WHAT THIS FUND CAN OWN

The thesis is partly a statement about fund size

A thesis that ignores the fund's cheque size is a market view, not an investment strategy. This page is the bridge.



Capital raised or sought by layer, against a ₹25-50 Cr cheque (shaded). Compute and model layers are 50-500x the cheque; physical-interface companies are 2-5x. Sources: [S28, S29, S30, S16, S19, S09, S44, S46]. Neysa = equity plus intended debt; Yotta = IPO target.



Ownership for ₹40 Cr falls fast with price. Above ~₹800 Cr post-money, the fund drops below a 5% floor. Post-money values: Simplismart and Netrasemi are media estimates MEDIA-REPORTED; CynLr is our proposal ASSUMPTION.

What the two charts imply

1. **Ownership floor.** A fund of ~22 companies returning 30% gross needs positions large enough to matter. I use a 5% fully diluted floor at entry and expect ~65-75% retention through later rounds. ASSUMPTION
2. **Price ceiling follows.** With ₹40 Cr, the 5% floor means a post-money of ~₹800 Cr or less. That excludes most AI-infra Series B rounds and selects for physical-interface companies at Series A/B.
3. **Fund-returner arithmetic.** At 6-7% exit ownership, returning a ₹600 Cr fund needs a ~₹9,000 Cr (~$1B) exit (p.34). So the companies we back must have a credible path to $1B, which, given p.5, means global customers.

> **Corollary:** the most-funded Indian AI companies are excellent businesses that this fund should mostly not own. That is a statement about fit, not quality.

7 · CONSTRAINTS, FAILURE MODES AND ARCHETYPES

What would make the thesis wrong

Constraints

CONSTRAINT	WHY IT BITES	MITIGANT WE WOULD REQUIRE
Hardware working capital	Inventory and receivables consume cash as revenue grows; base model NWC reaches ₹93 Cr by FY32	Venture debt and milestone billing; NWC days in the diligence plan
Design-win cycles	18-36 months at OEMs; two-year pilots are normal COMPANY CLAIM [S11]	Price for slippage; tranche on POs
Supply chain	CynLr sources 400+ parts from 14 countries COMPANY CLAIM [S09]	Dual-source plan for long-lead items
Talent density	Senior analog/power and robotics talent is scarce	Track hiring velocity; Swiss/US centres
Exit valuations	Listed comps at 70-160x P/E are fragile [S41]	Underwrite on EV/revenue at strategic-sale multiples (2.4-7.7x), not listed P/E

The bear case for the whole thesis

Physical-interface companies remain sub-scale component vendors. Integrators and OEMs take the margin, revenue stalls at ₹50-150 Cr, and exits are acqui-hires at 1-2x revenue. In that world, the compute layer's winners, though capital-heavy, generate the only Indian deep-tech outcomes above $1B. **I put roughly 30% on this world.** INFERENCE

Archetypes to back

- **Interface owner with global design-ins:** sells a component or cell that a global OEM designs into a product line (C2i, CynLr).
- **Sovereign-first, export-ready:** first revenue from defence or space, with a product exportable within three years (Sisir Radar, GalaxEye).
- **Capital-efficient silicon with filed revenue:** real invoices before Series B (Mindgrove).

Archetypes to avoid

- **Compute resellers** and GPU-capex businesses at venture prices.
- **Horizontal AI tooling** whose value proposition is a discount on a deflating token.
- **Research labs with no customer,** e.g. companies founded a decade or more ago still pre-pilot.
- **Single-sovereign-customer companies** without an export roadmap (lumpy, as ideaForge's −85% then 5x quarters show FACT [S41]).
- **Pre-hardware companies priced like post-hardware companies** (6x step-ups before first flight or silicon).

> **What a fund must believe to run this thesis:** (1) at least 3-4 Indian physical-interface companies reach $100M+ revenue by 2033; (2) strategic acquirers pay 3-8x revenue for them; (3) we can enter at ≤₹800 Cr post-money with ≥5% ownership. All three are testable with each deal.

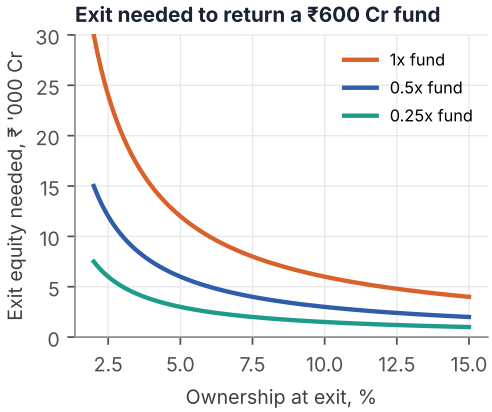
8 · FUND STRATEGY IMPLIED BY THE THESIS

Stage, cheque, ownership, reserves, exits

DECISION	RECOMMENDATION	REASONING
Stage	Series A, first commercial revenue; select Series B where a production PO exists	Technology risk mostly retired; commercial risk still priced in
Initial cheque	₹25-45 Cr, tranched where conversion is unproven	The tranche is worth more than the haggle (p.33)
Entry ownership	≥5% fully diluted; 7-12% target	Fund-returner arithmetic (p.34)
Reserves	~0.6-0.8x of initial cheque, concentrated in the top third	Pro-rata only where the step-up is earned by POs, not hype
Position cap	~10-12% of investable capital including follow-ons	Concentration discipline for a ~22-company fund
Board	Seat when leading; observer plus information rights otherwise	Value-add is customer introductions and supply chain, not a governance title
Exit paths	Strategic sale to global automation, semis or defence primes; Indian IPO for sovereign-led winners; secondaries to growth funds at Series C/D	Comps: ABB Robotics at ~2.4x revenue FACT [S23]; Saankhya to Tejas [S54]; Tonbo DRHP [S42]

Portfolio construction (illustrative)

Of ~22 companies: ~8 power and silicon, ~7 perception, robotics and automation, ~5 space, RF and defence sensing, and ~2 fintech infrastructure (outside this thesis but in the mandate). Expected distribution: ~40% return under 1x, ~35% return 1-3x, ~20% return 3-10x, and 1-2 outliers expected to return the fund. ASSUMPTION



The exit needed falls fast as ownership rises. At 6.6% exit ownership (our CynLr base case), returning 1x of a ₹600 Cr fund needs a ~₹9,100 Cr exit. INFERENCE

> **Secondary transactions** are an underused exit for Indian deep tech: buying founder or angel secondaries at a discount alongside primary rounds can lift ownership without raising the headline valuation. I would test this in every Series B where the seed cap table is crowded. INFERENCE

9 · RESEARCH UNIVERSE / DEAL MAP

48 companies mapped from public sources

This is research coverage, not proprietary deal flow: no company here was contacted. Key: IC memo screened interesting mapped context / comparable · ⚑ = existing Piper Serica Fund I company, flagged as a conflict and not screened [S05, S52]. Full fields (funding, investors, revenue, founders, concerns, confidence) are in the universe workbook.

AI INFRASTRUCTURE / COMPUTE (7)

Neysa
GPU cloud · Growth/control

Yotta Data Services
Sovereign GPU cloud / DC · Pre-IPO

Ziroh Labs
CPU inference runtime · Undisclosed

Simplismart
Inference optimisation · Series B

TrueFoundry
LLMOps PaaS · Series A

Sarvam AI
Sovereign LLMs · Series B

QpiAI
Quantum + AI · Series A + debt

SEMICONDUCTORS (14)

C2i Semiconductors
AI data-centre power delivery · Series A

Mindgrove Technologies
RISC-V MCU / IoT SoC · Series A

Netrasemi
Edge-AI vision SoC · Series A

InCore Semiconductors
RISC-V processor IP · Seed

AGNIT Semiconductors
GaN RF / power · Seed ext.

BigEndian Semiconductors
Surveillance vision SoC · Pre-Series A

Vervesemi
Analog / mixed-signal · Series A

Morphing Machines
Reconfigurable compute · Series A

HrdWyr
Edge AI SoC · Series A

Calligo Technologies
Posit HPC/AI chips · Raising Series A

FermionIC Design
RF beamformers / SerDes · First round

Sensesemi Technologies ⚑
Edge-AI sensor SoC · Seed

Saankhya Labs
SDR / broadcast SoC · Acquired

Kaynes Semicon
OSAT · Listed subsidiary

ROBOTICS & AUTOMATION (9)

CynLr (Vyuti Systems)
Vision-led dexterous manipulation · Series A (raising B)

Ati Motors
Autonomous industrial vehicles · Series B

Perceptyne
Industrial dual-arm / humanoid · Seed

Addverb
Warehouse automation · Strategic-owned

Niqo Robotics
Agri spot-spraying · Series B

Haber
Autonomous process control · Series C

Coratia Technologies ⚑
Underwater ROVs · Seed / pre-A

GreyOrange
Warehouse robotics · Series D

Sastra Robotics
Robotic device testing · UNKNOWN

SPACE & DEFENCE (18)

Sisir Radar
L/P-band SAR · Series A

GalaxEye
OptoSAR EO · Series A + debt

Pixxel
Hyperspectral EO · Series C

Digantara
Space domain awareness · Series B

Dhruva Space
Satellite platforms · Pre-Series B

Bellatrix Aerospace
Satellite propulsion · Pre-Series B

EtherealX
Reusable launch · Series A

Skyroot Aerospace
Small-lift launch · Late stage

Agnikul Cosmos
3D-printed engines / launch · Late stage

Astrogate Labs ⚑
Laser comms terminals · Pre-Series A

Xovian Aerospace ⚑
RF-intelligence satellites · Seed

Big Bang Boom Solutions
Counter-drone, armour · Growth

NewSpace Research & Tech.
Swarm UAS / HAPS · Growth

Raphe mPhibr
Military UAVs · Series B

Zuppa Geo Navigation
Drone autopilots · Bridge

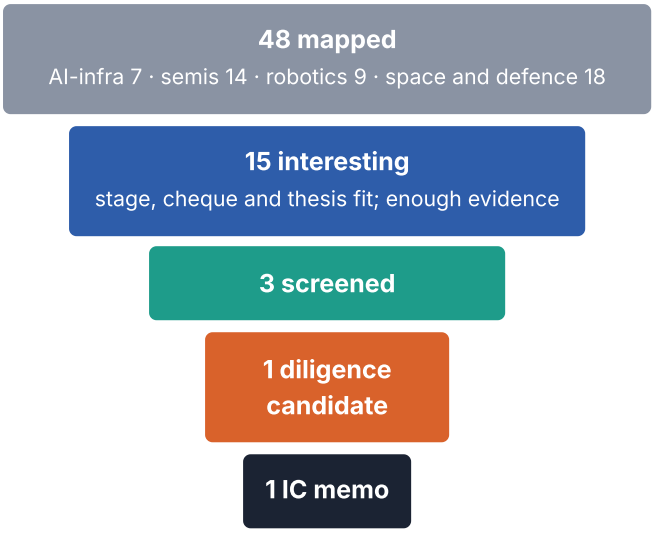
Tonbo Imaging
Thermal imaging, seekers · Pre-IPO

Data Patterns (listed)
Radar, EW, avionics · Listed

ideaForge (listed)
Military / enterprise UAV · Listed

10 · FROM 48 TO 1

The funnel, and why each company dropped out



Why these three screens

One per sector. Each tests a different part of the thesis: **Simplismart** tests the "avoid compute tooling" claim against a strong company; **C2i** tests the "power is the binding constraint" claim pre-revenue; **CynLr** tests the perception claim with three years of filed revenue. I deliberately did not pick the easiest companies to like.

What removed companies at each stage

FILTER	REMOVED	EXAMPLES
Not a venture deal or too late	14	Yotta, Neysa, Sarvam, Skyroot, Pixxel, Raphe, listed comps
Existing portfolio of the likely employer	4	Sensesemi, Coratia, Astrogate, Xovian
Too early or opaque for a ₹25-50 Cr cheque	9	InCore, FermionIC, Zuppa, Sastra, Perceptyne*
Thesis mismatch or long gestation	6	Morphing Machines (est. 2006, no customers), QpiAI, EtherealX (6x step-up pre-hardware)
Interesting, not screened this round	12	Mindgrove, Netrasemi, HrdWyr, Calligo, Ati Motors, Sisir Radar, GalaxEye, Bellatrix...

*Perceptyne kept on the watchlist as a seed-to-A candidate. Counts are approximate where a company met more than one filter.

Next five I would screen

1. **Sisir Radar:** low-frequency SAR niche, IAF iDEX wins, founder ex-Director SAC ISRO [S46]. Kill question: is there demand beyond one founder's reputation?
2. **Ati Motors:** revenue doubling, Bosch/Hyundai/TVS [S45]. A Series C co-invest candidate.
3. **Mindgrove:** the only RISC-V MCU company with filed revenue (₹3.0 Cr FY25) [S44].
4. **Bellatrix:** flat 2026 round may offer an attractive entry price, if revenue checks out.
5. **Netrasemi:** strong anchor (Zoho) but a ~₹630 Cr post on services revenue [S43]; screen for price.

11 · SCREENING FRAMEWORK

Gates first, then a sector-weighted score

A high score cannot rescue a failed gate. Weights differ by sector because what kills companies differs by sector.

DIMENSION (SCORE 1-5)	W: AI	W: SEMIS	W: ROBOTICS	SIMPLISMART	C2I	CYNLR
Interface ownership (thesis fit) Does the company own a hard interface between intelligence and the physical world (perception, power, RF) rather than reselling compute?	5	10	10	1	5	5
Technical differentiation & IP Is the hard part genuinely hard, proprietary and verifiable?	10	20	15	3	4	4
Commercial evidence Paying customers, POs, repeat orders: not LOIs or logos.	15	10	15	3	1	2
Path to gross margin Credible route to the margin structure that venture exits require.	15	5	15	2	4	3
Capital intensity (5 = light) Capital needed to reach the next value-inflection, relative to the fund.	5	10	10	4	3	2
Team & execution record Founder-market fit, hiring, and the gap between past guidance and delivery.	10	15	10	4	4	3
Market size & timing Bottom-up market and why it opens now.	10	10	5	4	5	4
Exit paths Named strategic buyers and listing route for this category.	10	10	5	3	4	3
Price & ownership Can we own ≥5% at a price that clears our return hurdle?	10	5	10	1	3	3
Platform / incumbent risk (5 = low) What happens if NVIDIA, a hyperscaler or an incumbent ships it?	10	5	5	1	3	3
Weighted score (0-100), own-sector weights				**52**	**74**	**63**

HARD GATES

GATE	SIMPL.	C2I	CYNLR
G1 Fund fit: Series A/B, ₹25-50 Cr, ≥5% ownership achievable	FAIL	PASS	PASS*
G2 No conflict with an existing portfolio company	PASS	PASS	PASS
G3 At least one paying customer or binding PO	PASS	FAIL	PASS
G4 No integrity or governance red flag in public record	PASS	PASS	PASS
Decision	**PASS**	**WATCH**	**PROCEED**

*CynLr passes G1 only at ≤₹400 Cr pre-money with a tranche (p.33).

WHY THE WEIGHTS DIFFER

- **AI infra:** commercial evidence, margin path and platform risk carry 40% combined, because the technology is rarely the failure point; pricing power is.
- **Semis:** technical differentiation and team carry 35%, because first silicon and the team's tape-out record decide everything; early revenue is weighted lightly because it is structurally late.
- **Robotics:** commercial evidence and margin path carry 30%, because deployment, uptime and service burden kill robotics companies more often than algorithms do.

Note what the score says. C2i scores highest (74) but is a WATCH, because it fails the paying-customer gate at a fund that invests at first revenue. The gates are the fund's mandate; the score is the company's quality. Keeping them separate stops a good story from overriding the mandate.

12 · SCREEN 1 · AI INFRASTRUCTURE / ADVANCED COMPUTING

Simplismart: a strong company this fund should not own

$14M
SERIES B, OCT-2026
MEDIA-REPORTED

~$100M
POST-MONEY (MEDIA EST.)
MEDIA-REPORTED

~4.5%
WHAT ₹40 CR WOULD BUY
INFERENCE

280x
TOKEN PRICE FALL, 2022-24
FACT

Company

Bengaluru/San Francisco inference-optimisation platform: a serving engine plus orchestration that lets enterprises run open and fine-tuned models on cloud or on-prem GPUs faster and cheaper. Founded by Amritanshu Jain (ex-Oracle Cloud) and Devansh Ghatak (ex-Google Search) MEDIA-REPORTED [S16].

Traction and evidence

- $7M Series A (Accel, Oct-2024) → $14M Series B (Dallas VC lead; NVIDIA, Accel, Shastra) in Oct-2026; ~$21M raised MEDIA-REPORTED [S16, S17].
- Says it cut InVideo's GPU serving cost by 56% and powers the serverless layer of Yotta's Shakti Studio COMPANY CLAIM [S16].
- Revenue, gross margin and retention: **UNKNOWN** publicly.

Why now / market

Every enterprise moving from API calls to self-hosted models needs serving infrastructure, and India's sovereign compute push (38,231 IndiaAI GPUs at ~₹65/GPU-hour FACT [S27]) creates domestic demand. The market is real and growing.

Technology and moat

Kernel-level and scheduling optimisations, batching, quantisation-aware serving and NVIDIA Dynamo integration. The engineering is good, but much of the technique is published or open-sourced (vLLM, SGLang, TensorRT-LLM) within months. The moat is speed of integration and enterprise trust, which are real but erodible. INFERENCE

The value-capture problem

Simplismart sells a *reduction* in the cost of a token. When token prices fall ~10x a year, a 56% saving on a shrinking base is a shrinking dollar pool, unless volume growth outruns price decline and the company prices on value rather than on cost. INFERENCE

WHO COULD SHIP THIS?	LIKELIHOOD	EFFECT
NVIDIA (NIM, Dynamo)	Already shipping components	Raises the floor that customers get for free
Hyperscalers (Bedrock, Vertex)	High	Captures cloud-native buyers
Open source (vLLM, SGLang)	Ongoing	Caps pricing for technical buyers
Indian GPU clouds (Yotta, Neysa)	Medium: partner today, builder tomorrow	Channel conflict

NVIDIA on the cap table is a credibility signal, but it also means the platform vendor has a front-row view of the roadmap. That cuts both ways. INFERENCE

Why this fund should not own it

- **Gate G1 fails:** ₹40 Cr buys ~4.5% at ~$100M post, only via an extension or secondary, with no lead position and no board influence.
- **Mandate fit is weak:** software margins exposed to platform pricing, not IP-bound physical technology.
- **Global competition** (Fireworks, Together, Baseten) is better capitalised by an order of magnitude.

12 · SCREEN 1 · CONTINUED

Economics, return math and decision

What a return would require

At a $100M entry and ~4.5% ownership diluting to ~3.2% at exit ASSUMPTION, a 3x return on ₹40 Cr needs ~₹120 Cr of proceeds, i.e. a ~₹3,750 Cr (~$425M) exit. For an inference-software company that implies roughly $45-60M ARR at 7-9x revenue by ~2031, or ~20x growth from an undisclosed base, against a deflating unit price. Possible, but the probability-weighted outcome is below 3x unless margins prove durable. INFERENCE

Business model questions I could not answer publicly

- Pricing unit: per GPU-hour managed, per token, or platform licence? Value-based pricing would weaken my objection.
- Gross margin when Simplismart resells GPU capacity versus software-only deployments.
- Net revenue retention and how much revenue comes from three customers.
- Share of revenue from on-prem sovereign deployments, which are stickier and less exposed to hyperscalers.

Scorecard summary

Weighted score (AI-infra weights)	**52 / 100**
Gate G1 (ownership at a price)	**FAIL**
Strongest dimensions	team, market, capital efficiency
Weakest dimensions	price, platform risk, interface ownership

> **The discipline being shown:** I like this company. If I were its founder, I would be proud of the Series B. The pass is about this fund's cheque, this round's price and this layer's pricing power, not about the team.

PASS

Confidence in this decision: 70%

Probability I would make the same call with full diligence access, given what is public today.

TOP 3 REASONS

1. Fails the ownership gate: ~4.5% for ₹40 Cr at ~$100M post, via extension or secondary only.
2. Value capture in inference tooling deflates with token prices; platforms and open source set the free floor.
3. Outside the IP-led physical-technology mandate the fund was raised for.

TOP 3 RISKS TO THIS DECISION

1. Sovereign and on-prem demand in India proves sticky and high-margin, a niche hyperscalers do not serve.
2. NVIDIA partnership becomes distribution, not a threat.
3. Fund misses an AI-infra outcome that would have anchored its AI story with LPs.

KILL QUESTION

> What share of revenue comes from customers who could move to an open-source serving stack or a hyperscaler endpoint within one quarter, without re-engineering?

NEXT DILIGENCE STEP

Not pursued for investment. Keep a relationship. Request a 15-minute conversation on sovereign on-prem demand, which is useful market intelligence for our semis and power companies.

WHAT WOULD CHANGE MY MIND

- Disclosed gross margin above 70% with net revenue retention above 130% across 20+ customers.
- Multi-year on-prem contracts with Indian sovereign or BFSI buyers that hyperscalers cannot serve.
- An allocation of ≥5% at ≤$60M post (e.g., a structured secondary).

13 · SCREEN 2 · SEMICONDUCTORS

C2i Semiconductors: the right problem, one step too early

$16.7M
SERIES A (PEAK XV, TDK)
MEDIA-REPORTED

$1.5→3.8B
AI-DC POWER SEMIS 2025→28
MEDIA-REPORTED

0
NAMED PAYING CUSTOMERS

2026
FIRST SILICON TAPE-OUT
COMPANY CLAIM

Company and team

Bengaluru fabless start-up building power-delivery silicon for AI servers: a digital controller ("Manas") and smart power stages ("Sarayu") for the last conversion step before the GPU. Six co-founders, presented as senior power-IC engineers; Preetam Tadeparthy is CTO COMPANY CLAIM [S20]. Prior employers are reported as a major US analog house but **not verified** in sources I could read. That is a first diligence item.

Problem

Each conversion stage between the grid and a GPU core loses energy as heat. At AI rack densities, a few points of efficiency are megawatts across a data centre and directly limit how many GPUs fit within a site's power envelope. Power is the binding constraint on AI build-outs, not chips. INFERENCE

Technology claims

- >96% conversion efficiency vs ~94% incumbents; 10-15% system-level performance per watt; "self-tuning" modular units COMPANY CLAIM [S20, S19].
- Roadmap to direct 800V-to-sub-1V conversion with GaN partners for high-voltage stages COMPANY CLAIM [S20].
- Pre-silicon validation only; tape-out scheduled 2026 COMPANY CLAIM. **No independent data.**

Market

UBS projects AI-DC power semiconductors at ~$1.5B (2025) rising to ~$3.8B (2028) MEDIA-REPORTED [S21]. That is a single broker estimate with a loose definition, so I treat it as order-of-magnitude only. Bottom-up, the market is accelerators shipped × power-stage content per accelerator. I could not source content per accelerator reliably, so it is a diligence item, not a guess. INFERENCE

Competitive map

PLAYER	POSITION
Infineon, MPS, Renesas, TI, Vicor	Incumbent power-stage and controller suppliers on GPU reference boards
Hyperscaler in-house teams	Custom power architectures for their own racks
GaN players	Partners for high-voltage stages; competitors at the system level
C2i	Software-defined control + power stage, targeting server OEMs and ODMs

Why the thesis likes it

- Owns a physics-bound interface (watts into the GPU). Platforms cannot simply ship it as software.
- Global buyers from day one; India's cost base stretches the tape-out budget.
- Credible syndicate: Peak XV leading, TDK Ventures (a power-electronics strategic) participating MEDIA-REPORTED [S18, S19].

Why not now

- **Gate G3 fails:** no paying customer or binding PO; first silicon not yet characterised.
- Design-in at a server OEM or ODM typically takes 18-36 months after working silicon. Incumbents sit on reference designs, so a first-time vendor must be qualified board by board. INFERENCE
- Its Series A was just priced and extended by top-tier investors; there is no round for us to lead until 2027.

Capital intensity: fabless with external foundries, but multiple tape-outs at advanced analog nodes plus reliability qualification. Tape-outs can exceed $3M each COMPANY CLAIM [S34]. Expect $40-60M before meaningful revenue. INFERENCE

13 · SCREEN 2 · CONTINUED

What has to be true, and the re-entry trigger

Return logic at a plausible Series B

If silicon works and one OEM design-in is visible, I would expect a 2027 Series B at $120-180M post. ₹45 Cr would buy ~3-4%, below our floor. That pushes us to one of three routes.
INFERENCE

1. **A larger co-lead allocation** secured through a relationship built now, before the round is competitive.
2. **An insider-led extension** if first silicon is late. Late silicon often makes the round cheaper, but only if it is late for reasons that do not touch the core idea.
3. **Accept a below-floor position** only if the expected multiple exceeds 5x, given the global market.

What must be true

ASSUMPTION	TEST	PASS THRESHOLD
Efficiency claim holds on silicon	Independent bench data at full load and transients	≥95.5% peak, ≥2 pts better than the reference design
A buyer will qualify a new vendor	Calls with 2 server OEM / ODM power engineers	Named programme with a qualification timeline
Team has shipped power silicon before	Employment and tape-out references	≥2 founders with volume-production power ICs
Gross margin of 55%+ at volume	BOM and package cost model	Credible without heroic yields

Scorecard summary

Weighted score (semis weights)	**74 / 100** (highest of three)
Gate G3 (paying customer)	**FAIL**
Strongest	thesis fit, market timing, team
Weakest	commercial evidence

> **Why WATCH is an active decision.** A watch without triggers is a polite pass. This one has three dated triggers and an owner (me): meet the founders this quarter, ask for silicon results the month they arrive, and talk to two ODM power engineers before the Series B opens.

Platform risk

NVIDIA defines GPU reference power designs. If NVIDIA standardises a vendor list that excludes start-ups, C2i's route runs through hyperscalers' custom racks. That is slower but larger per win. Not a kill, but it reshapes the go-to-market.
INFERENCE

WATCH

Confidence in this decision: 65%
Probability I would make the same call with full diligence access, given what is public today.

TOP 3 REASONS

1. Best thesis fit in the universe: a physics-bound constraint with global buyers.
2. Strong syndicate and team signals (Peak XV lead; TDK strategic).
3. Fails our first-revenue gate; the next priced round is ~2027, after silicon.

TOP 3 RISKS TO THIS DECISION

1. The Series B is pre-empted by insiders at a price we cannot meet: watching costs us the deal.
2. Silicon underdelivers on efficiency, and the moat becomes the control software alone.
3. Qualification at OEMs takes longer than the company's runway.

KILL QUESTION

> Will a server OEM or ODM qualify a first-time vendor's power stage onto a GPU board within 24 months of working silicon?

NEXT DILIGENCE STEP

Founder meeting this quarter; review first-silicon characterisation when available; two calls with ODM power-design leads (Taiwan) through TDK or our network.

WHAT WOULD CHANGE MY MIND

- Silicon data matching claims plus a named OEM or ODM design-in → move to diligence immediately.
- An insider-led bridge at or below the Series A price → would consider investing early.
- Evidence that founders have not shipped volume power silicon → downgrade to Pass.

14 · SCREEN 3 · ROBOTICS / INDUSTRIAL AUTOMATION

CynLr: real revenue, real logos, and a credibility gap to price

₹1.1→5.9 Cr
REVENUE FY23→FY25 MEDIA-REPORTED

~$1M
ANNUALISED, MAR-2026 COMPANY CLAIM

$22M
2027 REVENUE GUIDED IN 2024 COMPANY CLAIM

$40M
FIRST TRANCHE BEING SOUGHT MEDIA-REPORTED

Company

Vyuti Systems Pvt Ltd (CIN U72200KA2015PTC079783) FACT [S15], trading as CynLr, builds a vision-led robotic platform (the CLX vision stack plus the three-armed CyRo manipulator) that picks and handles *unseen, variable* parts without per-part training. Founded by Gokul N A and Nikhil Ramaswamy, both ex-National Instruments machine-vision engineers. Bengaluru lab plus a Swiss R&D centre (EPFL and CSEM collaborations) MEDIA-REPORTED [S09].

Problem

Most industrial robots are blind: parts must arrive fixtured in known positions. In high-mix production (semiconductor equipment, auto Tier-1), part variety makes fixturing and re-programming the dominant cost, so the work stays manual. GM is cited as needing one platform for 22,000+ parts MEDIA-REPORTED [S09].

Commercial evidence

- Three years of filed revenue growth (₹1.1 → 3.1 → ~5.9 Cr) MEDIA-REPORTED [S14].
- Pilots / co-development: GM and DENSO MEDIA-REPORTED; two of the largest semiconductor-equipment makers, one with ~6,000 machines in the field, with 300-400 deployments a year possible after 2028 COMPANY CLAIM [S11].
- ~70% of revenue from semicap customers; customers entirely outside India COMPANY CLAIM [S11].
- Price $150-180k per system plus annual fees COMPANY CLAIM [S12].

Why the thesis likes it

- Owns perception and manipulation, which is the hardest part of the stack and the part platforms are least able to commoditise in brownfield factories.
- Sells to global industrial buyers from India's cost base.
- Fits a ₹25-50 Cr cheque in a round we can help shape.

The problem I cannot explain away

In November 2024, the company guided **$22M of revenue for 2027** COMPANY CLAIM [S10]. In March 2026 it reported **~$1M annualised** COMPANY CLAIM [S11]. A twenty-fold gap 15 months into a 36-month forecast is the most important fact in this file. It does not mean the technology fails; two-year pilots in semicap are normal. It means **management forecasts cannot be used as a base case**, and price must absorb the slippage. INFERENCE

Other risks

- Burn ~₹20 Cr a year against ~$15M raised; the $75M plan is large relative to traction MEDIA-REPORTED [S13, S14].
- No patents surfaced in open searches; IP position unknown.
- Public financial data is internally inconsistent (FY25 revenue ₹5.5 vs ₹6.3 Cr; PAT −₹14 vs −₹17.5 Cr) MEDIA-REPORTED [S14].
- Valuation undisclosed. The founders' ask may sit well above our walk-away price.

14 · SCREEN 3 · CONTINUED

Why it advances to IC despite a 63/100 score

Scorecard summary

Weighted score (robotics weights)	**63 / 100**
Gates	all pass (G1 conditional on price)
Strongest	interface ownership, technology
Weakest	capital intensity, commercial proof, team record on guidance

CynLr scores lower than C2i but passes every gate. It has revenue, named global pilots and an open round. The score says "good, not great". The memo exists to find out whether a structure can make a good company a great *investment*.

Initial return sketch (before the model)

If the semicap pilot converts at even 150 units a year by FY30 at ~$150k, that one account is ~₹190 Cr of revenue. At 5x revenue and 6-7% exit ownership, that alone returns ~₹60-65 Cr on ₹45 Cr, which is not enough unconditionally. The upside lives in a second account converting and in a platform sale. The model on pp.28-34 tests this properly. INFERENCE

What has to be true

BELIEF	HOW WE WOULD KNOW
Semicap pilot converts to production	PO for ≥50 systems within 12 months
Hardware margin reaches ≥45% by ~FY30	BOM roadmap; contract manufacturing quotes
Deployment time falls to days, not months	Customer-site logs for the last 5 installs
The founders can hire a COO-grade operator	Named candidate or a hire before Tranche 2

> Advancing to IC is not a recommendation to invest at any price. It is a recommendation to spend diligence time, because the question can be answered and the answer matters.

PROCEED

Confidence in this decision: 60%

Probability I would make the same call with full diligence access, given what is public today.

TOP 3 REASONS

1. Owns the perception-plus-manipulation interface; sells to global industrial buyers (thesis core).
2. Three years of filed revenue growth and named Fortune-500 pilots: rare at Series B in Indian robotics.
3. An open round where a ₹45 Cr tranched cheque can be meaningful and shape terms.

TOP 3 RISKS TO THIS DECISION

1. Guidance credibility: a ~20x miss against the 2024 forecast.
2. Hardware economics (BOM, field service, inventory) may cap margins below 45%.
3. Price: the founders' ask may exceed our ₹400 Cr walk-away level.

KILL QUESTION

> Will either semicap pilot issue a production purchase order for ≥50 systems within 12 months of our investment?

NEXT DILIGENCE STEP

Reference calls with the semicap OEM's automation lead and DENSO; MCA AOC-4 reconciliation; patent and IP-assignment schedule; last five installation logs.

WHAT WOULD CHANGE MY MIND

- Evidence the semicap pilot is stalled or de-prioritised → Pass.
- Hardware gross margin on current units below 20% with no BOM path → Pass.
- A signed production PO before close → invest unconditionally, at a higher price.

15 · INVESTMENT COMMITTEE MEMO · CYNLR (VYUTI SYSTEMS PVT LTD)

Invest subject to conditions: a tranched ₹45 Cr into the Series B, with a walk-away price

INVEST SUBJECT TO CONDITIONS

Do not write an unconditional cheque at any plausible market price: expected MOIC is 2.45x against a 3.0x hurdle. A tranche of **₹15 Cr at close plus ₹30 Cr on a production PO** lifts expected MOIC to **3.17x** and cuts capital at risk in the failure state from ₹45 Cr to ₹15 Cr. Walk away above **₹400 Cr pre-money**.

Proposed transaction

Instrument	Series B CCPS, 1x non-participating preference, broad-based weighted-average anti-dilution
Round	₹180 Cr (~$20M) equity, plus ~₹60 Cr venture debt for working capital ASSUMPTION
Our cheque	**₹45 Cr:** T1 ₹15 Cr at close; T2 ₹30 Cr at the same price on milestone
Milestone (T2)	Production PO for ≥50 systems from a semicap or auto customer within 12 months
Pre-money	**₹375 Cr** (~$43M); walk-away ₹400 Cr ASSUMPTION
Post-money incl. 5% pool	₹584 Cr
Ownership	7.7% fully diluted after T2 (~2.6% after T1 only)
ESOP	Top-up to a 5% unallocated pool, post-money
Governance	Board seat if we co-lead with ≥₹45 Cr committed; otherwise observer plus full information rights
Pro-rata	Contractual right; exercised at Series C only if the PO milestone is met

Returns at a glance (model, gross)

	BEAR	BASE	BULL
FY32 revenue	₹75 Cr	₹304 Cr	₹677 Cr
Exit equity value	₹208 Cr	₹1,690 Cr	₹5,283 Cr
Ownership at exit	4.9%	6.6%	6.7%
Invested (incl. pro-rata)	₹45 Cr	₹64 Cr	₹76 Cr
Gross MOIC	**0.63x**	**1.73x**	**4.69x**
Gross IRR	-9.6%	13.9%	45.1%

PROBABILITY-WEIGHTED	E[MOIC]	CAPITAL AT RISK
Unconditional ₹45 Cr	2.45x	₹45 Cr
Tranched (recommended)	**3.17x**	**₹15 Cr**

Bear MOIC reflects the 1x preference recovering part of the capital. Exit valued 31-Mar-2032 on FY32 revenue. All valuation figures are proposals or model outputs ASSUMPTION; CynLr's actual round terms are not public.

> **Recommendation in one line:** buy an option on CynLr's semicap conversion, not the company's plan. Pay for the plan only after the PO.

15 · EXECUTIVE SUMMARY

The answers the IC will ask for first

What is the company?

A Bengaluru-built robotic platform that gives robots the ability to see and handle parts they have never encountered, in factories where part variety makes conventional automation uneconomic. Revenue of ~₹5.9 Cr in FY25 MEDIA-REPORTED [S14]; ~$1M annualised by March 2026 COMPANY CLAIM [S11]; customers entirely outside India.

Why now?

Perception compute became cheap enough (inference ~280x cheaper since 2022 FACT [S26]) to run general-purpose vision in a factory cell. Meanwhile, high-mix manufacturers (semicap equipment, auto Tier-1) face labour constraints that fixtured robots cannot address.

Why could it be large?

One semicap customer claims potential for 300-400 deployments a year after 2028 COMPANY CLAIM [S11]. At ~$140k each, that is $42-56M a year from *one* account, more than our entire FY32 base case. A general-purpose manipulation platform proven in semicap is the kind of asset global automation majors buy (ABB Robotics sold for $5.4B FACT [S23]).

Why us?

A ₹45 Cr tranched cheque is meaningful in a ~₹180 Cr equity round. Our likely value-add is Indian supply-chain localisation to lower BOM, introductions to Indian electronics manufacturing for a second market, and structure discipline. We should not claim robotics expertise we do not have. INFERENCE

What can go wrong?

- Pilots never become production (the guidance record argues caution).
- Hardware economics cap margins; field service eats the recurring line.
- Foundation-model robotics players with billions commoditise perception.
- Semicap cyclicality: ~70% of revenue from one industry COMPANY CLAIM.

Why is the price acceptable?

At ₹375 Cr pre-money (~$43M) the tranched structure clears the 3.0x hurdle (3.17x); at ₹450 Cr it does not (2.8x). The unconditional value to this fund is only ~₹260 Cr pre-money, so **the structure, not the headline price, is what makes the deal work**. Ask why the founders would accept: a credible answer is a smaller, faster equity round paired with venture debt; a weak answer is that no one else will lead. That distinction is a diligence item.

What must be true?

1	A semicap or auto pilot converts to a ≥50-unit PO within 12 months **(P = 40%)** ASSUMPTION
2	Hardware gross margin climbs from ~25% to ~50% by FY31 (base)
3	The company reaches ~₹300 Cr revenue by FY32 on ~₹430 Cr of total equity (B + C)
4	A strategic buyer pays ≥5x revenue in 2031-32

> **Highest-conviction view:** the perception problem is real and expensive for customers. **Lowest-conviction assumption:** the 40% conversion probability. Everything on pp.33-34 is most sensitive to it.

15 · COMPANY HISTORY AND FOUNDERS

Eleven years to a commercial platform

2015	Founders leave National Instruments; Vyuti Systems incorporated FACT [S15]
Aug-2019	Seed ₹5.5 Cr (Speciale, Arali, CIIE) MEDIA-REPORTED [S09]
2021	Operations at scale begin (company's account) COMPANY CLAIM [S11]
Apr-2022	$4.5M pre-Series A (Speciale, growX, Info Edge) MEDIA-REPORTED [S09]
Sep-2024	Swiss design & research centre opens (EPFL, CSEM ties) FACT [S09]
Nov-2024	$10M Series A (Pavestone, Athera). Guides $22M revenue by 2027 COMPANY CLAIM [S10]
FY25	Revenue ~₹5.9 Cr; loss ₹14-17.5 Cr MEDIA-REPORTED [S14]
Feb-2026	Object Intelligence platform launched; seeks $75M ($40M first) MEDIA-REPORTED [S12]
Mar-2026	~$1M annualised; ~80 staff; semicap pilots COMPANY CLAIM [S11]
Nov-2026	Round status unknown; this memo INFERENCE

The long arc is not itself a red flag for deep tech: the founders spent 2015-2019 on the vision approach before productising. But it shows that time-to-revenue has been consistently longer than planned, a pattern that runs through the guidance record. INFERENCE

Founders: observable evidence only

DIMENSION	EVIDENCE	READ
Founder-market fit	Both ex-National Instruments machine-vision specialists; Gokul led NI vision work across India, Saudi Arabia and Russia COMPANY CLAIM [S11]	Strong
Technical ability	Own vision stack and multi-arm manipulator; EPFL/CSEM partnerships	Strong
Commercial ability	GM, DENSO and semicap pilots won from India; but no conversion yet	Unproven at scale
Recruiting	60 → ~80 staff across 3 countries; Swiss centre COMPANY CLAIM	Good
Learning velocity	Shift from GM-style auto to semicap where the pull is stronger	Positive signal
Integrity signals	None negative found. Forecasts have been very optimistic (not the same as dishonest)	Test in references
Complementarity	Gokul: design and product; Nikhil: GTM and fundraising MEDIA-REPORTED	Good; no CFO/COO visible
Key-person risk	High: the vision IP and customer relationships are founder-centred	Mitigate with ESOP, insurance, documentation

> I deliberately make no judgment about personality. The only founder judgment I am confident in from public data is **calibration**: the team forecasts ambitiously. The memo prices that rather than penalising the founders for it.

Nikhil Ramaswamy's education and earlier roles: UNKNOWN from public sources. Diligence item.

15 · PROBLEM AND PRODUCT

Why "blind" robots leave high-mix work manual

Who has the problem

Manufacturers with many part variants per line: semiconductor-equipment makers (machine tending, wafer-handling sub-assemblies), auto Tier-1s (kitting, assembly feeding) and electronics. GM is cited as wanting one platform for 22,000+ parts MEDIA-REPORTED [S09].

How painful, how often

Every new part means new fixtures and robot re-programming, or a human. For high-mix lines this happens weekly, so the economics of fixed automation never close. INFERENCE

What customers use today, and its cost

ALTERNATIVE	COST DRIVER
Manual labour	Wages, shifts, quality variance, scarcity in US/EU/Japan
Fixtured robots + 2D/3D vision	Fixtures and integrator programming per part
Bin-picking vision kits	Work for known CAD parts; fail on transparent or reflective parts

Customer ROI is UNKNOWN publicly. Illustration ASSUMPTION: a cell replacing 1.5 operators across 3 shifts in a US plant (~4.5 FTE at ~$60k loaded = ~$270k a year) against a $165k system plus a ~12% annual fee pays back in under 1 year. To be verified with the semicap customer's own business case.

Product architecture



Supply chain dependency
400+ parts from 14 countries (claim): actuators, optics, compute modules (GPU/edge AI from global vendors), sensors. Bengaluru lab with 25 robo

Revenue attaches at three points
System sale ($150-180k, claim) · annual maintenance and development fee on installed base · paid pilots / NRE for new applications

The investment case rests on blocks 1-2. If perception is genuinely general (handles unseen, reflective and transparent parts with no per-part training), CynLr can sell a platform. If it is not, CynLr is a capable custom integrator, and integrators earn integrator multiples. Technology diligence (next page) is designed to tell these apart. INFERENCE

15 · TECHNOLOGY DILIGENCE AND IP

What is genuinely hard, what is commodity, what needs proof

QUESTION	CURRENT EVIDENCE	STATUS
What is technically difficult?	Generalised grasping of unseen, reflective, transparent parts in uncontrolled lighting; hand-eye coordination across 3 arms	Credible problem
What is proprietary?	CLX vision approach (vergence, event imaging); object-learning software; manipulator design COMPANY CLAIM	Plausible; unverified
What is commodity?	Actuators, compute modules, sensors, safety PLCs, integration	Known
Supplier / platform dependence	Global compute and actuator vendors; 400+ parts COMPANY CLAIM	Moderate risk
Can an incumbent replicate?	Fanuc/ABB/Keyence have vision but optimise for structured cells; Alphabet's Intrinsic and NVIDIA's robotics stack target general manipulation as a platform INFERENCE	The main medium-term threat
Benchmarks	Learn unseen object in 10-15 s; new task in ~1 hr COMPANY CLAIM [S12]. No third-party benchmark found	**Needs proof**
TRL	Paid pilots at customer sites; commercial platform launched Feb-2026: TRL 7, approaching 8 INFERENCE	Reasonable
Scaling constraints	Per-site integration time; safety certification 2-3 years COMPANY CLAIM [S11]; field service	Key cost risk

Claims requiring independent verification

1. **First-pass grasp success rate** on a customer-chosen set of 50 unseen parts, including reflective and transparent ones, under site lighting. Pass: ≥95% with no per-part training.
2. **Cycle time** against the incumbent manual process at the semicap pilot.
3. **Mean time between interventions** over the last 3 months of pilot logs.
4. **Deployment time** for the last 5 installations, from site survey to production.

Who verifies: an independent robotics expert (e.g., a manipulation researcher from IISc or the EPFL network the company already works with, with conflicts declared), plus a customer-site visit. Budget: ~₹8-12 lakh and two weeks. ASSUMPTION

Intellectual property

ITEM	FINDING
Patents	**None surfaced** in open searches for Vyuti Systems or CynLr. Request a patent schedule (India, PCT, Swiss filings)
Ownership / assignments	Confirm founder IP from 2015-19 is assigned to the company; confirm Swiss-centre IP is owned by the Indian parent
Trade secrets and datasets	Likely the real moat: deployment data across part types. Request data rights in customer contracts
Licensing / open source	Review OSS licences in the perception stack (copyleft exposure)
Freedom to operate	Screen against the vision-guided manipulation patent families of major incumbents

This is a list of diligence questions, not legal advice. IP counsel would run the FTO.

15 · MARKET

Bottom-up, and anchored on one account

I do not use third-party "$X billion by 2030" robotics reports. The market below is built from robot installations and one customer's own fleet.

Top-down frame from installation data

STEP	VALUE	LABEL
Global industrial robot installations, 2025 forecast	575,000	FACT [S22]
× share of applications blocked by part variety or unseen objects	10%	ASSUMPTION
= addressable cells per year	~57,500	INFERENCE
× system value	$150k	COMPANY CLAIM [S12]
TAM (annual flow)	**~$8.6B**	INFERENCE
× share in CynLr's target verticals and geographies (semicap, auto Tier-1, electronics; US/EU/Japan/Korea)	25%	ASSUMPTION
SAM	**~$2.2B / yr**	INFERENCE
Base-case FY32 revenue	$34.5M	Model
SOM (FY32 share of SAM)	**~1.6%**	INFERENCE

The 10% blocked-application share is the weakest number on the page. It counts applications that are manual today because fixtured automation fails, so the TAM partly reflects demand that does not yet exist, which is penetration risk. Halving it halves TAM and SAM but does not change the SOM conclusion.

Account-level bottom-up (more useful)

Semicap pilot customer's installed machines	~6,000	COMPANY CLAIM [S11]
Deployments a year possible after 2028	300-400	COMPANY CLAIM
× ASP in FY30 (base)	~$146k	Model
Revenue potential, one account	**$44-58M / yr**	INFERENCE

> **Implication:** the base case (35 $M in FY32) does not need the market to be large. It needs one customer to convert partially. That makes the diligence question narrow and answerable, which is why the tranche is tied to it.

Why the market opens now

- Cheap perception compute [S26]; labour scarcity in high-wage geographies.
- Semicap equipment demand linked to AI chip capacity (a tailwind, and also the source of cyclicality).
- India itself is not the market: 9,100 installations in 2024 FACT [S22]. CynLr's export focus is the right shape.

15 · CUSTOMERS AND COMPETITION

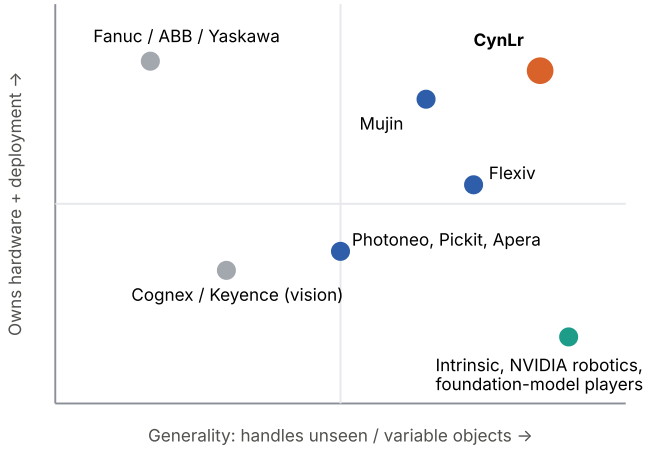
Logos are pilots; competition is coming from above

Customer analysis

ITEM	WHAT IS KNOWN
Named accounts	GM, DENSO (co-development) MEDIA-REPORTED; two large semicap OEMs (unnamed pilots) COMPANY CLAIM
ICP	High-mix manufacturers with labour scarcity: semicap equipment, auto Tier-1, electronics
Procurement	Phased two-year pilots, then production; safety certification 2-3 years COMPANY CLAIM
Sales cycle	~24-36 months to production (inferred) INFERENCE
ACV	$150-180k system + annual fee COMPANY CLAIM; ~12% assumed ASSUMPTION
Concentration	~70% of revenue semicap COMPANY CLAIM; likely 2-3 customers = most revenue INFERENCE
Retention / expansion	**UNKNOWN.** The key diligence metric: systems per site over time
References	Required: semicap automation lead, DENSO, and one lost or stalled pilot

Ask for the lost pilot. The most informative reference is a customer who stopped. It reveals failure modes the winning customers will not mention. INFERENCE

Competitive map



Positions are my inference from public product descriptions, not measured benchmarks.

CATEGORY	PLAYERS	THREAT
Direct	Mujin, Flexiv, Photoneo/Brightpick, Pickit, Apera AI	Medium-high
Incumbents	Fanuc, ABB (now SoftBank), Keyence, Cognex	Medium: they could acquire rather than build
Platform	Alphabet Intrinsic, NVIDIA robotics stack, foundation-model robotics labs	**High, 3-5 yrs**
Build internally	Automakers' and semicap OEMs' automation teams	Medium
Do nothing	Manual labour plus fixtures	The real incumbent

If a platform ships general grasping: CynLr's software differentiation shrinks, and its value shifts to the deployed fleet, the data and the integration. A platform shipping a model does not make it work in a semicap cleanroom. That buys time, not immunity. INFERENCE

15 · BUSINESS MODEL AND UNIT ECONOMICS

A hardware sale with a recurring tail, priced on our assumptions

Who pays, for what

Who pays	Manufacturing operations / automation budget of the OEM
For what	A working cell: system + integration + support
Pricing	$150-180k per system COMPANY CLAIM; annual maintenance and development fees COMPANY CLAIM
Recurring vs one-off	One-off hardware (~85% of base revenue) plus recurring fees (15% of FY32 base revenue)
Services / NRE	Paid pilots and new-application engineering, ₹3-5 Cr a year ASSUMPTION
Usage-based?	Not today. A per-pick or per-uptime model is an upside option to test

Per-system economics, base case, FY30

ASSUMPTION

ASP	₹1.29 Cr (~$146k)
Hardware gross margin	46%
Hardware gross profit per system	₹0.59 Cr
Annual fee (12% of ASP) at 70% GM	₹0.11 Cr / yr
Lifetime gross profit (5-yr service life)	**~₹1.13 Cr**
Customer acquisition cost	UNKNOWN

Customer acquisition cost is dominated by 2-year pilots. A proper CAC is (pilot engineering + sales) / systems eventually ordered. With a 300-unit account it is trivially recovered; with a 5-unit account it is not. That is why account depth is the metric.

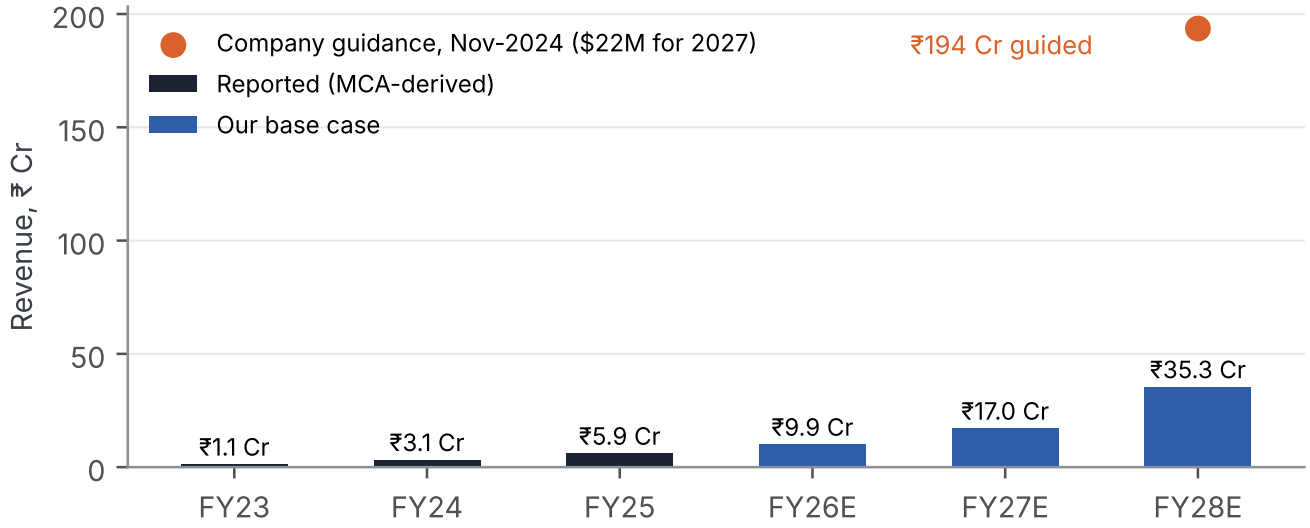
Hardware-specific economics to diligence

LINE	OUR ASSUMPTION	WHY IT MATTERS
BOM and margin	25% → 52% hardware GM, FY26 → FY32 (base)	Low-volume BOM is the largest margin drag
Manufacturing yield and rework	Implicit in GM	Multi-arm calibration at volume
Inventory	120 days of hardware COGS	400+ parts from 14 countries COMPANY CLAIM
Receivables	90 days	Milestone billing at OEMs
Warranty and field service	Inside recurring COGS (30%) and variable opex (8% of revenue)	Can quietly consume the recurring margin
Installation	Inside NRE / pilot revenue	Deployment time is the scaling constraint
Capex	₹3 Cr → ₹18 Cr a year	Demo cells, application labs, tooling
Supplier risk	Single-source actuators and compute	Ask for the top-10 BOM items and alternates
Production ramp	4 → 210 systems a year (base)	Company targets 1 a day by 2028 COMPANY CLAIM; we assume ~3 years later

Where an Indian fund adds real value: localising sub-assemblies (machining, harnesses, enclosures) through Indian EMS partners could take several points off the BOM. That is a value-creation plan, not a slogan. Target: 5 points of hardware margin by FY29. ASSUMPTION

16 · FINANCIAL ANALYSIS · HISTORICAL

Reported numbers, and a guidance record that sets the base case



Revenue: reported versus guided. Reported FY23-FY25 from MCA-derived data MEDIA-REPORTED [S14]; FY25 is the mid-point of two inconsistent figures (₹5.5 and ₹6.3 Cr). Guidance: $22M for 2027 given Nov-2024 COMPANY CLAIM [S10], shown at ₹88/$ ASSUMPTION.

₹ CR	FY23	FY24	FY25	LABEL
Revenue	1.1	3.1	5.5-6.3	MEDIA-REPORTED
Growth	n/a	182%	~90%	INFERENCE
Total expenses	n/a	n/a	20.4	MEDIA-REPORTED
EBITDA	n/a	n/a	−16.5	MEDIA-REPORTED
Net loss	−5.0	n/a	−14.0 / −17.5	MEDIA-REPORTED
Expense / revenue			~3.5x	INFERENCE

No audited statements were available. Condition precedent: reconcile to AOC-4 filings and audited FY26 accounts.

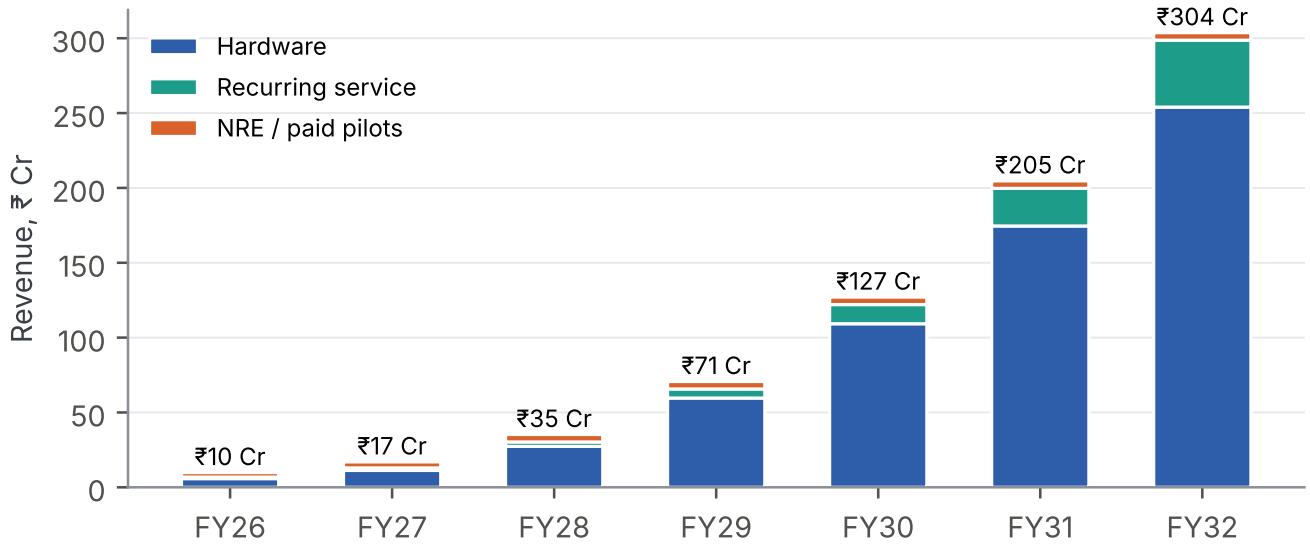
How the guidance record shapes the model

- **Revenue is growing** (~2-3x a year from a small base), so the business is real.
- **Forecasts have been 10-20x too high** over 1-3 year horizons. I therefore start the base case from observed run-rate (~$1M in Mar-2026) and assume the company's 300-units-a-year target arrives ~3 years late (FY31-32), not in 2028.
- **Burn has been modest relative to ambition** (~₹20 Cr a year): capital-efficient by robotics standards. That is a genuine positive that the guidance gap can obscure.

> The founders' plan is not my base case; it is close to my bull case (918 cumulative units by FY32 against 518 in base).

16 · FORECAST

Revenue is built from units, prices and installed base, not a growth curve



Base-case revenue build. Hardware = systems shipped × ASP; recurring = opening installed base × 12% fee × ASP; NRE from paid pilots. All drivers are on the Assumptions sheet.

BASE CASE	FY26	FY27	FY28	FY29	FY30	FY31	FY32
Systems shipped	4	8	20	45	85	140	210
Revenue, ₹ Cr	9.9	17.0	35.3	70.7	127.1	204.7	303.7
Gross margin	37%	39%	41%	45%	49%	52%	55%
EBITDA, ₹ Cr	-20.6	-24.7	-29.2	-30.1	-23.9	-6.4	19.1
Op.+inv. cash flow, ₹ Cr	-25.1	-33.0	-43.3	-51.5	-53.4	-44.6	-29.3
Closing cash, ₹ Cr	30	233	181	371	291	222	171

Every major assumption, and why

DRIVER	BASE	WHY
Units FY26 → FY32	4 → 210	FY26 anchored to the ~$1M run-rate; 300/yr company target delayed ~3 yrs
ASP	$165k, −3%/yr	Mid of claimed range; erosion from competition and volume
Recurring fee	12% of ASP	Typical industrial service contract
Hardware GM	25% → 52%	Company targets 60%; reached only in bull
Headcount	80 → 320	~80 in Mar-2026 COMPANY CLAIM; ~₹0.23 Cr loaded cost, +8%/yr

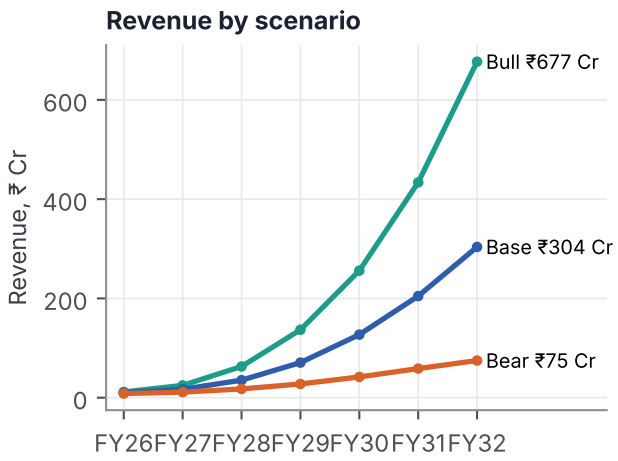
Cash and funding needs

Base cumulative pre-financing cash burn FY26-FY32 is ~₹280 Cr, including ~₹90 Cr of working-capital build. It is funded by the ₹180 Cr Series B, ₹60 Cr of venture debt (repaid FY30-32) and a ₹250 Cr Series C in FY29 at ₹2,000 Cr pre-money ASSUMPTION. Minimum year-end cash is ₹30 Cr (FY26, before the Series B closes), so **the round is needed within months**. EBITDA turns positive in FY32 in base, FY30 in bull, and never within the horizon in bear.

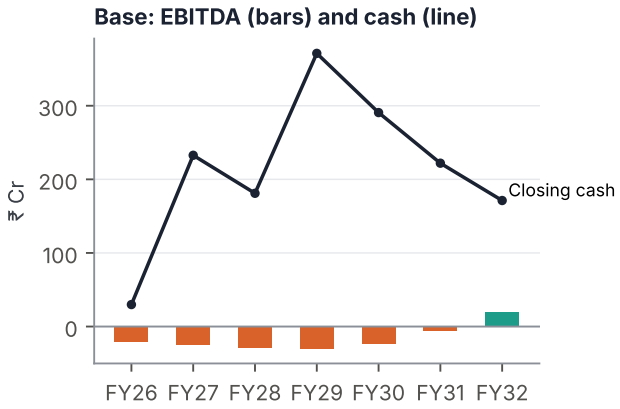
Total capital before profitability in base: ~₹430 Cr equity + ₹60 Cr debt (~$56M), close to the company's own $75M plan in total but spread over more time.

16 · SCENARIOS, BURN AND RUNWAY

Three worlds, one structure



Bear assumes neither pilot converts at scale; bull assumes the semicap account converts (380 units in FY32).



Base EBITDA and closing cash. The Series B (FY27) and Series C (FY29) are visible as cash steps.

METRIC	BEAR	BASE	BULL
FY27 revenue	₹11 Cr	₹17 Cr	₹25 Cr
FY29 revenue	₹28 Cr	₹71 Cr	₹137 Cr
FY32 revenue	₹75 Cr	₹304 Cr	₹677 Cr
Revenue CAGR FY26→FY32	44%	77%	97%
FY32 gross margin	43%	55%	63%
FY32 EBITDA margin	-71%	6%	31%
Cumulative units FY26-FY32	177	518	918
Lowest closing cash FY26-FY32	₹21 Cr	₹30 Cr	₹36 Cr

What separates the scenarios

DRIVER	BEAR	BASE	BULL
FY32 units	60	210	380
ASP change / yr	−6%	−3%	0%
FY32 hardware GM	40%	52%	60%
Series C	₹150 Cr down round	₹250 Cr at ₹2,000 Cr	₹400 Cr at ₹4,000 Cr
Exit EV / revenue	2.5x	5.0x	7.0x

Runway reading

In every scenario the company stays cash-positive only because the Series B closes in FY27. Bear burns roughly ₹45-55 Cr a year with revenue of ₹10-75 Cr. By FY30 it would need either a recapitalisation or a sale, which is what the 1x preference protects against. Monthly pre-financing burn in base peaks at ~₹3.9 Cr in FY30. INFERENCE

> The bear case is not a zero: the preference recovers ~63% of capital in a ₹208 Cr sale. A true zero (a failed company) is the separate "failure" outcome on p.33.

16 · VALUATION

Three methods, one range, one proposed price

1 · Exit comparables (used to set exit multiples)

COMPARABLE	EV / REVENUE	USE
ABB Robotics → SoftBank (2025)	2.4x	Floor: mature arm maker FACT [S23]
Cognex (Q1-2026)	7.7x	Vision-led automation MEDIA-REPORTED [S24]
Teradyne (Q1-2026)	12.2x	Excluded: semi-test cycle [S25]
Indian listed electronics / defence	70-160x P/E	Excluded: profit-based, fragile [S41]

Base 5.0x sits between a strategic sale of a mature business and listed machine vision, for a company still growing ~50% with ~55% gross margin in FY32. Bear 2.5x; bull 7.0x. INFERENCE

2 · Venture method: what is it worth *to this fund*?

Probability-weighted exit equity (blended)	~₹1,830 Cr
Retention of ownership to exit (no pro-rata)	75.9%
Target expected MOIC	3.0x
Maximum pre-money, unconditional	**₹260 Cr (~$30M)**
Maximum pre-money if the milestone is already met	**₹561 Cr**

3 · Scenario-weighted structure (the deciding method)

The tranche lets us pay a conditional price for conditional value. The walk-away grid on p.33 finds the highest pre-money at which the tranched structure still clears 3.0x: **₹400 Cr**.

Valuation range and proposal

REFERENCE POINT	PRE-MONEY
Unconditional value to this fund	₹260 Cr
Proposed (tranched)	**₹375 Cr**
Walk-away (tranched)	₹400 Cr
Value if the milestone is already met	₹561 Cr
Inferred Series A post-money (not disclosed)	~₹350-400 Cr ASSUMPTION

₹375 Cr is roughly a flat round on my *assumed* Series A price. Given the guidance gap, I think a flat round is fair, but founders and existing investors may disagree. If the market price is higher, we decline and re-engage after the PO, paying more for less risk.

Why no DCF

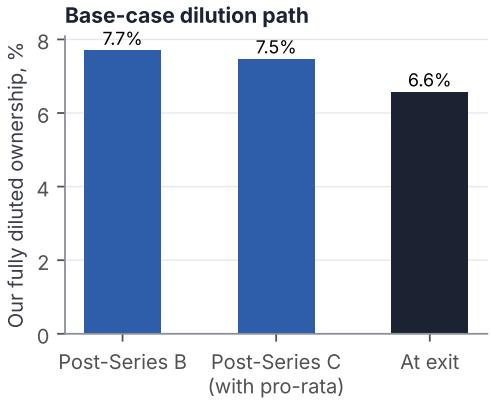
Over 90% of value lies beyond the forecast and hinges on one binary event: pilot to production. A DCF would hide that binary inside a terminal growth rate and present it as precision. INFERENCE

16 · CAP TABLE, OWNERSHIP AND DILUTION

From 7.7% to 6.6%: following our shares to exit

HOLDER (BASE CASE)	PRE-B	POST-B	POST-C	EXIT
Founders	52.0%	33.4%	28.8%	25.3%
Seed / pre-A investors	18.0%	11.6%	10.0%	8.8%
Series A investors	20.0%	12.8%	11.1%	9.7%
ESOP (existing)	10.0%	6.4%	5.5%	4.9%
Series B: our fund	0.0%	7.7%	6.6%	5.8%
Series B: other investors	0.0%	23.1%	19.9%	17.5%
ESOP top-up at B	0.0%	5.0%	4.3%	3.8%
Series C: new investors	0.0%	0.0%	9.9%	8.8%
Series C: our pro-rata	0.0%	0.0%	0.8%	0.7%
ESOP top-up at C	0.0%	0.0%	3.0%	2.6%
Pre-exit round (Series D / IPO)	0.0%	0.0%	0.0%	12.0%
Total	100.0%	100.0%	100.0%	100.0%

Pre-round splits (founders 52%, seed/pre-A 18%, Series A 20%, ESOP 10%) are ASSUMPTIONS: CynLr's cap table is not public. Ownership maths for new money is exact given the price; absolute share counts in the model are indexed (1,000,000 pre-B fully diluted shares).



Our position through the base case: Series B, Series C with ₹19 Cr pro-rata, a 12% pre-exit round and 3% ESOP top-up at C.

Mechanics

- **Series B price** = ₹375 Cr / 1,000,000 shares; ₹180 Cr raised; 5% post-money pool top-up → post-money ₹584 Cr.
- **Series C (FY29, base):** ₹250 Cr at ₹2,000 Cr pre-money = 3.4x step-up on our price. We take full pro-rata (₹19 Cr) because the milestone is assumed met in base.
- **Bear Series C** is a down round (0.86x our price). The model applies no anti-dilution adjustment, which is conservative; broad-based weighted-average protection would lift bear ownership modestly.

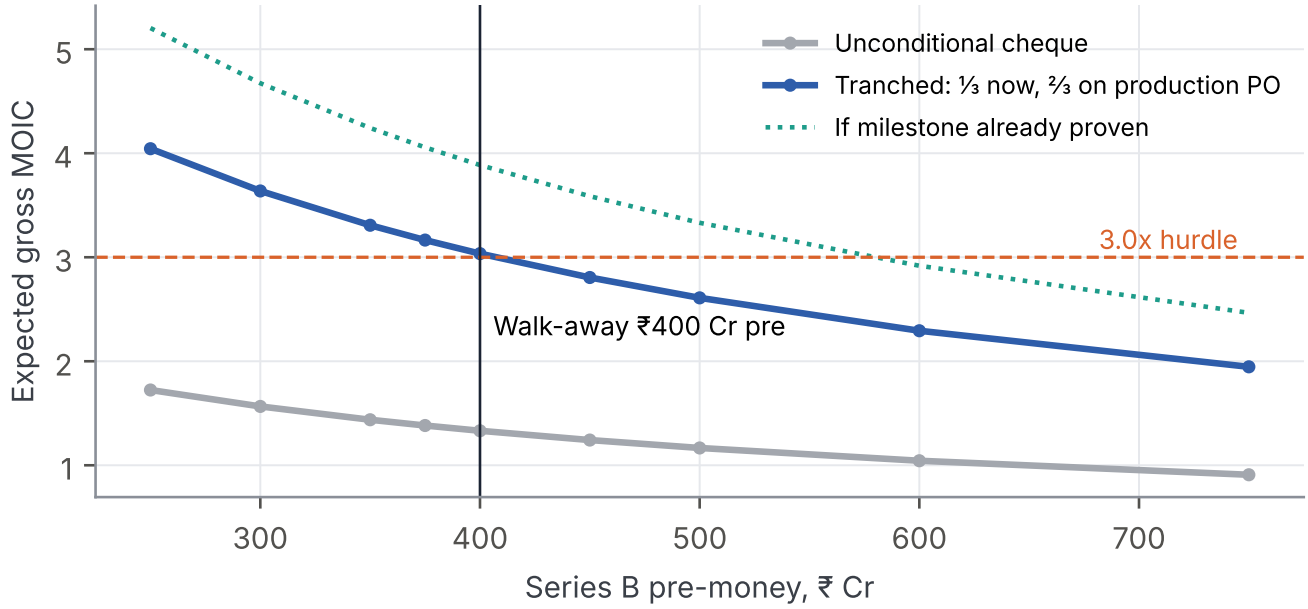
Founders after dilution

Founders hold ~33% post-B and ~25% at exit in base, which is enough to stay motivated. I would resist any structure that pushes founders below ~20% before Series C, because founder incentive is the real collateral in a two-year pilot cycle. INFERENCE

> Pro-rata is a *separate* investment decision taken at Series C on the evidence then. It is not a promise made today.

16 · RETURNS AND EXPECTED VALUE

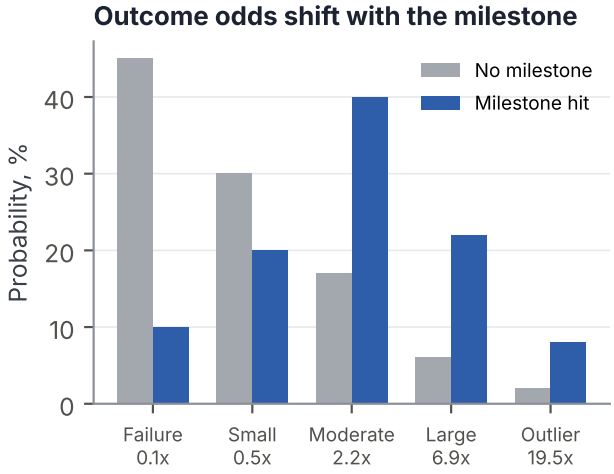
The tranche, not the haggle, makes this investable



Expected gross MOIC against Series B pre-money. The unconditional cheque never clears 3.0x in the plausible price range. The tranched structure clears it up to ₹400 Cr. Probabilities and outcomes below; every point is recomputed in the model's Expected_Value sheet.

OUTCOME	EXIT EQUITY	P NO MS	P MS	P BLEND	MOIC
Failure / asset sale	₹60 Cr	45%	10%	31%	0.14x
Small (bear case)	₹208 Cr	30%	20%	26%	0.48x
Moderate (base case)	₹1,690 Cr	17%	40%	26%	2.19x
Large (bull case)	₹5,283 Cr	6%	22%	12%	6.86x
Outlier (strategic platform sale)	₹15,000 Cr	2%	8%	4%	19.48x
Expected MOIC		**1.38x**	**4.06x**	**2.45x**	

MS = milestone (production PO ≥50 units within 12 months), P = 40% ASSUMPTION. MOIC per ₹ of Series B cheque, including the 1x preference floor; preference stack held at the base-case ₹430 Cr (conservative vs the bear sheet). Failure = ₹60 Cr asset sale; outlier = ₹15,000 Cr (~$1.7B) strategic platform sale ASSUMPTION.



STRUCTURE	E[INVESTED]	E[PROCEEDS]	E[MOIC]
Unconditional ₹45 Cr	₹45 Cr	₹110 Cr	2.45x
₹15 Cr + ₹30 Cr on MS	**₹27 Cr**	**₹85 Cr**	**3.17x**

16 · FUND-LEVEL MATH AND SENSITIVITIES

Can this position matter to the fund?

TO RETURN...	₹600 CR FUND	₹800 CR FUND	VS BASE EXIT
0.25x fund	₹2,281 Cr	₹3,042 Cr	1.4x
0.5x fund	₹4,563 Cr	₹6,084 Cr	2.7x
1.0x fund	**₹9,125 Cr**	**₹12,167 Cr**	**5.4x**

At 6.6% exit ownership, CynLr returns the fund only at a ~$1,037M exit: the outlier case, roughly 5x our base exit. A 0.25x-fund outcome (~$259M) sits between base and bull. **This is an option position, not a fund-maker by plan.** INFERENCE

Concentration check

Full T1 + T2 + base pro-rata = ₹64 Cr, 13.4% of ~₹480 Cr investable capital: **above my 10-12% cap**. Policy: cap total exposure at ~₹55 Cr, so pro-rata at Series C is partial unless the bull case is visible. I flag this rather than hide it inside the base case. INFERENCE

MOIC: entry pre-money × exit equity (₹ Cr)

PRE ↓ / EXIT →	₹1,000	₹2,000	₹3,500	₹6,000	₹10,000	₹15,000
₹300 Cr	1.5x	3.0x	5.3x	9.0x	15.0x	22.5x
₹450 Cr	1.1x	2.3x	4.0x	6.9x	11.4x	17.2x
₹600 Cr	0.9x	1.8x	3.2x	5.5x	9.2x	13.9x
₹750 Cr	0.8x	1.6x	2.7x	4.7x	7.8x	11.6x
₹900 Cr	0.7x	1.3x	2.3x	4.0x	6.7x	10.0x

MOIC: FY32 revenue × exit multiple (at ₹375 Cr pre)

REV ↓ / EV/REV →	2.5X	4.0X	5.0X	7.0X	10.0X
0.50x base	0.5x	0.8x	1.0x	1.4x	2.0x
0.75x base	0.7x	1.2x	1.5x	2.1x	3.0x
1.00x base	1.0x	1.6x	2.0x	2.8x	3.9x
1.50x base	1.5x	2.4x	3.0x	4.1x	5.9x
2.00x base	2.0x	3.2x	3.9x	5.5x	7.9x

Series B cheque only, base dilution, conversion case. Green ≥3x, amber 1.5-3x, red <1.5x. A ≥3x outcome needs roughly 1.5x base revenue at a 7x+ multiple: the bull world.

17 · PROPOSED TERMS

Terms that matter economically, and those that are hygiene

Terms that change the outcome

TERM	PROPOSAL	WHY IT MATTERS
Tranche	₹15 Cr at close; ₹30 Cr at the same price on a ≥50-unit production PO within 12 months; our option to waive the condition	Moves E[MOIC] from 2.45x to 3.17x; the single most valuable term
Price	₹375 Cr pre; walk away above ₹400 Cr	Derived from fund math (p.33)
Liquidation preference	1x, non-participating, pari passu with other preferred	Recovers ~63% in the bear sale. Participating would be off-market and damage alignment
Anti-dilution	Broad-based weighted average	Protects against a bear-case down round without punishing founders
Pro-rata	Contractual right through Series C	Optionality to concentrate if the bull case emerges
ESOP	5% post-money top-up	Hiring a COO and US sales leads; post-money pool avoids pricing friction

Hygiene terms

Information rights	Monthly MIS (cash, pipeline, units, installs), quarterly board pack, audited annuals within 120 days, annual budget
Board	Seat if co-leading; otherwise observer
Reserved matters	Standard: new securities, debt above a threshold, related-party transactions, change of business, IP transfer, Swiss/US subsidiary capital moves
Founder terms	Updated employment and non-compete; 4-year vesting on any new founder grants; no re-vesting of earned stock
ROFR / co-sale / tag	Standard for preferred holders
Drag-along	After year 6, at ≥2x the Series B price, with majority-preferred consent

Conditions precedent

1. FY24-FY26 financials reconciled to MCA AOC-4 filings and audited FY26 accounts; one FY25 revenue figure confirmed.
2. Written confirmation from the semicap pilot customer of pilot status and decision timeline (or a reference call with their automation lead).
3. IP schedule; founder and employee invention assignments; Swiss-centre IP owned by the Indian parent.
4. FEMA / ODI compliance for Swiss and US subsidiaries.
5. Export-control review for semicap end-use customers.
6. Key-person insurance on both founders.
7. Agreed 18-month operating plan with headcount gates tied to the PO.

> **What I would concede:** board seat (if we are not co-lead), drag threshold timing and observer access, in exchange for keeping the tranche and the price cap. **What I would not concede:** the tranche condition and the walk-away price.

Why the founders might say yes

It closes the round faster with a credible co-investor, the tranche is at the same price (no ratchet), and venture debt reduces equity dilution. If they say yes too easily, we should ask whether better-informed investors have declined. Adverse selection is part of reading a deal. INFERENCE

17 · DILIGENCE PLAN

Twelve questions, each with an owner and a threshold

AREA	QUESTION	WHY IT MATTERS	EVIDENCE NEEDED	WHO VERIFIES	PASS / FAIL THRESHOLD
Technical	Does grasping generalise to unseen, reflective and transparent parts without training?	Platform value vs integrator	Blind test on 50 customer-chosen parts; pilot logs	Independent robotics expert + site visit	≥95% first-pass; no per-part training
Technical	How long does deployment take?	Scaling constraint	Last 5 installs: survey → production dates	Customer + company logs	≤8 weeks by FY27
Commercial	Is the semicap pilot heading to production?	The kill question	Pilot plan, success criteria, budget owner	Reference call with automation lead	Named decision date ≤12 months
Customer	Why did any pilot stop?	Reveals failure modes	Lost or stalled pilot reference	Sathvik + partner	Failure not perception-related
Financial	What are true FY25-26 numbers and burn?	Data conflicts in public sources	AOC-4, audited FY26, bank statements	Fund CFO / auditor	Reconciles within 5%
Financial	Unit economics per system	Margin path	BOM by system, service cost per site	Fund + external hardware cost expert	Current HW GM ≥20%; credible path to 45%
Legal	Corporate hygiene and subsidiaries	Structure risk	Charter, SHA, FEMA/ODI filings, Swiss entity docs	Legal counsel	No material defects
IP	Who owns the IP; any FTO issues?	Moat and exit value	Patent schedule, assignments, OSS audit, FTO screen	IP counsel	IP in the Indian company; no blocking patents
Founder	Calibration and integrity	Guidance gap	Ex-colleague, investor and customer references; past board decks vs outcomes	Partner + Sathvik	Optimism, not misrepresentation
Market	Is 300-400 units/yr at one OEM real?	Underpins the bull case	OEM fleet data; automation budget	Customer + industry expert	Supported by the OEM's own plans
Cyber / data	Data rights and security in customer fabs	Moat (data) and adoption	Contract data clauses; security architecture	Security reviewer	Company retains model-improvement rights
Regulatory	Safety certification and export controls	Time-to-revenue; end-use	ISO 10218 / regional certification plan; export classification	Specialist counsel	Plan within 24 months; no end-use block

Sequencing (four weeks)

Week 1: financial reconciliation and the semicap reference call (cheapest kills first). Week 2: technical blind test and site visit. Week 3: IP, legal, FEMA. Week 4: unit economics and terms. Stop at the first failed threshold. ASSUMPTION

Budget and people

~₹20-30 lakh external (robotics expert, IP and legal counsel, hardware-cost advisor), plus ~3 weeks of my time and ~5 days of a partner. Small relative to a ₹45 Cr decision. ASSUMPTION

17 · RED TEAM

The case against the investment

Written as the most sceptical partner would argue it, followed by my response. No strawmen.

OBJECTION	THE SCEPTIC'S ARGUMENT	MY RESPONSE
The founders sell a vision, not a business.	A ~20x miss against a 3-year revenue forecast, made at the Series A, is not a rounding error. Investors who priced that forecast have lost money on paper; we would be next.	Agreed on the facts, which is why the base case ignores the forecast and starts from run-rate, and why ⅔ of our capital waits for a PO. We are not buying the forecast.
It is a hardware company dressed as a platform.	ABB's robotics business sold for ~2.4x revenue. Hardware margins, service costs and inventory will make CynLr look like ABB, not Cognex.	Possibly. The base uses 5x, bear 2.5x. Even at 2.5x on base revenue the 1x preference limits the loss. The sensitivity table shows the bull world (7x+) is needed for 3x: priced, not assumed.
Better-funded platforms will commoditise perception.	Foundation-model robotics players and Big Tech robotics stacks have 100x CynLr's capital. In three years, general grasping is an API.	The strongest objection. My answer: deployment in semicap cleanrooms is the moat, not the model, and CynLr may become an acquisition target for exactly those platforms. If a platform achieves ≥95% general grasping in brownfield cells, I would revisit (falsifier on p.5).
Semicap concentration is cyclical.	~70% of revenue from one cyclical industry; a capex downturn freezes pilots for 18 months.	Real. The tranche protects ⅔ of capital; we would want a second vertical (auto Tier-1 via DENSO) in the 18-month plan.
The position is too big for an option.	₹64 Cr with pro-rata is 13% of investable capital for a base case of 1.7x.	Correct, and the memo caps total exposure at ~₹55 Cr. Pro-rata is a Series C decision on Series C evidence.
If they accept our price, we are being adversely selected.	A strong company does not take a flat round with a tranche. Acceptance itself is a negative signal.	A fair point, and a diligence step: ask who else was approached and why they declined. A credible reason is speed and debt-equity mix; an incredible one ends the process.

> **Verdict of the red team:** objections 1, 3 and 5 are not fully answered. They are *contained* by the tranche, the price cap and the exposure cap. If the IC believes objection 3 will play out within five years, the right decision is to pass, and I would not argue hard against that.

17 · PRE-MORTEM

It is 2031. The investment returned 0.3x. What happened?

FAILURE PATH	THE STORY	PROBABILITY
Pilot never converts	Semicap capex downturn in 2027; pilots extended; the PO never arrives. T2 not drawn, T1 marked down.	15%
Platform commoditisation	A foundation-model robotics stack plus commodity arms reaches 'good enough' grasping; CynLr's pricing collapses to integrator levels.	10%
Hardware margin trap	Field failures and service visits eat the recurring line; hardware GM plateaus near 30%; every unit sold deepens burn.	10%
Capital crunch	Series C fails in FY29; insiders recapitalise at a deep discount; preferences wiped by new senior money.	8%
Exit market closes	Company reaches ~₹150 Cr revenue but no strategic buyer pays; listed small-cap multiples compress; sold at ~1.5x revenue.	5%
Key person / team	A founder departs; the US sales team does not scale; customer relationships were founder-held.	6%
IP / legal	Core IP sits partly in founder names or the Swiss entity; FTO dispute with an incumbent during the exit process.	3%
Total probability of a sub-1x outcome		**57%**

The total (57%) is consistent with the blended "failure + small" probability in the expected-value table (57%). The paths overlap, so the split is judgmental. INFERENCE

Early-warning indicators to monitor

- No written decision date from the semicap OEM by month 6.
- Hardware GM on new units not improving by 3+ points per half-year.
- Installs taking longer than 12 weeks.
- Senior hires in US sales missing plan by two quarters.
- A platform announcement with brownfield deployment claims at a named OEM.

What we would do

- Do not draw T2; preserve the ₹30 Cr.
- Push a pivot to the strongest vertical (whichever pilot is closest to production).
- Explore an early strategic sale while the team and IP are intact. The preference makes a ₹200-400 Cr sale recover most of T1.

18 · INVESTMENT COMMITTEE RECOMMENDATION

CynLr Series B: invest subject to conditions

Decision	INVEST SUBJECT TO CONDITIONS
Cheque	₹45 Cr: ₹15 Cr at close, ₹30 Cr on milestone; total exposure capped at ~₹55 Cr including pro-rata
Price	₹375 Cr pre-money (~$43M); walk away above ₹400 Cr
Ownership	7.7% post-B fully diluted (both tranches); ~6.6% at exit in base
Key thesis	Owns the perception-plus-manipulation interface; global high-mix buyers; one account can carry the base case
Key risks	Guidance credibility; conversion; hardware margins; platform commoditisation; semicap cyclicality
Required diligence	12 items (p.36); kill items first: semicap reference and financial reconciliation
Conditions precedent	MCA reconciliation; customer confirmation; IP assignments; FEMA/ODI; export-control review; key-person cover
Expected return	E[MOIC] 3.17x tranched; base 1.73x / 13.9% IRR; bull 4.69x / 45.1%
Downside	Bear 0.63x via preference; failure ~0.1x on T1 only (₹15 Cr at risk)

What must be true

1. The semicap (or auto) pilot converts to a ≥50-unit PO within 12 months.
2. Hardware gross margin reaches ~45% by FY30.
3. A strategic buyer pays ≥5x revenue in 2031-32.
4. The founders hire an operator-grade COO before T2.

What would make us walk away

- Pre-money above ₹400 Cr, or refusal of the tranche structure.
- Semicap reference indicates pilot de-prioritised or no budget owner.
- Financials do not reconcile within 5%, or misstatement is found.
- Core IP not owned by the Indian company.
- Blind technical test below 90% first-pass success.

> **Vote requested:** approval to spend four weeks and ~₹25 lakh on diligence and to issue a non-binding term sheet on the terms above. Final approval to follow diligence.

19 · DECISION REVIEW

How the view changed, and why

Initial view

CynLr looked like the cleanest expression of the thesis: proprietary perception, a Fortune-500 logo list, export revenue and capital efficiency. I expected to recommend a straightforward participation at whatever the market price turned out to be.

Initial hypothesis

"Indian physical-AI company with global pilots, early in a large market. Underwrite at ~$60-80M pre-money; base case 3-4x."

Evidence that strengthened the case

- Three consecutive years of filed revenue growth (rare in Indian robotics at this stage).
- The semicap pilot's scale: one account could exceed our FY32 base revenue.
- Capital efficiency: a multi-arm platform plus vision stack on ~$15M.
- Customers entirely export; sales to sophisticated buyers.

Evidence that weakened the case

- **The guidance gap** ($22M guided vs ~$1M actual run-rate).
- No patents found in open searches.
- Inconsistent public financial data.
- Logos turned out to be pilots and co-development, not purchase orders.
- Semicap concentration (~70%) in a cyclical industry.

What surprised me

That the base case returns the money but not a venture multiple even at a flat price. I had assumed price negotiation would fix the returns. It does not: the unconditional value to this fund is only ~₹260 Cr pre-money, below any realistic founder ask.

What I almost missed

Two things. First, the structure: I spent a day refining exit multiples before realising the tranche moved expected MOIC more (2.45x → 3.17x) than any plausible multiple change. Second, the portfolio conflict: four companies in my universe are likely already owned by the employer's Fund I. I had initially shortlisted one (Coratia) on its Navy contract.

What changed the decision

The walk-away grid. It turned "proceed at market price" into "proceed only with a tranche, at ≤₹400 Cr". That is a different recommendation, and a more defensible one.

What we still don't know

- The real valuation ask and who is leading.
- Current hardware gross margin.
- Whether the semicap pilot has a budget owner and a decision date.
- IP ownership and patents.

Where we may be wrong

The 40% conversion probability. Raise it to 60% and the unconditional cheque nearly clears; drop it to 25% and even the tranche fails. That one number deserves most of the diligence budget.

19 · DECISION REVIEW · CONTINUED

What would change my mind, and what this teaches

What would change our mind

EVIDENCE	NEW DECISION
Signed production PO before close	Invest unconditionally; accept up to ~₹550 Cr pre
Pilot de-prioritised or no budget owner	Pass
Blind grasp test <90%	Pass: integrator, not platform
Current hardware GM >35%	Raise base margins; price cap rises
Platform player shows brownfield general grasping at a named OEM	Pass or re-underwrite as an acquisition target only

What this teaches for future deals

1. **Build the walk-away grid first.** It tells you whether to spend time on price or on structure.
2. **Treat guidance history as a data series,** and give it a column in every screen.
3. **Check the employer's portfolio before shortlisting.** Conflicts are a gate, not a footnote.
4. **Separate "good company" from "good investment" with gates,** or the score will win the argument.

My bias check

I come from the founder side. These are the places that background most likely pushed my judgment, and what I did about each.

BIAS	HOW IT SHOWED UP	CORRECTION
Founder empathy	My first instinct was "two-year pilots are normal; the founders will get there"	Priced the guidance gap; tranche on POs, not on time
Overvaluing product velocity	Impressed by the 10-15 s learning claim	Made it a blind test with a pass threshold
"They'll figure it out"	Assumed margins improve with scale	Base hardware GM 52% vs the company's 60%; bear 40%
Underweighting capital structure	First draft ignored the preference stack	Modelled 1x non-participating, pari passu; bear recovers via the preference
Ignoring fund math	First draft sized the cheque at ₹50 Cr with full pro-rata	Fund_Math sheet; exposure cap ~₹55 Cr
Charisma by proxy	Strong founder interviews in the press	Used interviews only as COMPANY CLAIM; references required
Liking what I understand	As an AI-product builder, I find the perception software intuitive	Scored commercial evidence and hardware economics heavily (robotics weights)

Bias I may not have corrected: a preference for interesting technology over boring, cash-generative businesses. A partner from a PE background would probably weight Ati Motors' revenue more heavily than I did.

20 · SOURCES AND METHODOLOGY

What this rests on

ID	CLAIM (ABRIDGED)	SOURCE	DATE	LABEL
S04	Bharat Tech Fund: ₹600 Cr + ₹200 Cr greenshoe Cat II AIF; Series A/B; ₹25-50 Cr cheques; 30% gr...	Inc42	2026-05-20	FACT
S10	CynLr guided $22M revenue by 2027 (Nov 2024)	Inc42	2024-11-05	COMPANY CLAIM
S11	CynLr ~ $1M annualised revenue Mar-2026; ~$600k in prior 6 months; ~70% from semicap customers;...	Inc42	2026-03-12	COMPANY CLAIM
S12	CynLr pricing $150-180k per system + annual maintenance/development fees; Object Intelligence p...	Business Today	2026-02-13	COMPANY CLAIM
S14	CynLr (Vyuti Systems) revenue FY23 ₹1.1 Cr, FY24 ₹3.1 Cr, FY25 ₹5.5-6.3 Cr; FY25 loss ₹14-17.5 ...	Inc42 company page (MCA-derived)	accessed 2026-11-02	MEDIA-REPORTED
S16	Simplismart $14M Series B led by Dallas VC with NVIDIA, Accel; ~$100M post; InVideo 56% GPU-cos...	YourStory	2026-10	MEDIA-REPORTED
S18	C2i $15M Series A led by Peak XV with Yali, TDK Ventures	Entrackr	2026-02-16	MEDIA-REPORTED
S21	AI data-centre power semiconductor market ~$1.5B (2025) → ~$2.5B (2026) → ~$3.8B (2028)	UBS via Rallies.ai	2026-03-08	MEDIA-REPORTED
S22	Global industrial robot installs 542k (2024); China 295k (54%); India 9,100 (+7%); 575k forecas...	IFR World Robotics 2025 press release	2025-09-25	FACT
S23	ABB Robotics sale to SoftBank $5.375B; 2024 revenue $2.28B; ~7,000 staff	TechRepublic	2025-10	FACT
S26	Inference cost of GPT-3.5-level output fell ~280x (Nov-2022 to Oct-2024)	Stanford HAI AI Index 2025 (via Tom's Hardware)	2025-04	FACT
S28	Neysa: Blackstone-led up to $600M equity + $600M intended debt; majority stake	Blackstone press release	2026-02	FACT
S32	ISM 1.0 ₹76,000 Cr; ISM 2.0 ₹1.27 lakh Cr approved Jul-2026	Angel One (Cabinet/PIB)	2026-07-15	FACT
S37	MoD capital acquisition FY27 BE ₹1,85,467 Cr (+24.6%), 75% earmarked domestic	MP-IDSA budget analysis	2026-02-13	FACT

Key sources only. All 54 sources with URLs, access dates, reliability scores and evidence are in the source-log workbook.

Method

- **Research window:** 2 Nov 2026, public sources only. No management access, no paid databases (Tracxn, PitchBook and Tofler Pro were unavailable).
- **Source hierarchy:** filings and government data first, then company statements, then media. Anything traceable to a single founder interview is logged once, as a company claim.
- **Model:** 1,583 live formulas, zero errors after recalculation; cap tables reconcile to the share; inputs labelled.
- **Employer:** identified by inference (~60-65% confidence). The analysis is calibrated to the mandate archetype, not to private knowledge of any firm.

Limitations

Valuations are proposals; cap tables are assumed; probabilities are judgments. The value of this document is the reasoning it makes checkable, not the precision of its numbers.

> **Thank you for reading.** I would welcome the chance to defend, and be corrected on, any page of this.