



One AI layer built once and reused across the portfolio, technical due diligence on AI claims, and evidence that holds up at exit

Vikram Jha · Principal FDE · AI Architect · Builder · 15+ years in production AI · Co-Founder at a US-based native-AI startup in stealth mode, working on AI governance, agentic security and context graphs.

WHY NOW

A 2026 survey of 200 fund and operating leaders finds only 7 percent of portfolio companies running AI at enterprise scale, with talent the top barrier ([FTI Consulting survey](#)). Regulators are also acting on inflated AI claims ([FTC](#)), which makes diligence on a target's AI a valuation question.

WHAT I DO

I work inside your team and build the whole stack with your engineers (context, models, agents, evaluation, security, deployment), prove it against your own baseline, and leave it running in your environment, owned by your people. If AI is not the answer, I say so in writing.

THE FOUR DELTAS, EACH WITH THE NUMBER THAT PROVES OR DISPROVES IT

1 One owner, full stack

Architecture, build and governance in one accountable person; no handoff loss

MEASURED BY Handoffs removed; days from question to decision

2 An evaluated slice in weeks

A working, measured slice of one workflow in 3-4 weeks of data access

MEASURED BY Weeks to a slice that meets its written threshold

3 Governance built in

Authority bands, permission receipts, red-team and a tested kill path before go-live

MEASURED BY Share of agent actions with an authority receipt

4 Open-source cost and sovereignty

Open-weight models wherever blind evaluation shows parity; data in-region

MEASURED BY Cost per resolved case against the frontier baseline

WHERE TO START

A FIRST SLICE I WOULD PROPOSE	MEASURED BY	SHAPE
AI technical due diligence on a live target	Findings by severity	Production slice sprint (4-8 weeks)
Portfolio AI baseline sprint	Companies scored	Production slice sprint (4-8 weeks)
Shared layer plus two company slices	Time to the second slice	Production slice sprint (4-8 weeks)

WHAT YOU CAN HOLD ME TO

48 h

From your reply to a 30-minute scoping call

1 week

From a signed scope note to the first working session

3-4 wks

From data access to a first evaluated, working slice

2 days

Working days from any session to its written summary and next decision

WHAT OTHERS HAVE PUBLISHED ABOUT THEIR OWN RESULTS

- **FTI survey of 200 fund leaders:** [only 7% of portfolio companies run AI at enterprise scale](#)
- **FTI survey of 200 fund leaders:** [talent is the primary barrier \(35%\)](#)

Self-reported or vendor-published, linked; they show what is achievable, not a prediction for you.

WHY ME

- Pucho, an Indic voice assistant (2014-2020): 40,000 hours of speech data augmented to about 600,000
- Generative drug design (2020): three COVID-19 candidates approved by NIV Pune for pre-clinical trials
- Twelve engagement shapes, from a 10-day diagnostic to a retained design authority
- Published reasoning on production agents at vikramjha.work/writing

NEXT STEP

A 30-minute scoping call, within 48 hours of your reply. Tell me the workflow or the stalled pilot; the written summary follows within two working days.

Book 30 minutes: calendly.com/vikram01jha/30-mins

Email: vikram@vikramjha.work

Writing: vikramjha.work/writing

WhatsApp / phone: [+91 73034 70111](https://wa.me/917303470111)

LinkedIn: linkedin.com/in/invinciblejha

