

EXECUTIVE BRIEF

Native AI in production, owned by your team

Agents that clean your catalog, answer your customers and protect your margin, measured outlet by outlet and order by order

PREPARED FOR

Retailers, consumer brands, marketplaces and D2C companies

PREPARED BY

Vikram Jha · Principal FDE · AI Architect · Builder

DATE

7 October 2026

Edition 1.0 · sources checked 6-7 October 2026

EXECUTIVE BRIEF · THE ANSWER

The answer in one minute

I am a senior forward-deployed engineer and native-AI architect. I build the whole stack with your team (context, models, agents, evaluation, security, deployment), prove it against your baseline and leave it running in your environment, owned by your people.

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

THE DELTA

MEASURED BY

One owner across the full stack

Handoffs removed; days from question to decision

An evaluated slice in weeks

Weeks from data access to a slice that meets its threshold

Governance built in

Share of agent actions carrying an authority receipt

Open-source cost and sovereignty

Cost per resolved case against the frontier baseline

WHY NOW

What changed, and where to start

The largest retailer reports generative AI doing catalog work that would have needed nearly 100 times the headcount ([Walmart](#)), and regulators have said plainly there is no AI exemption from consumer protection law ([FTC](#)). [\[PUBLIC FACT\]](#)

A FIRST SLICE I WOULD PROPOSE [\[PROPOSED\]](#)

MEASURED BY

Catalog quality agent

Attribute completeness

Returns and refunds resolution

Resolution time

Store associate assistant

Time to answer

PUBLISHED RESULTS

What organizations report about their own work

ORGANIZATION

WHAT THEY REPORTED

Walmart

[850 million catalog data points created or improved with generative AI](#)

FTC

[no AI exemption from consumer protection law \(Operation AI Comply\)](#)

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

ENGAGEMENT SHAPES

Twelve ways to work together, each ending in a written decision

SHAPE	LENGTH	SHAPE	LENGTH
1 · Workflow opportunity and control diagnostic	10 business days	7 · Production slice sprint	4-8 weeks
2 · Pilot-to-production rescue	2-6 weeks	8 · Open-weight and sovereign migration	4-10 weeks
3 · Architecture review and second opinion	1-2 weeks	9 · Agentic software-delivery uplift	3-6 weeks
4 · Go-live readiness and red-team	2-3 weeks	10 · Platform and program build	8-20 weeks
5 · CXO and board AI portfolio sprint	1-2 weeks	11 · Fractional principal FDE and architect	Monthly, recurring
6 · Build-vs-buy and technical due diligence	1-3 weeks	12 · Retained design authority or co-delivery	Monthly retainer

WHY ME

Fifteen-plus years of production AI, and the layer most programs miss

Pucho, an Indic voice assistant (2014-2020: 40,000 hours of speech data augmented to about 600,000; telecom and device partners); generative drug design in 2020 (three COVID-19 candidates approved by NIV Pune for pre-clinical trials); and today, co-founder at a US-based native-AI startup in stealth mode, working on AI governance, agentic security and context graphs.

NEXT STEP

A 30-minute scoping call, within 48 hours of your reply

Tell me the workflow or the stalled pilot. Within two working days you receive a written summary and a proposed starting shape; if AI is not the answer, I will say so.

BOOK A 30-MINUTE SCOPING CALL calendly.com/vikram01jha/30-mins

EMAIL vikram@vikramjha.work

PHONE AND WHATSAPP (INDIA) [+91 73034 70111](tel:+917303470111)

LINKEDIN linkedin.com/in/invinciblejha

WRITING AND SELF-DIAGNOSTIC vikramjha.work

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