

AI HEPHS

The Engineering Context Layer for AI-Native Enterprise Software

From iterative requirements discovery to governed, production-grade systems.

Investor Pitch | Raising INR 5 Crore


INR 30 Crore Pre-Money Valuation



AI can generate code. It still cannot engineer an **enterprise system**.

What current AI tools **optimise**

01 Prompt-to-code speed



02 Developer productivity



03 Isolated task completion



04 Rapid prototyping



What enterprises actually **require**

01 Complete requirements


02 Architectural consistency

03 Traceability

04 Security and validation

05 Governance

06 Long-term maintainability



The gap between code generation and enterprise engineering is now the **primary adoption barrier**.

WHY THE MARKET IS OPENING NOW

AI is moving from assistance to **autonomy**.

01

AI coding → **autonomy**

AI is moving beyond code suggestions toward systems that can reason, execute, and complete increasingly complex engineering workflows.

CODING

AUTONOMY



02

GenAI → **agentic AI**

Models are evolving from conversational copilots into agents capable of planning, coordinating actions, and operating across enterprise workflows.

GENAI

AGENTIC AI



03

Governance → **infrastructure**

As AI becomes autonomous, governance, traceability, security, and control become foundational infrastructure rather than optional safeguards.

GOVERNANCE

INFRASTRUCTURE



The faster AI becomes autonomous, the more valuable the **control layer** becomes.

The missing ingredient is **Engineering Context**.

01

Industry Context

Regulations, terminology, domain workflows and industry conventions.



02

Enterprise Context

Systems, policies, users, data models, integrations and organisational constraints.



03

Engineering Context

Requirements dependencies, architecture decisions, interfaces, validation rules, lifecycle state and traceability.



Without engineering context, AI generates fragments. With **engineering context**, AI Helps engineers systems.


AI Hephs:

An **AI-Native** Enterprise Software Engineering Platform



One persistent engineering context. One lifecycle. From requirements to production and beyond.

Users answer questions. They do not engineer prompts.



Each answer enriches a **persistent engineering model** reused across requirements, architecture, code, testing and change.

HOW AI HEPHS WORKS

From requirement to production — with control built in.



ENGINEERING CONTEXT

Architecture

Systems

Policies

Data

CONTROL & VALIDATION LAYER

Agent
Execution

Generation
Execution

Testing

Security

Approvals

Audit Trail

Every generated artifact remains connected to the requirements, architecture, validation and governance that produced it.

Start where the cost of AI error is highest.

01

Lending & Underwriting

Complex decisions, structured evidence and regulated workflows.



02

Payments & Modernisation

Mission-critical systems, legacy complexity and controlled change.



03

Regulated Software Delivery

Engineering workflows where traceability, validation and governance matter.



04

High-Consequence AI Agents

Autonomous systems operating where errors carry significant business or operational consequences.



COMPLEXITY



RISK



ENGINEERING CONTROL



AI DEPLOYMENT

Initial beachhead: complex, regulated and high-consequence enterprise workflows.

Real Enterprise Problems. **Real Delivery.**

Four active engagements validating AI-native engineering across critical business workflows.

01 | PAYMENT GATEWAY

UK PAYMENT GATEWAY MODERNISATION

Signed contract to reverse-engineer and migrate
a .NET Framework 4.8 payment gateway to .NET 10.

.NET 4.8 → .NET 10

02 | INTELLIGENCE

VALINOR INTELLIGENCE

Delivering an intelligence platform for complex, expert-driven
intelligence workflows.

INTELLIGENCE PLATFORM

03 | CONSTRUCTION ERP

AI-NATIVE CONSTRUCTION ERP

Delivering an ERP for a construction company
operating across 40 sites in India.

40 SITES → INDIA

04 | FINTECH

VEDIKA CREDIT CAPITAL

Delivered a small-business loan assessment tool reported as significantly
more accurate than the client's prior OpenAI-based implementation.

CREDIT ASSESSMENT → AI

Four days instead of four to six months

A functioning enterprise application delivered from requirements to beta in days — not months.

04 DAYS

Actual delivery time

20-30x

Implied acceleration

4-6 MONTHS

Conventional estimate

CASE — STEADY HABBIT

ENTERPRISE HEALTH-ASSESSMENT TOOL

Enterprise-grade assessment workflow

SERVICENOW

Runs on ServiceNow

BETA DELIVERY

Completed beta delivery

FUNCTIONING APPLICATION

Delivered as a functioning application

Demonstrates requirements-to-application delivery beyond coding alone.

Defensibility compounds with every engineering lifecycle.

— REQUIREMENTS — • ARCHITECTURE — • IMPLEMENTATION — • TESTING — • CHANGE

01

Persistent Engineering Memory

Decisions and dependencies survive beyond prompts and model sessions.

02

Structured Engineering Grammars

Requirements, architecture, interfaces and validation are represented as reusable structures.

03

Lifecycle-Aware Orchestration

Context travels across requirements, design, implementation, testing and change.

04

Validation Infrastructure

Outputs are checked for consistency, traceability and architectural alignment.

05

Enterprise Workflow Intelligence

The platform learns how complex workflows are specified and delivered.

06

Foundation-Model Independence

Value compounds above the underlying model layer.

AI HephS operates above the coding-copilot layer.


From code and workflows to persistent, lifecycle-aware enterprise systems.

Capability	Coding CoPilots	Low-Code / App Builders	 AI HEPHS
Primary Unit	Code / task	Screen / Workflow	Enterprise system
requirements	Prompt - led	Configuration- led	Iterative discovery
context	Session - Bound	Platform-specific	Persistent
Architecture	Partial	Template - constrained	Lifecycle-aware
Governance	Limited	Variable	Built-in
Modernisation	Developer - Assisted	Limited	Reverse engineering + migration

The difference is not better code generation. It is a higher-order engineering system.

A new category is forming.

AI Hephs sits at the intersection of AI-enabled Enterprise Engineering, Agentic AI, and AI Governance.



IDEAL CUSTOMER PROFILE

- 01 Mid-market and enterprise organisations
- 02 Complex or regulated workflows
- 03 Legacy estates with high migration costs
- 04 Multiple systems, sites or business units
- 05 Pressure to modernise without increasing headcount

ECONOMIC BUYER

CIO / CTO

HEAD OF ENGINEERING

CHIEF DIGITAL OFFICER

TRANSFORMATION LEADER

Enterprise platform revenue with services-assisted adoption.



Land with a mission-critical use case. Expand into the engineering stack.

01

LAND

Target urgent and expensive engineering problems.

Start with a high-value, mission-critical enterprise problem where the cost of inefficiency is immediately visible.

02

PROVE

Deliver a measurable outcome quickly.

Demonstrate clear ROI, operational improvement, or engineering productivity through a focused deployment.

03

EXPAND

Extend into additional workflows, applications and teams.

Expand from the initial use case into adjacent engineering workflows and enterprise teams.

04

SCALE

Partner with consulting firms and system integrators.

Accelerate enterprise-wide deployment through strategic implementation partners and system integrators.

EXPANSION PATH

Enterprise Assessment

Pilot

Production

Enterprise Platform

TEAM

Built by enterprise software operators.

01 —
33 Years

Samir Gupta's enterprise software
engineering experience

02 —
16 Years

Experience as CTO

03 —
7 Years

Experience as CEO

04 —
13 Years

Collaboration with CTO Abhinav Jain

05 —
12

Current team size

06 —
10

Technical team members

The leadership team combines deep enterprise software delivery experience with product vision and technical execution capability.

From governed delivery to autonomous enterprise engineering.

01 —

GOVERNED AI-NATIVE SDLC

Requirements-to-code traceability

Architecture-aware generation

Validation and governance

Enterprise deployment hardening

02 —

AUTONOMOUS ENTERPRISE ENGINEERING

Multi-agent orchestration

Autonomous QA and compliance

Migration automation

Deployment orchestration

03 —

SELF-EVOLVING ENTERPRISE SYSTEMS

Autonomous defect remediation

Self-healing architectures

Continuous optimisation

Persistent engineering memory

GOVERNED

AUTONOMOUS

SELF-EVOLVING

Increasing autonomy, increasing engineering leverage.

Raising INR 5 crore to convert validation into repeatable enterprise scale.

CAPITAL SOUGHT

INR 5 Crore

PRE-MONEY VALUATION

INR 30 Crore

POST-MONEY VALUATION

INR 35 CroreIMPLIED INVESTOR
OWNERSHIP**14.29%**

USE OF FUNDS

- 01 Product hardening and platform engineering
- 02 Enterprise delivery capacity
- 03 Sales and GTM
- 04 Security and governance readiness
- 05 Benchmarking and measurable proof points

₹5 CRORE
CAPITAL

PRODUCT + DELIVERY + GTM + GOVERNANCE

REPEATABLE ENTERPRISE SCALE**|** From validated capability to repeatable enterprise scale.

18-MONTH MILESTONES

What this round is designed to achieve

01

PRODUCT

Production hardening and
repeatable deployment

02

COMMERCIAL

Convert active engagements into
references, renewals and expansion

03

EVIDENCE

Publish defined benchmarks

04

GTM

Build enterprise pipeline and
S1 partnerships

05

READINESS

Advance security and
compliance readiness

06

SCALE

Move toward a repeatable platform and
delivery model

The operating system for **AI-native enterprise software engineering.**

Engineering context. Governed delivery.
Enterprise outcomes.

RAISING

INR 5 CRORE

INR 30 CRORE

PRE-MONEY VALUATION