



NOCTRODEX

(OPC) PRIVATE LIMITED

Two equal product stages under one company.

AI AppSec software now · sovereign on-prem hardware next

STAGE 1 // SOFTWARE

AI SAST & DAST

AppSec software assist
Code & apps stay under customer control

STAGE 2 // HARDWARE

On-prem AI hardware

Appliances you own or manage
No per-query rent for inference



Architecture: EDHTA (paper v3) — in Stage 1 apps · designed for Stage 2 appliances

One thesis, two stages

Software wedge first · sovereign compute next — full objects always include both

01

Prove value in AppSec

AI SAST + DAST assist where code and apps stay under customer control.

02

Own the inference path

Same cyber AI delivered on customer premises — no token rent.

03

Scale the stack

EDHTA architecture IP + software + appliance SKUs + managed path as capital allows.

Go-to-market is staged — both stages are real

Commercial sequence, not abandonment of hardware (SBI / business overview still hold)

STAGE 1 — NOW

AI AppSec software

- AI SAST — static analysis assist on owned code
- AI DAST — authorised dynamic testing assist
- Ghostmaster AI core (in development)
- Operator / analyst software surface
- Faster pilots · software IP · lower capital

STAGE 2 — NEXT

On-prem AI hardware

- Purpose-configured AI appliances
- Local inference · no per-query metering
- Outright sale or managed (ISP-router)
- Models co-designed to hardware envelope
- Sovereign compute for cyber + industrial

Stage 1 problem — AppSec needs better AI


Cloud copilots leak code · one model window cannot do everything well

The pain

- Proprietary code into public metered AI
- Scanner noise vs late, expensive vulns
- SAST misses runtime · DAST misses root
- One model config cannot max long context + deep judgment + tools

AppSec wants AI on SAST + DAST — with ownership of source and scope.

The single-core trilemma — one config cannot max all four



Why one model window isn't enough — EDHTA splits reader vs judge.

Stage 1 products — AI SAST + AI DAST

Active development / pilot path • not public GA

AI SAST

Static Application Security Testing — AI-augmented

- Ingest customer-owned source / builds
- AI-assisted finding prioritisation
- Analyst-friendly review workflows
- Built for AppSec + engineering SDLC
- Keep code off public metered APIs


AI DAST

Dynamic Application Security Testing — AI-augmented

- Authorised testing of running applications
- AI-guided triage of dynamic findings
- Pairs with SAST for root-cause linkage
- Customer policy & scope always apply
- No unauthorised scanning — ever

Stage 1 buyers

AppSec-first positioning (not revenue share)



■ AppSec / engineering ■ Security consultancies ■ Product cos (SDLC)
■ Research / labs ■ Industrial path→S2


Outcomes they want

- Faster SAST triage
- Authorised DAST, less noise
- Source stays controlled
- Later path to on-prem AI


Stage 1 workflows

Abstract paths — no internal tooling inventory

SAST path



DAST path (authorised scope only)



Ghostmaster — AI core (public view)

Cybersecurity-oriented model line · under development · not GA

⚡ EDHTA · paper v3

What it is

AI for research and defensive workflows on infrastructure the customer controls.

Powers analyst assist for AppSec — SAST and DAST triage — without putting sensitive code on a public token API.

Powered by EDHTA: long-context reading + high-precision judgment on your path.

Status: active iteration · not public GA.

Product layers

1. Ghostmaster (model)
2. Apps & software
 - Operator console
 - SAST / DAST focus
3. EDHTA architecture (paper v3 · company IP)
4. On-prem hardware (Stage 2)

EDHTA — the architecture we ship

⚡ EDHTA · paper v3

Formal ML paper (v3) · company IP · Stage 1 active · Stage 2 designed to host



How AppSec work becomes a brief — not a raw dump into a single window.

Stage 1 — active

In softwares & apps: Ghostmaster path · AppSec assist · reader / judge / memory.

Stage 2 — designed

Same architecture on owned on-prem appliances — local inference.

Company IP


EDHTA-v3 paper · copyright package · not “patented” · not public GA.

Full paper under company IP / NDA package · Document ID EDHTA-v3 · public view only

Stage 1 architecture — EDHTA (public)

⚡ EDHTA · paper v3

Shape map only · no ports · same design Stage 2 hosts on-prem



Same architecture Stage 1 software runs · Stage 2 appliances are designed to host it on customer-owned mind.

Long materials

Code, reports, large pastes → briefs

Sharp judgment


Prioritise & assist without mega-window tax

Admitted memory

Findings enter history only when gated

Stage 1 competitive frame

Illustrative design intent — not a lab benchmark



Position

Not “replace every scanner tomorrow.”

AI-native AppSec software that respects ownership of code — with a clear Stage 2 hardware upgrade path.

Scores = design intent.


STAGE 2 // SOVEREIGN COMPUTE

Own the compute

On-prem AI appliances · co-designed models · no per-query metering
Where EDHTA runs on metal you own — equal brief: problem, product, deployment, buyers, economics

Stage 2 problem — renting intelligence fails production

From business overview: economics · sovereignty · fit



The missing category

Turnkey AI appliances where the model and the machine are designed together, installed on the customer's premises, and fully operated by the customer.

Not raw GPUs. Not metered cloud.

Stage 2 solution — co-design model and machine

⚡ EDHTA · paper v3

Local inference · EDHTA on owned hardware · not rent per query

Hardware

Purpose-configured AI appliances — compute, storage, power, thermals.

Model stack

Domain-tuned models (incl. cyber) sized for the appliance envelope.

EDHTA architecture

Same paper as Stage 1 apps — reader / judge / memory on the mind.

Ops path

Design → assemble → install → customer operates; air-gap capable.

One architecture

Stage 1 software
EDHTA in apps



Stage 2 appliance
EDHTA on owned mind

Same paper. Your metal.

Stage 2 product surface — appliance classes

Thematic roadmap variants · not SKU quotes

Edge

Compact sovereign inference for constrained sites and offline cells.

Mid-range


Team / lab class stacks for AppSec + research workloads.

High-capacity

Facility-class appliances for heavier local model + tool loops.

Stage 2 deployment — outright vs managed

ISP-router analogy from business overview / SBI profile



What never moves

Compute stays on the customer's premises. Data does not leave for inference. Nothing is billed per query.

What changes: title, capex vs opex, who maintains the metal.

*Managed = Noctrodex maintains.
Outright = customer or contract.

Illustrative scores.

Stage 2 engagement path

How an appliance programme runs

01

Discovery

Workload, latency,
privacy, environment

02

Co-design

Compute class,
memory, model
architecture

03

Build

Assemble, validate,
load model bundle

04

Deploy

Install on premises,
hand over control


05

Support

Integration &
maintenance — not
access rent

Stage 2 buyers

India first · industrial & institutional sovereignty (business overview §6)




■ Manufacturing / OT ■ Healthcare on-site ■ Research & education
■ SME / mid-market ■ Air-gap / critical

Why they buy appliances

- Data residency / offline
- Predictable cost envelope
- Own the metal and models
- Graduate from Stage 1 software

Stage 2 economics — planning view

BOM buckets illustrative • no supplier SKUs • re-quote before purchase



■ GPU compute ■ CPU / memory / board
■ Storage & network ■ Power / chassis / other

Unit economics (public)


Margin from BOM + assembly + integrated model value — not infinite cloud COGS per query.

Managed deployments only when fee/advance covers hardware outlay (bank profile honesty).

Early company: capital discipline first.

Stage 2 competitive frame

Business overview alternatives — illustrative axes



Differentiation


Not a cloud API reseller.
Not a white-box GPU dump.
Not a pure body-shop SI.

Co-designed model + machine for
sovereign cyber and industrial AI.

Scores = design intent.

Company revenue mix — staged (illustrative %)

Pre-revenue • not a forecast of contracted business



Reading the chart


Y1* software-led (Stage 1).

Y2–Y3* hardware sale + managed rise (Stage 2).

Axis = illustrative mix % of revenue composition, not rupees.

Company snapshot

Public legal facts · street address omitted



Legal facts

NOCTRODEX (OPC) Pvt Ltd

Incorp: 27 Jul 2026

RoC: Kolkata, India

CIN: U62010WB2026OPC288993

NIC: 62010 · 62020 · 26201

Pre-seed · founder-funded

HQ: Kolkata, India

noctrodex-hq.com

Founder

Public site bio · education + professional background



Soumabha Majumdar

Founder & CEO · AI & Cyber Security Engineer

Dual prior career tracks: corporate software developer and infosec / security engineer.

Cyber forensics-trained AI security engineer. Professional experience spans corporate product/engineering development and information-security engineering — including security research, healthtech infosec, offensive research, and full-stack delivery — before founding NOCTRODEX (OPC) PRIVATE LIMITED in 2026.

Corporate path

Prior corporate developer / engineering experience building software products and full-stack systems.

Infosec path

Infosec engineer track: security research, healthtech infosec, authorised testing craft.

Education · certs

MSc Digital Forensics (NFSU)
BTech · CEH · eWPTX
Kolkata, India

Stage honesty + roadmap themes

Build velocity · S1 then S2 · priority scores are illustrative (1-5)

₹0

External equity

0

Named clients


Pre

Revenue

S1→S2

GTM path

Near-term priority (1-5 illustrative)



The ask

⚡ EDHTA · paper v3

Stage 1 pilots now · Stage 2 owned mind CAPEX (planning ceiling · upper +5%)

SAST pilot

Authorised code-review assist

DAST pilot

Authorised dynamic-test assist

Design partners

Industrial / air-gap talks

Stage 2 capital — owned Ghostmaster v11 headless mind (planning ceiling)

Mind: dual 96 GB class GPUs · headless Linux appliance (owned only — not cloud rent)

Operator: Linux laptop plane (System76 Lemur Pro 14 class) · does not host the dual-GPU mind

Upper BOM bound × 5% contingency.
Soak one GPU first. Re-quote before purchase.

Mind

\$40k
₹39L

Laptop

\$3.2k
₹3.0L

System

\$43k
₹42L

High + import

\$61k
₹59L

Soumabha Majumdar · Founder & CEO · prior corporate developer + infosec engineer
MSc Digital Forensics (NFSU) · BTech · CEH · eWPTX · soumabha@noctrodex-hq.com · noctrodex-hq.com