

Sensor reliability becomes software.

NADIR turns episodic sensor calibration into a continuous software loop:
detect drift from existing signals, correct supported outputs, route
physical service when required, and preserve signed evidence.


- See the product →
- Export PDF
- Contact founders

THE 30-SECOND PITCH

Continuous reliability infrastructure for
sensor-equipped vehicles and machines.


BUYER	Fleet operations, repair quality, autonomy engineering, and platform teams.
INPUT	Existing telemetry, sensor streams, route context, and service history.
OUTPUT	Health state, correction eligibility, action queue, forecast, and evidence.
WEDGE	Four-week commercial-fleet deployment on 50-200 vehicles.

- 5
system layers
- 3
operating modes
- 1
evidence graph



Calibration is an event. Drift is continuous.

Vehicles leave a bay aligned, then vibration, impact, thermal change, glass replacement, and ordinary operation move sensors out of position.



THE GAP

No continuous layer connects sensor health, software correction, physical service, and evidence.

Operators see warning lights or outcomes. Engineering teams see raw telemetry. Neither gets a durable, fleet-wide reliability record.



Safety requirements are rising. Reliability infrastructure is lagging.

The vehicle is becoming more sensor-dependent while calibration remains largely event-driven and fragmented across operations, repair, engineering, and compliance.

MAY 2024

FMVSS 127 finalized

NHTSA establishes automatic emergency braking and pedestrian AEB requirements for light vehicles.

2026

Deployment window

Fleets, repair networks, and suppliers have time to establish continuous health and evidence workflows.

2029

Compliance horizon

New light vehicles up to 10,000 lb GVWR move toward the rule's performance requirements.

01

More safety functions depend on perception.

Camera, radar, LiDAR, and fused outputs increasingly sit upstream of critical decisions.

02

Service records do not show continuous condition.

A completed calibration cannot prove what happened weeks later on the road.

One reliability layer. Four operational jobs.

NADIR converts existing sensor and fleet signals into a clear operating decision—and preserves the reason behind it.

TRACK01 Know current health. Continuously score camera, radar, LIDAR, and cross-modal consistency.	CORRECT02 Compensate in software. Apply versioned transforms in real time for supported, validated integrations.
PRIORITIZE03 Act on what matters. Separate software-correctable drift from conditions that warrant physical inspection.	DOCUMENT04 Preserve the lineage. Seal source signals, model versions, transforms, decisions, and outcomes.

BOUNDARY: NADIR corrects supported sensor outputs. It does not directly command steering or braking.


From raw signal to reliable output.

NADIR / STACK COMPOSABLE VERSIONED TRACEABLE API · SDK · EDGE	01 INGEST Signal fabric Sensor streams, telemetry, route context, and service history.
	02 INFER Reliability models Residual geometry, changepoints, uncertainty, and trajectory.
	03 CORRECT Correction plane Eligibility, versioned compensation, explicit validation, and fallback.
	04 OPERATE HORIZON + Console Fleet queue, forecast, dispatch, validation, and closeout.
	05 PROVE Evidence graph Sources, models, transforms, actions, service, and outcomes.

The system is tangible today.


Real product surfaces cover fleet operation, technical investigation, correction integration, simulation, and evidence review.

NADIR HORIZON — forecasting engine, live run (A1 → A2 → A3)



Large installed base. Fragmented reliability spend.

NADIR's planning model spans ADAS hardware, calibration services, calibration tools, and the software and data layer connecting them.



Start where drift already costs money.

Commercial fleets offer accessible telemetry, concentrated uptime pain, repeatable workflows, and a short path to measurable deployment value.

CONNECT	OPERATE	VALIDATE	DECIDE
50–200 vehicles	4 weeks	Monitor + correct	Go / revise / stop
Existing exports or live data	Health queue + weekly digest	Explicit integration boundary	Signed closeout package

PUBLIC ENTRY OFFER

\$750

four-week micro deployment

✓ Baseline and health queue

Sell a measured decision, not an indefinite proof of concept.


The entry motion is intentionally small: qualify one operational question, validate one data path, run a bounded deployment, and close with evidence.

QUALIFY 1 owner + 1 question Sensor scope, data path, operating mode, and decision criteria.	VALIDATE Representative sample Identity, timestamps, units, quality, context, and manifest.
DEPLOY 50–200 vehicles · 4 weeks · \$750 Baseline, health queue, correction-fit review, and evidence samples.	EXPAND \$750/mo → \$4,500+/mo Ongoing operations, larger scope, HORIZON, APIs, correction, and exports.


CURRENT COMMERCIAL READINESS		JULY 2026
PRODUCT	Available	
HORIZON, Console, Calibration Lab, API Theater		
INTEGRATION	Available	
Python SDK, REST API, edge packages, sample schemas		

One core system. Four market surfaces.

The data contract and evidence graph stay constant while each buyer receives the workflow they already own.



Land with operations. Expand through infrastructure.



Current public pricing. Final enterprise scope varies by vehicle count, runtime, model depth, correction integration, and support.

A software-led planning scenario.

The prior deck modeled fleet ARR first, followed by repair, evidence, risk data, and later OEM licensing. This visual preserves that model without presenting it as guidance.



CAC / VEHICLE-EQ \$480 legacy acquisition assumption	MODELED ARPV \$220 annual revenue / vehicle
---	--

The moat is the longitudinal evidence graph.

Every deployment can connect sensor behavior to correction, service, and outcome—creating a technical history that point solutions do not retain.

SOURCE

Signals

Frames · telemetry · context



MODEL

Inference

Residual · confidence · tier



ACTION

Transform

Correction · fallback · policy



Built enough to test. Honest about what remains.

SHIPPED

Product system

- HORIZON + Console
- Calibration Lab + demos
- SDK, API, and edge packages
- Evidence bundle workflow
- Pilot packaging and pricing

NOW

Deployment validation

- Qualified fleet data paths
- AV and robotics integrations
- Correction envelope validation
- Operator workflow learning
- Evidence closeout reviews

NEXT

Repeatability

- Reference deployments
- Hardened segment manifests
- Expanded correction modules
- Partner distribution
- Longitudinal outcome data

Built close to the mobility problem.

A small founder-led technical team working across product, perception reliability, fleet workflows, APIs, and applied sensor math in Ann Arbor.



CO-FOUNDER & CEO

Dhruv Hegde

Product, platform, deployments, and translating technical depth into an operating workflow.

[Read profile →](#)



CO-FOUNDER & CTO

Sri Balaji

Algorithms, Console, HORIZON, and the engineering systems behind inference and correction.

[Read profile →](#)

Fund deployment learning and product hardening.

The prior deck's pre-seed scenario focused capital on a small technical team, deployment hardware, field programs, and the infrastructure required to turn pilots into a repeatable product.

PROPOSED PRE-SEED SCENARIO

\$200–400K

\$4–6M modeled pre-money range

18 months

to deployment, product, and seed-preparation milestones

USE OF FUNDS

\$200K–\$400K SCENARIO

Technical team	40%	\$80–160K
Hardware + manufacturing	30%	\$60–120K
Pilot deployments	20%	\$40–80K
Infrastructure + operations	10%	\$20–40K

M1–6

Core hires + first deployment

Harden connectors, correction boundaries, and evidence closeout.

M7–12

Repeatable pilot motion

Multiple deployments, segment manifests, and workflow proof.

The risks are real. The sequence is deliberate.

LONG OEM CYCLES

18-36 months

Build fleet and repair revenue first; use deployments to earn OEM and Tier-1 conversations.

MITIGATION / WEDGE FIRST

IN-HOUSE COMPETITION

Tier-1 + OEM

Stay cross-platform and neutral, with aftermarket evidence and fleet workflow depth.

MITIGATION / EVIDENCE GRAPH

SAFETY + CERTIFICATION

High burden

Keep current correction boundaries explicit and sequence formal pathways with deployment maturity.

MITIGATION / EXPLICIT SCOPE

EARLY ECONOMICS

Unproven

Treat model assumptions as hypotheses; instrument conversion, support, retention, and outcome value.

MITIGATION / MEASURE

Help make sensor reliability continuous infrastructure.

We are looking for qualified fleet and autonomy design partners, correction-integration collaborators, and aligned early-stage investors.

CAPITAL SCENARIO

\$200–400K

18-month planning horizon

DEPLOYMENT SCOPE

50–200

vehicles in a four-week entry program

SUCCESS CONDITION

Repeatability

reference outcomes, manifests, and expansion

DEPLOY

Run NADIR on real data.

Start a deployment →

INTEGRATE

Build with correction and evidence APIs.

Open developers →

PARTNER

Talk directly with the founders.

founders@nadirai.net →

