



5G / 6G AUTONOMOUS SENSING PLATFORM

InfraSense

Autonomous networks that perceive, reason and act —
one AI control center for the entire 5G/6G network.

Powered by NVIDIA · Built by DataSigns.ai

Investor Briefing · June 2026 · Confidential



THE PROBLEM

Today's networks are blind, manual and expensive

▶ Reactive operations

Faults are found in hours-to-days and fixed by hand.

▶ Manual, siloed optimization

Each domain is tuned separately — OPEX keeps climbing.

▶ Wasted sensing

The radio already 'sees' the physical world, but that signal data is thrown away.

▶ 6G demands autonomy now

ITU-R IMT-2030 targets zero-touch operations today's stacks can't deliver.

Hours → days

to detect & fix faults today

Manual

config & optimization

100% wasted

of the network's sensing signal

WHY NOW

Four fast-growing markets converge in one platform

\$30B+

Open RAN market by 2030
industry estimate

\$200B+

AI-RAN infrastructure TAM
industry estimate

2,000+

ISAC patent families since
2020
a core 6G pillar

\$25B+

Network digital-twin segment
industry estimate

InfraSense sits at the intersection of all four — Open RAN, AI-native networks, integrated sensing, and the network digital twin — instead of competing in any single one.

Tailwinds: O-RAN disaggregation breaks vendor lock-in, NVIDIA's AI-RAN push makes GPU-in-the-RAN mainstream, and 3GPP / ITU-R are standardizing sensing (ISAC) as a first-class 6G capability.

Market figures are illustrative industry estimates — sources to be confirmed before external use.

OUR VISION

Telecom as a living sensing fabric

01

PERCEIVE

Every base station becomes a sensor. Integrated Sensing & Communications (ISAC) reads range, velocity, position, presence and weather from the same radio that carries traffic.

02

REASON

AI agents and a natural-language orchestrator analyze telemetry + sensing, run digital-twin what-ifs, and decide the best action — with safety scoring built in.

03

ACT

Closed-loop xApps apply changes autonomously across the RAN — self-deploying, self-optimizing, self-healing toward zero-touch operations.

Target: **ITU-R IMT-2030 (6G)** — zero-touch, zero-wait, zero-fault operations.

THE PRODUCT

One control center, six integrated layers

AI ORCHESTRATION

Natural-language intents → agent + RAN actions · NeMo / NIM · Tree-of-Thoughts · LangGraph

DIGITAL TWIN

NVIDIA Omniverse + Aerial AODT · real-time twin · what-if planning

SIMULATION

NVIDIA Sionna RT ray tracing · Aerial 5G NR PHY · Modulus surrogate

NETWORK

O-RAN Near-RT RIC · xApp / agent marketplace · NTN (3GPP Rel-18)

SENSING

ISAC — cm-level range / velocity / position / presence / weather

GPU COMPUTE

NVIDIA Brev L40S cloud — real GPU-accelerated workloads

KEY ATTRIBUTES

Six pillars that make the network autonomous

01

ISAC sensing

Turn every base station into a sensor — cm-level range, velocity, 3D position, presence, gesture and RF weather.

02

Real-time digital twin

NVIDIA AODT + Omniverse + Sionna RT — a live, GPU-accelerated replica for what-if planning and training.

03

AI-native orchestration

Natural language → structured agent + RAN actions, with Tree-of-Thoughts and LangGraph planners.

04

Autonomous 6G RAN

Self-deployment, -optimization, -healing, -evolution across autonomy levels L0 → L5.

05

Real GPU compute

Sionna RT, Aerial, Modulus, Isaac Sim run live on NVIDIA Brev L40S — not slideware.

06

xApp / agent marketplace

Lifecycle-managed xApps and AI agents — deploy, stop and monetize closed-loop apps.

ISAC: every base station becomes a sensor

± 0.3 m

Range estimation

± 0.5 m/s

Velocity detection

± 0.5 m

3D positioning

± 1.5°

Angle of arrival

99.2%

Presence detection

96.8%

Gesture recognition

96.2%

Fall detection (CSI)

94.8%

Activity recognition

CSI sensing: 512 subcarriers @ 1 kHz · breathing 98.5% **Waveforms:** OFDM · FMCW hybrid · PMCW coded · pulsed radar
Multi-target Kalman tracking · cm-level 3D mapping — sensing and comms share one radio, one spectrum slice.

ATTRIBUTE · DIGITAL TWIN

A live, GPU-accelerated replica of the network

► NVIDIA Aerial AODT + Omniverse

Photoreal 3D digital twin with collaborative USD hosting (Enterprise Nucleus, live).

► GPU ray tracing (Sionna RT)

Real propagation paths over real city geometry — ~5 s per scene on an L40S.

► Sub-100 ms previews (Modulus)

Physics-informed surrogate for instant what-if drags; full Sionna one click away.

► Sionna ↔ AODT sync + USD export

Bidirectional, with Blender 4.2 USD export of real environments.

53 paths

real multipath, 1 TX × 1 RX

135 buildings

real OSM city geometry

~5 s / scene

GPU ray trace on L40S

< 100 ms

Modulus preview latency

AI-native orchestration, not bolt-on automation

► Natural-language command center

Plain-English business intents → structured agent + RAN-layer actions via NVIDIA NIM (llama-3.3-70b).

► Tree-of-Thoughts planner

Fans out 3 candidate plans, a critic scores each on safety / effectiveness / cost — best wins.

► LangGraph state machine

Production-grade graph orchestration, fully LangSmith-traced for observability.

► Persistent agent memory

Agents remember context across sessions (episodic / semantic / procedural recall).

► Hybrid RAG over ops docs

Dense + keyword retrieval (Reciprocal Rank Fusion) grounds answers with citations.

► xApp / AI agent marketplace

Deploy, stop and monetize closed-loop xApps and specialized agents.

Self-driving RAN, L0 → L5

Level	Capability	Status
L0	Manual operations	Legacy
L1	Assisted automation	Current
L2	Partial automation	Demo-ready
L3	Conditional automation	In development
L4	High automation	Roadmap
L5	Full autonomy	Target

SELF-X CAPABILITIES

- ▶ Self-deployment
- ▶ Self-optimization
- ▶ Self-healing
- ▶ Self-evolution

ITU-R IMT-2030 (6G) TARGETS

- ▶ 1 Tbps peak throughput
- ▶ 99.99999% reliability (HRLLC)
- ▶ 10M devices / km² (mMTC)
- ▶ cm-level ISAC sensing

This is running infrastructure, not slideware



NVIDIA Brev GPU cloud

Live 495-instance catalog + provisioning on L40S-48GB. All GPU work runs here.



Sionna RT (GPU)

Real ray tracing on CUDA — Sionna RT 2.0.1, Mitsuba 3.8, ~5 s/scene on L40S.



Aerial 5G NR PHY

End-to-end PUSCH chain: CDL channel → MMSE-IRC → SINR/RSRP → MCS throughput.



Modulus FNO (trained)

Physics-informed surrogate trained on the L40S (torch 2.6 + CUDA), GPU-verified.



Isaac Sim 4.5

Headless Physical-AI / robotics sim streaming over WebRTC from the L40S.



Omniverse Nucleus

Full 12-service Enterprise stack live for collaborative USD hosting.

GPU policy: every GPU workload runs on NVIDIA Brev — no other GPU vendor or local-GPU path.

DIFFERENTIATION

A defensible position, not a feature

Conventional RAN	InfraSense
Add-on AI	AI native in every layer
Static planning	Real-time digital twin
Manual optimization	Autonomous closed loop
Separate sensing	Integrated ISAC
Vendor lock-in	Open O-RAN architecture

DEFENSIBLE IP — 4 PATENT FILINGS

Multi-agent ISAC orchestration	HIGH
ISAC ↔ digital-twin fusion	HIGH
RF-based weather sensing	MED-HIGH
Disaster-recovery orchestration	MED-HIGH

WHERE IT LANDS

Five verticals, one platform

Mobile operators

Cut OPEX with autonomous optimization, self-healing and energy savings across the live RAN.

Private 5G & enterprise

Campus, factory and industrial networks with built-in sensing for automation and safety.

Smart venues & live events

Live Event ISAC: per-zone crowd density, movement, surge and fall detection from one radio.

Public safety & disaster

NTN satellite failover, survivor location and priority traffic — autonomous emergency recovery.

Dynamic spectrum sharing

CBRS incumbent-radar protection and NR↔LTE sharing via lifecycle-managed xApps.

TRACTION

Live and verified in the product today

► O-RAN SC RIC on GKE

E2 / A1 / O1 interfaces running on a production Kubernetes cluster.

► 5+ lifecycle xApps

Video-aware MAC, spectrum-sense, DSS, interference-mgmt, traffic-steer — deploy / stop live.

► GPU twin on L40S

Sionna RT ray tracing + AODT EM simulation verified end-to-end on NVIDIA Brev.

► Live AI orchestrator

Natural-language → actions via NIM, plus Tree-of-Thoughts and LangGraph engines.

► Federation GraphQL layer

One schema federating every agent backend — cross-domain queries in one round trip.

► Physical-AI stack live

Isaac Sim 4.5 + Omniverse Enterprise Nucleus running on the L40S box.

Stage: working platform, pilot-ready — pre-revenue. Seeking design partners and capital to reach GA.

ROADMAP

From working platform to 6G trials

NOW · H1 2026

- ▶ O-RAN RIC + 5 xApps live
- ▶ GPU digital twin
- ▶ ISAC demos
- ▶ AI orchestrator

NEXT · H2 2026

- ▶ ISAC platform GA
- ▶ NTN integration
- ▶ First paid pilots
- ▶ Design partners

2027

- ▶ L3 autonomy
- ▶ Agent marketplace GA
- ▶ Multi-operator deploys
- ▶ Patent grants

2028+

- ▶ 6G early trials
- ▶ L4 autonomy
- ▶ Scale GTM
- ▶ Ecosystem partners

Roadmap reflects current plans and targets, not commitments.

BUSINESS MODEL

Platform revenue, with operator ROI to match

30-40%

OPEX reduction

75% faster

Mean time to repair

+25%

Network capacity

20%

Energy savings

HOW WE CHARGE

- ▶ Platform license — the AI control plane (per network / per site)
- ▶ Usage — GPU compute for twin, sim & training (NVIDIA Brev)
- ▶ Marketplace — revenue share on third-party xApps & agents

DEPLOYMENT OPTIONS

- ▶ Cloud — fully managed SaaS
- ▶ Hybrid — cloud control plane, on-prem data plane
- ▶ On-prem / edge — in the operator's data center

Operator ROI figures are modeled targets — payback ~12-18 months (illustrative, to be validated in pilots).

THE ASK

Fuel the path to GA and first revenue

RAISING

\$[__]M [Seed / Series A]

Placeholder — fill in target raise & round.

18-MONTH MILESTONES

- ▶ ISAC platform & digital twin to GA
- ▶ 3-5 paid operator / enterprise pilots
- ▶ L3 autonomy in production
- ▶ File the 4-patent portfolio

USE OF FUNDS

Product & GA engineering [%]

ISAC / 6G R&D [%]

Pilots & go-to-market [%]

IP & patent portfolio [%]

THANK YOU

Let's build the autonomous network.

InfraSense · Powered by NVIDIA · Built by DataSigns.ai

[your name] · [email] · datasigns.ai