

LIGHTCONE | MULTI-SENSOR GROUND-TRUTH CAPTURE

Production-grade ground truth for autonomy, **at scale.**

LightCone is ROVR's second-generation data-capture unit. Automotive-grade LiDAR, a 4K calibrated camera, high-precision IMU, and tri-band GNSS with RTK, fused on-device and delivered as geo-referenced training data from a globally distributed fleet you don't have to operate.



LIGHTCONE · ROOF-MOUNTED MULTI-SENSOR RIG

WHAT THIS BRIEF COVERS

The sensor architecture behind LightCone, the raw and processed data products it produces, who buys that data today, and how automakers, AV technology companies, robotics firms, simulation platforms, and data-labeling vendors can engage with ROVR. Written for technical evaluators and data buyers.

2.8M+ km
DATA COLLECTED

<5 cm
POSITIONAL ACCURACY

200 m
LIDAR RANGE

3 LOIs
+ 1 POC CONTRACT

POSITIONING

Why LightCone.

The ground-truth problem every perception team has, and nobody wants to build around.

01 · THE MISREAD

Vision-only at inference is not vision-only at training.

It's common to hear that Tesla "doesn't use LiDAR," and that's true for inference. The production vehicle runs on cameras. But training is a different story. Public footage shows Tesla test vehicles with LiDAR rigs mounted in almost the exact position LightCone sits, and that's because LiDAR depth gets used as ground truth to supervise the vision signal during training. Without it, vision models can't learn reliable 3D geometry.

02 · THE BUILD-VS-BUY

Every perception team needs this data. Few want to collect it themselves.

Running a LiDAR collection fleet means buying sensors, hiring drivers, building calibration pipelines, maintaining vehicles, and managing geographic coverage. Most teams either don't have the budget for that or would rather spend engineering on their stack than on data ops. LightCone gives them a second path.

03 · WHY ROVR

We're already in the field.

Since September 2024, ROVR contributors have captured over 2.8 million kilometers of driving data through our first-generation TarantulaX units. LightCone is the multi-sensor successor, and it pairs that deployed network with the sensor stack perception teams actually want: LiDAR, calibrated camera, high-precision IMU, and RTK, all on a unit engineered for automotive-grade stability.

"For any modern training pipeline that uses vision as a primary modality, LiDAR ground truth isn't optional. It's how the model learns depth."

INDUSTRY SNAPSHOT

Who uses LiDAR at inference

Robotaxis. Waymo, Baidu Apollo Go, and others run LiDAR live. Operators assume liability, so sensor redundancy matters.

Premium ADAS. NIO, Xpeng, Li Auto, and several others ship LiDAR on production vehicles.

Pure-vision OEMs. Tesla and similar ship cameras only at inference, but still train models against LiDAR-captured depth.

CURRENT TRACTION

Commercial signal

3 Letters of Intent signed with automakers and autonomy partners.


1 active POC contract in delivery.

Additional partnership letters pending disclosure. LightCone pilot fleet begins Q1 2025 with 50 units, and the 1,000-unit production run starts Q2 2025.

ARCHITECTURE

Four sensor modalities, one synchronized data product.

LightCone fuses LiDAR, camera, IMU, and RTK-corrected GNSS on-device. The output is aligned, geo-referenced, and ready for training pipelines or HD map processing.



HARDWARE

Specifications.

Full sensor, compute, and environmental profile for the LightCone unit.

LIDAR

Detection range	200 m @ 10% NIST
Angular resolution	0.2° × 0.2°
Ranging accuracy (1σ)	< 2 cm
Positional accuracy	< 5 cm
Grade	Automotive

CAMERA

Resolution	4K HD
Grade	Automotive
Calibration	Factory, to LiDAR
Function	Point-cloud color

COMPUTE

AI chip	Rockchip RK3588
Workload	Fusion · calibration

IMU

Gyro bias (XYZ)	0.13 °/√h
Accel bias (XYZ)	0.050 m/√h
Class	Tactical-grade

GNSS / POSITIONING

Antenna	Tri-band GNSS
RTK service	GEODNET
Positioning accuracy	Centimeter-level

ENVIRONMENTAL

Operating temp.	−45 to +85 °C
Storage temp.	−65 to +105 °C
Ingress protection	IP67

DESIGN NOTES

Why these specs matter for training data

200m LiDAR range at 10% NIST. This captures low-reflectivity objects (dark vehicles, pedestrians in dark clothing) at highway stopping distances, which is the exact set of edge cases camera-only systems struggle with.

IMU bias at 0.13°/√h. This keeps vehicle pose accurate through GNSS dropouts in tunnels, urban canyons, and underpasses. Pose drift directly degrades point-cloud quality over time, and this class of IMU keeps that error inside the centimeter envelope for typical drive durations.

Factory LiDAR-to-camera calibration. Every point cloud arrives with pixel-accurate color. No field recalibration, no customer-side alignment pipeline needed before data enters training.

DATA PRODUCTS

What LightCone delivers.

Two data products, five customer types. Raw sensor data for model training, processed HD map data for autonomy stacks.

RAW DATA

Synchronized multi-sensor capture

Colorized LiDAR point clouds with IMU-compensated pose, RTK-corrected geo-location, and time-aligned 4K camera frames. Delivered in training-ready formats.

- Perception model training (3D detection, segmentation, depth)
- End-to-end autonomy model supervision
- Synthetic data generation seed for simulation platforms
- Ground truth for labeling vendor QA pipelines
- Derived products: secondary data sets curated by third parties

PROCESSED DATA

HD map data · ROVR-processed

ROVR's internal pipeline converts raw captures into HD map layers with semantic features, lane geometry, and localization primitives ready for autonomy stacks.

- HD map generation for autonomous driving
- Fleet localization priors
- Continuous map freshness via the deployed contributor network
- Geographic coverage expansion on demand
- Per-region HD map licensing, by geography or vertical

CUSTOMER TYPES

Who buys LightCone data.



Global automakers

Training data for ADAS and L2+ systems, and HD maps for production deployment.



Autonomy tech firms

Perception and end-to-end training sets, plus long-tail edge-case coverage.



Robotics companies

Outdoor navigation data for delivery bots, humanoids, and inspection robots.



Simulation platforms

Real-world scene seeds for photoreal and geometric sim environments.



Data labeling vendors

Ground-truth references for QA on camera-only labeling contracts.

SCALE · ENVIRONMENT · ENGAGEMENT

How to work with us.

Fleet scale today, operating conditions in the field, and the four ways customers engage commercially.

NETWORK SCALE

2.8M + km

CAPTURED DATA SINCE 9/2024

1,000 units

LC PRODUCTION RUN · Q2 2025

50 units

LC PILOT FLEET · Q1 2025

TARANTULAX DEPLOYED ACTIVE

LC PILOT Q1 2025

LC SMALL-BATCH Q2 2025

OPERATING ENVELOPE

Built to stay in the field.

Operating temp. -45 to +85 °C

Storage temp. -65 to +105 °C

Ingress rating IP67

Mount Roof-mounted

LiDAR grade Automotive

Camera grade Automotive

Environmental spec matches what perception teams expect from an automotive data pipeline. Units run through snow, desert, and tropical conditions without ruggedization changes, and contributors across regions feed the same data format back into the network.

ENGAGEMENT PATHS

Four ways to buy.

01 · Data licensing

Subscribe to raw or processed feeds by region, volume, or data type. Recurring commercial terms.

02 · Custom collection

Target routes, times of day, weather, or edge cases. ROVR tasks the contributor fleet to your spec.

03 · Proof of concept

Scoped POC with defined sample sizes and success criteria. Typical duration 4 to 8 weeks.

04 · Strategic partnership

Joint roadmap on data formats, sensor configs, or exclusive coverage. Case-by-case terms.

NEXT STEP

Request a sample data package or schedule a technical call.

rovrlabs.io
x.com/ROVR_Network
discord.gg/RjV3E3u4F2