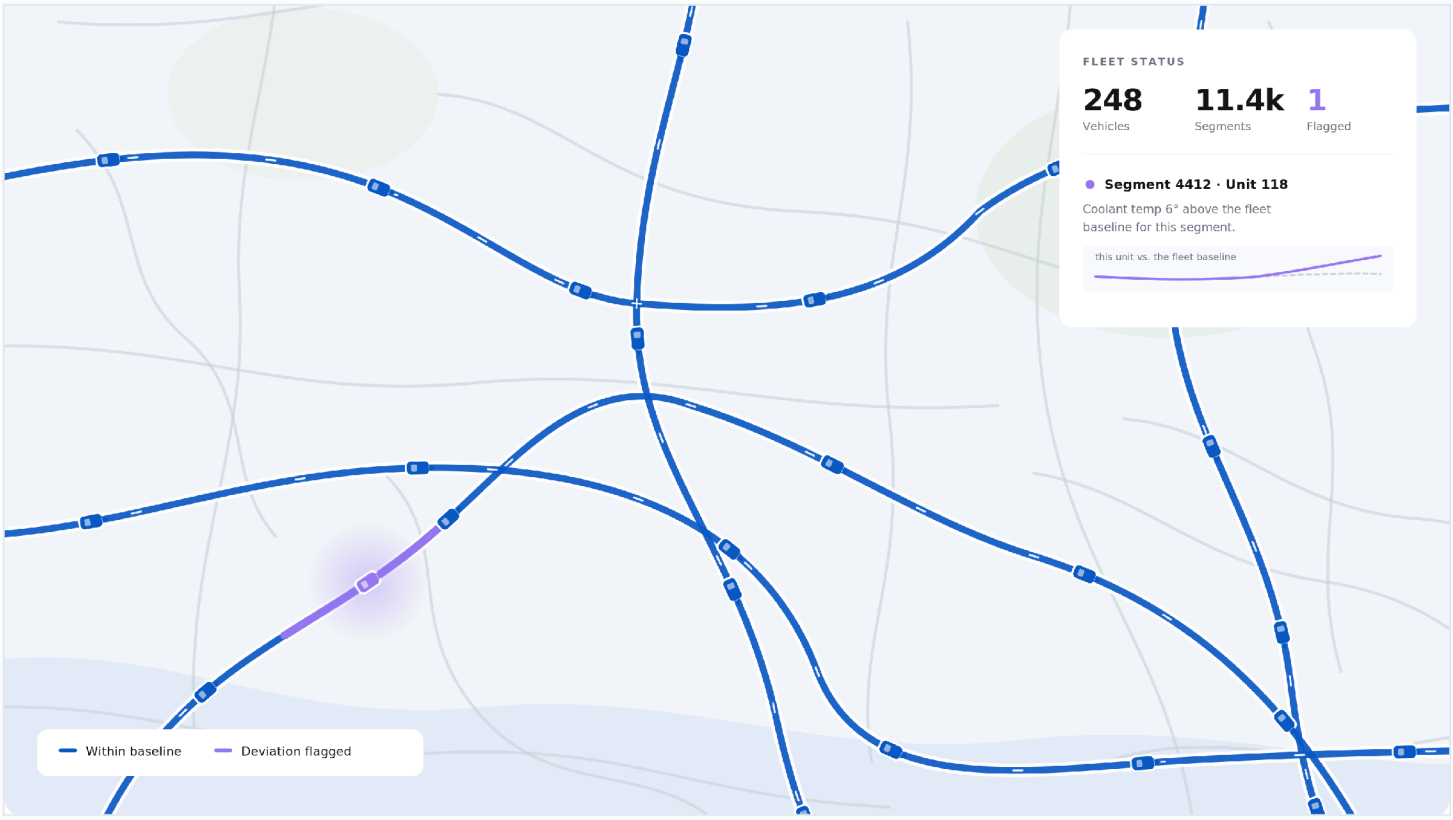
WHEROBOTS + ACTION ENGINE

The breakdown was in your data before the truck stopped



Your fleet operations team only catches issues when the driver calls it in, and by then, your truck could need major maintenance or is already broken down.

For autonomous vehicles, the stakes are even higher. Missing a camera or sensor slowly degrading can go from a stalled truck to an incorrect turn causing an accident in the middle of an intersection.

The issue is not a lack of technology; it's that traditional fleet observability platforms work as a Yes/No function. Something either works or doesn't. Current systems are built to rely strictly on static thresholds.

Aspen Fleet is a platform that moves beyond traditional alarm thresholds to catch issues as they approach failure.

01 THE PROBLEM

The spatial blind spot.

Every telemetry record has a location attached. Most platforms store it and stop there. Without the road segment, the grade, the ambient conditions, and the fleet's history on that stretch, there is no way to separate a genuine fault from normal operation.

Built for a different fleet.

Legacy platforms were designed around traditional vehicles and low-frequency ECU telemetry. Modern fleets carry cameras, perception stacks, and sensors reporting thousands of times faster, and the platform underneath was never shaped for it.

Reactive by design.

A threshold cannot predict. It marks a line already crossed. Managers hear about a failure after the repair bill exists, and after the vehicle is already down.

Alert fatigue.

Threshold alerts fire constantly and carry no explanation. Operators learn to dismiss the feed, which means the alert that mattered gets dismissed with the rest.

A fleet operations team running on fixed thresholds pays twice: first for the hours spent clearing alerts nobody could act on, and again for the breakdowns nobody saw coming.

Aspen Fleet starts from the road segment. →

02 THE SHIFT

Evidence, without the noise.

Aspen uses spatial and temporal context like elevation, road grade, and current temperature conditions to discard false positives and surface problems that threshold-based checks cannot detect. Every flagged issue carries the evidence behind it.

AI-first by design.

High-frequency fleet data is a native input. Aspen takes in the high-frequency camera and perception data that traditional platforms were never built for. Camera detections, perception outputs, and model confidence scores are validated like any other signal.

Predictive by construction.

Aspen flags failures when a vehicle deviates from the fleet norm for that segment, long before a fixed threshold would be crossed. Issues like motor degradation, sensor faults, and perception-stack decay surface while they are still repairable.

Issue to insight to remediation.

Aspen shows the root-cause of all issues found, and tells you how to solve them. By reporting each issue with the segment baseline, the pattern history, and the segment context behind it, teams see what went wrong and act on it.

03 HOW IT WORKS

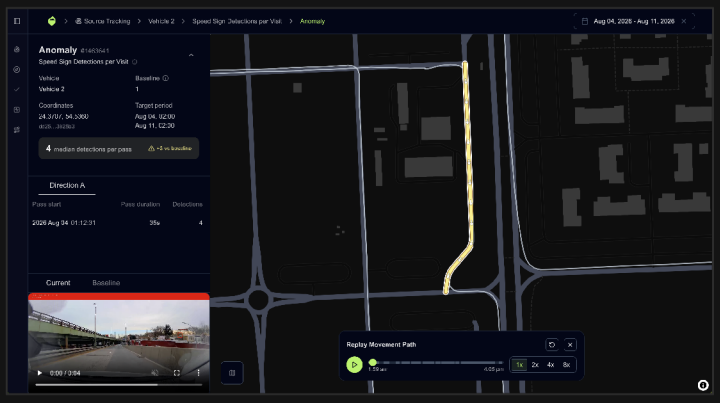
ASPEN

01 **Spatial referencing.**
Aspen maps every vehicle data point to a specific road segment.

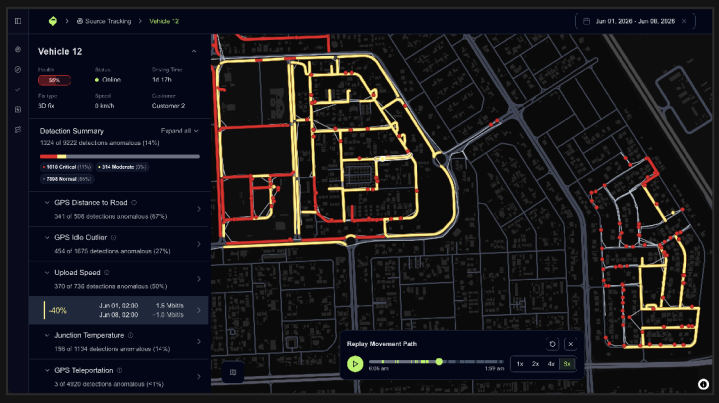
02 **Segment baselines.**
Aspen builds a behavioral baseline for every road segment from the fleet that drives it. A vehicle is compared against what normal looks like on that road, on that grade, in those conditions.

03 **Segment context.**
Aspen enriches every segment with contextual data, including road type, elevation, and ambient conditions, and analyzes it as part of every check.

Aspen Fleet includes over **100 pre-built checks** across location validity, signal integrity, operational behavior, sensor envelopes, and computer-vision output. None of them rely on a uniform threshold. Aspen runs them two ways: continuously against streaming telemetry, and retrospectively across years of stored history to build a clean baseline. Historical scans across billions of telemetry records complete in minutes.



SEGMENT BASELINES, IN PRODUCT
Speed sign detections per visit, Vehicle 2. Four detections on a pass where the fleet baseline for that segment is one, with the footage from that pass beside it.



EVIDENCE, WITHOUT THE NOISE
Vehicle 12, health 55 percent. Upload speed 40 percent below where it started, 2.5 down to 1.5 Mbit/s over a week, with the route colored by the severity of issues found on each segment.

> "Wherobots has been more than infrastructure for us. Their team helped us design a system that treats spatial context as a first-class signal, and their engine runs it reliably in production against the highest-volume fleets running with Aspen Fleet. That combination of product and expertise is why our fleet data is ready for the AI stack our customers depend on."

SERGEY SUKOV
CEO, ACTION ENGINE

The engine underneath, and the same method at network scale. →

04 THE ENGINE UNDERNEATH

Action Engine built Aspen Fleet on Wherobots, the AI Context Engine for the Physical World. Wherobots supplies the spatial compute: 300+ spatial functions covering vector and raster data, with native Spark SQL for tabular operations, built by the original creators of Apache Sedona and its 62M+ downloads.

Spatial engineering at this scale is why most fleet platforms do not have it. There are roughly 1.5 million spatial software developers against 47 million general software developers, and general-purpose infrastructure makes spatial data engineering cost prohibitive. Aspen Fleet gives fleet teams the capability without the build.

300+
SPATIAL FUNCTIONS, VECTOR AND RASTER

62M+
APACHE SEDONA DOWNLOADS

1.5M / 47M
SPATIAL VS. GENERAL SOFTWARE DEVELOPERS

05 ADJACENT PROOF

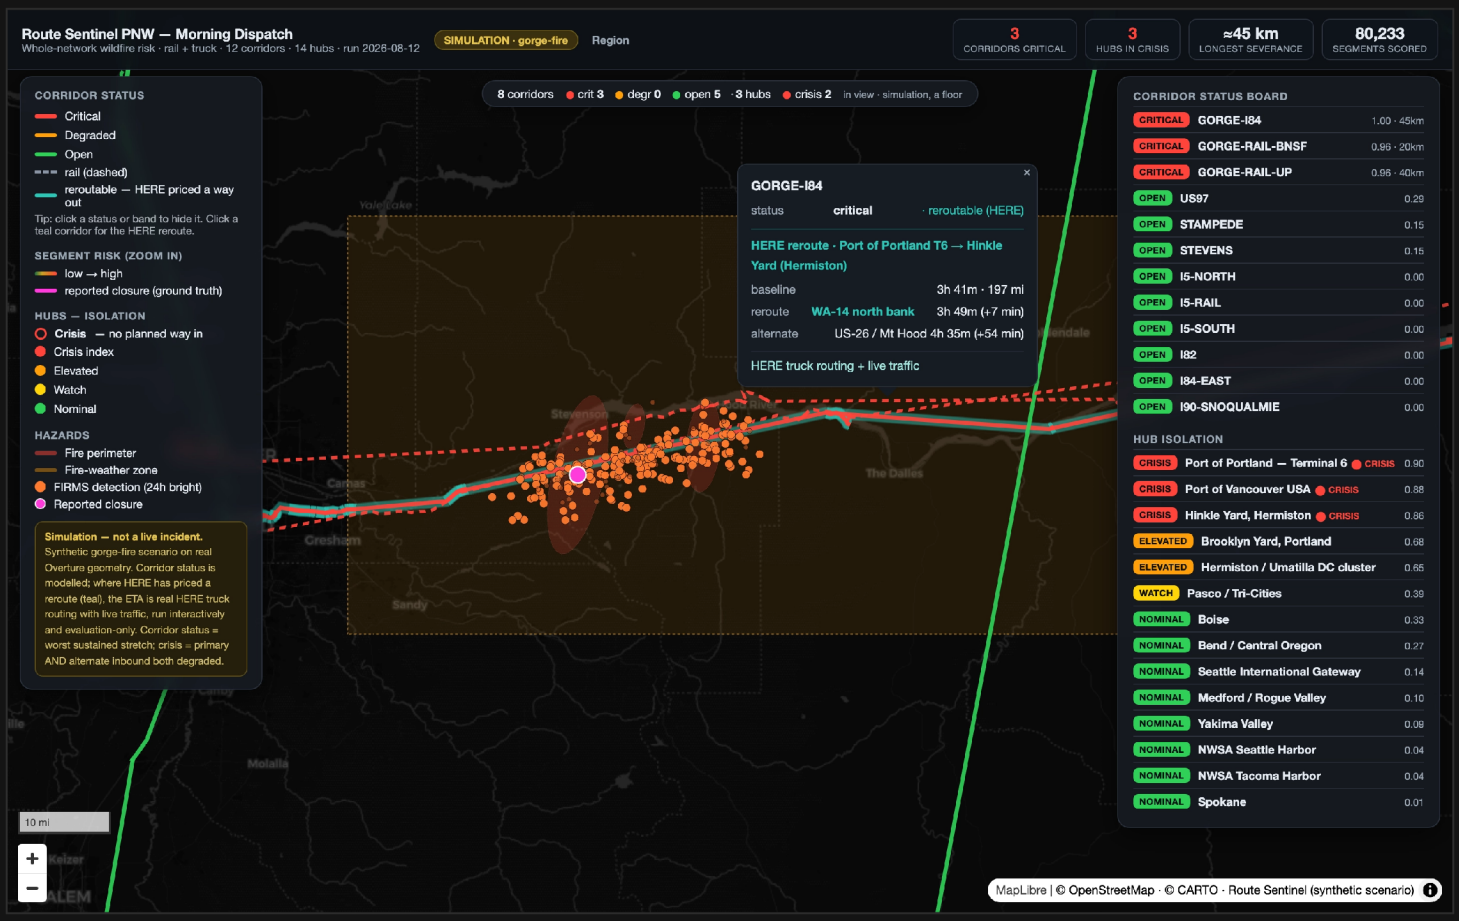
ROUTE SENTINEL · WHEROBOTS

The same engine, one step further out

Vehicle context and network context are the same problem at different scales. Route Sentinel scores an entire freight network under a wildfire scenario: 12 corridors and 14 hubs across real Pacific Northwest road and rail geometry, with each corridor scored open, degraded, or critical as conditions move. Route Sentinel flags a community as crisis only when its primary and alternate inbound routes both degrade.

In one run, a gorge fire severs I-84 at Cascade Locks. The reroute north on WA-14 costs 7 minutes. The Mount Hood alternate costs 54. Both look like detours on a map. Only one is a dispatch decision.

The wildfire is a simulated scenario played out on real road and rail geometry. Corridor scoring runs as it would for a live incident. Live app: route-sentinel-app.vercel.app



See it against your own telemetry.

Aspen Fleet connects through your existing telematics platform, a REST endpoint, batch files in cloud storage, or a live feed from onboard systems. Ingestion, normalization, and processing are handled. Run it against your historical data and see what threshold-based checks have not caught.

BEN PRUDEN
Head of GTM, Wherobots
ben@wherobots.com

KIRILL SUKOV
Business Development, Action Engine
kirill.sukov@actionengine.com

Wherobots | actionengine