Sovereign Edge AI Software

AI gets better with more data.
The data that matters most is data you can't move.

Scaleout Edge trains where the data lives.
No transfers. No exposure.

AI sovereignty by design.
Not by promise.

BAE SYSTEMS SAAB NATO OTAN FMV SCANIA

Why Edge AI fails with conventional ML

Standard ML assumes you can centralize data, train in the cloud, and deploy a static model.

At the edge (across defense, industrial, and regulated environments) none of those assumptions hold.

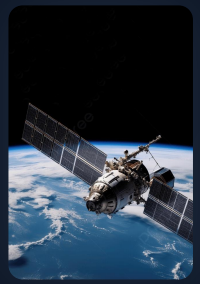

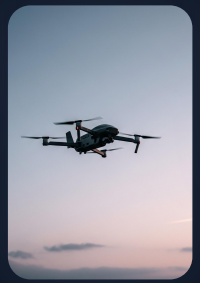

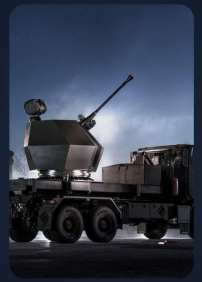

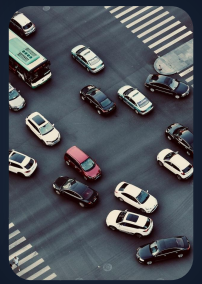

- **The data can't move.**
  Privacy law, classification boundaries, and bandwidth costs make centralization impractical or illegal. This isn't a policy preference, it's a constraint that disqualifies most ML architectures before they start.
- **The network can't be trusted.**
  Contested environments, remote sites, and mobile fleets operate where connectivity is intermittent or denied. Any AI system that requires a live cloud link to function isn't mission-critical, it's a single point of failure.
- **The fleet is a black box.**
  Thousands of nodes running models independently with no central record of which version ran where, what data shaped it, or when it last updated. In regulated or safety-critical systems, that's not an ops problem, it's an accountability void.
- **One model doesn't fit the fleet.**
  A model trained on centralized, averaged data washes out local patterns. Edge environments are genuinely different operating contexts, not noisy copies of a single distribution. And they shift faster than centralized retraining cycles can follow.

DIFFERENT MISSIONS, SAME CHALLENGE:

Distributed data that cannot move, and AI that must keep learning.

The model goes to the data. Not the other way around.

Scaleout Edge uses federated learning to train where the data lives. Models go to the data, train locally, and only updates return.

The platform is built on four principles shaped by edge deployment constraints.

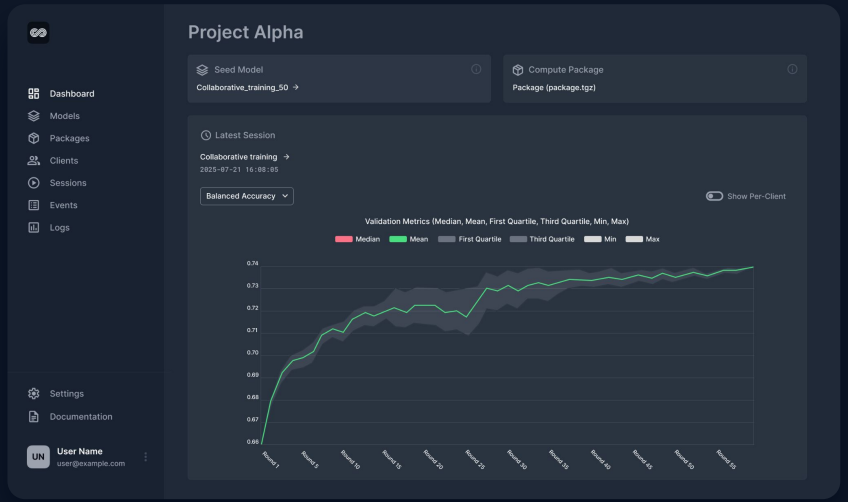

- **Raw data never leaves the device.**
  Only model updates travel to the aggregation layer, weight adjustments, not sensor feeds or telemetry. All communication is encrypted in transit via TLS, and at enterprise tier, models can be cryptographically signed and envelope-encrypted before distribution. Sovereignty is enforced by system design, not access control policy.

- **Designed for denied connectivity.**
  The platform never assumes a stable connection. Nodes train independently and sync updates when connectivity allows. A drone that loses signal mid-mission continues learning; when it reconnects, its updates merge automatically.

- **Immutable audit trail across the fleet.**
  The control layer captures training metrics, model provenance, and node health across every device. Every model version is recorded with its compute package hash and session lineage. The accountability chain required for regulated and safety-critical deployments.

- **Continuous training on live data.**
  Instead of centralized retraining cycles, edge nodes train directly on local data as conditions change. Any ML code is packaged into structured Compute Packages and distributed to the fleet for continuous, autonomous learning that adapts to each node's environment.

Not a replacement. An extension.

Scaleout Edge connects your ML environment to your edge fleet, preserving your existing AI stack while managing distributed execution.

ML FRAMEWORKS & LIBRARIES

PyTorch · Keras · TensorFlow · Scikit-learn · TFLite

ML PLATFORM & WORKFLOW

SageMaker · Azure ML · Vertex AI · KubeFlow · MLFlow

SCALEOUT EDGE

Server-side

REST API & Web UI: Manage sessions, models, and fleet configuration. No custom client required.

Model Registry: Version, stage, and distribute models.

Observability Backend: Fleet-wide dashboards for training metrics, node health, and model lineage.

CLOUD PLATFORMS & INFRASTRUCTURE

AWS · Oracle · Kubernetes · Azure · OpenStack · Google Cl

SCALEOUT EDGE

Client-side

Lightweight SDKs that run on constrained hardware. Three languages, minimal dependencies.

Python SDK: Full training and inference. For devices with standard compute, Jetson, x86 edge servers, on-prem nodes.

C++ SDK: For embedded systems and resource-constrained hardware where Python isn't viable.

Kotlin SDK: For Android-based edge devices and mobile platforms.

HARDWARE TARGETS

NVIDIA Jetson · Raspberry Pi · x86 Edge Servers · Android

Edge AI modules

Federated-ready edge AI modules with built-in models, training setups, and workflows.

Turning months of integration into days.

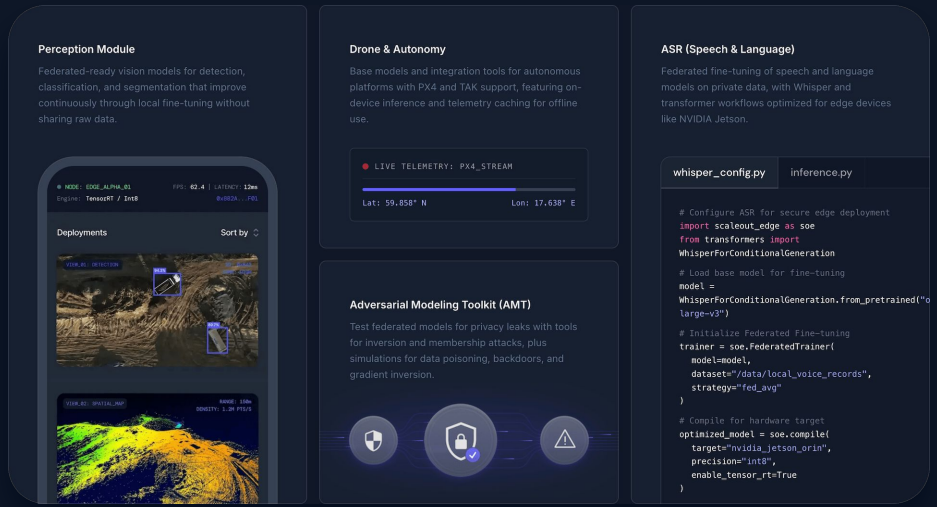

- **Computer Vision**
  Federated-ready vision models for detection, classification, and segmentation that improve continuously through local fine-tuning without sharing raw data.
- **Drone & Autonomy**
  Base models and integration tools for autonomous platforms with PX4 and TAK support, featuring on-device inference and telemetry caching for offline use.
- **Speech & Language**
  Federated fine-tuning of speech and language models on private data, with Whisper and transformer workflows optimized for edge devices like NVIDIA Jetson.
- **Adversarial Modeling (AMT)**
  Test federated models for privacy leaks with tools for inversion and membership attacks, plus simulations for data poisoning, backdoors, and gradient inversion.

Vision Ground Node

Deploy one node for immediate capability.

Add more and federated learning compounds the value across sites.

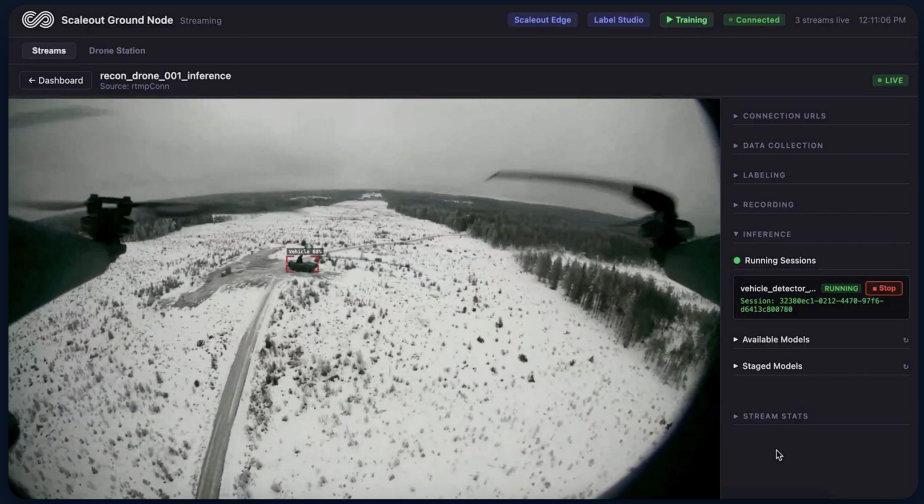

- **Full ML model lifecycle**
  Ingest, select, annotate, train, and deploy, all running locally on the edge node with no data leaving the site and no cloud dependency.
- **Field validated**
  Tested in live exercises with real drone feeds, annotation workflows, and federated training running across multiple nodes. Not a prototype.
- **Runs on standard hardware**
  Deploys on ruggedised edge servers, command vehicles, or laptops. No specialised infrastructure required.
- **C2 integration**
  Detections and video streams connect directly to OpenTAK and other C2 systems, distributing to end user devices on the network.

One platform. Three operational contexts.

Different missions, same challenge: distributed data that cannot move and AI that must keep improving.

Tactical Edge

Autonomous Intelligence in Denied Environments

Drones, unmanned systems, and forward-deployed sensors operating in contested or bandwidth-denied environments.

- On-device training directly on sensor streams.
- Sensitive raw data never crosses classification boundaries.

NATO DIANA · BAE Systems Bofors · FMV

Autonomous Transport

Fleet-Wide Learning Without Bottlenecks

Thousands of vehicles generating terabytes of telemetry daily — too much to centralize, too valuable to ignore.

- Build collective fleet intelligence without exposing individual routes.
- Compressed weight updates transmit instead of raw telemetry.

Traton / Scania · BMW

Cross-Jurisdictional

Federated Intelligence Without Data Exchange

High-value datasets in government and healthcare siloed by privacy law and classification.

- GDPR and HIPAA compliance enforced architecturally.
- No party sees another's raw data or proprietary information.

Banks · Government · Healthcare

Where we operate

Demonstrated across defence, automotive, space, and industrial domains.

About the company

AI systems are increasingly built on data that cannot be centralised. Medical records, industrial data, defence systems and connected devices all operate in environments where data must remain local.

Our mission is to enable AI training and operation across these environments without moving sensitive data. By enabling organisations to collaborate on machine learning while retaining full control over their data, we help build more secure, resilient and trustworthy AI systems.

Scaleout was founded in 2018 by researchers from Uppsala University working on large-scale distributed systems and scientific computing. The company was selected as a 2025 NATO DIANA Innovator.

We are a team of 20, with 7 PhDs and 2 Associate Professors.

Investors:

FAIRPOINT capital · Navigare Ventures · B50 Beijerstiftelsen · SEB · almi

Learn more

Detailed documentation for technical evaluation and integration planning.

Document

Platform overview

What the platform does, why it exists, and how federated learning enables adaptive AI at the edge. Start here.

To the platform overview →

Document

Technical brief

Architecture, deployment options, security model, device support, and integration patterns. For technical evaluators and systems integrators.

To the technical brief →

Document

Capability modules

Deep dive into the Vision, Drone & Autonomy, ASR, and Adversarial Modeling modules. Features, workflows, and supported hardware per domain.

To the capability modules →

Full technical documentation, API reference, and getting-started guides. The ground truth for implementation. To the documentation →

Sovereign Edge AI Infrastructure

scaleoutsystems.com/contact

AI sovereignty by design. Not by promise.