

The agent runtime that makes robots operational at scale

Hardware gives robots intelligence. mimOE gives them the ability to coordinate, self-organize, and operate autonomously across any silicon, any environment, at any scale.

NVIDIA Jetson/IGX

Intel Arc/Core Ultra

AMD Ryzen Embedded

Arm-based SoCs

x86 / RISC-V

THE GAP

Inference compute is solved. Inference operations are not.

Silicon vendors have solved acceleration. But running inference at fleet scale requires coordination, security, and operational resilience that no hardware vendor provides.

What hardware vendors leave to you

- Agent discovery across devices
- Fleet management at scale
- OTA updates per device context
- Offline and low-connectivity ops
- Zero-trust security at agent layer
- Fleet-wide observability
- Cross-vendor interoperability
- Zero-configuration agent setup

What mimOE provides out of the box

- Dynamic service mesh discovery
- Decentralized fleet coordination
- Context-based OTA updates
- Online, offline, multi-cloud ops
- Built-in zero-trust and API gateway
- Fleet-wide telemetry and tracing
- One runtime across all silicon
- Zero config: agents join seamlessly

THE POC TRAP

Most physical AI projects prove the concept.

Few make it to production.

The barrier is never the model. Across the robotics industry, teams demonstrate compelling prototypes and then stall. The missing layer is operational infrastructure: the ability to coordinate agents across a fleet, push updates safely, maintain security, and keep systems running when connectivity fails. This is not a research problem. It is an execution problem.

90%

of physical AI pilots never reach production
IDC / Gartner

No. 1


barrier: operational complexity, not compute
McKinsey, 2024

mimOE closes the gap

Same runtime from PoC to full fleet. No rebuild. No delay.

CORE CAPABILITIES

What mimOE brings to physical AI



Dynamic agent discovery

Agents join a live service mesh without static configuration. Peers found automatically.



Decentralized coordination

Agents self-organize without a central controller. Resilient to individual node failures.



Context-based OTA updates

Agents update based on their environment and workload, not a central push schedule.




Zero configuration

Agents discover, join, and connect with no manual setup. Deploy and run from day one.



Built-in API and MCP gateway

Inference services become discoverable and routable without extra infrastructure.



Zero-trust security

Baked into the agent layer, not bolted on. Every interaction authenticated and authorized.



Fleet-wide observability

Telemetry, logs, and tracing at every point across all devices from one control plane.



Seamless agent updates

Staged rollouts, automatic rollback, and peer-to-peer delivery across the fleet.



Offline-first operation

Full agent operation when disconnected. Cloud resumes seamlessly when available.



Heterogeneous hardware

One runtime across NVIDIA, Intel, AMD, Arm, and x86. Write once, deploy anywhere.



Unified inference API

REST APIs for ONNX and GGUF models. Dynamic routing between edge and cloud.

PLATFORM COMPATIBILITY

Runs on the silicon you already use

mimOE is hardware-agnostic by design, leveraging hardware acceleration automatically where available.

NVIDIA

Jetson Orin / Thor
CUDA acceleration

Intel

Arc / Core Ultra
OpenVINO

AMD

Ryzen Embedded
ROCm compatible

Arm

Cortex-A / Neoverse
CMSIS-NN

x86 / RISC-V

Industrial PCs
Cloud VMs

WHO IT IS FOR

Built for the people who ship robots

Robotics OEMs and hardware makers

Device providers, humanoid robot companies

- ▢ Ship devices agent-ready, mimOE pre-installed
- ▢ Managed OTA and fleet as recurring revenue
- ▢ Differentiate on software, not hardware specs
- ▢ Mixed silicon fleets, one runtime
- ▢ Same runtime from PoC to full deployment
- ▢ Zero config: agents connect out of the box

System integrators

Vertical solution builders, enterprise teams

- ▢ One runtime across all environments
- ▢ From project delivery to fleet operations
- ▢ Staged rollout, rollback, update-as-a-service
- ▢ Single observability plane, all customers
- ▢ Deliver production systems, not just pilots
- ▢ Seamless fleet updates, no manual steps

Solution providers

Manufacturing, healthcare, logistics, retail

- ▢ Focus on domain logic, not infrastructure
- ▢ Deliver turnkey solutions out of the box
- ▢ Cameras, sensors, legacy systems supported
- ▢ Cut undifferentiated engineering overhead

Developers

Agent developers, robotics software engineers

- ▢ Write agent logic once, any hardware
- ▢ No inference engine or serving setup
- ▢ Security, observability, coordination built in
- ▢ Prototype code runs in production at scale

CAPABILITY COMPARISON

What silicon vendors provide vs. what mimOE adds

Capability	Hardware vendor (typical)	mimOE
Inference acceleration	Requires manual integration	Agent-accessible via REST
PoC to production path	Not addressed	Same runtime, no rebuild
Agent runtime and discovery	Not provided	Dynamic service mesh
Zero configuration	Not provided	Automatic, out of the box
Fleet-scale management	Partner responsibility	Unified fleet management
OTA and seamless updates	Partner responsibility	Context-based, peer-to-peer
Zero-trust security	Not provided	Built-in
Offline operation	Not designed for	Full agent operation
Heterogeneous hardware	Vendor-specific only	Any compute platform

"Hardware makes robots powerful. mimOE makes them operational: discoverable, coordinated, observable, and resilient at fleet scale, across any silicon."

mimik Technology — Device-First Continuum AI