LF EDGE

One Control Plane for the Edge: Unifying Containers, VMs and AI Models With Open Horizon

How Mainsail Industries is putting everything its Starlight platform runs under one policy model, with no dependency on a live connection back to base.

mainsail

At a Glance

LF Edge Project
Open Horizon

Company
Mainsail Industries

Deployment Status
Standalone management hub and console running in Mainsail's lab, tested against real edge agents

Executive Summary

Running applications and AI at remote sites is difficult when connectivity is unreliable, hardware is physically exposed, and each workload requires a different management tool. For defense and critical infrastructure operators, that complexity can mean slow updates, security gaps and costly on-site maintenance.

Mainsail Industries is building on Open Horizon, an LF Edge project, to bring policy-driven deployment to its secure edge platform, Starlight. The goal is to manage containers, legacy virtual machines, microVMs, operating system images and AI models through one control plane, helping devices keep working when disconnected and reducing the tools operators must learn and maintain.

Mainsail has tested a hardened management hub and web console in its lab with real edge agents, demonstrating simplified onboarding, signed service publishing, fleet visibility and automated recovery. Planned extensions will add broader workload support, independent AI model delivery and confidential computing protections.

By building on a vendor-neutral, open source project and planning upstream contributions, Mainsail aims to make secure edge operations easier to automate while giving customers greater control over their infrastructure. The work offers a path toward deploying and updating edge AI with less manual effort, fewer operational risks and no required cloud service.

> "Open Horizon lets us manage containers, VMs and AI models on Starlight with one set of policies. Confidential computing keeps those workloads' data protected while in use."
>
> *Brad Sollar, Co-Founder and CTO, Mainsail Industries*

Challenge

Edge fleets running video analytics, sensor fusion, industrial control and AI inference have outgrown the tooling built to manage them.

- **Separate tools for separate workloads.** Containers, VMs and OS images are each deployed and updated through their own system, and many sites still depend on legacy applications that only run as full VMs.
- **Fleet updates are mostly manual.** Rolling out a new app, model or OS version across hundreds of devices is slow, hard to audit, and a failed rollout can require a physical visit to a site that isn't easy to reach.
- **AI adds a new kind of asset to manage.** Model weights are large, updated frequently and often the most valuable IP on the device, with no established way to version and ship them independently of the runtime that serves them.
- **Devices operate outside the owner's physical control.** Workloads, data and models need to stay protected even if a device is stolen, tampered with, or compromised at the host.

Solution

One control plane instead of five

Mainsail is making Open Horizon the single control plane for everything Starlight runs at the edge. Open Horizon already uses policy to deploy containerized services across large, intermittently connected fleets. Mainsail is extending that workflow to VMs, microVMs, OS images and AI models.

Operators declare what runs and where it may run. The Open Horizon agent on each Starlight node carries out deployments and updates, even while cut off from the hub.

- Mainsail has built a hardened, standalone management hub and web console and tested it in its lab against real edge agents. The hub is tuned for Starlight but doesn't require it.
- **Management hub.** Exchange, Agbot, Cloud Sync Service (CSS), FDO device onboarding and OpenBao for secrets. Runs as Podman containers on RHEL 10 with SELinux and fapolicyd enforcing.
- **Gateway.** Envoy terminates TLS, rate-limits the API, restricts admin endpoints to admin networks and keeps internal container networks private.

- **Web console.** Fleet overview, geographic fleet map, services, deployment policies, agreements and hub health. Operators publish a signed service by pasting its definition or importing it from a GitHub repository, then check the source for updates.
- **One-command onboarding.** The hub serves agent packages. Each node gets a one-time token, so nobody copies credentials by hand.
- **Self-healing.** Starts at boot, unseals secrets automatically and recovers from container failures on its own.

Roadmap

CAPABILITY	WHAT IT ADDS	WHY IT MATTERS
VMs as services	VM images are published, signed, matched to nodes by policy, and updated through the same agreement workflow as containers, running on Starlight's hypervisor.	Legacy and vendor applications that can't be containerized get automated lifecycle management.
MicroVMs	The same lifecycle extended to lightweight microVMs.	VM-level isolation for untrusted or multi-tenant workloads, at near-container startup time and density.
AI models	Model runtimes deploy as services while weights ship and version separately through the Model Management System (MMS) and CSS; policy places models only on nodes with the right GPU, NPU and memory.	New model versions ship without a new container build or software release.

Why LF Edge's Open Horizon

Mainsail evaluated several orchestration approaches first. Most assume a Kubernetes cluster at every site, a constant link to a central control plane, or a specific cloud. None of that holds at the edge. Open Horizon was built for large fleets of individual devices that go offline for long stretches. Its agents negotiate and maintain their own deployments, so a node keeps working when the network doesn't.

Policy expresses intent. "This model runs only on GPU nodes in secure facilities." Operators stop scripting individual devices, and the same policies place AI and confidential workloads.

The trust model fits. Signed service definitions, per-node identity and zero-touch onboarding through FIDO Device Onboard (FDO) match Starlight's security posture.

One workload model extends. Services and agreements generalize from containers to VMs, microVMs, OS images and models. Mainsail extends one control plane instead of building five.

Model delivery is already separate. MMS and CSS split model and data delivery from app deployment. Large weights need that split.

Why Open Source

An edge control plane is long-lived, mission-critical infrastructure. It has to outlast any vendor's roadmap, pricing or acquisition. A vendor-neutral project under LF Edge gives customers that. They can audit every line of code that controls their fleet and self-host it in air-gapped or sovereign environments, with no required cloud service. For public sector and regulated buyers, that transparency is usually a requirement.

Open governance also makes Mainsail's own work count for more. Mainsail builds on Open Horizon and contributes upstream instead of carrying a private fork. Its VM, microVM, bootc, AI model and confidential computing extensions become shared ground that other vendors, integrators and users can adopt and improve. The community brings use cases and review Mainsail couldn't get alone.

Results to Date

The standalone hub and console are running in Mainsail's lab today, managing real edge agents.

- **Onboarding** — a single command replaces a manual, multi-step agent setup.
- **Publishing** — signed services can be pushed from the console or pulled directly from a Git repository.
- **Fleet visibility** — the full fleet is visible on a live map.
- **Resilience** — the hub recovers from container failures and restarts on its own, verified through automated chaos testing.
- **Hardening** — SELinux and fapolicyd enforce, internal networks stay private, a TLS gateway rate-limits traffic, and no secrets are stored in configuration files.

Together, these give Mainsail one operating model for containers, legacy VMs, microVMs, the host OS and AI models, under the same policies, signing chain and console, rather than a separate one for each.

What's Next

Mainsail is folding the hub into Starlight and working through the roadmap in order:

1. Add VM and microVM services to the Open Horizon agent on Starlight, with signed images and policy-driven updates.
2. Distribute models and weights through MMS with hardware-aware placement.
3. Run confidential containers and VMs with attestation-gated key release.
4. Contribute these extensions upstream and develop the designs in the open with the LF Edge community, starting with secure AI at the edge.

Get Started with Open Horizon

Open Horizon is a platform for managing the service software lifecycle of containerized workloads and related machine learning assets. It enables autonomous management of applications deployed to distributed webscale fleets of edge computing nodes and devices without requiring on-premise administrators.

GitHub: github.com/open-horizon

Web: lfedge.org/projects/open-horizon