

v3.4-EDG-001 – Edge-Native Deployment Logic & Resource-Aware Scaling

Document Title	Edge-Native Deployment Logic & Resource-Aware Scaling
Version	v3.4
Document ID	v3.4-EDG-001
Date	2025-03-22
Author	Take Back Your Data – Edge Architecture Group
Document Type	Public / Certification / Internal

1. Purpose & Scope

This document defines the logic for edge-native deployment and adaptive scaling in MaxOneOpen forks. It ensures efficient execution across heterogeneous devices while maintaining full sovereignty, performance, and verifiability.

2. Edge-Native Deployment Logic

- Deployments must support schema-driven configuration without central orchestration
- Twin lifecycles must adapt to local context, hardware profile, and policy state
- Local compute nodes must operate independently of cloud signaling
- Deployment triggers include identity presence, resource match and mobility anchor

3. Resource-Aware Scaling Logic

Scaling Trigger	Fork-Level Action	ZK-Proven State
CPU/GPU Saturation	Activate adjacent twin	Load parity + init trace
Thermal Drift	Route to low-heat node	Snapshot + thermal hash
Bandwidth Spike	Limit queue depth	Policy replay anchor
Idle State	Suspend non-critical twin	TTL audit + sealed freeze

4. Certification Hooks

- Forks must support autonomous scaling and rebalancing without centralized control
- All scaling actions must be provable, policy-driven, and schema-aligned
- Local resource use must remain within certified operational bounds

5. Certification Triggers

- Cloud fallback, external orchestration or missing scaling proofs disqualify fork
- Runtime misalignment with policy-defined thresholds invalidates certification

6. Certification Relevance

MaxOneOpen forks must demonstrate edge-native behavior with adaptive, resource-aware deployment logic. Only forks with sovereign scaling mechanisms qualify for certification in the Edge context.