

v3.4-AI-001 – Distributed AI Runtime & Edge-Inference Control

Document Title	Distributed AI Runtime & Edge-Inference Control
Version	v3.4
Document ID	v3.4-AI-001
Date	2025-03-22
Author	Take Back Your Data – Edge AI Systems
Document Type	Public / Certification / Internal

1. Purpose & Scope

This document defines the MaxOneOpen AI runtime architecture and logic for edge-based model execution. It ensures AI tasks can be executed in a distributed, secure and sovereign manner without centralized inference dependencies.

2. Runtime Architecture Design

- AI runtimes are modular, isolated, and container-bound
- Inference happens locally or in federated edge clusters
- Control logic enforces per-model trust and version validation
- Models run on demand, under token-governed execution contracts

3. Edge Inference Flow

Phase	Action	Security Control
Model Load	Schema + signature check	Hash-verified source only
Runtime Init	Container scoped boot	Resource-bound + tokenized
Inference Call	Data pull + compute	Context-isolated execution
Result Return	Post-token validate	ZK-notarized proof

4. Isolation & Failover Mechanism

- Inference containers operate independently
- Each failure triggers twin-internal recovery logic
- Failed models may not cascade or retry without token renewal
- No centralized fallback or cloud redirect allowed

5. Certification Logic

- Forks must expose inference entrypoints + audit traces
- Certification checks include: model integrity, runtime validation, token logics
- Edge-only and offline compatibility must be demonstrable

6. Certification Relevance

Distributed and sovereign AI inference is a core certification condition. No model may run outside the validated edge-controlled runtime architecture.