

MaxOneOpen – Management Summary

Status: March 2025

Vision

MaxOneOpen is the first fully auditable, sovereign, and decentralized AI platform, designed to fundamentally diverge from all BigTech LLM systems in both architecture and economics. It returns control, scalability, and security to operators—without central authorities, proprietary dependencies, or computational compromise.

Core Principles

- **Sovereign:** No centralized infrastructure, no external control, no API lock-in
- **Auditable:** Full transparency, zero-knowledge secured, forkable and verifiable
- **Efficient:** Edge-optimized, ultra-low CAPEX/OPEX, high inference performance
- **Secure:** No tokens, no gateways, no cloud-based attack surfaces
- **Modular:** Each component (training stack, inference runtime, benchmarking) is independently usable

Technological Highlights

- Capsule-based inference with < 50 ms latency
- < 250 MB RAM per model instance
- Local replicability with under €500 hardware budget
- Full SHA-256 encryption of all twin states
- Real-time inference twins with audit-trigger logs
- Seamless interoperability with zero-trust networks, decentralized identities, and sovereign storage systems

Technical Benchmark (CTO-Validated)

Criteria	MaxOneOpen	GPT-4
FLOPS per Inference	5.4 TeraFLOPS	4.4 TeraFLOPS
Tokens per Second	27.8 tokens/s	20 tokens/s
Training Energy	466,172 kWh	1,287,000 kWh
Inference Energy Usage	< 1 Joule	Up to 5 Joules
Architecture	Specialist, modular	Monolithic, centralized
Data Privacy	Zero-Knowledge, UDUH	Centralized storage & profiling
Security	98% PenTest success, 680 min TTB	85%, 320 min TTB

Criteria	MaxOneOpen	GPT-4
Scalability	Up to 2 billion users	Unknown
AI Inference	Standard hardware (32 GB RAM)	Specialized AI chips required
Cloud Dependency	None (Edge & On-Premise)	Yes (Cloud-only)

Economic Metrics (from official documentation)

- **CAPEX:** < €500 per node (local setup)
 - **OPEX:** < €1/month per edge module (power & maintenance)
 - **TCO:** Up to 99% savings over centralized LLM services
 - **ROI cycle:** < 2 months under standard usage conditions
-

Differentiation from BigTech

Category	MaxOneOpen	BigTech LLMs
Control	Edge-operated by users	Platform-operated
Infrastructure	Local, decentralized	Centralized, cloud-based
Transparency	Fully auditable	Proprietary & opaque
Forking	Allowed and cryptographically enforced	Prohibited or limited
Runtime Cost	Minimal, localized	High, volume-dependent
Data Sovereignty	Fully user-owned	Externally controlled

Readiness

MaxOneOpen is fully documented, verifiable, buildable, and production-ready. All modules have passed final validation. Any competent team can implement the system within 12 weeks—without relying on centralized infrastructure, cloud services, or license restrictions.

Target Stakeholders

- Governments & public institutions
 - Enterprises with critical security or data compliance needs
 - NGOs, educational networks, and independent media
 - Edge providers and sovereignty-first digital infrastructures
-

Next Steps

- Public release on April 1, 2025
- Reference architecture distribution via Git & IPFS

- Onboarding for forks, audits, and deployment partners

Contact for evaluation, certification or fork access:

info@take-back-your-data.com