



Sovereign AI Engine

Open-source private AI infrastructure for when SaaS AI isn't an option



What It Is

QXLI is an open sovereign agentic AI engine — modular by design, built on best-in-class open-source components, and engineered for organizations that require production-grade private AI on infrastructure they own, operate, and govern.

Unlike public AI tools and SaaS platforms that route your data through external systems and foreign jurisdictions, QXLI keeps every prompt, model, embedding, and output securely inside your environment. No external model calls. No vendor access. No data leaving your boundaries.

Originally assembled to meet our own needs, we are now making QXLI available to help other organizations unlock the power of secure, sovereign AI — without the complexity, compromises, and dependence on external platforms.

What's Inside

QXLI integrates the following into a single, cohesive sovereign AI engine:



Private LLM Runtime — Run and fine-tune leading open-source language models locally. Model-agnostic. No API dependency. Switch or pin models on your terms.



Workflow Automation — n8n-powered automation layer connecting AI outputs to real business processes and enterprise systems.



Agentic Orchestration — Deploy intelligent AI agents that operate autonomously across workflows — retrieving knowledge, making decisions, and executing tasks inside your environment.



Security — Granular role-based access, identity-aware control, and optional Zero Trust security. Air-gap capable for the most regulated environments.



RAG Pipeline — Vector search and document ingestion powered by pgvector. Your organization's knowledge, made instantly queryable by AI.



Observability & Governance — Full operational visibility. Every model interaction, agent action, and system event — logged, traceable, and under your control.

Who It's For

QXLI is built for organizations where data sovereignty is non-negotiable:



Built on Trusted Foundations

QXLI is assembled from leading open-source infrastructure — Langchain, Ollama, n8n, pgvector, NGINX, and more. No proprietary lock-in. Full auditability and extensibility by design.

As part of the NVIDIA Inception Program, QXLI also leverages NVIDIA-optimized models, deployment blueprints, and agentic AI toolkits to help organizations run high-performance private AI workloads on infrastructure they control.

The QXLI Difference

	Public / SaaS AI	QXLI Sovereign AI Engine
Data location	Vendor's cloud, foreign jurisdiction	Your infrastructure, your jurisdiction
Model control	Vendor's roadmap, vendor's schedule	Your choice, pinned, fine-tuned, swappable
Governance	Vendor's audit trail	Your audit trail, your lifecycle
Compliance	Best-effort under SaaS terms	Aligned with data residency and sector mandates
Air-gap capability	No	Yes

Get Started

Self-Deploy — Access the QXLI AI engine via GitHub. Deploy it on your own infrastructure with your own team.

Request a demo — See the QXLI AI engine running in a sovereign deployment environment.