



Sovereign AI Engine

The open sovereign agentic AI engine for organizations where SaaS AI isn't an option



Zero Data Egress
Architecture



Deployed on your
Infrastructure



Engine available on
GitHub



Inception
Program




Google Cloud
Partner



The Enterprise AI Dilemma


CIOs face an impossible choice when deploying AI

01  **Option 1: Public Cloud AI**

- Sensitive data leaves your infrastructure
- Unpredictable costs that scale with usage
- Vendor lock-in to proprietary services
- Compliance teams block deployment

— // ————— // —

Great for general tasks, but regulated industries can't afford the risk.

02  **Option 2: Build From Scratch**

- 18-24 months development timeline
- Requires scarce, expensive AI/ML talent
- \$2M-5M upfront investment
- Constant maintenance as models evolve

— // ————— // —

Fewer than 10% of enterprise AI projects advance beyond pilot." — McKinsey



The Core Problem:

Current options force you to choose between power and privacy, or between speed and control.

QXLI: The Open-Source Third Way

QXLI gives teams a faster path to sovereign AI: a modular AI engine, based on open-source foundations, that runs on infrastructure you control — with optional enterprise support to take it from GitHub to production.

Open by Design



Inspectable, extensible, GitHub-available engine built from leading open-source components.

Governed by Design



Policy-controlled access, auditability, secure deployment patterns, and enterprise-ready controls.

Owned by You



Runs 100% on your infrastructure (cloud or on-premises). No vendor lock-in. True strategic independence.



2-4 Weeks

Rapid Deployment

From pilot to production



0%

Data Egress

Information stays internal



100%

Ownership


Full platform control

Origin Use Case: Space Armour



Space Tech Modular, integrated, space-optimized AI Engine Deployment aboard Satellites / Space Stations in-orbit

The Challenge

-  **Data Overload**
Satellites generate high volumes of data but lack the ability to process or act on it in real time
-  **Process Bottleneck**
Reliance on ground-based processing creates latency, bandwidth dependency, and limits true autonomy
-  **Lack Of Trusted Operations**
Space operations lack robust security, and in-orbit autonomous decisions have no verifiable trail



The Solution

-  **Secure Autonomous In-Orbit AI**
QXLI's modular AI Engine, integrated with cryptographic integrity, and Zero Trust security
-  **NVIDIA Inception Powered**
QXLI AI Engine Deployed on the NVIDIA Jetson Orin Platform with Jetpack OS and NVIDIA tooling
-  **Accelerated Deployment**
A productized, production-ready AI stack—integrated, pre-configured, and ready for launch in under 30 days

The Impact

- 100%**
Autonomous Edge AI
On-ground → In-orbit
- NVIDIA**
Powered
Hardware and Software
- 30 Days**
Rapid Deployment
Ready-to-go

Engineered for Space. Optimized for Enterprise.

Originally developed for Space Armour, QXLI is now being made available as an open-source sovereign AI engine for enterprises, developers, and partners.

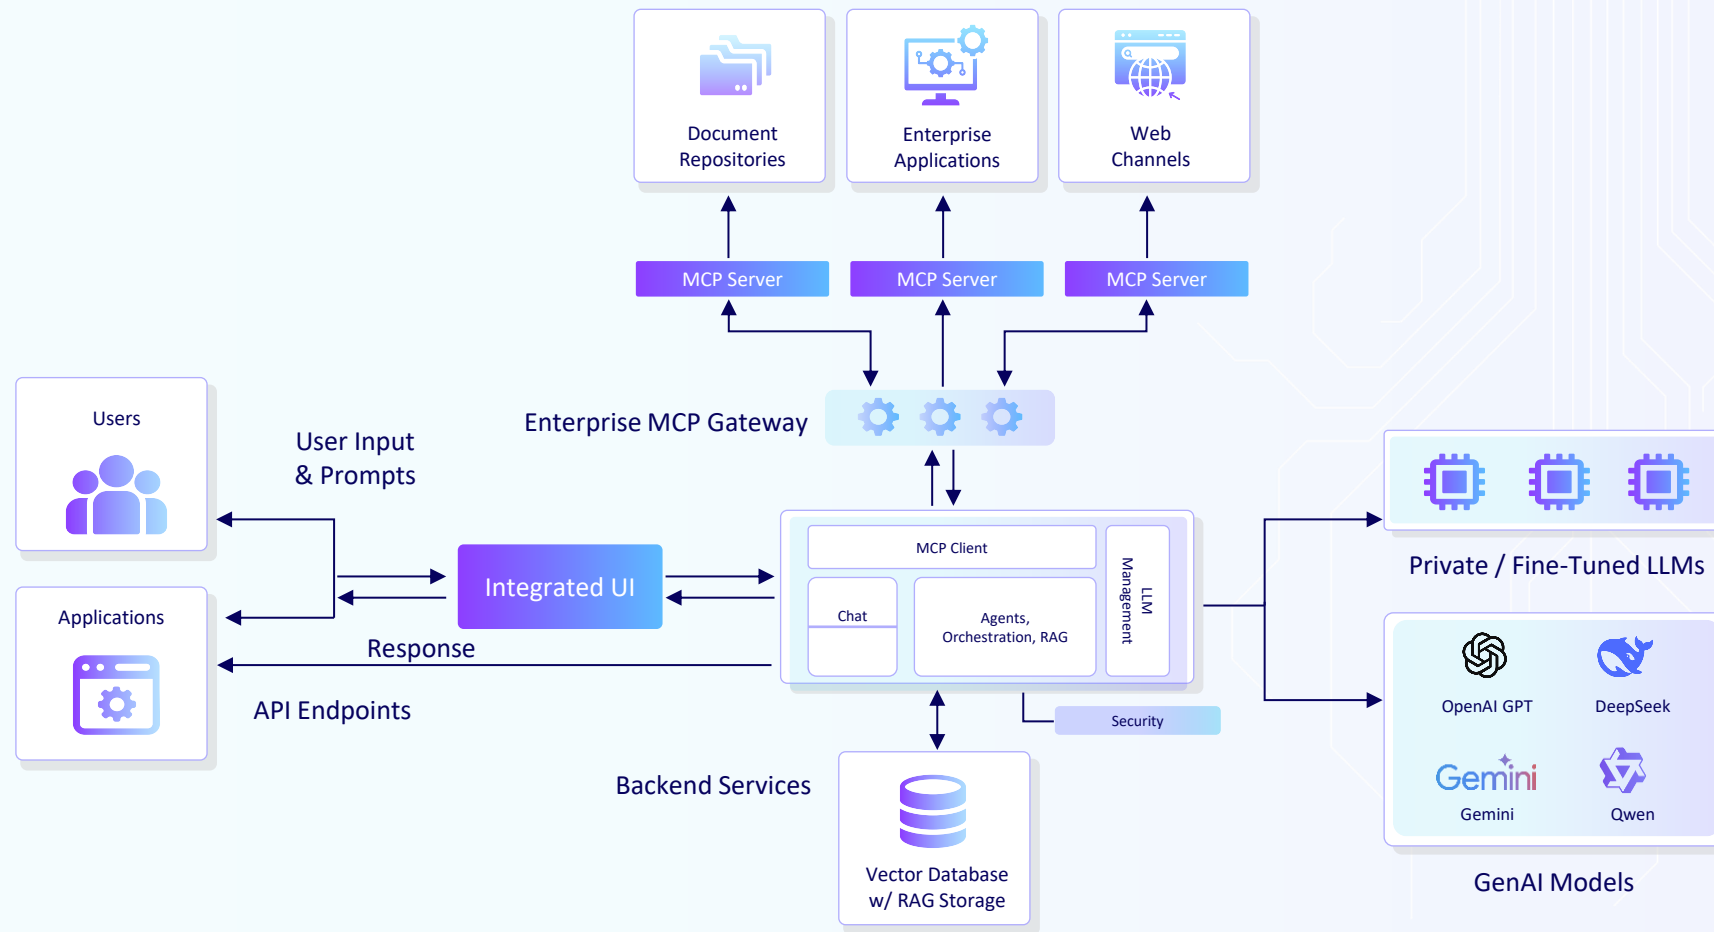
QXLI Architecture – Open, Modular, Best-of-Breed AI

Leading open-source components, optimized private LLMs, and advanced agent tools – Integrated, configured, ready-to-deploy



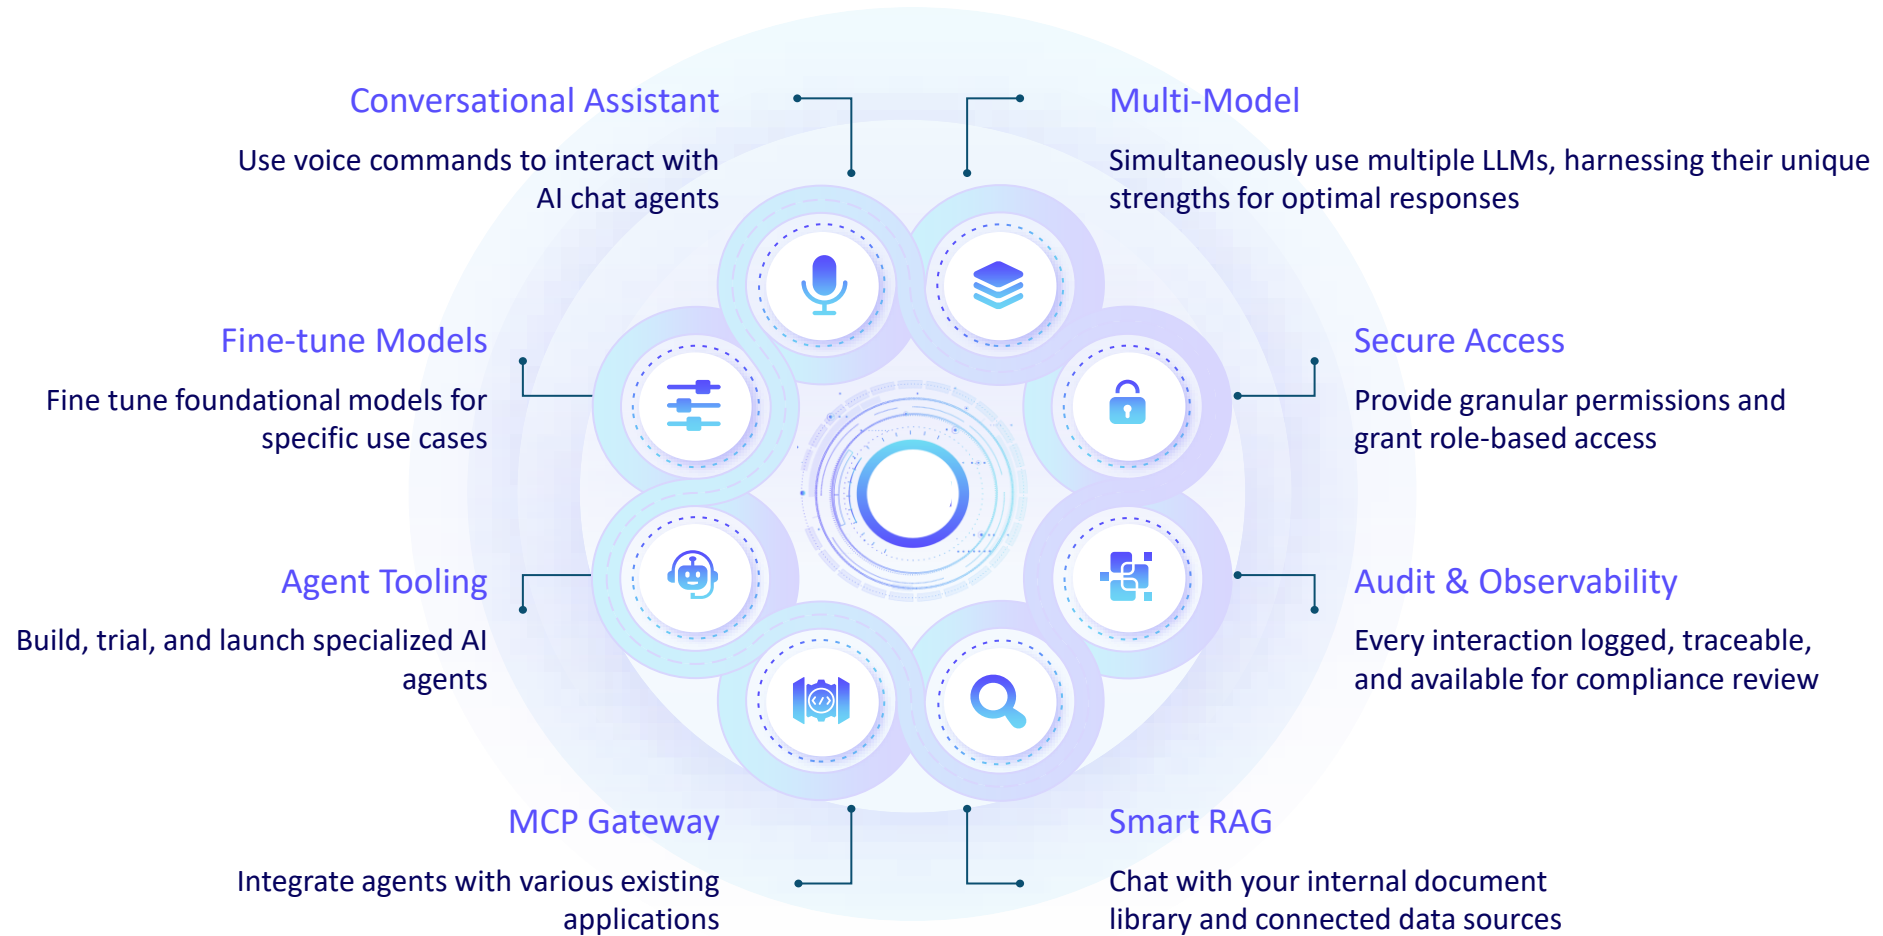
The QXLI 'MCP-Powered Agentic RAG' Blueprint

The Technical Foundation for Enterprise AI



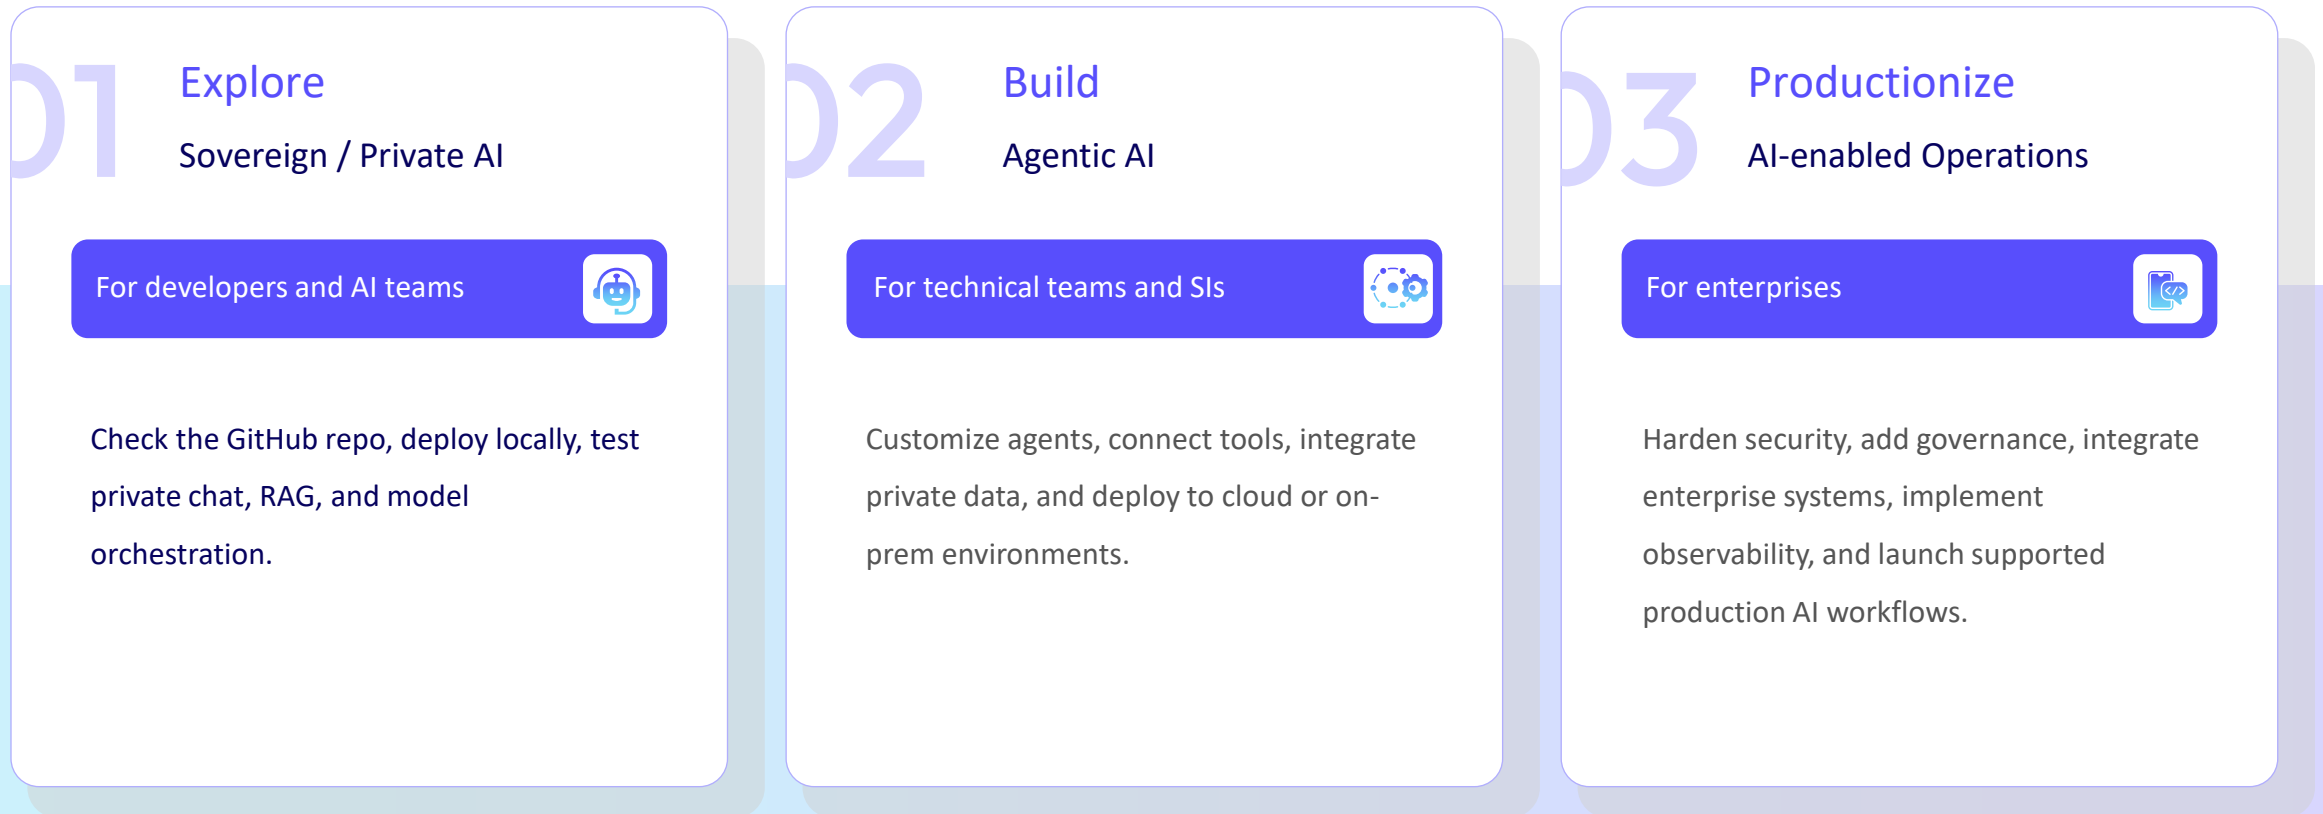
Unified Capabilities for Enterprise-Grade Private AI

From conversational interfaces to agent orchestration, secure access, and smart RAG — all the critical components required to deploy governed AI at scale.



One Engine, Three Ways To Engage

Progressive Capabilities for Every Stage of Enterprise AI Adoption



QXLI vs. SaaS AI

When Data Sovereignty and Cost Predictability Matter

Feature

Data Sovereignty

Cost Model

Vendor Lock-in


Agent Source Code


Deployment

SaaS AI

 Data leaves enterprise boundaries






 Unpredictable (Per-token)

 High (Proprietary APIs)

 Black Box

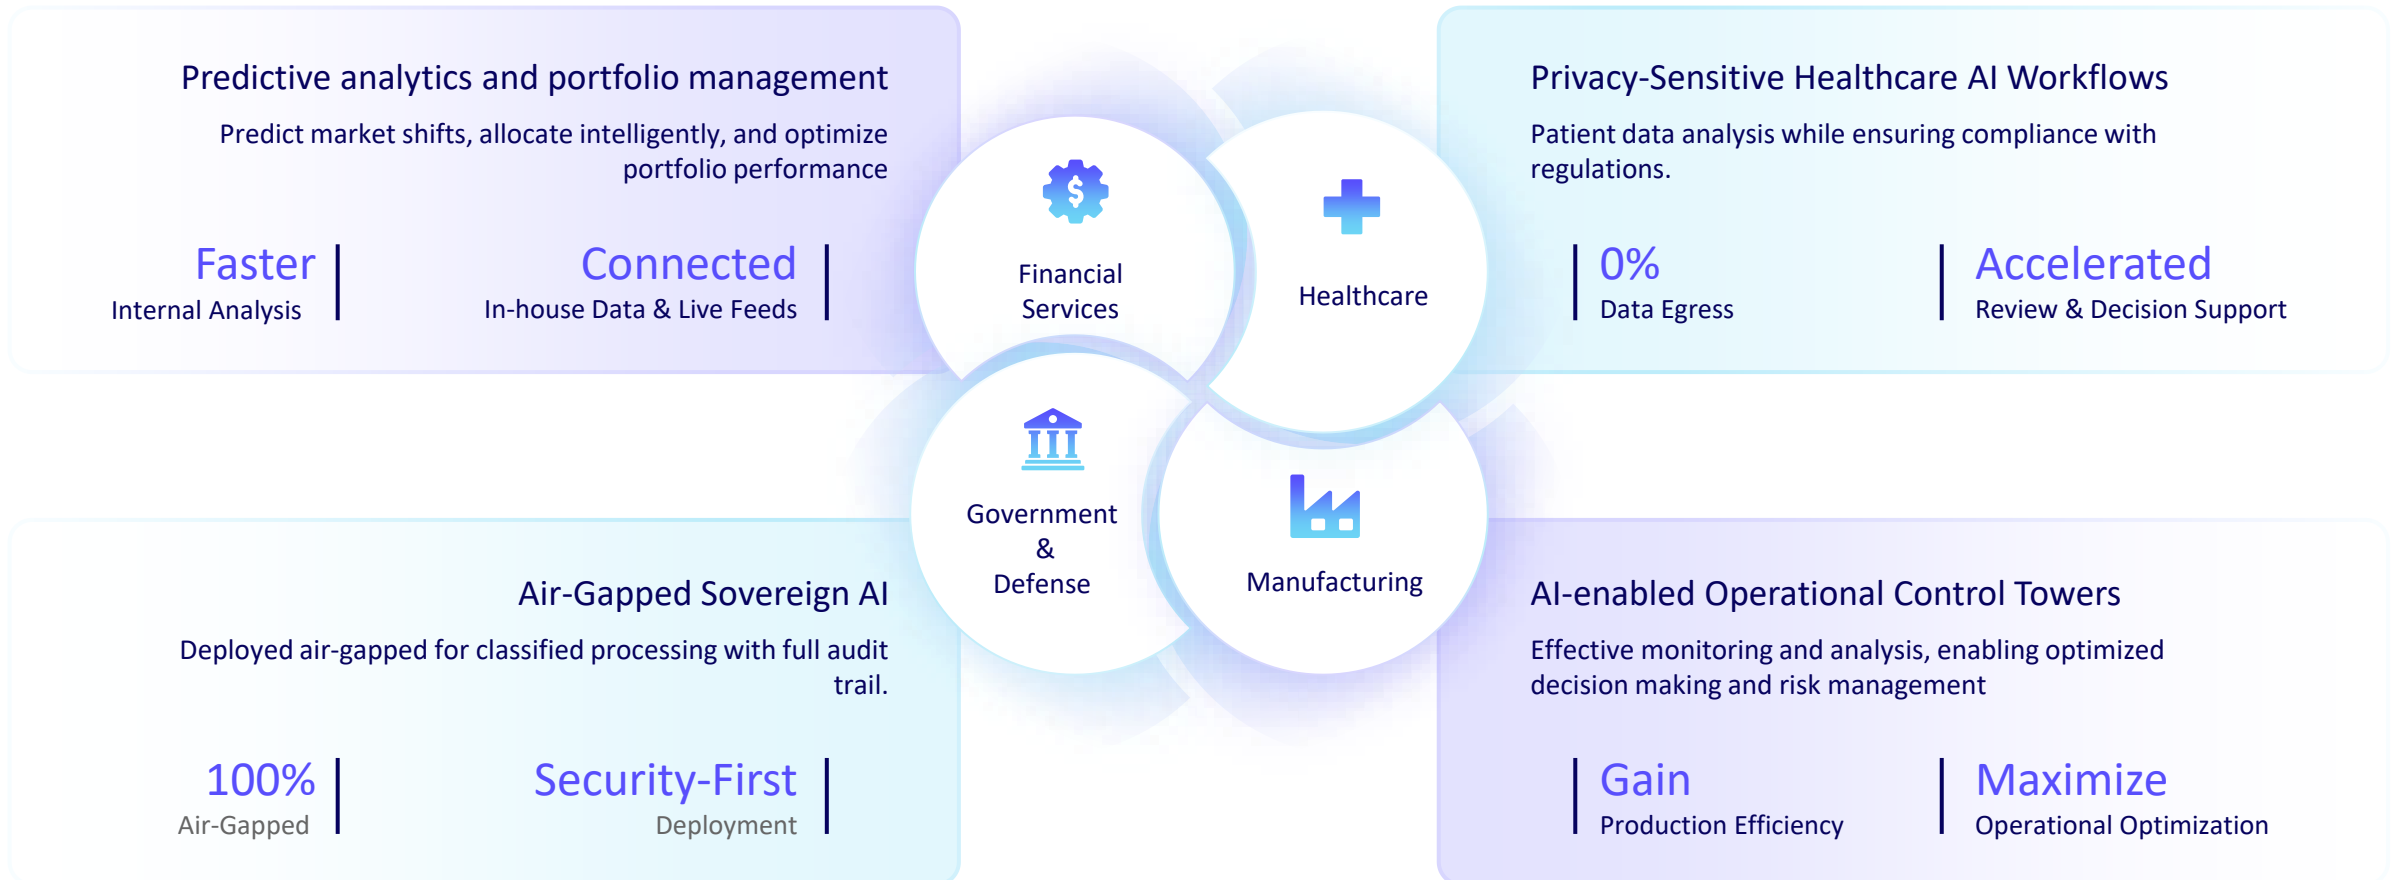
 Public Cloud Only

QXLI Open Sovereign AI

-  0% Data Egress
-  Predictable (Fixed Infra)
-  Zero (Open Standards)
-  Full Access
-  Air-gap / On-prem / Cloud

Where Sovereign AI Matters Most

QXLI's emphasis on privacy and security makes it ideal for organizations that require AI operations to be conducted within a controlled and protected environment



Launch Your Own Private Enterprise AI Factory

From Concept to Production in 4 Weeks

WEEK
01



Discovery & Design

Define scope, review architecture, and hardware sizing.



WEEK
02



Infrastructure Setup

Deploy infrastructure, setup QXLI AI stack, and configure components.



WEEK
03



Core AI Engine Activation

Ingest data/documents, setup RAG pipelines, and customize agents.



WEEK
04



Production Launch

User acceptance testing (UAT), team training, and go-live.

Zero Trust Ready



Edge Capable Deployment



Modular & Extendable



Data Sovereignty & Residency



Get The QXLI AI Advantage



Try it yourself — Deploy the QXLI AI engine on your own infrastructure via GitHub.



Request implementation support — Get help with deployment, integrations, custom agents, and production rollout.



Website

www.qxli.com



GitHub

github.com/SpaceArmour/qxli-sovereign-ai-engine/



Email

info@qxli.com