The platform of last resort for data security.

On-premise Sovereign AI
engineered specifically for
high-trust enterprises.

Intelligence is most powerful when it is entirely owned.

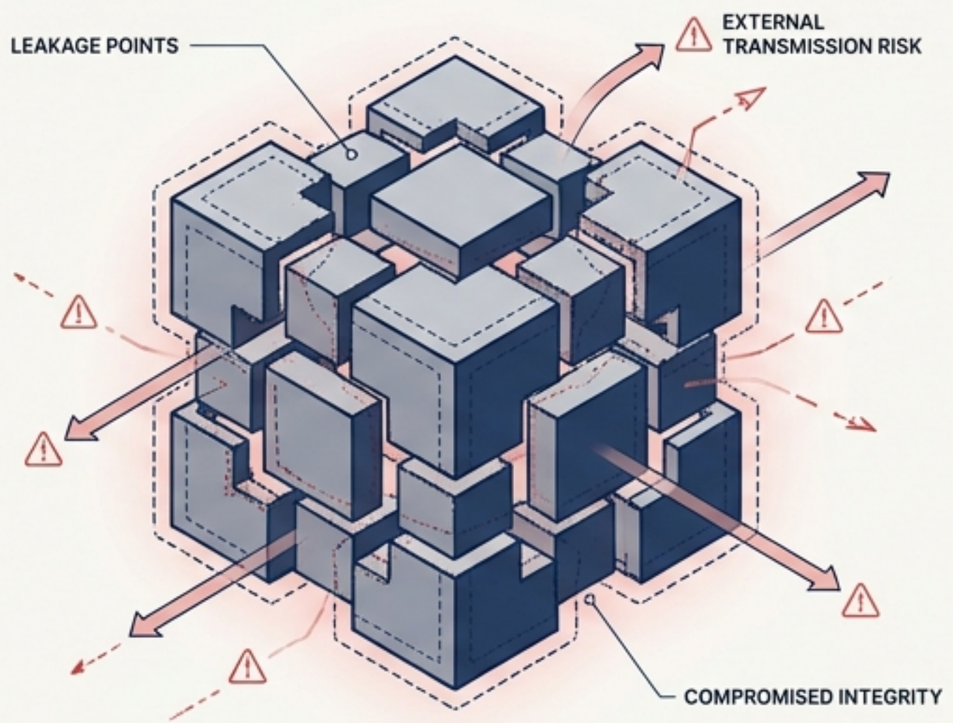

The Vulnerability

Public LLMs force high-trust enterprises into an impossible choice: sacrifice advanced AI capabilities or surrender mission-critical data sovereignty. For organizations facing rigid regulatory requirements, external data transmission is an unacceptable risk.

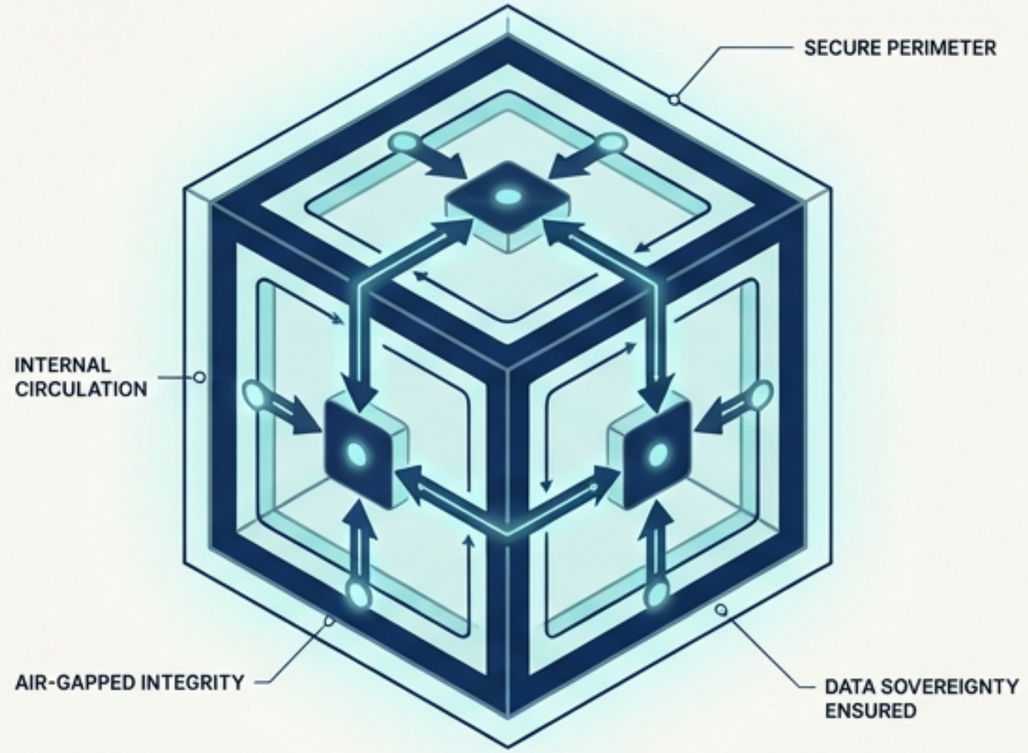

The Sanctuary

OmniNStack acts as the central nervous system for organizations where data sovereignty is absolute. It guarantees organizational autonomy by bringing the intelligence to the data, rather than sending the data to the intelligence.

High-trust operations require structural guarantees, not just privacy promises.

	Public/Cloud AI	OmniNStack Sovereign AI
Network Reliance	Dependent on public internet.	Engineered for total air-gapped isolation; zero public internet connection required.
Vendor Control	Strict vendor lock-in to specific proprietary models.	Multi-Model Gateway; seamlessly toggle **between local private LLMs** and public APIs.
Security & Governance	Shared responsibility model, opaque auditing.	Deterministic Zero-Trust via Policy-as-Code (OPA) and immutable identity-aware auditing.
Hardware Strategy	Fully dependent on external provider infrastructure.	Self-hosted, hardware-agnostic foundation optimized for NVIDIA TensorRT and Apple MLX.

Three core pillars engineer an environment of absolute autonomy.

1. Absolute Sovereignty

Function seamlessly in total air-gap environments. Provide absolute privacy for the most sensitive operational data without compromising processing power.

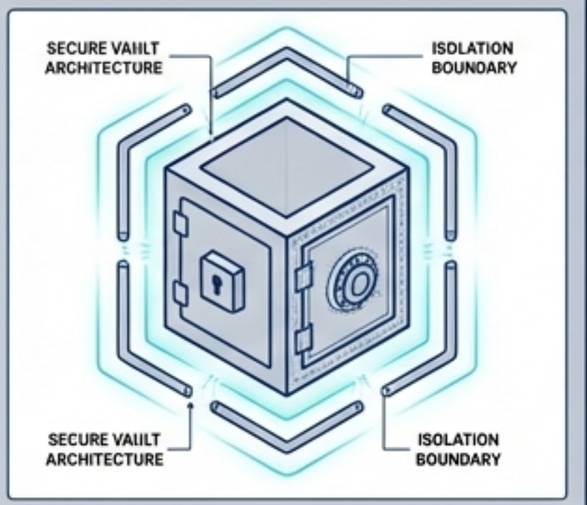

INTEGRATED SYSTEM FLOW

DATA PATHWAY

2. Multi-Model Gateway

Eliminate vendor lock-in permanently. A secure configuration dashboard allows instant toggling between private local LLMs and public APIs based on real-time security requirements.

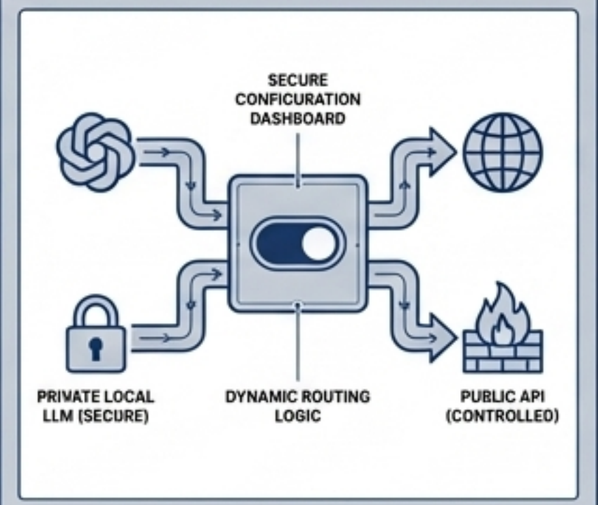

INTEGRATED SYSTEM FLOW

DATA PATHWAY

3. Agentic Knowledge

Deploy autonomous AI agents that act as tireless human researchers. Bridge highly secure local documents with global intelligence via secure, citation-backed Retrieval-Augmented Generation (RAG).

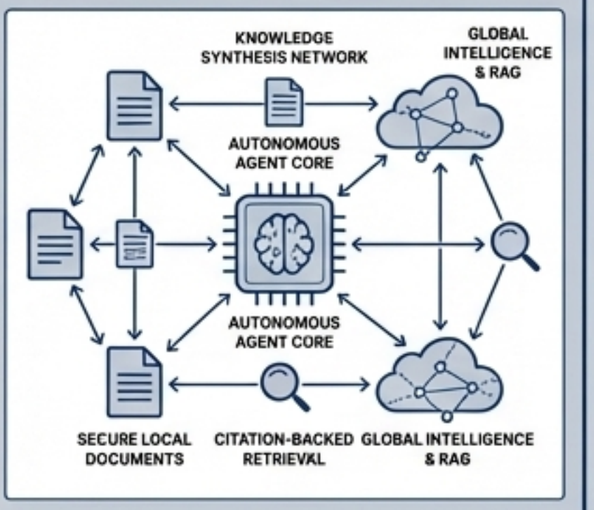

A modular, five-layer architecture guarantees scalability without compromising security.

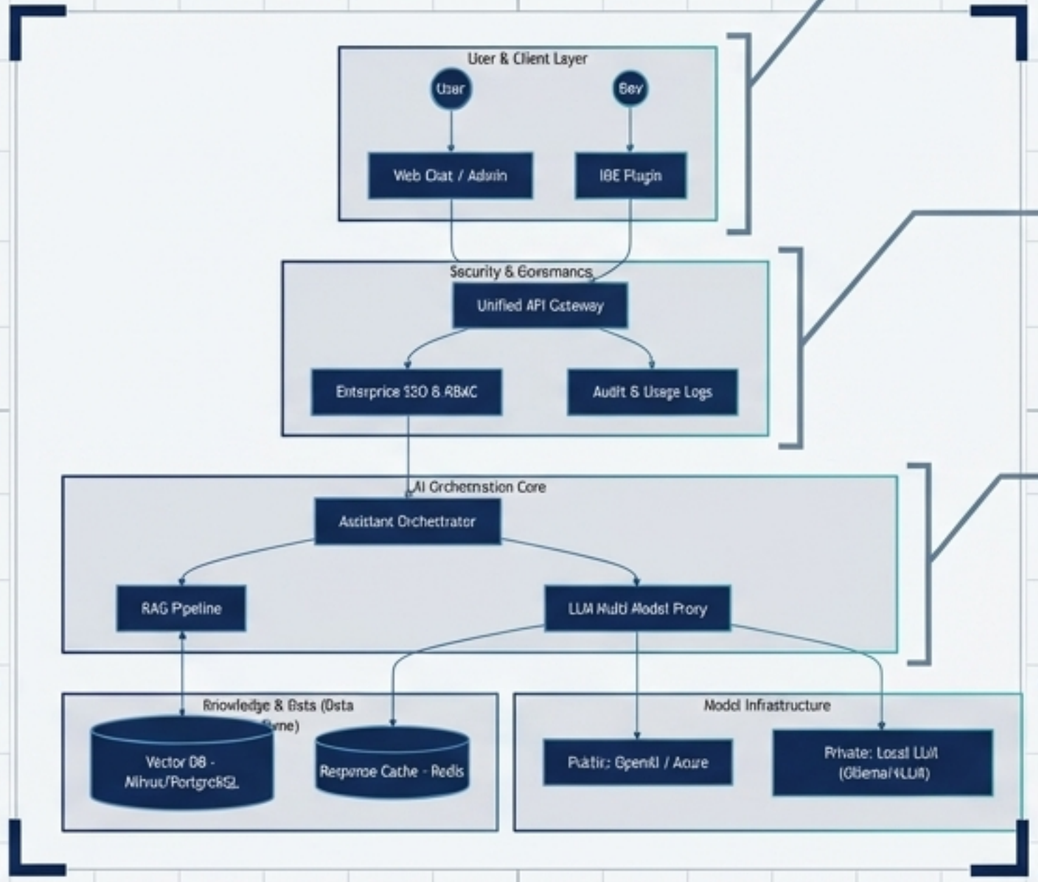

Presentation Layer
(User & Client Interface)

Governance Layer
(Security & Unified API)

Orchestration Layer
(AI Core & Assistant Routing)

Data Layer
(RAG, Vector DB, Cache)

Infrastructure Layer
(Public/Private Model Hosting)

All layers interact safely within a unified boundary, capable of executing complex intelligence routing without ever accessing the public internet.

The human-AI interface is protected by a deterministic security barrier.

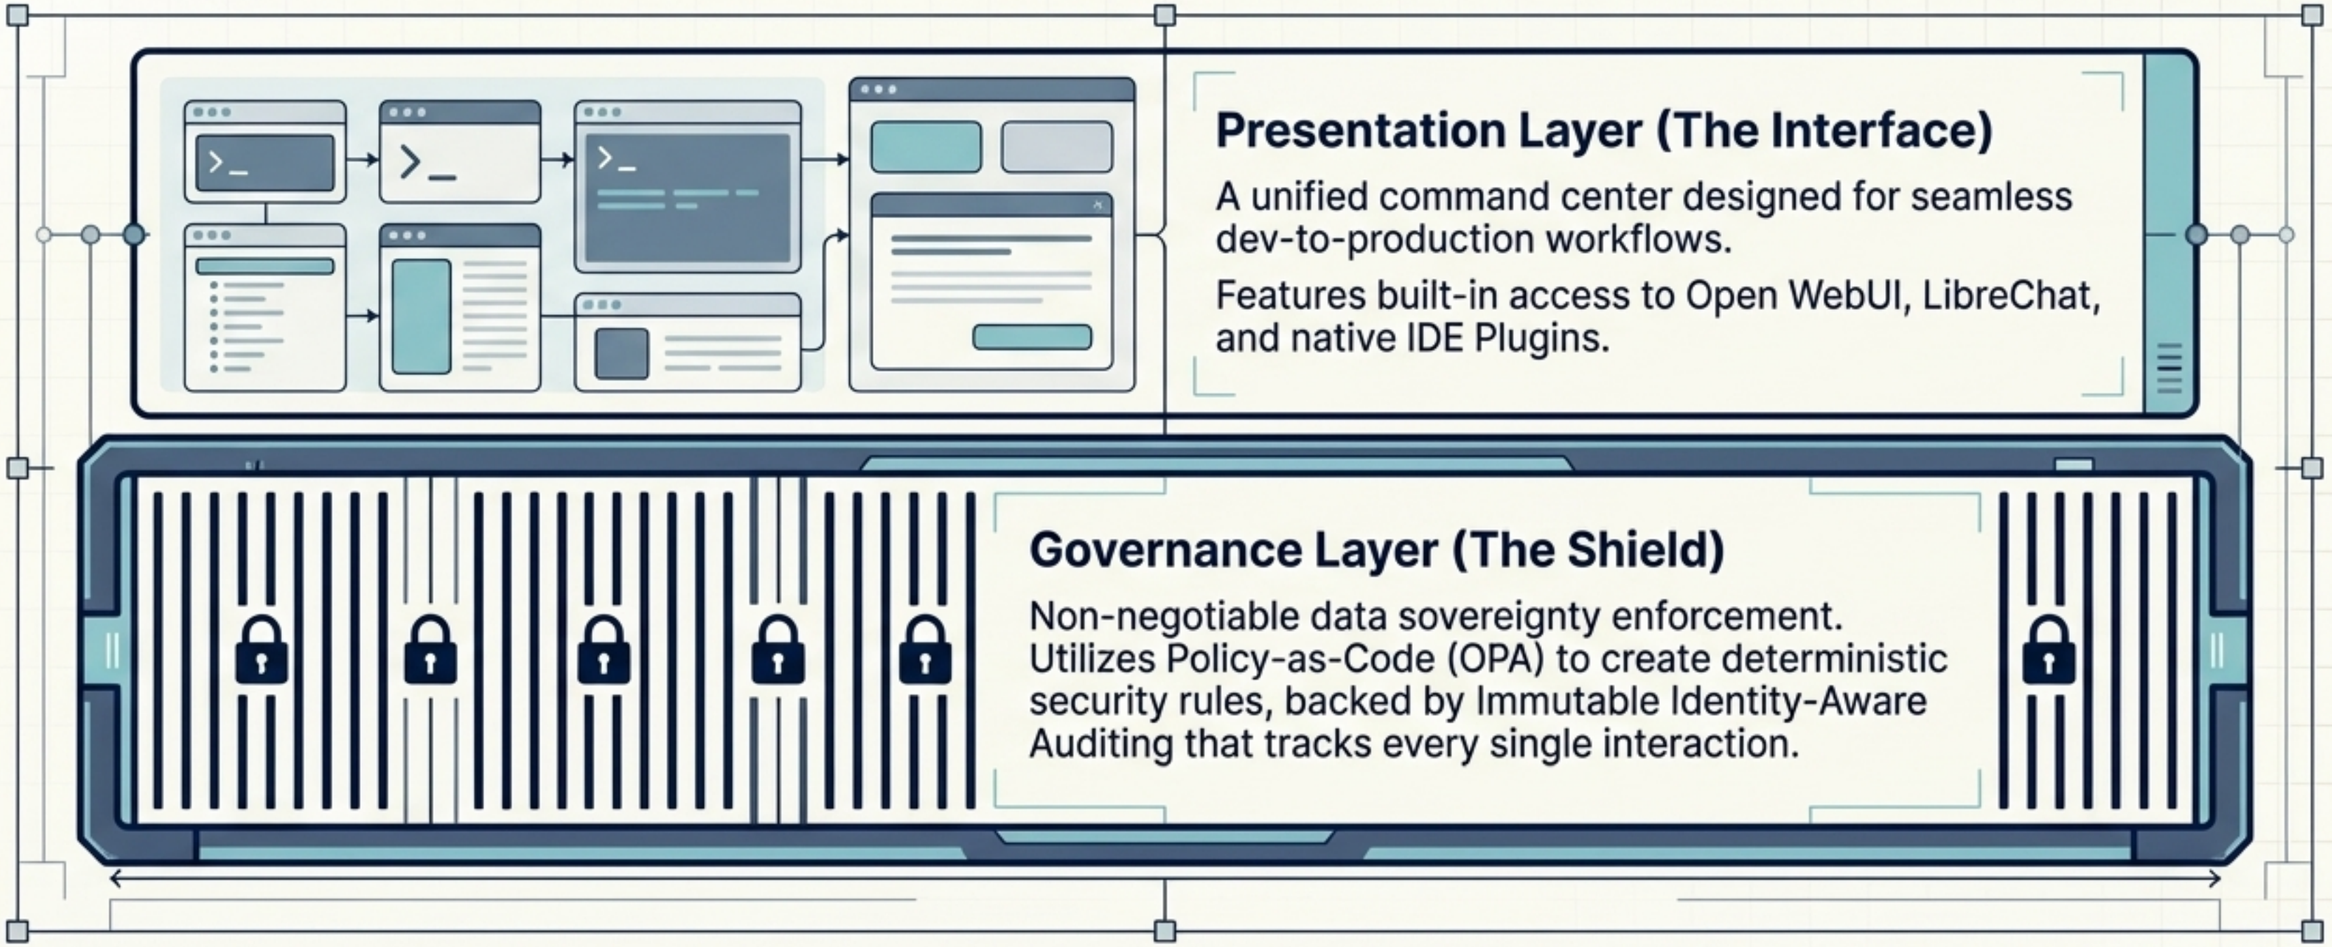

Presentation Layer (The Interface)

A unified command center designed for seamless dev-to-production workflows.

Features built-in access to Open WebUI, LibreChat, and native IDE Plugins.

Governance Layer (The Shield)

Non-negotiable data sovereignty enforcement. Utilizes Policy-as-Code (OPA) to create deterministic security rules, backed by Immutable Identity-Aware Auditing that tracks every single interaction.

The cognitive nervous system manages autonomous reasoning entirely on-premise.

Orchestration Layer (The Brain)

Manages complex Agentic Reasoning pathways.

Uniquely capable of Industrial AIoT Bridging—integrating real-time telemetry from physical assets directly into AI decision loops.

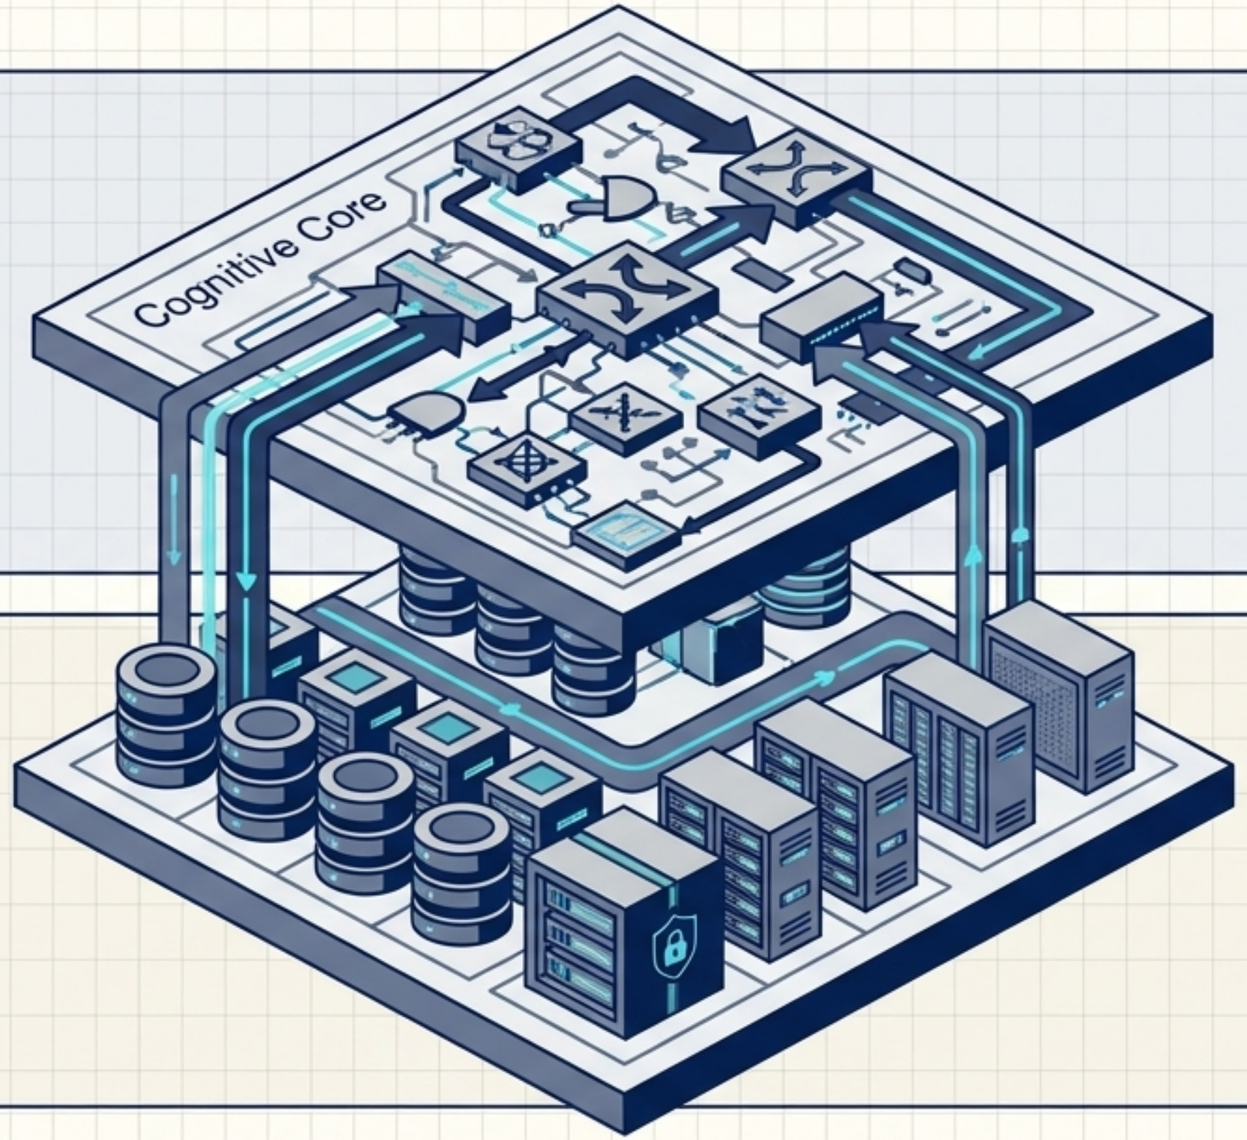

Data Layer (The Memory)

A sovereign knowledge base powered by Air-Gapped RAG.
Utilizes highly secure Vector Stores and a private object repository equipped with continuous semantic consistency checking.

A self-hosted, hardware-agnostic foundation optimized for immediate deployment.

Infrastructure Layer

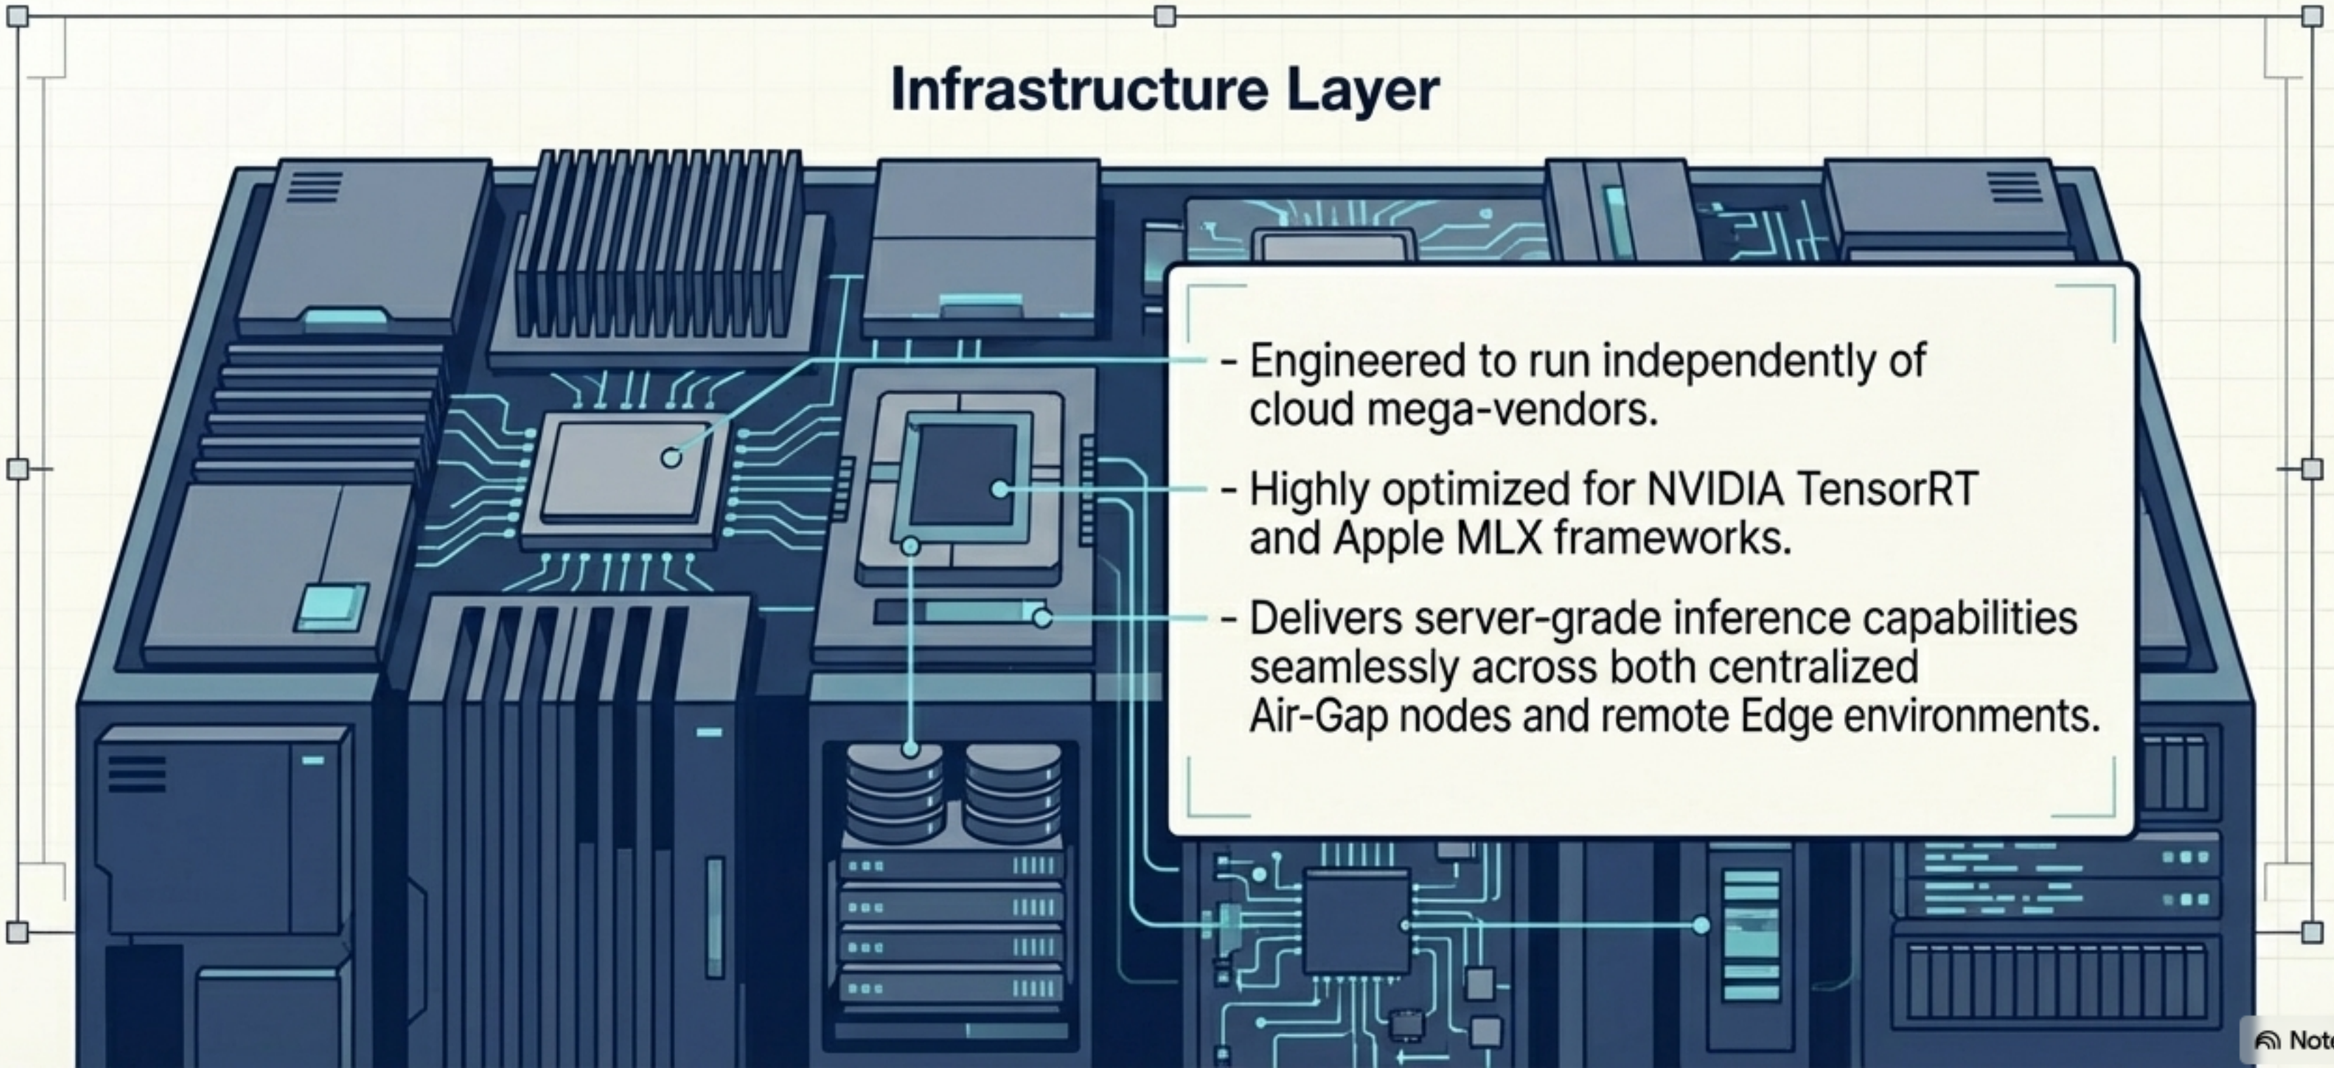

- Engineered to run independently of cloud mega-vendors.
- Highly optimized for NVIDIA TensorRT and Apple MLX frameworks.
- Delivers server-grade inference capabilities seamlessly across both centralized Air-Gap nodes and remote Edge environments.

Engineered specifically for environments where failure is not an option.

Gov & Public Sector

Challenge:
Synthesizing massive, highly sensitive legislative and public records.

OmniNStack Capability:
Automated zoning and legislative analysis delivered with pixel-perfect, secure citations.

Healthcare Systems

Challenge:
Utilizing AI on patient data without violating HIPAA or data residency laws.

OmniNStack Capability:
Complete administrative automation and de-identified clinical inference where PHI never leaves the internal network.

Physical AI & Industrial

Challenge:
Managing real-time analytics for remote resource sectors and edge robotics.

OmniNStack Capability:
Deploying localized cognitive layers for physical assets directly at the edge, untethered from cloud latency.

A leadership team uniting cyber-physical architecture with advanced AI orchestration.

June Hong

CEO

20+ years in embedded systems and cyber-physical platforms.
Architected mission-critical systems for Samsung, NTT, and NEC.
Leads the vision for bridging software orchestration with Edge and Physical AI.

Hyunseung Lee

Principal Software

Cross-disciplinary AI Scientist.

Specializes in building production-grade Sovereign AI stacks using LangGraph, RAG, and LangChain for high-trust enterprise environments.

Joseph S. Yu, Ph.D.

Principal Hardware

Physical AI Architect.

Drives cross-border technology alignment and market expansion the South Korea–North America innovation corridor.

Transition from bare metal to autonomous AI rapidly and transparently.

One-Click Deployment

Radically simplified orchestration allows for rapid deployment across any existing on-premise infrastructure or private cloud environment.

Unified Observability

Total, real-time monitoring of the entire stack. Security teams maintain granular visibility over model performance, operational health, and strict data lineage tracking.

OmniNStack

Own your intelligence. Secure your future.

OmniNStack is the sanctuary where organizational autonomy is guaranteed. Break free from vendor lock-in and embrace a future where AI innovation thrives in a secure, self-governed ecosystem.

Contact our Solutions Architects to schedule a deployment demo or request the complete technical whitepaper.

info@omninstack.com | LinkedIn