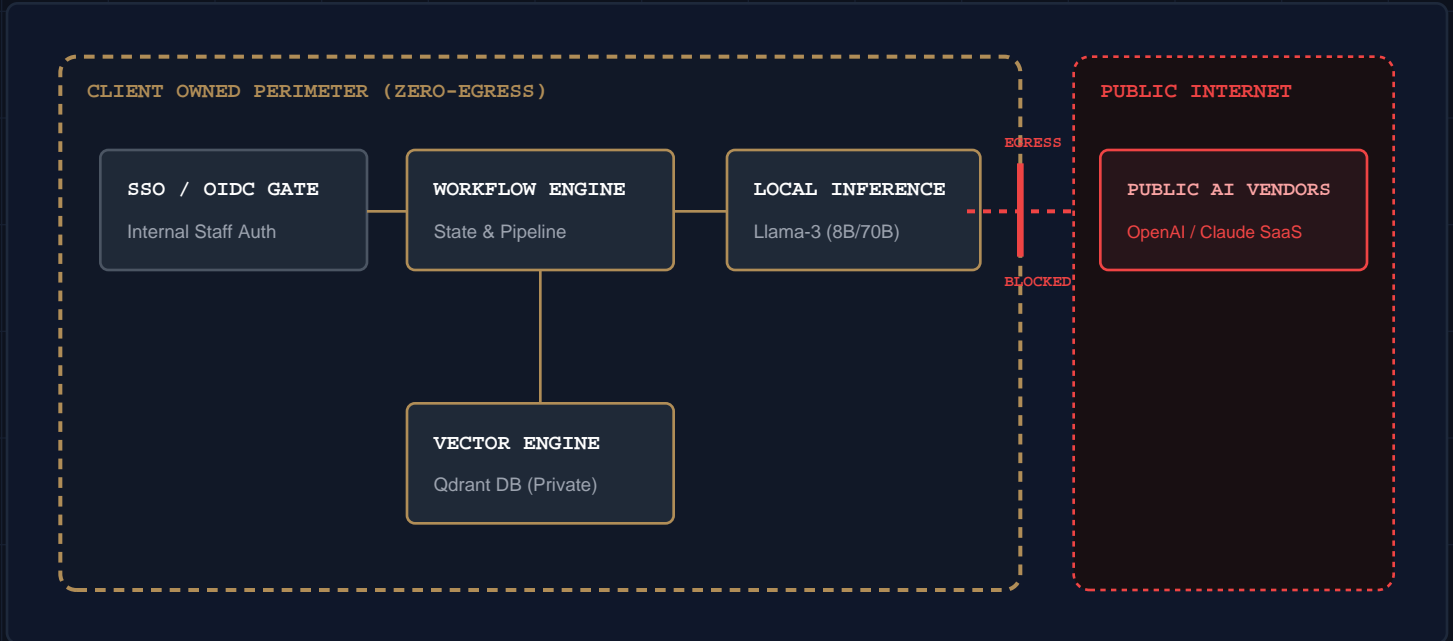


AETHONEX // COGNITIVE INFRASTRUCTURE

TECHNICAL SPECIFICATION: ZERO-EGRESS AI DEPLOYMENT MAP v2.6

SYSTEM TOPOLOGY SCHEMATIC



CORE ARCHITECTURAL PILLARS

01 // LOCAL DATA PARSING & ZERO LEAKS

All input payloads, client briefs, and raw operational metrics remain strictly within isolated NVMe storage. No external API queries are transmitted, and zero network endpoints exit the security boundary to commercial AI APIs.

03 // ENTERPRISE PRIVACY COMPLIANCE

Bypasses standard vendor terms of service that allow model training on user payloads. Enables agencies to guarantee strict client data residency and GDPR compliance to enterprise clients who demand audit-ready infrastructure.

02 // FIXED ECONOMIC INFRASTRUCTURE

Traditional SaaS licensing taxes scaling with seat costs and volume caps. By transitioning to owned, local inference nodes (using quantized open weights), marginal processing costs drop to zero, leaving only fixed hardware maintenance.

04 // PERMANENT OPERATIONAL ASSET

Unlike rented SaaS platforms which can modify pricing or terminate accounts, the Aethonex deployment forms a permanent software asset under your corporate ownership, run on client-hosted virtual private servers.

SECURITY ISOLATION CHECKLIST & STANDARDS

STANDARD	IMPLEMENTATION DETAIL	STATUS
Identity / Access	OpenID Connect (OIDC) via Authentik or Keycloak with strict MFA rules.	VERIFIED
Inference Engine	vLLM / Ollama executing Llama 3.1 8B / 70B in FP16 or GGUF quantization.	VERIFIED
Vector Storage	Qdrant running inside Docker on isolated private bridging networks.	VERIFIED
Egress Firewalls	iptables rules preventing outbound connections on ports 80/443 for inference nodes.	VERIFIED

