

TEKIZ.AI

Private AI routing for regulated businesses

Earl Vanze Co · Founder · hello@tekiz.ai

The problem

- Every AI tool sends prompts, customer data, and credentials to a different provider.
- No shared policy, routing evidence, or fail-over rule across providers.
- Governance, data-residency, and audit pressure outpace adoption.

Why now

- AI adoption accelerating across healthcare, finance, legal, and SMB.
- Enterprise buyers need proof of custody and residency before deploying.
- Model supply is fragmenting across providers, local hosts, and edge devices.

The solution

- One OpenAI-compatible API that routes requests to approved local models, customer provider accounts, and private servers.
- TEKIZ.AI adds policy enforcement, routing evidence, and fail-over rules as a layer above every model.

How it works

- 1. Define policy: which providers, regions, models, and data rules apply.
- 2. Connect sources: local Ollama, customer accounts, private servers.
- 3. Route and record: TEKIZ.AI picks the eligible path and logs evidence.

Physical infrastructure angle

- TEKIZ.AI coordinates a network of customer-owned edge appliances, community-hosted GPU/CPU nodes, and customer-controlled routers.
- The routing and settlement layer is software, but the supply and buyer value depend on physical compute.

Traction

- Hosted Lite, Pro, and Max plans are live.
- Provider integrations: OpenAI, Anthropic, Azure, AWS Bedrock, Gemini, Groq, Mistral, DeepSeek, xAI, OpenRouter, NVIDIA NIM, and Ollama.
- Community node network and self-service enrollment are running today.
- Early usage across healthcare, finance, legal, and small-business workflows.

Market

- Every organization adopting AI while facing data residency, vendor lock-in, and audit pressure.
- Primary beachhead: regulated SMB and mid-market teams that need governed AI routing now.

Business model

- Self-service subscriptions (Lite, Pro, Max) with usage-based routing.
- Marketplace fee on verified inference tokens between buyers and independent hosts.
- Enterprise contracts for custom policy, residency, and audit requirements.

Strategic fit

- Ollama: expand local-model reach into regulated business workflows through governed routing.
- NVIDIA NIM: enterprise customers need credential separation and audit trails before adopting NIM.
- We amplify both partners without reselling their services.

The ask

- Seed funding and strategic partnerships to scale supply, enterprise readiness, and go-to-market.
- Primary targets: AI infrastructure funds, dev-tools angels, sovereign-cloud operators, Ollama / NVIDIA partnerships.

Contact

- Earl Vanze Co · Founder · TEKIZ.AI
- hello@tekiz.ai · <https://tekiz.ai/funding>