


Rethink Energy


Persistent operational intelligence for machines and infrastructure. Built for organizations whose most valuable data can never leave their own walls.



investment brief

Parham Shariat, Co-founder / CEO

Mohammad Salamati, Co-founder / CTO



Rethink Energy

A “Private Language Model” called The Brain. This Artificial Intelligent model is built from scratch by our team and it is designed for companies that **can not** use general purpose AI model due to security and cost.

A model whose cost grows in a straight line rather than a square, with a persistent memory for every tracked entity, running entirely inside the customer’s infrastructure, and small enough to deploy on a device.

The production server is written in Rust programming language, known for performance and memory safety and is already built and tested.


What We Built

Why you can't use ChatGPT or Lama?

Standard transformer models get more expensive, quadratically, as a record gets longer. That's fine for a chat message. It breaks down when the record is five years of sensor readings on one machine.

Most tools solve this by throwing history away: a short window, a rolling summary, a truncated log. That's exactly where the pattern that would have predicted the failure gets lost.

Where AI Breaks



Rethink Energy

The Brain gives every tracked entity, a pump, a vehicle, a server, its own persistent memory that accumulates over its lifetime instead of starting from nothing every session.

Our architecture targets linear cost growth in sequence length, instead of the quadratic cost of standard transformer attention. That is an architectural claim, and we say plainly on the next slide where it stands today.

Not A Private ChatGPT

Regulation is creating a market for private language models like our Brain. Regulated industries generate huge volumes of operational data every day. Most of it can't leave their own infrastructure for legal and security reasons, and conventional AI gets too expensive to run once records get long.

Regulators are now pushing companies toward on-premise AI. Sovereign cloud spending reached \$80 billion globally in 2026, up 35.6% in a single year, and the EU AI Act's obligations are already being phased in.

Why Now

Where We Really Are

BUILT AND TESTED

- Production server, written in Rust
- Industry-standard interface, existing tools work unmodified
- Deployment and monitoring tooling

IN PROGRESS, NEXT 3 MONTHS


- Confirm the speed fix on production hardware
- Train the foundation model
- Prove the persistent memory system

We're being direct: the serving infrastructure is done. The core model claims, linear scaling and persistent memory, are the next three months of work, not finished results.

How We Compare

	Public LLM	Local LLM	RAG	The Brain
On-premise	No	Yes	Yes	Yes
Persistent entity memory	No	Limited	Limited	Yes
Long operational histories	Expensive	Expensive	Partial	Targeted
Air-gapped capable	No	Yes	Yes	Yes
Edge deployment	No	Limited	Limited	Yes
Linear cost scaling	No	No	Unclear	In validation*

*Linear scaling is our core architectural claim. It is being validated on production hardware this quarter; see Where We Really Are.



Rethink Energy

Picture an operator with 50,000 pieces of equipment. Compressor C-74 gets its own persistent memory: installation history, operating parameters, vibration trends, maintenance events, technician notes.

Instead of asking what an alarm means, the operator can ask why C-74's vibration profile has changed over six months, and whether that pattern preceded a failure elsewhere in the fleet. The data never leaves their infrastructure.

Our Beachhead: Oil & Gas

Expansion Path



Phase 1

- Oil and gas, industrial infrastructure. One beachhead, one paying design partner.

Phase 2

- Aviation, transportation, manufacturing. Same entity-memory model, adjacent operational data.

Phase 3

- Government and regulated enterprise, once the architecture is proven at scale.

Who We Don't Serve

Deliberately narrow. If a workload is mostly free text rather than sequential, entity-tracked data, a general-purpose model is the better tool, and we'd rather say so.

- News and general journalism
- Social and platform content
- E-commerce catalogues
- Scientific literature search
- General business correspondence
- Medical imaging

Compute vs. Intelligence

The customer buys and owns the hardware, roughly \$70,000 to \$100,000, giving them full control of their infrastructure and their data.

We charge for the intelligence layer: a one-time deployment fee of \$25,000 to \$50,000, then a recurring platform license starting at \$1,500 a month per site, scaling with entities and data volume.

Customer owns the compute. Rethink owns the intelligence. That's a recurring, high-margin software relationship, not a one-time integration job.

Team

PARHAM SHARIAT, CO-FOUNDER / CEO

- Senior systems engineer at a major hospital network, overseeing mission-critical infrastructure in a tightly regulated environment
- Later led U.S. operations and international trade for an industrial company

MOHAMMAD SALAMATI, CO-FOUNDER / CTO

- Software engineer and architect with 10 years building and scaling production systems
- Led cloud security, DevOps and infrastructure at prior startups

Part of this raise funds our first applied-ML hire to lead foundation model training and validation.

Rethink Energy

We're raising \$750,000 for 25% of a new company built to hold The Brain: its technology, IP, customers, and revenue going forward. IP transfer terms and cap table detail are in the data room.

This raise funds production-hardware validation of the core technical claims and our first paying design partner in oil and gas. That validated customer is the milestone for our next round.

The Ask

Roadmap & Milestones

Next 3 months

- Confirm the speed fix on production hardware, train the foundation model, and prove the memory system.

3 to 6 months

- Sign our first paying customer in oil and gas, the milestone that unlocks the next round.

6 to 18 months

- Expand to a second and third vertical once the first customer renews.

Rethink Energy

We'd love to walk you through the architecture, the validation plan, and the numbers in more detail.



Thank You

parham@rethinkenergy.com

mohammad@rethinkenergy.com

+1 619 626 7545

22112 Summit Hill Drive, Lake Forest, California 92630