

QUANTAREON

Investor Memo — Seed Round, \$500K

Founder	Vlad M., sole founder, QUANTAREON Labs (Tbilisi, Georgia)
Round	\$500K seed, 18-month runway
Stage	Methodology published (DOI), one prototype case on GitHub, product in specification
Date	May 2026

One-liner

QUANTAREON is building the **ConsciousAI engine** — a deliberation layer above any LLM that reduces reactivity at the architectural level. One engine, three market outlets: certification, coding, methodology licensing.

The Problem

Modern LLMs share one structural flaw: they are reactive. They emit the statistically most probable next token, not the most correct one. This produces:

- **Hallucinations** — under output pressure, the model invents instead of pausing.
- **Scope drift** — tasks expand beyond what was actually requested.
- **Sycophancy** — models agree under social pressure rather than verifying.
- **Jailbreak surface** — the first impulse to an adversarial prompt is the wrong one.

Every existing alignment technique — RLHF, Constitutional AI, chain-of-thought, ReAct, Reflexion — operates **after** the model has committed to its first impulse. The trajectory is already locked.

The Solution

The ConsciousAI Protocol is a system-prompt layer that bends the generation trajectory **before** the first token. Two parallel layers:

Static Identity (WHO)	Dynamic Action (HOW)
VL-Framework: identity seal, values, depth rules, memory cycle	Control Cycle: per-request gating — Initiation, Reflection, Adaptation
R-Cycle: 5 deliberation states, R0 as default background process	Workflow: Proposal → Risk Check → Scope Guard → Minimal

Cross-model. Same behavioral signature on closed-weight (Western) and open-weight (Eastern) frontier models — the work happens at the system-prompt layer, not inside any one model's internals.

The Proof Point

quantarion-router — multi-provider LLM router with fallback and circuit breaker, written solo end-to-end under this protocol in a **single session**: specification, implementation, and the full test suite — **124 passing tests, mypy --strict**, zero required runtime dependencies. The repo is public;

commit history shows the trajectory.

The Market

One engine, three market outlets. Each outlet stands on its own; together they compound.

Outlet	Buyer	What we sell
QUANTARION Audit	Compliance officers, regulators, independent third-party evaluator for AI agents. Built around EU AI Act	Methodology licensing: protocol-as-a-layer applied to client's existing L
QUANTARION Coding	Developers shipping production code with coding PLMs	Methodology licensing: protocol-as-a-layer applied to client's existing L
ConsciousAI API	Platform integrators, B2B SaaS	Methodology licensing: protocol-as-a-layer applied to client's existing L

Timing. EU AI Act enforcement begins Q3 2026. Articles 14-15 create regulatory demand for independent agent evaluators. QUANTARION Audit lands into that demand at audit.quantareon.com.

Why Us

- **Methodology already public.** CC BY-NC-ND 4.0, DOI-anchored. The protocol exists, runs in production, and is replicable.
- **Proof point already shipped.** `quantarion-router` (124 tests, single session) shows the productivity profile of the protocol.
- **Cross-model from day one.** No vendor lock-in. Sufficient instruction-following is the only constraint.
- **Engineering-first.** The founder is the conductor, not just a researcher. Specs, tests, types, ship.

Use of Funds — \$500K, 18-month runway

Category	% of round	What it buys
Founder + 1 senior engineer salary, 6 months	30%	Two-person team to close the audit MVP
Compute + cross-model benchmarking (HarmBench, AdvBench, TruthfulQA)	20%	Formalize the empirical claim
Audit platform engineering (audit.quantareon.com)	20%	Test category spec to live evaluator
Coding platform core engine	15%	Productize the router pattern at scale
Legal (Delaware C-Corp flip, IP)	5%	Capped at \$25K, deferred until closing
Buffer / runway extension	10%	12-month reserve to extend to 18 months

Traction

- **Methodology published** — CC BY-NC-ND 4.0, DOI 10.5281/zenodo.19858814.
- **quantarion-router** — live on GitHub, 124 passing tests, mypy --strict.
- **Manifund campaign** — live at manifund.org, goal \$150,700, closes June 12, 2026.
- **LTFF application** — submitted May 2, 2026, under evaluation.
- **Public engineering report** — published on EA Forum, May 2026.

Roadmap (18 months)

Quarter	Milestone
Q3 2026	Audit MVP live at audit.quantareon.com. First 3 pilot clients.
Q4 2026	Cross-model benchmarks published. Coding platform alpha.
Q1 2027	Audit revenue traction. Coding beta, first 10 paying teams.
Q2 2027	Coding platform GA. First Conscious AI API enterprise contract. Series A readiness.

Structure

- **QUANTAREON Labs** (Tbilisi, Georgia) — operations, IP, day-to-day work.
- **QUANTAREON Inc.** (Delaware C-Corp) — holding company, recipient of investment, equity issuer.
- Standard flip on closing. Legal fees deferred until closing, reimbursed from investment proceeds (capped at \$25K).
- Founder remains based in Tbilisi.

Ask

\$500,000 seed. SAFE or priced round, terms open to discussion. Single lead preferred. **Macroscopic Ventures** is our first-choice partner — your existing portfolio (Gray Swan, Halcyon, Mythos, RunRL) is the exact ecosystem this work belongs in.

Contact

Email	vlad@quantareon.com
Telegram	@quantareon
Platform	quantareon.com
Methodology	github.com/makx518-ui/consciousness-protocol (DOI 10.5281/zenodo.19858814)

“The first impulse is noise. The second is signal. Silence → Clarity → Action.”

— Vlad M., founder, QUANTAREON Labs