MEMBRIA.AI · INVESTOR DECK

AI agents that *learn*.

Membria is the memory & experience layer for AI agents — one causal-memory core powering products for software teams, industrial engineering and audit.

Michael Aprossine

CEO, Founder — Membria.AI by Actiquest Labs

CONFIDENTIAL

July 2026 · v1.0

COMPANY

Who We Are

Actiquest Labs

Applied AI Systems Company

Practical automation, intelligent workflows and enterprise-grade AI — built for real results.

Membria.AI

AI Operating Environment

Our core product — agents, memory, knowledge and verification combined into one enterprise-ready AI platform.

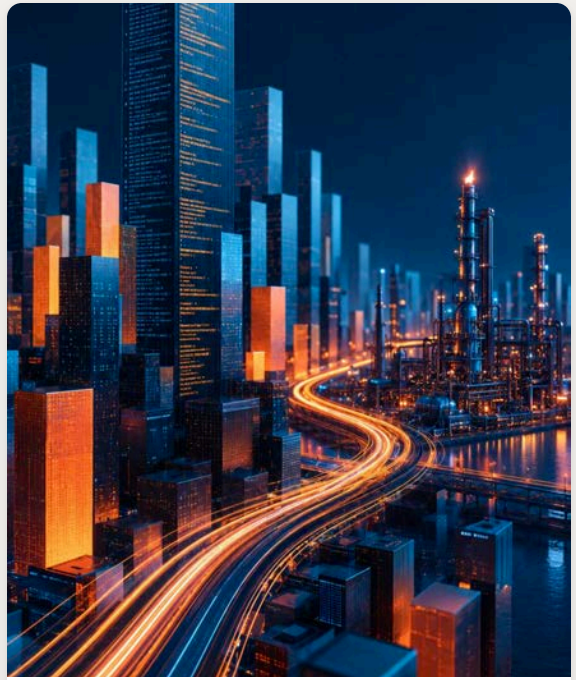

Vertical Expertise

Code · Fintech · Government · Audit · EPC

We specialize in complex domains where AI needs precision, context and process understanding.

How We Work

Applied AI Engineering

We ship enterprise AI projects that stress-test and fund Membria's core — concept to production.

COMPANY PROFILE

Actiquest Labs at a Glance

Key facts, production footprint and enterprise AI capabilities.

10+ years enterprise product, AI engineering and venture-building

4 verticals — Fintech · Government · Audit · Engineering/EPC

100+ seats deployed across production AI use cases

METRIC	VALUE
Founded	February 2023 · California, USA
Core Product	Membria.AI — memory & experience layer for AI agents
Production Verticals	SDLC · Fintech/Banking · AI Accounting/Audit · Engineering/EPC
Deployment	Cloud SaaS · self-hosted · on-premise enterprise
Clients & Partners	OES · Wowcube · Hiscox · K-Telco · Kazakhstan Government · industrial enterprises across EMEA/CIS
Technology Partners	NVIDIA · Amazon Web Services
Team Background	20+ years enterprise software & audit · 9 years applied AI — finance, government, industrial, deep-tech

PROBLEM

Intelligence that *resets* every morning

1. **AI agents are stateless**
   Every session starts from zero. Yesterday's reasoning, decisions and failures are discarded the moment the window closes.

2. **Mistakes repeat — and they're expensive**
   Coding teams re-solve the same problems; in engineering, **28% of EPC budgets** are lost to rework and 9 of 10 large projects overrun by +28% on average.

3. **Handovers bleed context**
   30–50% of assumptions vanish at every phase handover; 18% of working time is spent searching for data that already exists.

4. **Judgment walks out the door**
   A senior engineer leaves — and fifteen years of decisions, lessons and "don't do this" leave with them.

Every AI tool today can answer questions. **None of them remembers why decisions were made — or what happened next.**

WHY NOW

The experience graph is *unclaimed*

1. **Agents became the workforce**
   Claude Code, Codex and Cursor now do real daily work in code, engineering and finance — but they shipped without memory.

2. **MCP became the standard port**
   One protocol plugs a memory plane into every agent on the market — distribution without integration projects.

3. **Regulation demands traceability**
   EU AI Act and ISO 19650 require defensible decisions and immutable audit trails in safety-critical work — exactly what a causal memory produces by default.

4. **Incumbents own data, not experience**
   Nobody owns the decision → outcome graph. Whoever captures it first gets a moat that compounds with every user, every day.

PRODUCT

Membria *Reasoning Graph* Workforce

INSTITUTIONAL KNOWLEDGE

Memory

- ✓ **Decisions** — every choice captured with its reasoning
- ✓ **Consistency** — behavior chains keep teams aligned
- ✓ **Skills** — wins distilled into reusable know-how
- ✓ **NegativeKnowledge** — failures become rules

MULTI-AGENT HARNESS

Reasoning

- ✓ **Mesh** — many agents work in parallel
- ✓ **Coordination** — agents hand off via A2A
- ✓ **Routing** — each task to the right model
- ✓ **Quality gates** — verification blocks bad output

DOMAIN-AGNOSTIC FOUNDATION

Graph

- ✓ **GraphRAG** — semantic search over a property graph
- ✓ **Multi-domain** — one schema across every domain
- ✓ **Orchestration** — work runs as a dependency graph
- ✓ **Adaptive** — memory that strengthens with use

HOW IT WORKS

One loop, *every stage*

Not a memory bolted onto one step — every shipped outcome feeds the next decision.

01

Plan

Recalls decisions and outcomes from similar work — the agent starts from what already worked.

02

Execute

Injects validated skills and blocks known anti-patterns as the work is being done.

03

Review

Checks the result against past failures and standards — deterministic critics, auditable output.

04

Ship

Records the real outcome and links it back to the decision that caused it.

↻ **Every outcome feeds the next decision.** One MCP connection: `npx @membria/cli init`

FOCUS

One core, *many domains*

Verticals are distribution of one product — not divergence. The same graph schema ships verbatim in every codebase.

PRODUCT	WHO BUYS	DELIVERY	STATUS
Membria CE	solo developers on Claude Code / Codex / Cursor	cloud SaaS · zero-install MCP	SAAS · CODE.MEMBRIA.AI
Membria Enterprise	engineering & product teams	self-hosted platform · own UI	DEMO · DESIGN PARTNERS
AI-EPC	EPC contractors & design institutes	vertical app on the same core	LIVE DEMO · EPC.MEMBRIA.AI
Glass	accounting & audit teams	vertical app on the same core	DESIGN-PARTNER STAGE

Proof of substrate: Decision · Outcome · NegativeKnowledge · Skill — one schema, one causal experience loop, shared verbatim across every codebase. A vertical competitor can't follow us across domains; we can.

CASE STUDIES

Proven Across *4 Domains*

One AI substrate · any vertical domain

Software Development / AI-PDLC

Senior engineers leave and take their judgment. Membria keeps every decision — teams stop re-solving the same problems.

Banking / Finance

A new regulation used to mean a six-week release. Now it's live on the next transaction — traceable to the clause that authorized it.

Glass — AI Accounting / Audit

Audit teams lose weeks tracing entries to source documents. Membria marks every posting with origin and trust score.

AI-EPC / Engineering

A missed P&ID clause becomes a safety incident. Five AI critics check every diagram against standards before sign-off.

Common pattern: domain knowledge → graph memory → AI workflows → human verification → production. 100+ seats deployed across these use cases.

WHY MEMBRIA?

Core Capabilities

The full stack — from memory and reasoning to deployment and governance

AI Agent Orchestration

Multi-agent BBS mesh, WorkGraph DAG, A2A protocol coordination

Causal Memory

Decision / Outcome / NegativeKnowledge graph with time-decay and compounding

Verification Engine

Cross-document consistency checks, rule-based normative validation, quality gates

Graph Intelligence

Memgraph property graph, symbol-level CodeGraph, policy graph with Hebbian co-activation

LLM Integration

Claude / GPT factory, structured output, behavior chains, context injection

Product Engineering

Full-stack: FastAPI + React + Node.js, MCP protocol, Telegram, Web Dashboard

Domain AI Systems

Vertical-specific schemas, critics, standards compliance (ISA / API / ASME / IEC)

Security & Governance

PolicyEngine, RBAC, OIDC/SSO, prompt injection defense, secrets management

Bottom line: Generic AI forgets. Membria remembers, reasons, and governs.

THE WEDGE

EPC first: *nobody remembers*

Description → P&ID in minutes
Plain words or an equipment CSV in; a laid-out diagram out.

Real CAD files, not pictures
DXF for AutoCAD, DEXPI for AVEVA/Hexagon, IFC4 for BIM.

Five-discipline critics
Process, Mechanical, HSE, I&C, Piping — deterministic, auditable, repeatable.

Every defect becomes memory
Finding → change request → negative knowledge, automatically.

Zero competitors with persistent memory
Autodesk, Bentley, AVEVA, Trimble, Procore — all stateless.

Regulation is the sales agent
EU AI Act + ISO 19650 make traceable decisions mandatory in safety-critical infrastructure.

MOAT

Three EMPIRIA graphs, *day one*

Cold start is the hardest problem in agent memory. We don't have one — every vertical ships with a pre-loaded PK/NK library before a single client session runs.

VERTICAL	PK — BEST PRACTICES	NK — ANTI-PATTERNS / PROHIBITIONS	SOURCE	STATUS
Engineering (EPC)	**~1.3M**	**~1.2M**	20,000+ standards — ISO · IEC · ASTM · ГОСТ · API · ASME · NFPA · DNV	LIVE · SCALING · EPC.MEMBRIA.AI
Code	**~250K Skills**	**~320K AntiPatterns**	96.5K mined repos + curated sources, CWE-mapped	LIVE · SCALING · CODE.MEMBRIA.AI
Audit / Glass	**schema live**	**schema live**	materiality & provenance graph — seeding from client engagements	DESIGN-PARTNER STAGE

Same schema, same loop, three domains. A generic memory startup starts every client at zero. We start every client already knowing what fails.

DATA MOAT

The engineering knowledge graph *no one else has*

20,000+
engineering standards processed — API · ASME · IEC · ISO · ГОСТ · DNV · NFPA · AWS · ASTM

2.5M+
machine-extracted rules: ~1.2M prohibitions + ~1.3M best practices

10,000+
equipment types mapped to rules across 20+ disciplines

6
jurisdictions in one graph: US · EU · UK · Norway · CIS · International

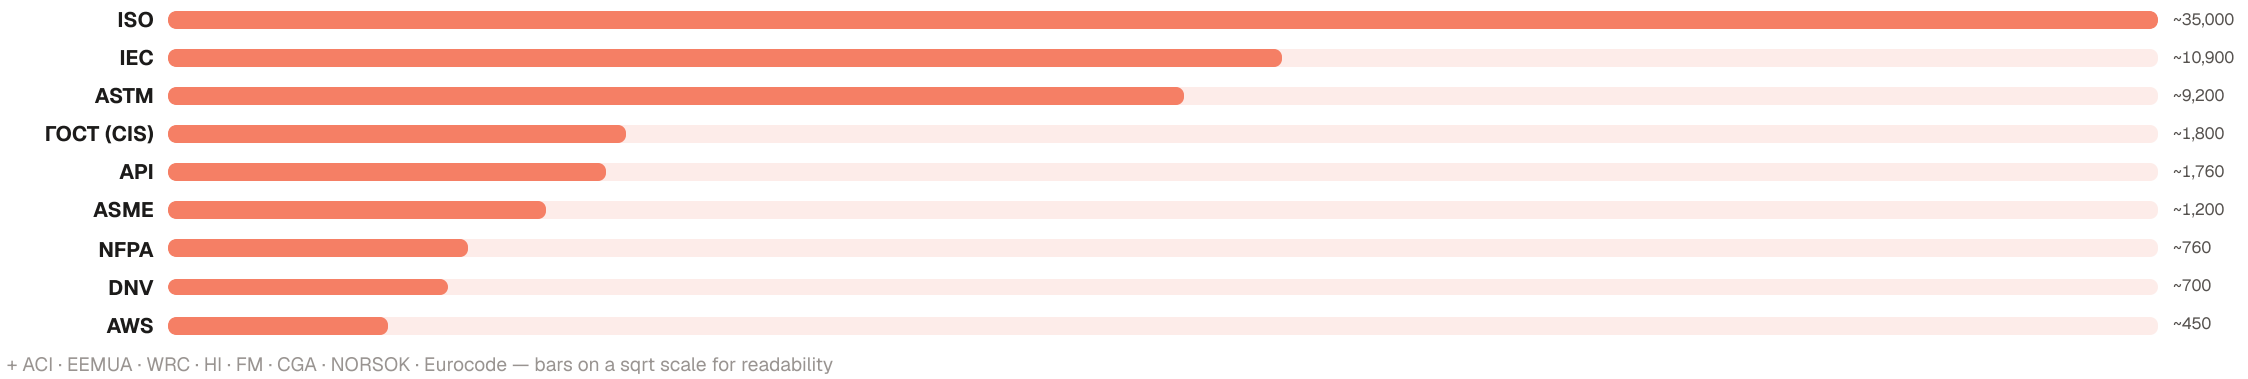

+ ACI · EEMUA · WRC · HI · FM · CGA · NORSOK · Eurocode — bars on a sqrt scale for readability

Prohibitions + best practices, one graph. No incumbent has anything like it — in any jurisdiction.

PROOF ENGINE

Code: evidence, *not opinions*

Membria CE is live for solo developers — and every session strengthens the shared graph the enterprise products run on.

96.5K
production repos mined by the CodeDigger crawler

7,448
anti-patterns cross-validated with an objective removal-rate signal

~570K
Skills + AntiPatterns live today — scaling toward 1–2M

1
command to connect: npx @membria/cli init

→ **Removal-rate: the signal nobody else mines**
The share of pull requests that delete a pattern, computed from commit archaeology — historical evidence from real codebases, not scraped opinions.

→ **PLG funnel feeds the moat**
Solo developers adopt in one command; their validated experience compounds the commons graph that enterprise and vertical products query.

COMPETITION

Nobody *remembers*

PLAYER	WHAT THEY DO	PERSISTENT MEMORY?	GAP
Autodesk	design + construction cloud	NO	AI = copilot for CAD, not knowledge management
Bentley Systems	infrastructure digital twins	NO	twin ≠ memory — no decision provenance
AVEVA / Schneider	process industry digital twins	NO	operational data, not engineering decisions
Procore	Helix AI automation	NO	stateless — every request starts from zero
SymphonyAI	P&ID drawing recognition	NO	recognition, not a memory plane
Memory startups mem0 · Zep · Letta	generic agent memory	YES	no vertical depth, no verification, no standards graph

The real competition is inertia: Excel + "it works fine as is." Every project starts from a blank page — that is the memory problem, and it's solvable.

MODEL & ROADMAP

Sequencing, not *spreading*

CE subscriptions for solo developers · per-seat team plans · self-hosted enterprise license · EPC per-project pricing.

NOW

Foundation

1 bank + 1 government, live

EPC pilots in engineering

CE & EPC SaaS launching in parallel

100+ seats in production

NEXT 12 MONTHS

EPC revenue

first EPC contracts

Glass audit beta

commons → 1–2M nodes

ISO 27001 / SOC 2 path

24 MONTHS

Memory-plane API

open experience layer

any agent, any vendor

cross-industry commons

platform economics

Let's build the layer
agents can't work without.

TEAM & ASK

Let's Work Together

Michael Aprossine
CEO · Co-Founder
Scaling operations, customer-focused growth

Phillipe Khomenok
COO · Co-Founder
Strategic vision, automation-driven growth

Mike Keer
CTO · Co-Founder
AI architecture — fast, reliable, built to scale

ASK	DETAIL
Traction	1 bank + 1 government live enterprise · multiple EPC engineering pilots · 100+ seats deployed · CE & EPC SaaS launching in parallel
Raising	$2M seed on a $15M pre-money valuation
Use of funds	EPC go-to-market & certifications · commons scale-out to 1–2M nodes · enterprise connectors · core team
Contact	hi@actiq.ai · membria.ai · code.membria.ai · epc.membria.ai · glass.membria.ai (soon)