Lumovra AI

INDEPENDENT EVIDENCE INFRASTRUCTURE FOR ENTERPRISE AI

The proof your AI can't fake.



Lumovra AI is building the independent verification layer for enterprise AI.

Starting with evidence captured directly from engineering workflows, Lumovra turns changes, tests and approvals into buyer-verifiable Proof Packs. Our Triple Lock™ architecture extends that evidence chain across model provenance and production reality.

CONNECT DECLARE CAPTURE PROVE PROOF PACK

Watch the launch video

Confidential

THE PROBLEM · WHY NOW

AI deals stall when evidence arrives after the questions.

Enterprise buyers ask how an AI system was built, changed, tested and approved. Teams have the raw evidence, but it is scattered across repos, tickets, logs and documents. Reconstructing it mid-deal costs time, engineering focus and trust.

01 · DEAL FRICTION

Questionnaire arrives mid-deal.

Governance and security reviews pull senior engineers into evidence collection instead of roadmap work.

Late-stage

Diligence lands after the deal is in motion

02 · ENGINEERING TAX

Evidence gets rebuilt under deadline.

Discovery conversations repeatedly surfaced 2–4 weeks of senior engineering effort to assemble diligence evidence.

2–4 wks

Reported per diligence cycle

03 · TRUST GAP

Self-assembled evidence has a credibility ceiling.

Policies and screenshots document intent, but buyers still need confidence that they reflect what actually happened.

Continuous

Provenance beats a reconstructed dossier

WHY NOW · AUTONOMY + PROCUREMENT + REGULATION

AI is moving from software that answers to software that acts.

As autonomy increases, enterprises need stronger evidence of what changed, what was approved, what controls existed and what happened in production. Regulation accelerates demand. Enterprise procurement is the durable driver.

Agentic systems raise the stakes: every action, tool call, override and approval creates a new evidence requirement.

THE SOLUTION

An independent, buyer-verifiable evidence trail for every AI system you ship.

Every AI system needs a record that buyers can examine.

Lumovra connects AI systems to their engineering evidence, linking changes, tests, approvals and findings into a traceable history. Teams generate Proof Packs for procurement, audits and governance reviews, without rebuilding evidence from scratch.

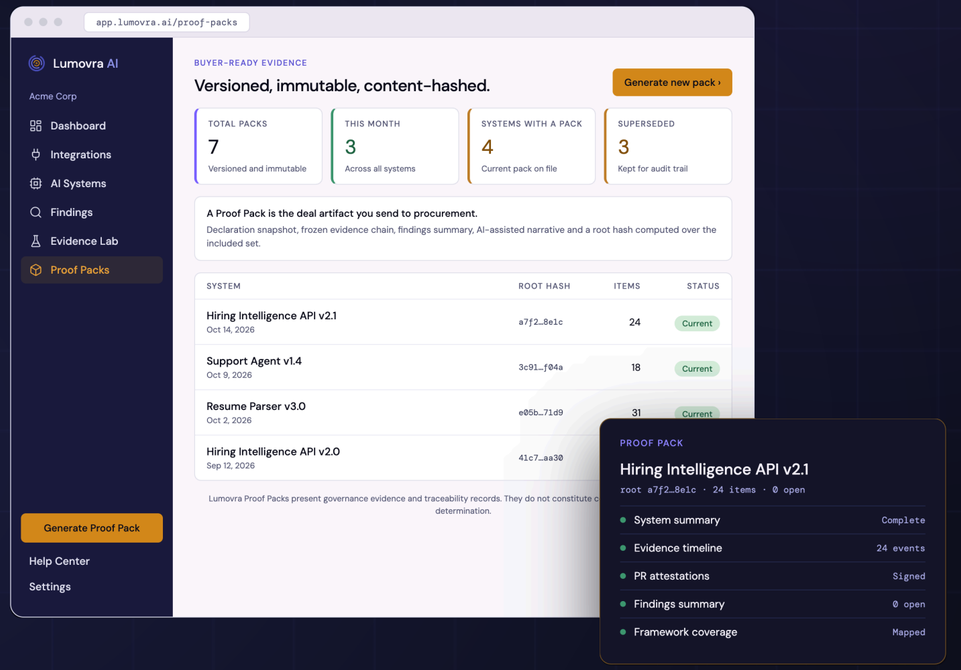
CODE · DATA · PRODUCTION

CAPTURED CONTINUOUSLY

INDEPENDENTLY VERIFIABLE

KEY DISTINCTION

Lumovra AI issues no compliance verdicts. It creates independent, attributable evidence that supports procurement, audits, governance reviews and incident reconstruction.



Minutes to inspect

Multi-framework mapping

Live GitHub evidence

HOW LUMOVRA AI WORKS

Evidence built into how you ship.
Proof ready before the buyer asks.

Four steps. Lives inside your CI. The evidence trail starts before procurement creates urgency.

01 · CONNECT

GitHub or GitLab in minutes

OAuth · minimum permissions · no installed agents · no infra

02 · DECLARE

Register your AI system or agent

What it does · who it affects · where the risk lives

03 · CAPTURE

Every AI-touching PR

Changes, approvals and findings enter the evidence graph automatically

04 · PROVE

Generate a Proof Pack

Buyer-ready evidence · framework mapped · inspectable in minutes

REGULATORY + FRAMEWORK MAPPING

Mapped in-product to the frameworks enterprise buyers already ask about.

EU AI Act
Risk + evidence mapping

NIST

AI RMF
Voluntary risk framework

ISO

ISO/IEC 42001
AI management system

US state frameworks

CO · CT · IL · NYC · CA · TX

US coverage includes Colorado · Connecticut · Illinois · NYC · California (FEHA / ADMT) · Texas TRAIGA

● LIVE Free tier

HISTORICAL BACKFILL

Lumovra AI can backfill Git and PR history to build the initial evidence record, then capture forward continuously.

THE MECHANISM · TRIPLE LOCK™

What was approved. What it was built on. What actually ran.

Three locks follow an AI system through its lifecycle. Each item carries its source, strength and capture time. Shows evidence coverage, not a compliance score.

LIVE

LOCK 1 · INTENT · DECISION + APPROVAL

Records who approved every change, under which policy.

Reviews by named people, the policy version in force at merge, and the CI evals that ran before release.

CONNECTORS

GitHub · GitLab · CI evals

TARGET LIVE · NOV 2026

LOCK 2 · PROVENANCE · DATA + MODEL

Records which data and model it was built on, and every change.

Training runs, dataset hashes and the model version the system actually calls, checked over time. Metadata only.

CONNECTORS

MLflow · OpenAI · Anthropic · Bedrock

TARGET LIVE · NOV 2026

LOCK 3 · REALITY · SERVING PATH

Records which controls were configured in production.

Guardrails, content filters and pinned versions, plus overrides, incidents and logged decisions.

CONNECTORS

Bedrock · Azure OpenAI · LiteLLM · Runtime events

CROSS-LOCK VERIFICATION · THE LUMOVRA DIFFERENCE

Triple Lock™ is designed to cross-check company claims against evidence from engineering, model and production systems. Discrepancies become visible in the Proof Pack.

WHAT CROSS-LOCK VERIFICATION WILL CATCH

Declared ≠ actual model
Code says Model A. Production calls Model B.

Data claim ≠ data evidence
“Trained on approved dataset.” Training records reference a different dataset version.

Control claim ≠ reality
Guardrail promised. Not configured in production.

Version drift
A provider changes the model version without a matching internal record.

EVERY SCAN STRENGTHENS THE EVIDENCE CHAIN

Each scan adds a timestamped evidence record. As Triple Lock™ expands, these records build a cross-provider history of what was approved, what changed and what actually ran.

THE ADVANTAGE

The independent evidence layer across the AI stack.

Triple Lock™ is designed to cross-check evidence across systems no single provider controls.

WHY INDEPENDENT

No single source system can verify the relationship between engineering intent, data provenance and production reality.



Ecosystem shown for context. Current and planned integrations are on slide 5.

WHY BUYERS CARE

One independent record across the AI lifecycle.

Buyers can inspect what was claimed against what the underlying systems actually show.

GOVERNANCE PLATFORMS — Credo AI · Holistic AI

Broad enterprise governance, policy and risk orchestration.

LUMOVRA AI · Engineering provenance + buyer-verifiable evidence.

RUNTIME / EXECUTION EVIDENCE — Glacis · Fiddler · Arize

Runtime controls, observability or verifiable execution evidence.

LUMOVRA AI · Building the link from pre-deployment decisions and provenance to runtime evidence.

PROVIDER COMPLIANCE — OpenAI · Anthropic

Deep evidence inside a single provider's own platform.

LUMOVRA AI · Building cross-provider, cross-stack verification.

MARKET

Start with enterprise AI vendors.
Expand into the evidence infrastructure for every AI system.

ESTIMATED INITIAL WEDGE

~4,000

Potential Series A–C enterprise AI vendors

PROPOSED ANNUAL PRICING

$24K+

Growth entry price, with higher-value Scale and Enterprise tiers

INITIAL OPPORTUNITY

$96M

4,000 potential accounts × $24K entry price

BROADER MARKET CONTEXT

$15.8B

Projected global commercial AI governance software spending by 2030

Forrester, Nov 2024.

ENTERPRISE EXPANSION

From AI vendors to enterprise-wide verification

WHY THIS CAN BECOME A LARGE MARKET

01 **AI vendors**
Initial wedge: enterprise procurement creates an immediate need for verifiable evidence.

02 **Internal enterprise AI**
Enterprises need evidence for AI they build and operate themselves.

03 **Agentic AI**
More autonomy creates new evidence requirements for permissions, actions, overrides and approvals.

04 **Assurance ecosystem**
Auditors, insurers and external evaluators need independently inspectable records.

05 **Verification network**
Buyer recognition of Proof Packs can strengthen distribution and create network effects.

BUSINESS MODEL

The buyer asks for proof.
The AI vendor upgrades to deliver it.

Land with evidence history. Convert when diligence creates urgency. Expand as customers govern more AI.

FREE

$0

Distribution layer. Evidence history begins before a buyer asks.

- 1 repo · limited findings
- Evidence history
- 1 Basic Proof Snapshot / month

GROWTH — CORE MOTION

$2K/mo

Series A / enterprise deal motion. Min ACV $24K.

- 3 AI systems
- GitHub + 1 connector
- Framework mapping · Proof Packs

SCALE

$5K/mo

Series B / full-stack AI. Base ACV $60K.

- 8 AI systems
- Multi-source evidence
- Dedicated CSM · V2 modules

ENTERPRISE

From $80K/yr

Regulated or complex enterprise deployments.

- 20+ AI systems · SSO
- Custom connectors · SLA
- Procurement support · bespoke MSA

EXPANSION ROUTES · USAGE-LED

+ AI systems
Each new model, agent or LLM feature.

+ Connectors
Data, runtime and agent evidence sources.

+ History
Backfill from the first AI commit.

+ Agentic evidence
Overrides, logged decisions, approvals and cross-system provenance.

+ RFP / Q&A
Answer questionnaires from the evidence graph.

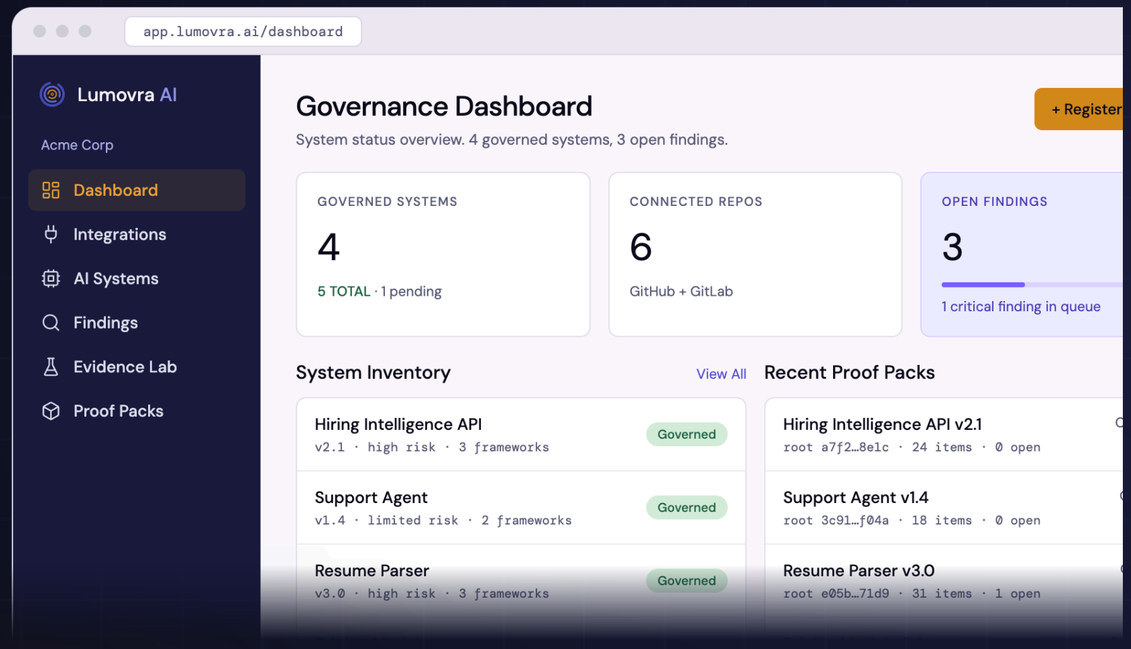
85–90%+ LONG-TERM MODEL TARGET · GROSS MARGIN

120%+ LONG-TERM MODEL TARGET · NRR

TRACTION + GO-TO-MARKET

Live product. Three design partners.

WHERE WE ARE



Live at app.lumovra.ai

GitHub + GitLab live · Free tier since Sept 2026

3 design partners onboarded

DESIGN PARTNER 01
Series A AI SaaS · Enterprise-facing

DESIGN PARTNER 02
Agentic AI platform · B2B

DESIGN PARTNER 03
Ed-tech AI platform

GO-TO-MARKET JOURNEY

NOW — Convert design partners

Get first Proof Packs into live diligence. Measure buyer acceptance and time saved.

NEXT — Repeat paid conversion

Questionnaires, RFPs, audits and buyer requests consistently turn evidence history into paid usage.

2027 — Expand distribution

Law firms, brokers, marketplaces and buyer-side verification become additional channels.

THE TEAM

Built by a team that understands both enterprise buying and the infrastructure behind AI.



Ezgi Turan

CEO & FOUNDER

LinkedIn ↗

- 15 years leading global growth and marketing for SaaS and tech across 6 continents.
- Enterprise GTM and partnerships involving Apple, Google and PayPal.
- Owns product vision, category, growth and fundraising.



Tom Haubner

CTO & CO-FOUNDER

LinkedIn ↗

- Founder with a successful technology exit.
- Architects Lumovra AI's CI-native evidence platform.
- Leads engineering, security, reliability and scale.



David Grant

CRO & CO-FOUNDER

LinkedIn ↗

- Commercial and operational leader focused on enterprise growth.
- Owns enterprise GTM, partnerships and design-partner conversion.
- Network across AI buyers, procurement and regulated verticals.

FOUNDING ENGINEER



Michael Sellmaier

LinkedIn ↗

- Previously worked on an AI governance product at IBM.
- Building V2 and the agentic and runtime evidence infrastructure.

WHY THIS TEAM

We came to Lumovra AI from different sides of the same problem: enterprise adoption, engineering infrastructure and commercial execution. What unites us is a shared conviction that as AI becomes more autonomous, trust has to be earned with evidence. This is personal for us. We care deeply about AI safety, and we believe building better ways to verify AI systems is one of the most important problems we could be working on.

REVENUE MILESTONE

A path to $500K+ ARR with 13 customers.

ANNUAL RECURRING REVENUE

9 GROWTH		3 SCALE		1 ENTERPRISE		BASE ARR
$216K	+	$180K	+	$120K	=	$516K
9 × $24K		3 × $60K		1 × $120K		13 paying customers

NEAR-TERM MILESTONES

1. **Paid conversion**
   Free and design-partner users convert to paid accounts.
2. **Buyer acceptance**
   Proof Packs are used and accepted in actual enterprise diligence.
3. **Account expansion**
   Customers add AI systems, connectors and evidence history.
4. **Repeatable buying trigger**
   Procurement, audits and buyer requests repeatedly trigger paid adoption.

PRICING ASSUMPTIONS

Growth	$24K
Scale	$60K
Enterprise	$120K

LONG-TERM OPERATING TARGETS

Gross margin	85-90%+
NRR	120%+

CAPITAL EFFICIENCY

Software-only product. Marginal cost should stay low as evidence volume scales.

THE ASK

Raising $1.5M to turn technical validation into repeatable enterprise revenue.

Taking Triple Lock™ V2 from first customers to enterprise-ready verification at commercial scale.

USE OF FUNDS · $1.5M PRE-SEED



45% — $675K

Engineering + Triple Lock™ V2

Harden and scale Triple Lock™ V2, expand provenance and runtime integrations, strengthen cross-lock verification and build enterprise-grade reliability.

35% — $525K

GTM + customer validation

Convert design partners, establish a repeatable founder-led sales motion and support the first paying customers.

15% — $225K

Security, legal + assurance

SOC 2 readiness, enterprise security requirements, contracting and regulatory counsel.

5% — $75K

Team + operations

Core company operations, founder compensation and essential operating costs.

WHAT THIS ROUND IS DESIGNED TO ACHIEVE

10–15
Paying customers

$500K+
ARR milestone

5+
Referenceable customers

15–18 mo
Target runway

Lumovra AI

THE INDEPENDENT EVIDENCE LAYER FOR ENTERPRISE AI

The next AI infrastructure giant won't build models. It will verify them.

The agentic AI economy will demand proof at scale. Lumovra AI is building the independent verification infrastructure to deliver it.

RAISING $1.5M PRE-SEED

lumovra.ai →

ezgi@lumovra.ai

The proof your AI can't fake.