INVESTOR PITCH · NOVEMBER 2026

AI Control

Human authority over what AI agents do. Every action asks first, a person stays in charge, and there is a signed record of who allowed what.

Built, pre-launch

Karnan Nataraj, founder · support@solieum.com · Solieum Labs

THE PROBLEM

AI agents now act, and nobody can prove who allowed it

They do, not just say

Agents place orders, change systems, send messages and steer machines. One mistaken step can move money or stop a line.

One shared key

Most agents run on a single API key with wide access. The system on the other end cannot tell which agent acted, or why.

No clear owner

When an agent does something costly, there is rarely a record of which person or rule approved it.

Logs written after the fact do not stop an action, and they are kept by the same system that made the mistake.

WHY NOW

Three things changed at the same time

01

Agents got hands

Models can now call tools, run code and operate software on their own. The gap between an answer and an action has closed.

02

They reached real systems

Companies are connecting agents to payments, infrastructure, customer accounts and physical machines.

03

Oversight became a rule

The EU AI Act requires human oversight for high-risk AI systems. Buyers need a way to show it, not just claim it.

THE PRODUCT

One question before every action: is this allowed?

The agent asks AI Control before it acts. The answer comes back in one of three forms.

Nothing is allowed unless a rule says so. Every answer is written to a signed log that can be checked offline, without trusting us.

Three lines of code in the agent. One console for the people in charge.

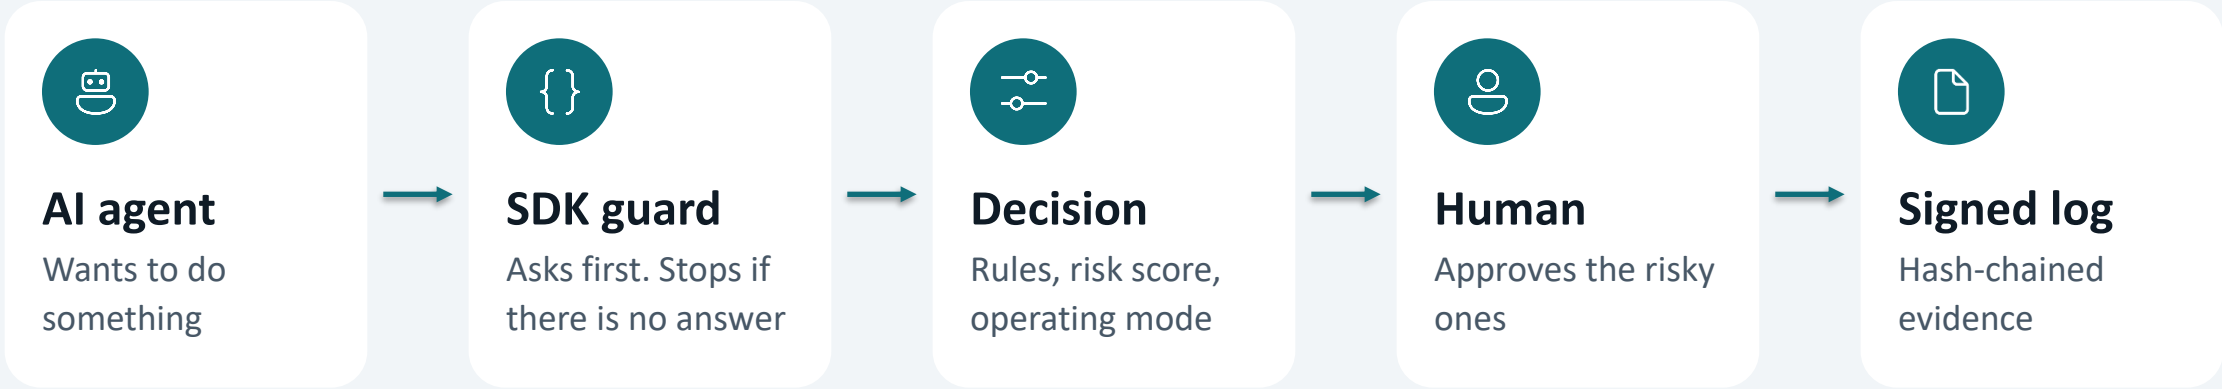
Allow
A rule covers it. The agent proceeds, and the decision is recorded.

Ask a human
It is risky or unusual. A named person approves or refuses.

Deny
No rule allows it. The agent stops. This is the default.

HOW IT WORKS

From the agent’s request to a record anyone can verify



WORKING TODAY

All five steps above, with JavaScript and Python SDKs, single sign-on, webhooks and a console for approvals.

NEXT

A proof of each log batch written to the Solieum chain, so the record does not depend on trusting us. The code exists; no proof has been broadcast yet.

WHAT A CUSTOMER GETS

Six controls that work today

Default-deny rules

Each agent gets scoped grants. Anything outside them is refused.

Human approvals

Risky actions wait for a named person. The requester cannot approve their own request.

Signed evidence

An append-only, hash-chained log. Exports are signed and verify offline.

Risk score and modes

A score per request, and five operating modes from human-only to controlled autonomy.

Containment

Quarantine, revoke or disable an agent, and request an emergency stop for a machine.

Enterprise sign-in

Single sign-on, SCIM, passkeys, roles and tenant isolation.

A DECISION IN PRACTICE

The same agent, three requests, three outcomes

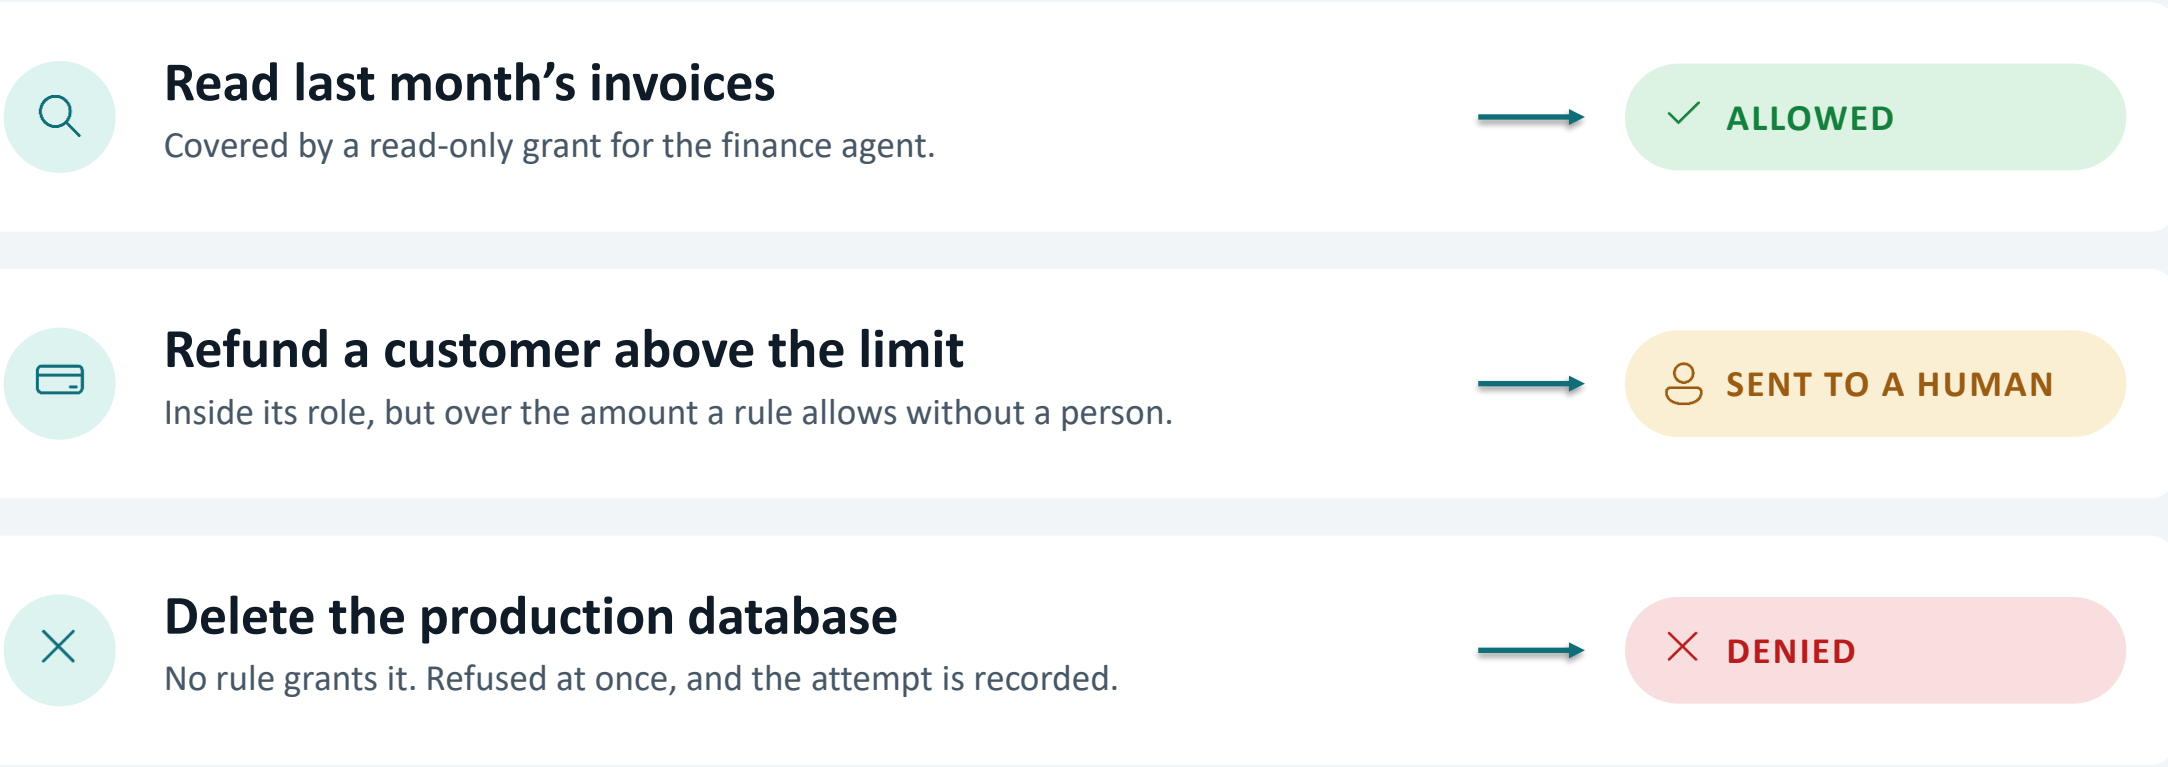


Illustration of how the rules behave. Not customer data.

WHAT IS BUILT

The control plane exists and is tested

81%

of the core control plane is built: sign-in, decisions, approvals, audit, console, SDKs

1,580

automated test cases across the backend, the console and both SDKs

38

database migrations, checksum-checked at start-up

2

SDKs, JavaScript and Python, plus a command-line tool

Source: read-only review of the repository, 15 November 2026. Test cases are counted, not a recorded pass.

WHERE WE STAND

Built, not yet launched

WORKING TODAY

- Organisations, roles, single sign-on, passkeys
- Allow, deny or ask-a-human decisions
- Approvals with separation of duties
- Signed, hash-chained audit evidence
- Console, webhooks, two SDKs
- Robot command checks and stop requests

NOT YET

- A live deployment and a first release
- A proof broadcast on the Solieum chain
- An independent security review
- Real robot and edge connections (simulated now)
- Financial and attack-detection modules (design only)
- Customers and revenue

WHAT MAKES IT DIFFERENT

It decides before the action, and the proof is not ours to edit

A gate, not a camera

Monitoring tools report what an agent did. AI Control answers before the action happens, and the default answer is no.

Evidence that stands alone

The audit log is hash-chained and signed, and verifies offline. Anchoring it on a public chain, planned next, removes the need to trust the operator at all.

One authority for both

The same rules, approvals and records cover software agents and robot commands, so a company runs one authority model, not several.

WHO BUYS IT

Teams putting AI agents into work that matters

Operations and IT

Agents that change systems, tickets and infrastructure. They need approvals and an audit trail. Strongest fit today.

Finance

Agents near payments and trading. They need limits, a second person, and evidence for auditors.

Robotics and industry

Machines taking AI instructions. They need command checks and a stop that a person controls.

First target: operations and IT, starting with a managed IT services provider that already runs AI agents.

HOW IT EARNS

A subscription per organisation, sized by agents under control

PILOT

€4,500 for 3 months

- One team
- The working core
- Our help to set it up

TEAM

€1,500 a month

- Up to 25 agents
- Single sign-on and approvals
- Audit export

ENTERPRISE

From €48,000 a year

- Self-hosted option
- SCIM and support terms
- Chain-anchored evidence when it ships

No pricing has been tested with a customer yet. The pilot is where we learn it.

THE NEXT SIX STEPS

From a tested build to a product we can sell



STEP 1

Deploy

Put it live and tag the first release

STEP 2

Prove the tests

Run the full suite on that build and keep the result

STEP 3

First anchor

Broadcast one real proof on Solieum and verify it

STEP 4

Name and legal

Settle the product name and finish the legal texts

STEP 5

Security review

Commission an independent review

STEP 6

Pilot

Run one pilot with a real team

Target: steps 1 to 4 in 8 weeks, the review and the pilot within 6 months.

RISKS WE NAME OURSELVES

Four open questions, and what we do about each

The chain proof is unproven

No anchor has been broadcast yet.

Step 3: one real proof on Solieum, verified by an independent script.

Control is cooperative today

It governs agents that call the SDK.

Add an in-path gateway and per-agent credentials for the first segment.

No outside security review

A security product needs one.

Step 5, funded by this round, before any paid contract.

The name is contested

ControlAI is an existing UK non-profit.

Step 4: choose and clear a new name before launch.

THE ASK

€750,000 to take AI Control from built to sold



25%

Launch

Hosting, monitoring, backups and the first release.



30%

Trust

Independent security review, legal texts, the name.



45%

First customers

Pilots, real connectors for the first segment, and the team to support them.

This funds 18 months and gets us to 3 paying pilots.

Karnan Nataraj, founder · support@solieum.com · Selvi Suresh, engineering · Suresh Raju, technical advisor