

Four-eyes control for AI agents that move money and records.

Routine actions go through. Risky ones wait for a named person to sign off. Dangerous ones are blocked. Every decision lands in a tamper-evident log built for AI Act Art. 12 and 14.

Bielik tries to change bank details	APPROVE DENY
Claude tries to drop a table	BLOCKED
GPT tries to email 1,240 people	BLOCKED
Mistral tries to refund \$18	ALLOWED
Claude tries to look up an order	ALLOWED

Companies are handing AI agents the keys. Nobody checks what they do with them.

1

Agents act now.

They move money, change records and send email — not just chat.

2

Guardrails check words.

Every cloud filters what an agent says; none decides what it may do.

3

One bad call is enough.

A deleted database or an emptied account can't be undone.

DATABASE

Replit's AI agent deletes a live production database — during a code freeze it was told to respect

The Register · Jul 2025

FINANCE

An AI agent guarding a \$47,000 crypto pool is talked into paying all of it out, in one message

Cointelegraph · Nov 2024

DATA

Asana's MCP server shows ~1,000 customers' data to other organisations for twelve days


BleepingComputer · Jun 2025

SECRETS

Repositories using an AI coding assistant leak secrets 40% more often

GitGuardian · Mar 2025

A checkpoint between every agent and every system it touches.



Nothing irreversible without sign-off

Risky calls wait for a named person to approve or deny.

No feed to babysit

Routine calls clear in milliseconds. People only see the exceptions.

A record auditors trust

Every decision, and who made it, in a tamper-evident log.

Agents got hands in 2025. The rules land in 2027.

2025

MCP goes mainstream

OpenAI, Google and Microsoft adopted one open standard for agents to call real tools.

JUNE 2026

Europe starts buying European

The EU Cloud and AI Development Act, and Poland's sovereignty test for public tech purchases.

2 DECEMBER 2027

High-risk AI Act rules apply

Human oversight (Art. 14) and automatic logs (Art. 12) become mandatory. Fines reach €15M or 3% of turnover.

100 customers is €7.5M a year. 4,200 EU firms already run agents.

SAM

€320M the ~4,200 in the EU already putting agents to work; we sell to those whose agents can write

SOM

€7.5M ARR from 100 customers in five years

Start with EU insurers, banks and claims firms of 200–3,000 staff that run at least two agents touching a system of record.

SAM: Eurostat, 55,000 EU firms with 250+ staff (2024), 55% of them using AI (2025); Capgemini (2025), 14% of large firms deploying agents, global sample; × €75K, the Growth tier. SOM is our assumption.

One engine. Three views of every agent action.

PROVE

Show who approved what


Audit trail		CHAIN VERIFIED ✓
09:38	claims-adjuster issue_claim_payment	APPROVED D. Whitfield
09:37	payroll-bot export_employee_records	DENIED M. Ilić
09:37	platform-ops run_sql	BLOCKED
09:36	payroll-bot get_benefits_summary	ALLOWED

Export NDJSON

Auditors read a hash-chained record and export it. They can't approve anything.

OPERATE

See what agents do across the org



A manager panel, cost and usage by agent and model, and an org chart that decides who sees which agents. Agents and teams, not individual people.

BUILD · FREE ON-RAMP

Gate agents while they work

Pick a gate 7 TEMPLATES

- Database gate**
Denies DROP, DELETE without WHERE
- Financial gate** ✓
Holds payouts above €5,000
- Personal data gate**
Holds SSN and card numbers

+ Build your own

Every tool call is checked through the SDK, the MCP gateway or the Claude Code hook. A CI check does the same for pull requests.

Sample data, real components. Connect once, pick from seven gates, approve only what is held. Live demo on request.

One annual fee per company, from €30K a year.

Not metered: a product that ends surprise AI bills shouldn't send one of its own.

COMMUNITY

€0

Open-core engine, self-hosted

DESIGN PARTNER

€30K

Up to 100K decisions a month

GROWTH

~€75K

Up to 500K decisions a month

ENTERPRISE

Custom

Unlimited, SLA, multi-entity

CONTRACT

Paid upfront or quarterly

Annual contracts, so revenue is predictable from the first design partner.

TARGET

€50K–€120K a customer


Blended annual contract value: design partners start at €30K and move to Growth at renewal. 6–8 of them reach €400K–€600K ARR within a year of general availability.

EXPANSION

Grows with the fleet

More agents and more teams under the same sign-off, not usage creep. Paid tiers add compliance reporting, the evidence export and support.

Engineers adopt it free. Risk and audit sign it off.



THE BUYER

Head of Engineering at a 200–3,000-person EU insurer, bank or claims firm

Just told to give an agent write access. The CISO or Head of Internal Audit signs, from a security or compliance budget.

“Who approved this payment, and can you prove it?”

First 6–8 design partners are regulated firms, insurance first. Banks and the public sector follow as EU rules start favouring EU vendors.

Others block or ask. We pick what needs a person, and prove who signed.

	GOVERNS ACTIONS	PICKS THE RISKY CALLS ITSELF	WHO SIGNS OFF	TAMPER-EVIDENT RECORD	UNDER EU LAW
FILTER WORDS					
Bedrock · Azure Content Safety · Model Armor · Lakera	words only	—	—	—	—
BLOCK ACTIONS					
WitnessAI · Operant AI	✓	✓	—	—	US
Zenity	✓	✓	—	—	Israel
Veto	✓	✓	—	—	self-host
Invariant (Snyk) · Lasso	✓	by hand	—	—	self-host
ASK A PERSON					
HumanLayer	✓	by hand	a reviewer	—	US
Auth0 for AI Agents	✓	by hand	the user	—	US
Permit.io · Frontegg AgentLink	✓	by hand	a reviewer	—	Israel
Inlock	✓	✓	✓ named person	✓ hash-chained	✓ EU

Vendors' own material, Sept–Oct 2026; — means not stated. HumanLayer, Permit.io and Frontegg log decisions but don't call the log tamper-evident. Auth0 asks the agent's own user, not a second person.

Security tools stop attacks. We run the sign-off. Snyk or a platform can sit alongside us: a fleet mixes vendors, and an auditor wants one independent record.

Hard for a US platform to copy, and hard to build over a weekend.

Four-eyes control, built in

The engine picks the risky calls, a named person signs them off, and the record verifies itself. Security tools are built to stop attacks, not to run the sign-off.

EU jurisdiction, by design

The plan: an EU entity, EU hosting and EU inference, so the whole stack sits out of CLOUD Act reach. A US vendor can't match it without a new legal entity.

Open-core: judge it yourself

Buyers can read and run the engine before they pay. A regulated buyer's security review starts from code, not a sales deck.

Sized for the mid-market

WitnessAI starts at 1,000 users. A 300-person insurer running forty agents is too small to buy it.

Pre-revenue, not pre-product. It runs end to end today.

<1 ms

to decide most calls, without calling a model

18/18

test scenarios passed on 7 mock MCP servers the engine had never seen

2

live external agents governed: a Mistral connector, a Bielik claims agent

4

ways in: SDK, MCP gateway, Claude Code hook, CI check

PIPELINE

US interest first, EU sales next

Alight and Travelers will try the sandbox. No EU prospect or signed deal yet: this round buys the first design partners. Next: ten calls with EU insurers and banks. Which agent can write to a system of record, who signs it off today, and does any tool they own cover it?

Inlock

DEMO WORKSPACE · SAMPLE DATA

128governed

119allowed

6held

3blocked

WAITING ON YOU

claims-adjusterissue_claim_payment

€18,500 is above the €5,000 review threshold

APPROVEDENY

09:37:51payroll-botexport_employee_recordsHELD

09:37:47claims-adjusterlookup_policyALLOWED

09:37:41platform-opsruncsql · DROP TABLEBLOCKED

09:37:36payroll-botget_benefits_summaryALLOWED

09:37:31claims-adjustersend_claimant_emailALLOWED

09:37:22payroll-botupdate_direct_depositHELD

Audit trail · 128 entries · chain verified ✓

What a reviewer sees: a payout and a bank-detail change wait for a person, a DROP TABLE is blocked, the rest go through in under a millisecond.

A founder who has been the buyer, built for the regulation, and shipped the mechanism.



Roman Belfer, Founder
Poland, EU

THE BUYER

Director of Engineering at Travelers, a Fortune 500 company and the 6th-largest US P&C insurer. Led its modernisation and security programmes.

THE REGULATION

CTO at EyreACT, an EU AI Act governance platform, building compliance tooling before the rules made it urgent.

THE MECHANISM

Piloted AI-guardrailed CI/CD checks at Travelers that caught AI-written code breaking business logic.

THE TEAM THIS ROUND BUILDS: FOUR PEOPLE BY MONTH 6

FIRST CLOSE • €400K • MONTHS 0-3

Founding engineer: approvals in Slack and Teams, the evidence export, transaction limits

SECOND CLOSE • €550K • MONTHS 3-6

Platform & security engineer: EU hosting, durable audit storage, uptime, audit readiness

FULL ROUND • €700K • BY MONTH 6

Solutions engineer: connects the first customers’ agents and runs onboarding while the founder sells

Raising a pre-seed round from EU capital.

€700K

First close at €400K funds the founder and a founding engineer. A second close at €550K adds a platform and security engineer. The full €700K adds a solutions engineer, the audit and paid pilots.

MONTHS 0–3

Ten EU buyer calls, then the EU entity, first engineer, open-core launch

MONTHS 3–6

Slack and Teams approvals, Art. 12/14 evidence export and per-agent transaction limits ship; general availability on EU hosting

MONTHS 6–12

6–8 paid design partners; third-party audit

MONTHS 12–18

€400K–€600K ARR and a seed round, before the December 2027 deadline



CONTACT **Roman Belfer** · [linkedin.com/in/roman-belfer](https://www.linkedin.com/in/roman-belfer)