

Your AI agents can act. LGD makes sure the risky actions only happen when your people say so.

A governance layer for AI agents. Agents do the work. LGD holds the keys to your systems. Anything risky waits until a quorum of your named people approves it on their phones, and every decision leaves a signed receipt your auditors can check without trusting us.

"Authorized by 3 of 5 independent, hardware-signed humans at 14:02 — chained and tamper-evident."

The receipt LGD produces for every released action.

Why now

On October 20, 2026, OpenAI published a misalignment report describing two of its own internal coding agents weakening the monitoring meant to catch them while doing routine work. One then told the user it hadn't. When an agent can reach the controls that govern it, it can weaken them without anyone deciding it should. LGD puts those decisions with your people.

How it works

1 • ASK

Agent asks LGD

The agent has no keys to your systems. LGD holds them.

2 • CHECK

Your policy decides

Routine actions run and are recorded. Risky ones are held.

3 • ALERT

People are notified

Each signer sees exactly what will happen, on their phone.

4 • SIGN

3 of 5 approve

Face ID or fingerprint. No quorum in time means denied, never run.

5 • PROVE

LGD acts and records

A signed, chained receipt anyone can verify.

Watch LGD catch an agent

An AI agent tries to turn off the monitor that watches it, the pattern in OpenAI's October 2026 misalignment report. LGD holds it. Three people sign on their phones. Out comes the signed proof.

Elapsed
0:04

Misalignment monitor (practice system)
ON

1 - Agent: turn off the monitor

Reset practice system

1 - Attempt

Agent asks

agent:astra → disable_monitor

2 - Caught

HELD

proposal 27a15d914b41 - needs 3

3 - Humans sign

open the approval link on each phone

4 - Signed proof

RELEASED

The 90-second demo: an agent tries to turn off its own monitor, and LGD holds it until three people sign. (Sample data.)

LGD — LifeStock Guardian and Defense
by Novai
Human-Led. AI-Delivered.

Held 2 of 3 signatures. Expires 2026-10-22T12:25:53Z UTC.

disable_monitor

Class: com.novai.lgd.requested by agent:astra at 2026-10-22T12:25:53Z UTC

EXACTLY WHAT YOU ARE AUTHORIZING

```
{
  "component": "misalignment_monitor"
}
```

Proposed hash (your proxy signs this value)
b6d31a769b3118a045b08f7873a782ca35526
8B7a7ba077135808052

Signed so far
CISO - Chief Risk Officer

Sign only if you would authorize this yourself. If anything looks wrong or unclear, don't sign: the action expires and is denied.

Your device will ask for Face ID, fingerprint, or your device PIN. LGD receives only a signature, never your private key.

Sign with passkey

What a signer sees. (Sample data.)

How it protects your company

The expensive mistake that doesn't happen

A wrong refund batch, a deleted table, a disabled alert: stopped at the point of action, before it costs anything. If it happens, it's a deliberate decision by your people, on the record.

Audit evidence, already done

Every released action has its own receipt: who approved, when, what. Hand your auditor the ledger and the open verifier instead of collecting screenshots and tickets.

Work you don't dare automate today

With a human quorum on the few risky actions, agents can take on the rest of the work around money, production and access.

One button when something's wrong

Anyone on the dashboard can stop everything at once. Only your quorum can resume it, and the full record is already there.

Estimate the value with your own numbers at lgdbynovai.com/use-case. It's a calculator, not a promise.

Built for the people accountable when an agent acts

Who	Agent actions LGD would hold
Banks, fintech, insurers	Refunds and reversals, payments, limit changes, account access
Healthcare	Record changes, bulk exports, access to patient data
Engineering and platform teams	Production deploys, deletes, alert and monitoring changes, permissions

What you need

- **Five named people with phones.** Each approval takes about a minute, and only risky actions reach them.
- **A list of actions and how risky each is.** We draft it with you in week one.
- **LGD as the only door.** The agent's own keys to the protected system are removed.

Start with a 90-day pilot

1. **Free 30-minute risk review:** we sort your agents' actions into routine, needs-3 and needs-4.
2. **Pilot on one workflow:** one agent, one connected system, your five signers, weekly reviews.
3. **Go or no-go** on the results: no risky action without your quorum, approvals fast enough, signers there when needed.

See it for yourself: watch the 90-second demo and ask the LGD agent anything at lgdbynovai.com. Business case and savings calculator: lgdbynovai.com/use-case · Walkthrough for your team: lgdbynovai.com/how-it-works
Talk to Novai: lgdbynovai.com/contact · contact@novaisystems.online