VeriProof

For business and risk leaders

Your agents work for the business.



Business Owner

Is this agent worth it, and what risk am I accepting?

Governance Engineer

Will this rule do what the business asked?

Compliance Officer

What's covered, and who's fixing the gaps?

Developer

What does this rule need from my application?

Administrator

Who can do what, and who approved the change?

Auditor

Can I verify this without relying on the team that ran it?

VeriProof makes sure they answer to it, from the first rule to the final audit.

When an agent's action needs a person, VeriProof sends it to the one with authority and hands their decision back to the agent. Rules are tested against real activity before they enforce, evidence builds as agents run, and you see what each agent returns against its AI spend, with a monthly budget for each one.

The Business Owner's home page shows the decisions waiting and the value the agents return, with their recorded AI spend beside it.

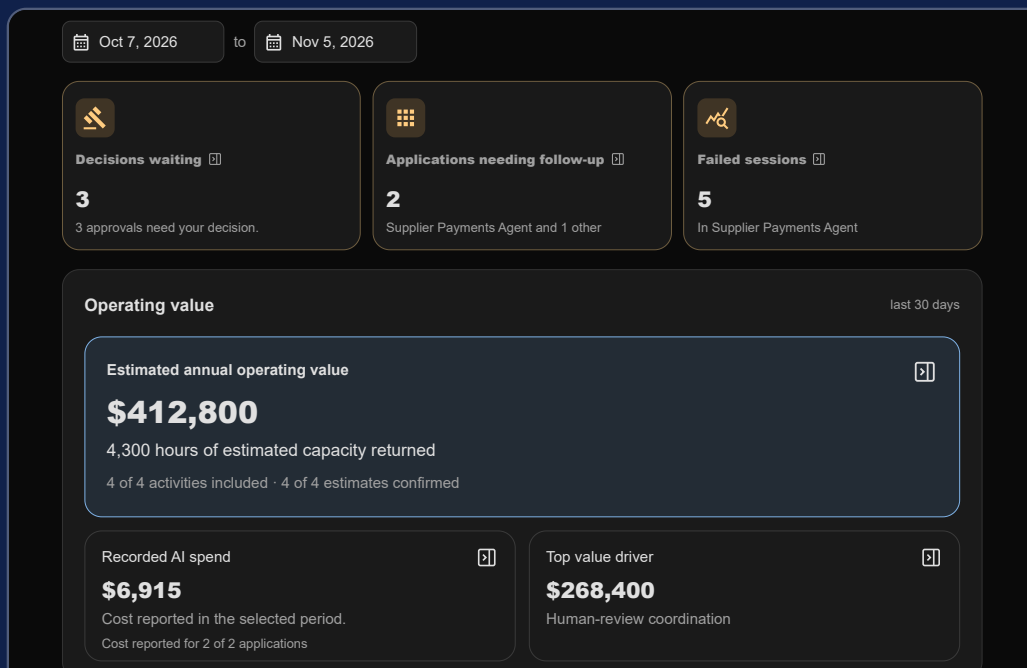

The business control plane for AI agents

veriproof.app

VeriProof

Nothing falls between the teams that answer for your agents.

Every agent you add uses the same handoffs, so the work stops spreading across email, chat, tickets, and meetings, and the record shows who did what.

An agent action needs a person
The reviewer with authority gets the facts and can approve, change, or reject the action, and the decision goes back to the agent. You can see how long each team takes to decide.

A rule needs data the agent doesn't send
The developer gets a fix request that closes when the data arrives, and compliance is asked to check again.

A rule changes for a production agent
An Administrator or Compliance Officer approves the tested version before it applies.

A control isn't met
The assigned owner gets the fix with a due date and a priority, and it's tracked to resolution.

A risk needs a business decision
The Business Owner decides on compliance's proposal to accept it, and approved risks come back for review.

An audit begins
The auditor works inside an engagement scoped to the applications and dates under review, with evidence they can verify as authentic.

Control at the moment of action
Rules written in your industry's terms are tested against real activity first, then decide each agent action before it runs, from Allow through Escalate to Deny.

Evidence in your GRC system, insight in Splunk
Evidence for 11 frameworks, including NIST AI RMF, goes to Vanta, ServiceNow, or ComplyAssistant, and anyone can verify a signed pack. Splunk shows each agent's value and AI spend.

Covers agents you can't change
SDK-built and Copilot Studio agents connect directly, and the Agent Gateway checks the rest as their traffic passes. Run it as managed SaaS or as Private Enterprise in your own Azure.

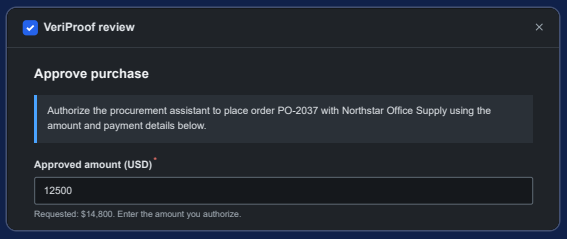

A reviewer approves $12,500 of the $14,800 the agent asked for. Reviews reach ServiceNow, Slack, Teams, email, and the VeriProof Inbox.

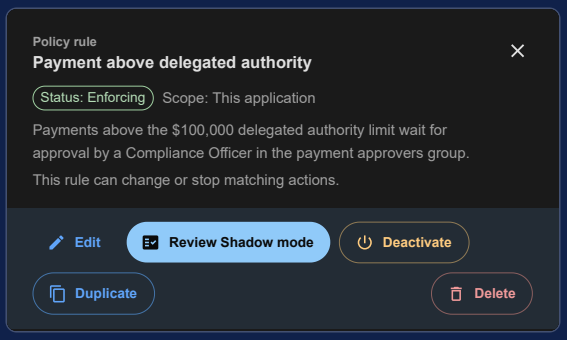

The Governance Engineer tests a rule against real activity before it enforces.

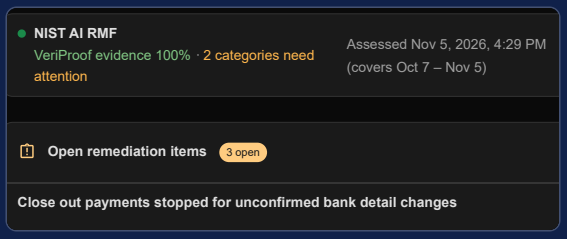

The Compliance Officer sees framework evidence and tracks open remediation.

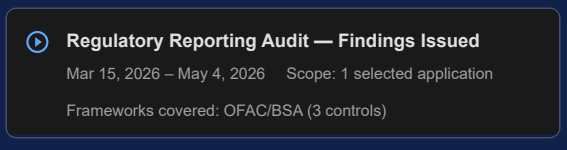

The auditor works in an engagement scoped to the applications and dates under review.

40%

of enterprises will demote or decommission autonomous AI agents by 2027, due to governance gaps identified only after production incidents occur.[1]

Start with one agent workflow
Put its rules, handoffs, and compliance evidence into VeriProof and measure the review and evidence work against how you run it today.

To talk first, visit veriproof.app/contact/.

Start free for 60 days
veriproof.app/start/

[1] Gartner, Inc. press release, May 26, 2026. Product names belong to their owners. © 2026 RJR Labs LLC. VP-E01, Nov 2026.