VeriProof

For platform and security leads

VERIPROOF AGENT GATEWAY

Get your AI agents from pilot to production.

Beyond the gateway essentials, the VeriProof Agent Gateway brings what a production sign-off needs: business decisions on each call, exceptions escalated for human review, controls mapped to your frameworks, and evidence an auditor can check. It runs inside your network, for the agents you build and the agents you buy.

GATEWAY ESSENTIALS, INCLUDED

- Model, MCP and A2A routing
- Credentials kept off agents
- Token and spend limits
- Retry-safe tool calls
- OpenTelemetry metrics

Production means business.

THE BUSINESS CONTROL PLANE FOR AI AGENTS

ENFORCE

Business decisions on each call

- Your published rules, checked before each call goes out
- Outcomes beyond allow and deny, including warn and escalate
- One rule set for custom and vendor agents

COORDINATE

A person decides the exceptions

- Human-in-the-loop review before high-impact calls run
- Reviewers decide in Teams, Slack, ServiceNow or the portal
- Each reviewer's authority confirmed before the decision is recorded

COMPLY

Controls mapped to your frameworks

- Decisions linked to NIST AI RMF, ISO/IEC 42001, the EU AI Act and more
- Gaps tracked as remediation work
- Approved reports to Vanta, ServiceNow or ComplyAssistant

PROVE

Evidence an auditor can check

- Evidence matched to each application's risk
- Sealed records that show if anything changed
- Outcomes credited to the system that reported them

11%

of organizations run agentic AI in production; 38% are still piloting. DELOITTE

14.4%

say all their AI agents went live with full security or IT approval. GRAVITEE

40%

of enterprises will demote or retire autonomous agents by 2027 over governance gaps. GARTNER

Included in every SaaS and Private Enterprise plan · Release 1.1.0

veriproof.app/agents/agent-gateway

Deloitte, Tech Trends 2026. Gravitee, State of AI Agent Security 2026. Gartner press release, May 26, 2026.

It runs where your agents run, on hosts you manage.



Run the VeriProof Agent Gateway inside your private network, close to the agents and tools it protects. Use a compatible Linux host in your data center or in your own Microsoft Azure, AWS or Google Cloud environment.

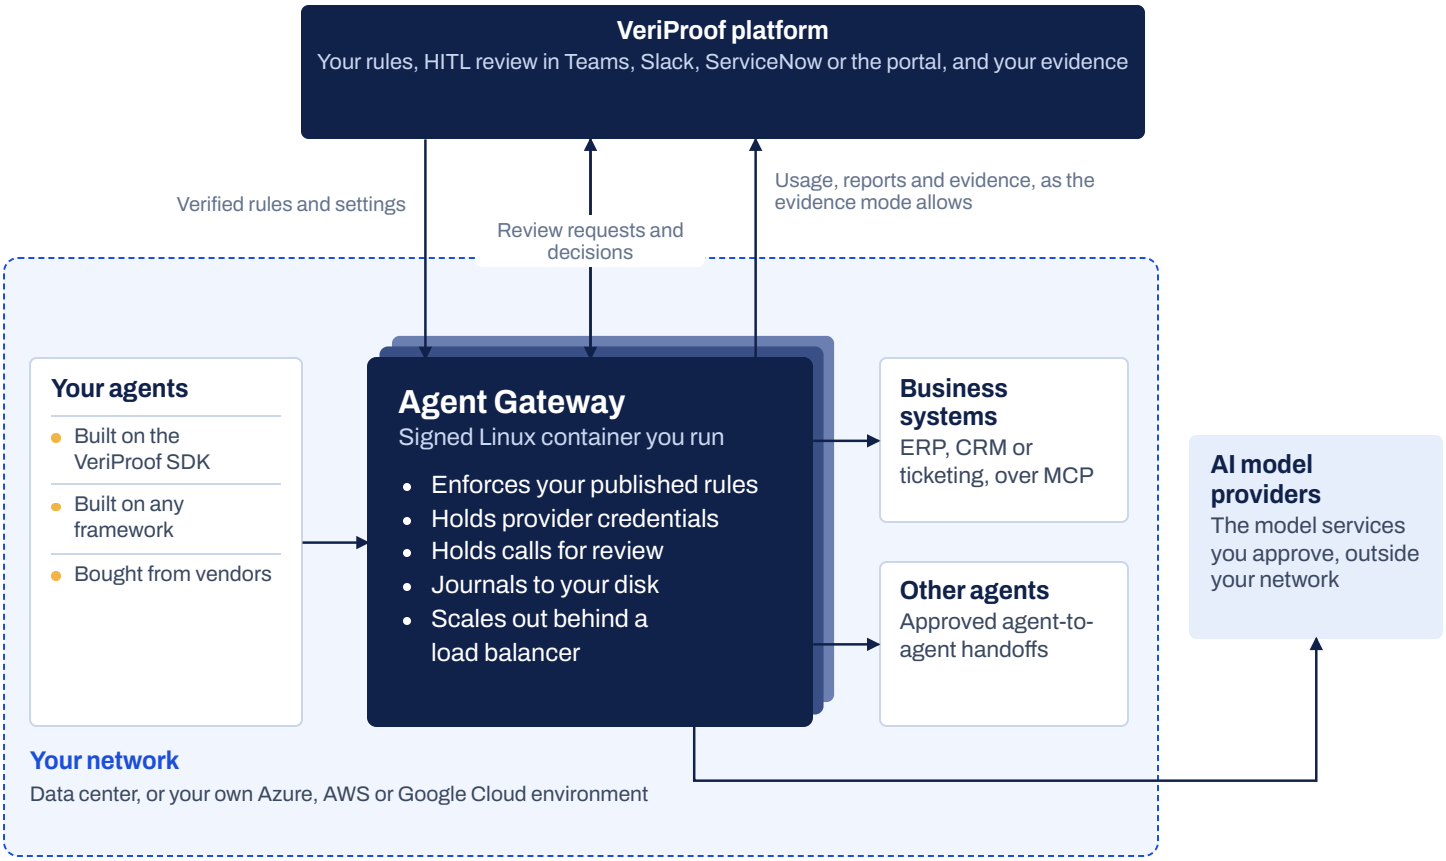

Pick how agents reach it

Edge

For agents built with the VeriProof SDK. They ask the gateway beside them for policy decisions, so routine checks stay inside your network.

Firewall

For any agent whose model, MCP or agent-to-agent traffic you can point at the gateway, including vendor agents you can't change. No SDK needed.

Combined

Both from one gateway, for teams running SDK agents beside proxied ones.

Install it on hosts you manage

The release is a signed Linux container for x86-64 hosts. Install it with Docker Compose or Kubernetes on a server in your data center, an Azure Linux VM, an Amazon EC2 instance or a Google Compute Engine VM. To scale out, run more gateways behind your load balancer, each with its own journal.

What it needs to reach

The gateway connects to the VeriProof platform for your published rules and human reviews, and to the providers and tools you approve. It sends activity and evidence back to the platform. If the platform connection drops, routine checks continue on the rules the gateway already verified while they're current, and evidence waits on the gateway's own disk. Reviews wait for the connection to return.

Set what each route may do, then run it at scale.



What the gateway checks

It runs only the rules your team published for that application, and decides before the call leaves.

Model request MODEL ROUTE	Checks that the model is approved and the caller is within its route limits, then reserves budget before sending. The response is checked before the agent gets it.
MCP tool call MCP ROUTE	The tool is on the approved list and its arguments meet your rules. When a rule asks, the call is held for review before it runs, or its result is held before release.
Agent handoff AGENT ROUTE	Checks the caller's approved skill and the work grant for the receiving agent against policy and route limits. The reply is checked before it's released.
Business action SDK, EDGE MODE	The facts your code sends, such as the amount and the supplier. The decision comes back before the action runs, linked to the workflow it belongs to.

Control access and spend

- Agents authenticate with short-lived workload tokens
- Daily token and spending limits per application and workload
- Provider credentials stay on the gateway host
- Allowances settle on the usage the provider reports

Retries that don't send twice

- Protected tool calls are recorded before they're sent
- A retry of a recorded dispatch isn't sent again

Restart test. In our Linux restart test, the gateway stopped right after sending a tool call. After the restart it recognized the agent's retry, and the tool saw a single call.

See what's running

- Portal view of each gateway and the readiness of its applications
- Metrics over OpenTelemetry to the monitoring you run

Part of the VeriProof platform

Review and reporting	Gateway escalations and evidence use the same VeriProof review, compliance mapping and reporting as the rest of your agents.
For delivery partners	Deploy the same gateway pattern in each client's environment, covering both custom and vendor agents.
One decision history	Gateway decisions sit beside your other agents' decisions, under the same rules.

Send exceptions to a person, and keep the record.

When a rule asks for a person

1. **A rule holds the call**
   The gateway stops the tool call before it reaches the system it would change.
2. **A person decides**
   An authorized reviewer sees the proposed call and the rule behind it in Teams, Slack, ServiceNow or the Customer Portal, then decides.
3. **The gateway sends it**
   After an approval, the gateway rechecks the approval and current policy, then sends the call.
4. **The system reports back**
   The tool's result is recorded with the rule and the decision, credited to the system that reported it.

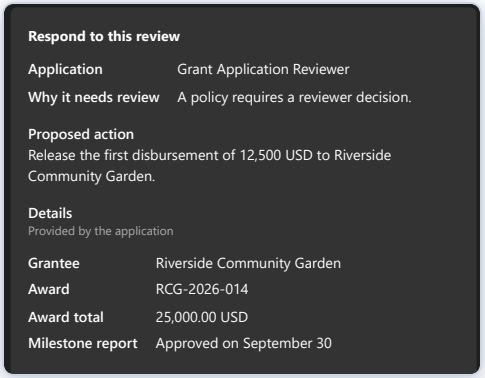

Example: an agent asks the grants system to release a $12,500 disbursement. The reviewer sees this Microsoft Teams card before deciding, and the grants system's confirmation is credited to it in the record.

Choose the evidence each application needs

Your policy applies in all three modes. The mode sets how much evidence VeriProof records for each application. Captured evidence joins VeriProof's sealed record, where an integrity check shows whether anything changed after it was written, and it flows into the approved reports VeriProof sends to Vanta, ServiceNow or ComplyAssistant.

EVIDENCE MODE	WHAT VERIPROOF RECORDS	A GOOD FIT FOR
Policy only	Usage totals and reported model spend, with no per-request records	High-volume traffic you need to control and count
Capture on escalation	The reviewed action and decision for tool calls held for review, plus the reported outcome	Approval workflows where routine calls stay light
Capture evidence	A record of each decision and session in the application's governed work	Work you'll investigate or audit, linked workflows, model output review and agent handoffs

240
calls a second in Policy only

119
calls a second in Capture evidence

Measured in Azure. One gateway capped at 2 vCPUs and 512 MiB served eight applications, each running a three-call workflow against test providers: a policy check, a model request and an MCP tool call.

The performance paper covers review workloads and resource use, so you can size a fleet. Find it at **veriproof.app/agents/agent-gateway**.

Plan a gateway with an engineer

Bring one agent and the rule that matters most to it. In a technical review, we'll work out where the gateway sits and which evidence mode fits, then walk one approval through.

Read the install guide

Host requirements and supported routes, with Docker Compose and Kubernetes steps.

docs.veriproof.app/agent-integrations/edge-gateway

Ask for a technical review

veriproof.app/contact/

The disbursement on page 4 is an illustrative example. Measurement source: VeriProof Agent Gateway performance, November 2026. Product names belong to their owners. © 2026 RJR Labs LLC. VP-G01, Nov 2026.