
USE CASE

See what each AI agent can reach and did

Aribot finds every AI agent that holds a credential and shows the access it was granted beside the access it used, with a verdict on every action.

THE PROBLEM

A team wired an AI agent into production with an API key, and nobody can say what it can reach or who signed it off.

Its permissions were set when the key was issued and nobody has compared them with what it does. When it acts at night, the log shows the action and says nothing about whether anything ruled on it.

WHY IT MATTERS NOW

82%

of organisations discovered shadow AI agents in the past year; 65% had at least one AI-agent-caused security incident in 12 months

Source: CSA & Token Security, Apr 2026

HOW ARIBOT HANDLES IT

1 **Discover credentialed agents**

The inventory starts from issued credentials, so agents nobody registered still appear.

2 **Compare granted with used**

Each agent shows the access it holds beside the access it has exercised.

3 **Read the decision trail**

Every action carries its verdict: allowed, refused with a reason, or errored before a decision.

4 **Rehearse a tighter boundary**

Replay recent activity against a narrower scope to see what it would have refused before you enforce it.

WHAT YOU GET

- An agent inventory that includes the unregistered ones
- A list of granted scopes never used, ready to revoke
- A recorded verdict on every gateway decision, refusals included
- An answer for the auditor on who did what and who agreed

SEE IT IN THE DEMO

- An agent inventory with labels such as REDUCIBLE, RIGHT-SIZED, UNREGISTERED and NOT ASSESSED
- A gateway decision panel where a request is refused with the reason `scope_not_granted`
- Granted against exercised scopes for a sample agent named `ci-deploy-bot`
- A step to set a boundary and rehearse its effect before it is enforced

See it for yourself

ayurak.com/demos/agent-assurance

€40 per Shadow AI scan · licence for agent governance