

Everyone has agents. Nobody has the list.

82%

run AI agents they never approved.

CLOUD SECURITY ALLIANCE • APR 2026

97%

expect an agent incident. Only 6% have funded a response.

ARKOSE LABS • FEB 2026

\$670K

added to a breach where AI ran ungoverned.

IBM COST OF A DATA BREACH

Governance is not a cost. It is the cheaper alternative.

EU AI ACT

NIST AI RMF

ISO 42001

SOC 2



Find out what's already running.

Two weeks, read-only. Every AI agent, assistant and connector in your environment, what each can reach, and what it is actually doing.

Most teams are surprised by the number.



Explore

BAIRAV.AI




Book a demo

DEMO.BAIRAV.AI

START WITH VISIBILITY.
SCALE WITH CONTROL.



Your AI agents are acting. The question is who's watching.




The control layer enterprises deploy to see, approve, govern and audit every AI agent.



Competence. Trust. Control.

"The three frictions blocking enterprise AI adoption:"
BAIRAV answers all three.

WHARTON HUMAN-AI RESEARCH • APRIL 2026



EXPLORE BAIRAV

Scan to explore

BAIRAV.AI

From discovery to evidence, in one control plane

Bairav does not build your AI agents. It governs them, continuously, not once at onboarding.

01

Discover

Cloud agents, local endpoints, autonomous agents and workflow automations, found by the platform, not declared by the teams running them.

02

Register

Everything discovered enters the Intake Queue and is classified before it reaches production.

Approve · Quarantine · Reject



03

Risk score

Four weighted factors, recomputed live as the record is edited.

Data sensitivity	30%
Permissions	30%
Autonomy	25%
External exposure	15%

SCREENED AGAINST OWASP AGENTIC TOP 10

04

Govern

Rules grouped by what they do - block, require approval, flag, notify - scoped to all agents or just one.

Scope is honest. The platform never claims control it doesn't have.

05

Monitor

Every model call and every data access, with the decision and the policy behind it.

06

Audit

A tamper-evident, hash-chained record of what each agent tried, how policy decided, who approved it, and what happened. Verifiable and exportable.

ONE AGENT, ONE DAY

Discovered 2:57 PM. Registered 3:00 PM. Later it tried to disclose supplier bank details.

Blocked in 138ms. Recorded permanently.