AYURAK, ON ITS OWN PLATFORM

Seventeen agent requests refused, each with a reason on record

Every AI agent that calls Aribot on our own platform goes through the same gateway our customers use. The gateway decides each request, writes the decision to a hash-chained ledger and refuses what the agent should not do. Of 40 governed requests, it refused 17, and each refusal names its reason.

23	17 of 40	268
agent identities registered with the gateway	agent requests refused, each with a recorded reason	entries in the hash-chained governance ledger, 28 July to 27 September 2026

THE SITUATION

We sell AI agent governance, so our own agents had better be governed. Agents reach Aribot through its API and its MCP server, and each one has to present a credential.

The question we wanted to answer for ourselves is the one customers ask: which agents exist, what may each one do, and when one does something it should not, who stopped it and why.

WHAT WE BUILT

Every agent gets an identity

An agent cannot call the platform anonymously. Each of the 23 identities is tied to a credential, so an agent nobody registered still shows up.

Decide every request

The gateway allows or refuses each request against the agent's scope, its threat status and any halt in force. Of 40 requests, 15 were allowed, 17 were refused and 8 failed before a decision.

Refusals carry their reason

Eleven refusals came from an agent under a threat restriction, five from a halt on execution and one from a forbidden capability. The reason is stored with the decision, not reconstructed later.

Enforce one step at a time

Most capabilities still run in observe mode, where the gateway records what it would have refused. On 18 September we switched one read-only capability to enforcement, chosen because it changes nothing and can only stop an agent that is already suspended or under a threat restriction. One command switches it back.

RESULTS

- Every request the gateway decided has its decision and its reason on record.
- The ledger is append-only and hash-chained, so a gap or an edit shows up when the chain is verified.
- One capability is now enforced rather than observed.
- Eight requests failed before a decision. They are counted as failures, not as allowed.

RELATED USE CASES

See what each AI agent can reach and did · Cut machine access down to what is used · An audit trail the auditor can verify

---

Customer zero: this is Ayurak running Aribot on its own platform. Every figure comes from our own workspace, counted on 1 October 2026. Governance runs in observe mode for most capabilities; one capability is enforced, as described below.

Products: AI Governance & Agent Security, Developer Portal, API & MCP Server