THE AUTHORITY LAYER FOR THE AGENTIC ECONOMY

Authority for your agentic workforce.

AI agents don't advise; they act. They move money, write to production, and delete records on their own, and the access that makes them useful is the access that makes them dangerous. **ValidMind is the independent authority that governs every agent like an employee**, at the point of action, on any runtime.

THE GAP EVERY AGENT PROGRAM HITS

They act, then you find out

Agents move money and change production in milliseconds. Review after the fact can't catch what already fired.

Your policy isn't binding

Rules live in docs and dashboards. Nothing stops an out-of-bounds action at the moment it happens.

No accountability to show

When an agent commits the firm, there's no manager to name, no charter to cite, no transcript to produce.

HOW IT WORKS · govern every agent like an employee

Give every AI agent a manager, a charter, a reporting line, and a record of what it did.

The org chart your AI workforce never had, enforced at the point of action, not just policy on paper. Here's the record one request produces:

1. **Intercept**
   Every agent tool call is caught **in the call path, before it executes**.
2. **Evaluate**
   Checked against policy. Routine work clears at **machine speed**; higher-stakes actions get a closer look.
3. **Decide**
   Allow, pause, or deny. Pauses **route to the right person**, the agent's reasoning attached.
4. **Record**
   Request, verdict, and who approved it, written to **an audit trail you own**.

MANAGEMENT TRANSCRIPT req #A7F2-1106 · logged ✓

AGENT	`data-ops-agent` · reporting line ↑ Data Platform
MANAGER	Priya N., Director, Data Platform
REQUESTED	**Delete production database** · `customers_prod` 1,206 records · irreversible
REASONING	*"Query returns empty, table looks unused. Dropping and rebuilding to clear the failed migration."*
CHARTER	**§4:** this agent may read & write production. It may not delete it.

Paused → escalated up the reporting line

Rejected by Priya N.

"The table isn't unused. ***Request denied.****"*

WHY VALIDMIND

The independent authority, by construction.

Three things a platform that builds and runs its own agents can't honestly claim.

Independent by design

It doesn't run the agents it governs, so it can truly say no. **No model grading its own work**, no platform refereeing the agents it also sells.

Runs on any stack

Runtime-agnostic enforcement in the call path, on any framework, harness, or cloud. **Govern the agents you didn't build**, not one vendor's walled garden.

Compliance that cascades

Write the charter once. Every agent, and every agent it spawns, inherits it, and a shared **rate & cost envelope** caps the loop a parent never authorized.

OPEN-SOURCE PROJECT

Atryum

The control plane, Apache 2.0, live and runnable today.

→

ENTERPRISE PRODUCT, BUILT ON ATRYUM

ValidMind Agent Authority

Design partners, now open

DEMONSTRABLE NOW

Watch ValidMind intercept, judge, and decide on a real agent action in a live demo, not slideware.

GROUNDED IN REGULATION

Aligned to SR 26-2, E-23, SS 1/23, effective challenge, DORA, and the EU AI Act.

GOVERNANCE HERITAGE

A decade of model-risk governance at G-SIB scale, now extended to the agentic workforce.

See ValidMind intercept, judge, and decide on a live agent action.

The authority layer for the agentic economy.

validmind.com/platform/agent-authority

info@validmind.ai