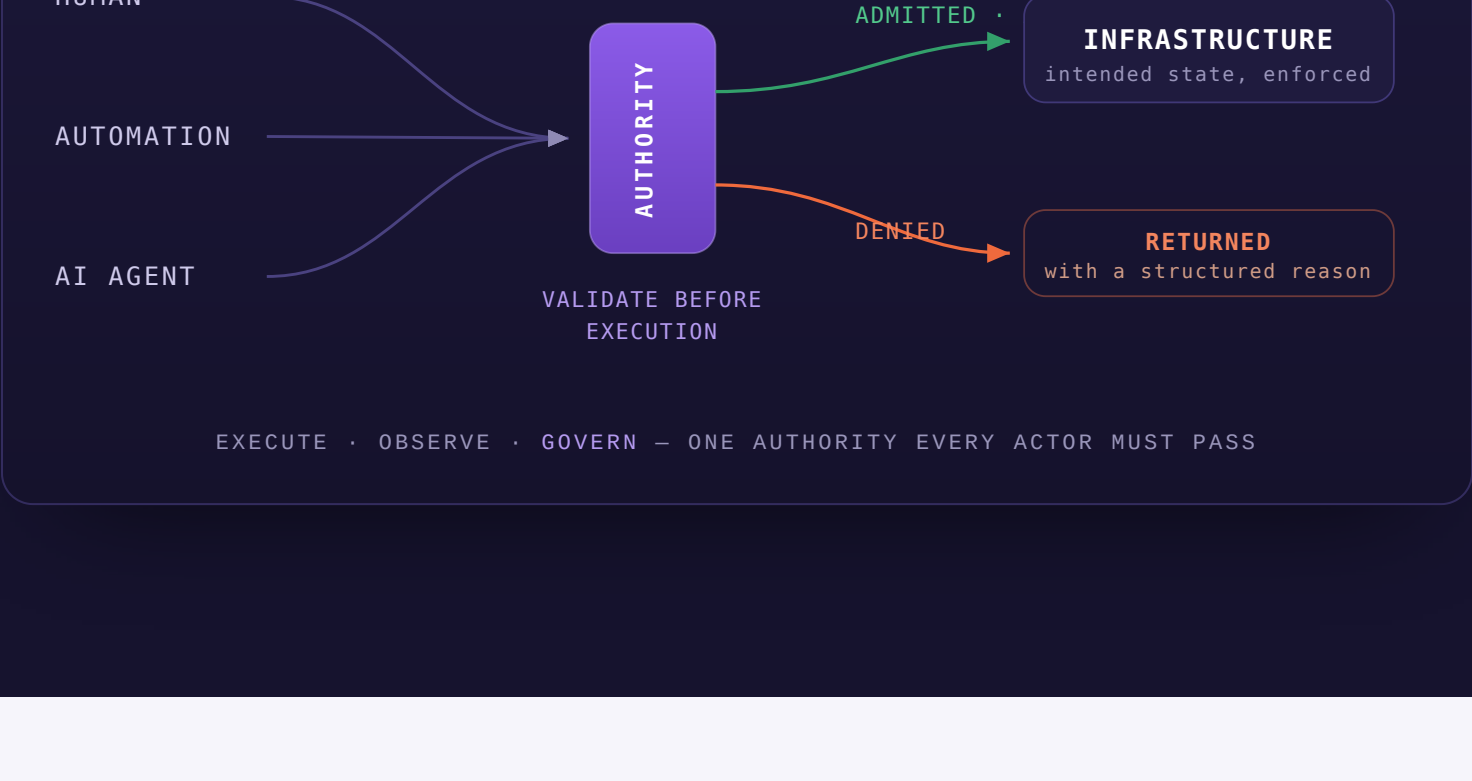


THE AUTHORITY LAYER FOR INFRASTRUCTURE CHANGE

Every change should answer to the same authority.

AuthorIOM is the Infrastructure Operating Model — the authority layer that validates every action, human, automated, or AI, against a living model of intent **before** it reaches your systems.

You do not have a tooling problem.
You have an operating model problem.



Two layers are mature. One is missing.

Enterprises can already make change and watch change. What no tool does is hold change to an authoritative model of what is intended and allowed — and decide, before execution, whether it is admissible.

EXECUTE
HAVE

Make change. Terraform, Ansible, CI/CD — pipelines that apply change quickly and repeatably.

OBSERVE
HAVE

Watch change. Datadog, Splunk, ServiceNow — visibility into what happened, after it happened.

GOVERN
MISSING

Validate change before it executes. Every action proves it is admissible against intent, dependencies, and policy — or it is denied with a reason. This is the authority layer.

Automation executes and forgets. An operating model **persists** a living model of intent and allowed states — and holds every change to it.

Two shifts turned the gap into exposure.

01 — Speed

Automation removed the human pause.
Change now happens at machine speed. A bad change propagates before anyone can intervene, so validation has to happen **before** execution — not in a review meeting after the incident.

02 — Autonomy

AI is now pointed at infrastructure.
Agents can propose and drive change directly. Letting them act and watching closely is not a safety model. The environment needs an authority every actor must pass — human, script, or agent.

AI SAFETY

Governed by design, not by oversight you hope holds.

There are two ways to let AI touch infrastructure. Only one is structural.

✗ OVERSIGHT

The agent acts. Humans watch.
Safety depends on catching the wrong change in time. At machine speed, review loses. The agent can still reach production by going around whatever is watching it.

✓ AUTHORITY

The agent proposes. The model decides.
A change either proves it is admissible against intent, dependencies, and policy — or it is denied with a reason. An AI agent can propose, but it **cannot route around the model**.

WHAT AUTHORIOM DOES

Four functions of an authority layer.

01

Hold the model
Maintain a living, reconciled model of assets, ownership, dependencies, and intended state — one source you trust without checking several places.

02

Validate before execution
Check every proposed change against the model first. Admissible changes proceed and are audited by default; the rest are denied with a structured reason.

03

Reason about blast radius
Evaluate what a change will touch against its dependencies before commit — not discover it in production by way of the incident.

04

Govern every actor
Hold humans, automation, and AI to one authority. No actor — and no agent — gets a path around the model.

THE INFRASTRUCTURE AUTHORITY MATURITY MODEL

Six levels. Most enterprises sit at two.

One axis, ascending: how authoritatively you can say what exists, what is intended, and what is permissible — and enforce it before change executes.

0	Tool-Driven	Knowledge lives in tools and a few people.
1	Documented	Someone wrote it down; it drifts from reality.
2	Observable	You can see what happens — after it happens. The top of the Observe ladder.
3	Modeled	A real, continuously reconciled model of what exists.
4	Governed	Change is validated against the model before it executes.
5	Authoritative	One authority every actor must clear — including autonomous AI.

Level 2 is not failure — it is the summit of observability. **The Govern ladder above it is a different climb**, and it is the one that decides whether automation and AI can act safely in your environment.

HOW IT DEPLOYS

Elevate your stack. Do not replace it.

AuthorIOM sits above what you already run. It reads first, models next, and governs on your terms.

01

Start agentless and read-only
It reads from the systems you already operate — no rip-and-replace, no agents in the critical path on day one.

02

Build the authoritative model
Assets, owners, dependencies, and intended state become one reconciled model you can trust without checking several places.

03

Move change through validation
Route human, automation, and AI changes through pre-execution validation — admitted with audit, denied with a reason — at the pace you choose.

FIND OUT WHERE YOU STAND

Score your environment from 0 to 5 in about a minute.

Get your authority score
A ONE-PAGE RESULT FOR YOUR ENVIRONMENT

Talk to the team
WALK YOUR STACK THROUGH THE MODEL