Cognous Open Control Stack

Open-source reference infrastructure for governable AI agents

Before enterprises can scale AI agents responsibly, they need more than prompts, logs, dashboards, and policy documents. They need a structured lifecycle: Declare → Control → Replay → Evidence.

The Problem

AI agents are beginning to call tools, access enterprise data, propose actions, draft communications, prepare exports, and trigger workflows. Existing governance systems were built for models, software, policies, dashboards, and human approvals — not for agents that propose actions through tools. A chatbot can be governed by reviewing what it said; an agent must be governed by reviewing what it can do, what it tried to do, and what stood behind its output.

Four gaps result:

1. **A declaration gap:** the agent's action surface is scattered across prompts, code, and institutional memory.
2. **A runtime control gap:** proposals, policy decisions, blocked actions, authority, and reliance are not captured as governance records.
3. **A replay gap:** specific runs cannot be reconstructed after an incident or audit question.
4. **An evidence gap:** business reviewers receive technical records but no coherent evidence package.

The Stack

Layer	Artifact	Question answered	Output
Declare	Agent Action Manifest	What may this agent propose?	Declared action surface
Control	Agent Control Plane	What did the agent propose, and what did policy decide?	Runtime control records
Replay	Agent Replay Bundle	Can we reconstruct this run?	Portable replay record
Evidence	Agent Governance Evidence Pack	What evidence supports deployment review?	Business-facing evidence package

The layers feed one another. Declaration gives runtime control a baseline. Runtime control generates records. Replay bundles package those records. Evidence packs translate them into business-facing review material. Every layer answers a question the next layer depends on.

Why It Matters

- Makes agent action surfaces explicit before deployment — tools, actions, authority, and review requirements.
- Records proposed and blocked actions at runtime, with policy decisions and evaluation traces.
- Preserves authority and source reliance context at the moment of decision.
- Makes agent runs reconstructable as portable, validated, redactable, signed records.
- Turns blocked actions into evidence that controls operated.
- Converts technical records into business-facing governance evidence — supporting audit, incident response, legal review, security signoff, customer assurance, and executive oversight.

What It is Not

It is not an agent framework, model runtime, hosted dashboard, complete AI governance platform, compliance certification product, security boundary, substitute for application authorization, or the ODES standard. It supplies structured artifacts for declaration, runtime control, replay, and evidence — turning agent governance from assertion into inspectable record.