Governance isn't a layer.
It's the execution model.

Governed execution for high-risk agent workflows: financial automation, healthcare, infrastructure. Policies don't check what happened. They gate what happens next.

THE PROBLEM

High-stakes workflows. Governance still on the Q3 roadmap.

- ✗ Policy checks added after the fact, after the agent has already acted
- ✗ No audit trail from the first step, only from when you remembered to add one
- ✗ Workflow state lost when infrastructure hiccups mid-run
- ✗ Human approvals are a Slack notification you might miss

```
# Three decorators. Full governance.
from waxell_sdk import agent, workflow, decision

@agent(name="financial-reconciliation")
class ReconciliationAgent:
  @workflow
  def run(self, ctx):
    result = ctx.call(self.validate)
    return ctx.call(self.execute, result)

  @decision
  def validate(self, ctx): ...
```

HOW IT WORKS: BUILD INSIDE THE GOVERNANCE LAYER, NOT ON TOP OF IT.

STEP 01

Build with Waxell SDK decorators

@agent, @workflow, @decision. No separate instrumentation step. Governance is native from line one. Agents on other Python frameworks are covered by Waxell Observe.

→ STEP 02

Every step is policy-gated pre-run

The same 50+ policy categories as Observe, cost, safety, PII, compliance, identity, enforced natively. There is no gap between execution and enforcement.

→ STEP 03

Durable, auditable, resumable

Workflows checkpoint at every step. Human approvals pause execution and wait. Infrastructure restarts, model timeouts, policy triggers, the run resumes from exactly where it stopped.

WHAT YOU GET

Isolated Execution

Every agent run in its own isolated environment. No shared state, no cross-contamination between workflows or tenants. The agent runs inside Runtime, not inside your application process.

Native Compliance

HIPAA, SOC 2, and PCI-DSS profiles out of the box. Data residency in US or EU. The lineage graph starts at run zero, with no post-hoc logging patched on after the fact.

Durable Workflows

Checkpointed at every step. Resumes from exactly where it stopped, after a network failure, model timeout, or human approval gate. Agents survive what production throws at them.

Kill Switches at Every Level

Stop any agent, workflow, or session, immediately. At the agent level, the workflow level, and the session level. No graceful shutdown. No waiting.

Runtime gates each step of the high-risk workflows you build on Waxell.

When the stakes require it: financial transactions, clinical decisions, infrastructure operations. Governance can't be a layer you add. It has to be the execution model.

3	50+	0
compliance profiles	policy categories	gap between run & govern

BUILT FOR: `Financial Automation` `Healthcare Workflows` `Infrastructure Ops`

BUILT WITH: `@agent` `@workflow` `@decision` `Waxell SDK`

WAXELL

waxell.ai/products/runtime