

Don't just observe. Govern what happens next.

Know what every agent did. Control what it does next, before it executes.

THE PROBLEM

A dashboard after the fact is not governance. It's an autopsy.

- ✗ Silent failures with no trace of what the agent actually decided
- ✗ Runaway spend you discover on Friday when you check the spreadsheet
- ✗ No way to stop a misbehaving agent mid-run
- ✗ Dashboards that show you what happened, not what's happening

```
# pip install waxell
# your existing agent code, unchanged
import waxell

waxell.init(
    api_key="wax_sk...",
    api_url="https://api.waxell.dev"
)
```

HOW IT WORKS: OBSERVE EVERYTHING. GOVERN WHAT HAPPENS NEXT.

STEP 01

Install & initialize in two lines

pip install, add two lines before your imports. Every LLM call, tool invocation, and agent decision is captured automatically from that point on.



STEP 02

See the full execution tree

Every span, cost, latency, and decision linked by session and lineage. Full parent-child traces for multi-agent systems, powered by OpenTelemetry.



STEP 03

Policies enforce. Agents respond.

50+ policy categories evaluate in real time, around every LLM call and tool invocation. When one triggers, it returns a clear action: warn, redact, throttle, block, or retry.

WHAT YOU GET

Auto-Instrumentation

200+ Python libraries instrumented at init: LangChain, LlamaIndex, CrewAI, OpenAI, Anthropic. No decorators required to start. No changes to your agent logic.

Full Execution Traces

Every LLM call, tool invocation, routing decision, token count, cost, and latency, captured and linked across multi-agent execution trees.

Governance at Execution Time

50+ policy categories, cost ceilings, content filters, rate limits, safety boundaries, evaluated as your agent runs. Not alerts. Enforceable rules.

Structured Feedback Loop

When a policy triggers, agents don't just stop. Each check returns a clear action: allow, warn, redact, throttle, block, skip, or retry.

Tracing tells you what happened. Waxell changes what happens next.

Your agents are in production. You're finding out what went wrong after it matters. Waxell makes observability a control system.

2

lines to
deploy

50+

policy
categories

200+

Python
libraries

WORKS WITH:

LangChain

LlamaIndex

CrewAI

OpenAI

Anthropic

+ more

WAXELLwaxell.ai/products/observe