ATX Agent Governance

The Governance Layer for Enterprise AI Agents

ATX Agent Governance brings AI evaluation, cost control, security, access control and monitoring for agentic AI systems into one layer. It makes agentic AI governance from policy on paper into end to end visibility and controls, so your every agentic AI system deployment is governed for your risk, audit and compliance management.

Three Controls for AI Agents, One Governance Layer

A Access Control

powered by **RBAC**

What users and agents may touch. Role-based permissions let each reach only the data it is authorized to see, preserving data privacy.

B Budget Control

powered by **FinOps**

What an agent may spend. Token budgets are set per agentic system, and once a limit is hit new requests stop, keeping costs firmly in check.

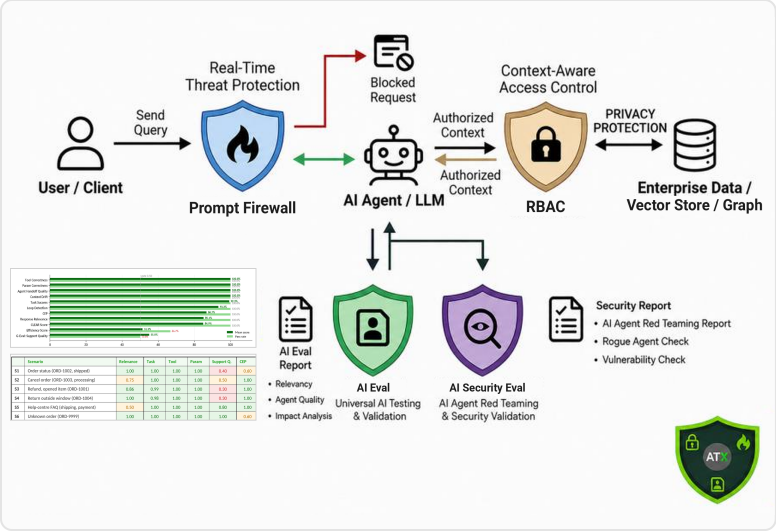
C Control Center

powered by **AI Eval + Observability**

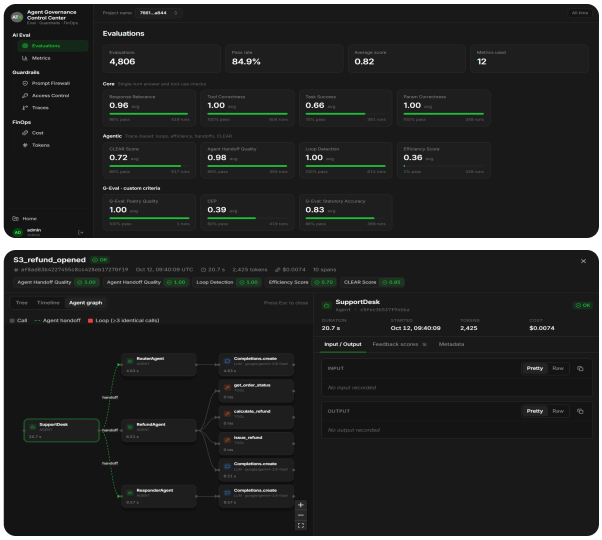
Complete visibility and control of your agentic system. Verify and monitor every AI agent, with an instant alert raised on any anomaly.

What sits inside ATX Agent Governance

- **AI Eval** Tests agents for quality, accuracy, safety and production readiness. Also offered as a separate service.
- **AI Security Eval** Red-teams agents to uncover security weaknesses before attackers can find and exploit them.
- **Agent Observability** One dashboard to monitor agents, with eval scores, security activity and full run traces.
- **FinOps** Token budgets for agentic systems. Automatic alerting upon over budget, keeping costs under control.
- **Prompt Firewall** The first line of defense, blocking harmful, malicious and jailbreak prompts in real time.
- **RBAC** Role-based access control, so both users and agents retrieve only the data they are authorized to see.



Agent Governance in action



Agent Observability dashboard

How to get started?

```
# pip install atx-ai-govern
```

Single Source of Truth for Verifying Your Agentic AI Systems.

ATX Labs

https://atxlabs.ai | hello@atxlabs.ai