

THE ENTERPRISE AI GOVERNANCE IMPERATIVE

Your AI POC Worked. Production Didn't. You're Not Alone.

AI is evolving from models to agents. This changes everything - economics, governance, risk. The enterprises that master AI governance now will define the next decade. The rest will spend it catching up.

THE INFLECTION POINT

Every enterprise is under pressure to deploy AI. POCs are easy - production is where most fail. Meanwhile, AI agents are proliferating across your organization faster than governance can keep up. AI regulation is inevitable. The enterprises that build governance now will be ready. The rest will scramble.

87%AI PROJECTS FAIL TO REACH
PRODUCTION**60x**

COST INCREASE POC TO PRODUCTION

47AVG. UNTRACKED AGENTS PER
ENTERPRISE**12+**

JURISDICTIONS WITH AI REGULATIONS

WHO WE ARE

Engineers with 15+ years building enterprise AI and cloud infrastructure at Google Cloud and Microsoft

OUR FOUNDATION

Peer-reviewed research and production-tested open source from Rotalabs (rotalabs.ai)

OUR FOCUS

Regulated industries where AI decisions have real consequences - banking, insurance, telecom

THE CHALLENGE

Three Crises Converging

Enterprise AI isn't failing because of technology. It's failing because of architecture, economics, and governance. These three crises are hitting simultaneously - and most enterprises aren't prepared.

Crisis 1**The Economics Don't Work**

Your POC cost \$500/month. Production costs \$30,000-50,000/month. Pure agentic AI uses expensive reasoning for every transaction. The business case that justified the project no longer works at scale.

Crisis 2**Shadow AI Is Everywhere**

While you built your official POC, departments deployed their own agents. No inventory. No audit trails. No policy enforcement. Regulators will ask questions you can't answer.

Crisis 3**Governance Is a Void**

How do your AI systems make decisions? Can you provide an audit trail? What policies govern agent behavior? If you can't answer these today, you won't be ready when regulators come knocking.

"We discovered 47 untracked AI agents across our organization. No one knew what models they were using, what data they accessed, or what decisions they were making. We were one regulatory inquiry away from a serious problem."

— Chief Risk Officer, Global Financial Institution

What Your Board Is Asking

- ✓ What is our total AI spend across the organization?
- ✓ How many AI systems are making decisions that affect customers?
- ✓ Can we demonstrate compliance with emerging AI regulations?
- ✓ What governance framework controls AI behavior at runtime?
- ✓ How do we explain AI decisions to regulators or in litigation?
- ✓ What is our AI risk exposure and how do we manage it?

The Competitive Reality

The enterprises that solve AI governance first will have a structural advantage. They'll deploy AI faster, at lower cost, with less risk. They'll satisfy regulators while competitors scramble. They'll attract talent that wants to work on production systems, not endless POCs. This isn't about being first to experiment - it's about being first to operationalize.

THE SOLUTION

Trust Intelligence Platform

Seven integrated products that solve AI governance at the architecture level. Not point solutions bolted on after the fact - infrastructure designed from the ground up for enterprise AI that works.

Guardian

AI reliability monitoring. Detect sandbagging, hallucination, drift before they impact operations.

Steer

Runtime behavior control. Adjust AI outputs without retraining using steering vectors.

Eval

LLM evaluation at scale. Rigorous testing that goes beyond benchmarks to production readiness.

Orchestrate

Multi-agent platform with Trust Cascade. Visual builder, audit logs, human-in-loop controls.

AgentOps

Enterprise agent governance. Registry, lifecycle, policy enforcement, flight recorder.

Context Engine

AI-ready data pipelines. ETL-C processing that preserves context for accurate reasoning.

AgentOps: The Governance Layer

Every AI agent gets a unique identity. Every decision is recorded. Every policy is enforced at runtime. AgentOps gives you complete visibility and control over your entire AI estate - the foundation for regulatory compliance and operational excellence.

- ✓ **Agent Registry** — URN-based identity across all environments
- ✓ **Flight Recorder** — Complete audit trail for every decision
- ✓ **Policy Engine** — Three-layer enforcement (gateway, sidecar, inline)
- ✓ **Composability Mapping** — See how agents interact

Trust Cascade: The Economics

Not "use less AI" - use AI where it matters. Route each decision to the cheapest processing layer that can handle it correctly. Expensive reasoning only when necessary.

86%

Cost reduction vs pure agentic

94%

Detection accuracy maintained

15x

Faster average response

Results from production deployment at global life insurer.

See the Platform in Action

Interactive demo of AgentOps and Trust Cascade.

[View Demo](#)

WHY ROTASCALE

Built by People Who've Done This Before

We're not another AI startup with a pitch deck. We're engineers who spent 15+ years deploying enterprise systems at Google Cloud and Microsoft, and industry veterans from banking who understand what regulated enterprises actually need.

Research-Backed

Every product built on peer-reviewed methods from Rotalabs. We don't just claim - we prove. Our research is published and open for verification.

Production-Tested

Built on open source that's already running at scale. Not untested ideas - battle-hardened systems you can verify before you buy.

Your Infrastructure

On-premise, private cloud, or hybrid. Your data never leaves your environment. Model flexibility with no vendor lock-in.

Engagement Options

ASSESSMENT**\$15K**

1-2 weeks. Current state analysis, agent inventory discovery, risk assessment, roadmap.

STRATEGY**\$25K**

2-4 weeks. AI trust roadmap, target architecture, governance framework design.

IMPLEMENTATION**\$30K+**

4-8 weeks. Product deployment, architecture design, team training, go-live support.

ADVISORY**\$8K/mo**

Ongoing. Monthly strategy sessions, architecture reviews, incident support.

The Path Forward

WEEK 1-2

Discovery call. Understand your current state, challenges, and objectives.

WEEK 3-4

→ Assessment. Agent inventory, governance gaps, risk analysis, quick wins.

WEEK 5-12

→ Pilot. Single use case implementation proving economics and governance model.

ONGOING

→ Scale. Expand to additional use cases with proven playbook and trained team.

Ready to close the AI governance gap?

30-minute executive briefing. No pitch - just an honest conversation about your challenges.

contact@rotascale.com

+1 (415) 524-0007
rotascale.com/get-started