

TRUST INTELLIGENCE FOR THE ENTERPRISE

The Trust Intelligence Platform

Monitor, evaluate, control, and orchestrate AI systems at enterprise scale. Production-grade reliability for regulated industries.

EXECUTIVE SUMMARY

Enterprise AI initiatives fail at production scale. POCs demonstrate capability; production reveals economics, governance gaps, and reliability failures. The Trust Intelligence Platform provides the complete infrastructure to deploy AI that actually works - with the monitoring, governance, and cost optimization enterprises require.

WHAT'S INSIDE

- 01 The Challenge: Why enterprise AI fails
- 02 Platform Architecture: Seven products
- 03 Agent Governance: The shadow AI crisis solved
- 04 Trust Cascade: Right-sizing intelligence
- 05 Getting Started: Engagement options

7

INTEGRATED PRODUCTS

86%

COST REDUCTION

94%

DETECTION ACCURACY

16

WEEKS TO PRODUCTION

Why Enterprise AI Fails at Scale

Your AI POC worked. Production didn't. You're not alone. The gap between demo and deployment isn't technology - it's architecture, economics, and governance.

"We spent 18 months and \$4 million on an AI initiative. The POC was impressive. Production cost 60x more than projected, and we had no visibility into what the AI was actually doing."

— VP of Digital Transformation, Fortune 500 Enterprise

The Economics Problem

Costs Explode at Scale

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

The Governance Gap

Shadow AI Everywhere

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

The Reliability Crisis

90% Isn't Enterprise-Grade

Hallucinations in customer communications. Edge cases everywhere. Decisions that don't match policy. "It works most of the time" isn't good enough for regulated industries.

The Operations Void

Can't Manage What You Can't See

No observability into AI behavior. No framework for when things break. When regulators or auditors ask how your AI makes decisions, you have no answer.

The Hard Truth

Enterprise AI isn't harder because of technology. It's harder because of context:

- ✓ Decisions have regulatory and legal consequences
- ✓ Costs must work at production volume
- ✓ Every AI action needs an audit trail

What enterprises actually need:

- ✓ Economics that scale (not 60x cost explosion)
- ✓ Governance over every AI agent
- ✓ Reliability monitoring and control

02 - PLATFORM ARCHITECTURE

Seven Products. Two Layers. Complete Coverage.

The Trust Intelligence Platform is organized into two complementary layers: a foundation for data and trust, and capabilities for agents and performance.

Data & Trust Layer

Context Engine	AI-ready data pipelines. ETL-C processing that preserves context for accurate AI reasoning.
Guardian	Continuous 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, reproducible testing with CI/CD integration.

Agent & Performance Layer

Orchestrate	Multi-agent platform with Trust Cascade. Visual builder, audit logs, human-in-loop controls.
AgentOps	Enterprise agent governance. Registry, lifecycle management, policy enforcement, flight recorder.
Accelerate	Inference optimization. 8x+ speedup with speculative decoding and custom kernels.

Common Use Cases

Fraud Detection

Trust Cascade routes decisions to cheapest sufficient layer. 94% detection at 86% lower cost than pure agentic approaches.

AI Assistants

Customer service, internal help desk, document Q&A. Guardian monitors reliability. Steer enforces compliance language.

Document Processing

Context Engine extracts meaning. Orchestrate coordinates multi-step workflows with full audit trail.

Risk & Underwriting

Multi-agent reasoning with complete explainability. Reasoning chains for regulatory documentation.

Compliance Monitoring

Guardian tracks AI behavior continuously. AgentOps enforces policies. Complete audit trails for regulators.

Claims Automation

Trust Cascade for adjudication. Simple claims in milliseconds. Complex cases routed appropriately.

The Shadow AI Crisis

While you focused on your AI initiative, agents proliferated across your organization. IT deployed chatbots. Marketing uses AI for content. Operations built automation. No one knows how many AI agents exist, what they're doing, or what decisions they're making.

"We discovered 47 untracked AI agents across our organization. No one knew what models they were using or what decisions they were making."

— Chief Risk Officer, Global Enterprise

Shadow Agents

Every department has their own AI experiments. No central registry. No governance. No visibility into what AI is actually doing.

No Decision Visibility

Regulators ask "How do your AI systems make decisions?" - you can't answer. No audit trail exists for any decision.

Policy Violations

Agents violating compliance rules. Bots making promises systems can't keep. You find out after the damage is done.

Untracked Spending

Unexplained API charges. Teams spin up expensive models without approval. No cost attribution or budget control.

The Questions You Can't Answer

- ✓ What AI systems are making decisions in your organization?
- ✓ How do you test for bias or compliance violations?
- ✓ Can you provide audit trails for any AI decision?
- ✓ What governance framework controls AI at runtime?
- ✓ How do agents interact with each other?
- ✓ What's your total AI spend across the organization?

AGENTOPS

Enterprise Agent Governance Platform

AgentOps gives you complete visibility and control over every AI agent in your organization. One platform to govern them all.

The Transformation

Capability	Before AgentOps	After AgentOps
Agent Inventory	"We think we have... some?"	Complete registry with URN-based identity
Decision Audit	Manual log aggregation	Flight Recorder captures every decision
Compliance	Post-hoc reviews	Real-time policy enforcement
Agent Relationships	Unknown dependencies	Composability mapping shows all connections
Cost Control	Surprise invoices	Budget enforcement and attribution

Core Capabilities

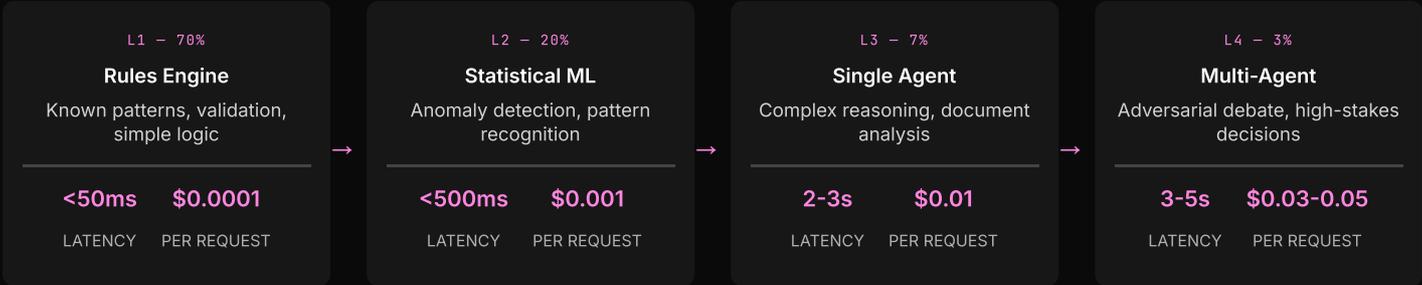
<p>Agent Registry</p> <p>Universal Resource Names for AI</p> <p>Every agent gets a unique identity: <code>urn:agent:fraud-detector:v2.1:prod</code>. Track across environments, versions, deployments.</p>	<p>Flight Recorder</p> <p>Complete Decision Capture</p> <p>Every agent interaction recorded: inputs, reasoning, outputs, tool calls. Compliance-ready audit trails. Litigation-ready documentation.</p>
<p>Policy Enforcement</p> <p>Governance at Runtime</p> <p>Three enforcement layers: gateway policies before requests, sidecar policies during execution, inline policies within agent code.</p>	<p>Composability Mapping</p> <p>See Agent Relationships</p> <p>Visualize how agents interact. Which agents call other agents? What data flows between them? Understand your AI ecosystem.</p>

See AgentOps in Action
Interactive demo of enterprise agent governance.

[View Demo](#)

Right-Sizing Intelligence

Not "use less AI" — use AI where it matters. The Trust Cascade routes each decision to the cheapest processing layer that can handle it correctly. Expensive reasoning only when necessary.



The Economics Difference

Metric	Pure Agentic AI	Trust Cascade	Improvement
Cost per 1M transactions	\$30,000 - \$50,000	\$2,300 - \$4,500	86-95% reduction
Average latency	3-5 seconds	200ms average	15-25x faster
Transactions needing AI	100%	~10%	90% handled cheaper
Decision accuracy	92%	94%	Better with less AI

"The question isn't 'How do we afford more AI?' It's 'How do we use AI only where it matters?'"

Your AI Journey

Whether you're starting fresh or rescuing a stalled initiative, we have engagement options designed for enterprise requirements.

AI READINESS ASSESSMENT	PLATFORM PILOT	ENTERPRISE DEPLOYMENT
\$15K	\$75K	Custom
1-2 weeks	6-8 weeks	3-6 months
<ul style="list-style-type: none">✓ Current state analysis✓ Agent inventory discovery✓ Risk assessment✓ Roadmap recommendations	<ul style="list-style-type: none">✓ Single use case implementation✓ Production deployment✓ Demonstrate economics✓ Prove governance model	<ul style="list-style-type: none">✓ Full platform deployment✓ Multi-use case implementation✓ Team training & enablement✓ Ongoing advisory support

Why Rotascale

- ✓ **Research-backed** — Every product built on peer-reviewed methods from Rotalabs. Published, verifiable research.
- ✓ **Production-tested** — Open source foundations running at scale. Battle-hardened, not experimental.
- ✓ **Your infrastructure** — Deploy on-premise, private cloud, or hybrid. Your data never leaves your environment.

Industries We Serve

Financial Services

Fraud detection, risk assessment, customer service, compliance monitoring, document processing.

Insurance

Claims automation, underwriting support, FWA detection, policyholder service.

Telecommunications

Network operations, customer experience, field service, revenue assurance.

Ready to deploy AI that actually works?

30-minute consultation to discuss your specific challenges and requirements.

Get
Started

View
Products

Contact

contact@rotascale.com · +1 (415) 524-0007
rotascale.com · San Francisco · Bengaluru

