CISCO Engage | Tech Day Miami

Optimizing Every Workplace with Agentic Ops

Forrest Burchell
Leader, Solutions Engineering
US Commercial & Latin America

February 19, 2026

I built an AI Application.

Mission Control

se-missioncontrol.com

An AI-powered platform built for our ThousandEyes SE teams. Used daily. In production. Real decisions. Real network calls.

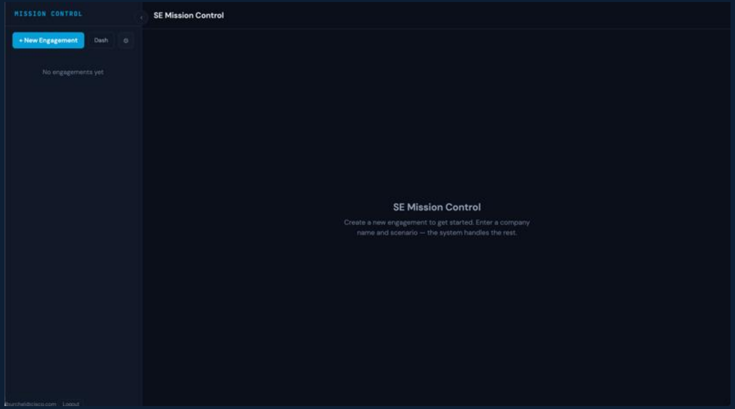

SE Mission Control

No engagements yet

SE Mission Control

Create a new engagement to get started. Enter a company name and scenario — the system handles the rest.

Automated Research

Transcript Analysis

Synthetic testing via TE API

Automated Runbooks / Playbooks

Work product deliverables

Then it broke.

Tuesday 2:14 PM

SILENT

AI gateway silent model downgrade

Anthropic API returns 529 (overloaded). LiteLLM triggers OpenAI failover automatically. Battle cards generated by gpt-4o-mini instead of Sonnet 4.5 -- shallower analysis, generic positioning, missed competitive nuances. SE walks into meeting with a C-grade card. No errors in logs, just a model field nobody checks.

Thursday 9:47 AM

SILENT

Guardrail Bypassed

Regional DNS for guardrail intermittently times out. AI Gateways pre-request guardrail hook fails open after 3s timeout. Prompts containing private data or adversarial content go directly to Anthropic unscreened. No alert fires because the request itself succeeds. Guardrail bypass rate climbs silently until someone audits the logs manually.

The following Monday

DISCOVERED

MCP tool unreachable — agent bypassed and kept going

ThousandEyes API server returned 503 during a POV setup. Agent skipped test creation step and told the SE it was complete. SE showed up to a customer meeting with zero tests deployed.

Under the Hood

One SE request to Mission Control. Seven network-dependent hops.

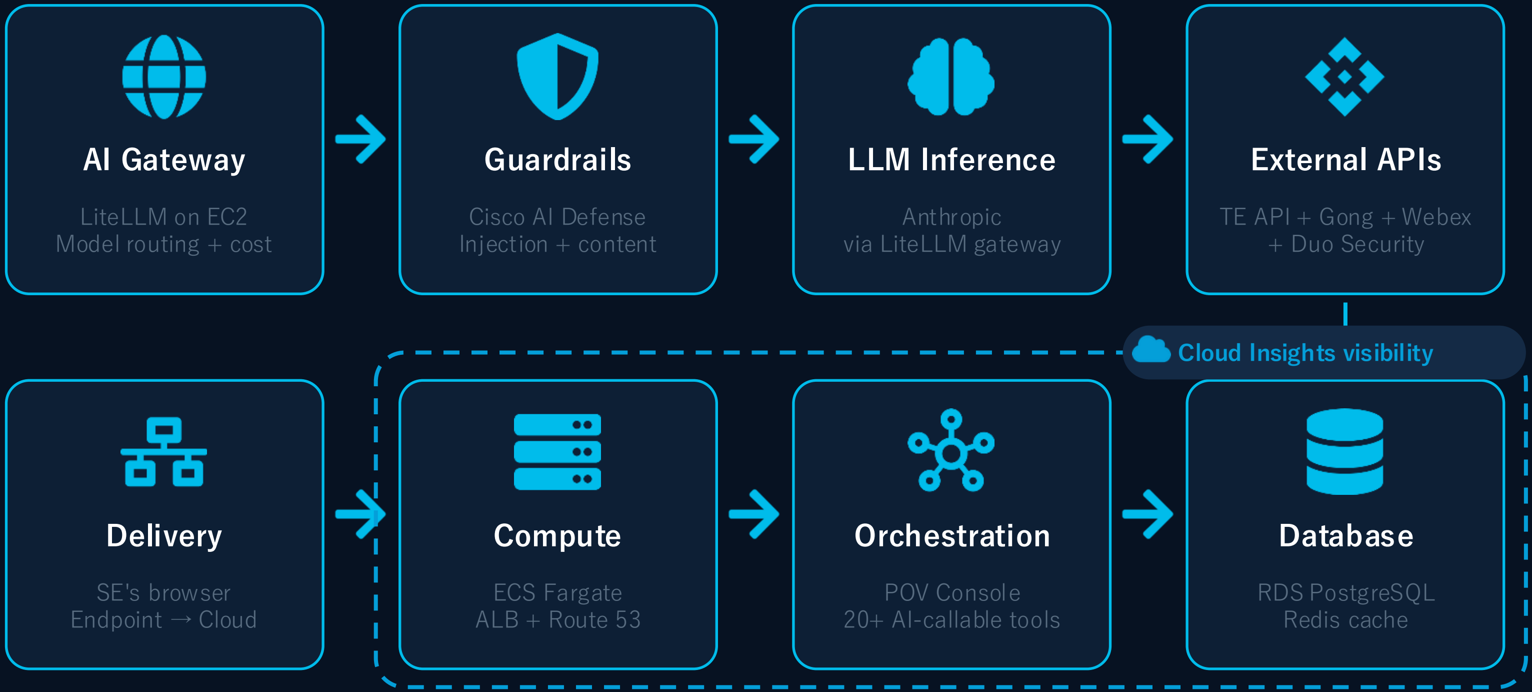

Every node is a network service. Every arrow is a network call. Every call can fail.

Mission Control has 12 AI endpoints.

How many do you have?

Do you know?

What silence costs

$67.4B

Losses attributed to AI hallucinations across enterprises in 2024. Unmonitored AI doesn't crash. It confidently gives wrong answers.

95%

of corporate AI initiatives fail despite $30-40B in investment. Only 5% of enterprise AI systems make it to production.

My agent gave wrong answers for days before I caught it. And I built it.

Your agents are doing it right now.

So I instrumented everything.

ThousandEyes tests running against Mission Control — right now

Foundation

- DNS: se-missioncontrol.com
- DNS: LiteLLM + LLM Providers
- HTTP: ALB + LiteLLM health
- BGP Reachability

Cloud Infra

- VPC flow logs: ECS ↔ RDS, ECS ↔ Redis traffic
- Config change correlation w/ synthetic test data
- Cloud topology: ALB → ECS → security groups → RDS
- Security group + route table change detection

AI-Specific

- Multi-step: prompt → LiteLLM → OpenAI → validate
- AI Test Template: baseline prompt + assertions
- API chain: Gong ingest → extract → store → analyze
- Guardrail latency: < 200ms threshold

Operational

- Internet Insights: OpenAI + Duo outage detection
- Endpoint agent: SE laptop → Mission Control path
- Alert → Webex notification
- Weekly AI dependency health summary

Demo

ThousandEyes AI Assurance

Your Turn

Day 1

Monday afternoon

DNS + HTTP tests for every AI endpoint. BGP monitoring for your top 3 AI provider prefixes. Network path visualization to your primary LLM.

Cloud Insights: connect your AWS/Azure account. Topology + flow logs start immediately.

30 minutes of setup. Zero excuses.

Day 30

Validation layer

Multi-step API tests mirroring your agent workflows. AI Test Templates with baseline prompts and assertions. MCP monitoring for tool integrity.

Cloud Insights: correlate config changes with synthetic test performance. Baseline your VPC traffic patterns.

Day 90

Operational maturity

Tune alerts from 60 days of baseline data. Integrate with incident management. Cloud Insights flow log anomaly detection tuned to your environment.

This is now normal ops.

One more thing.

ThousandEyes

- ✓ Hop-by-hop path to every AI endpoint
- ✓ DNS + BGP monitoring for AI providers
- ✓ Multi-step API chain validation
- ✓ Cloud Insights — topology + flow logs
- ✓ Internet Insights — collective intelligence

+

Cisco AI Defense

- ✓ Input guardrails — injection + jailbreak
- ✓ Output guardrails — PII + hallucination
- ✓ Action guardrails — tool permissions
- ✓ Continuous model evaluation
- ✓ Runtime policy enforcement

AI Defense is the seatbelt. ThousandEyes — from client to cloud — makes sure it's buckled.

The network is the AI runtime.

AI agents don't contain intelligence.
They call for it — across your network.
If you can't see every hop, you can't assure the outcome.

Monitor

Validate

Assure