Droma

FORMERLY GROKSTREAM + GROK (PRODUCT)

From Investigation to Delegation

Trusted, autonomous operations

Casey Kindiger

CEO & Founder, Droma

MORE TOOLS. MORE ALERTS. LESS CLARITY.

IT Operations Are Breaking

TODAY

- ✕ Alert storms bury the real signals
- ✕ Dedup and rules stop short of real compression
- ✕ Siloed tools and teams hide context
- ✕ LLMs alone can’t reason, predict or analyze deeply
- ✕ Faster routing ≠ incident prevention
- ✕ Manual effort doesn’t scale

WHAT MODERN IT OPS REQUIRES

- ✓ Predict issues before impact and human intervention
- ✓ Agentic AI that blends classical ML and generative AI to fit the problem
- ✓ Bring your own data
- ✓ One source of truth with real-time, integrated AI learning
- ✓ Closed-loop, human-in-the-loop automation

DELEGATION NEEDS A FUNDAMENTALLY DIFFERENT ARCHITECTURE

Why Most AIOps Can’t Get You to Agentic

- LLM copilots aren’t agentic AI
- Rule- and threshold-driven designs hit a ceiling
- Static automation can’t self-adjust
- No enterprise memory layer
- Prediction without action falls short
- Cross-domain blind spots
- Governance isn’t built in

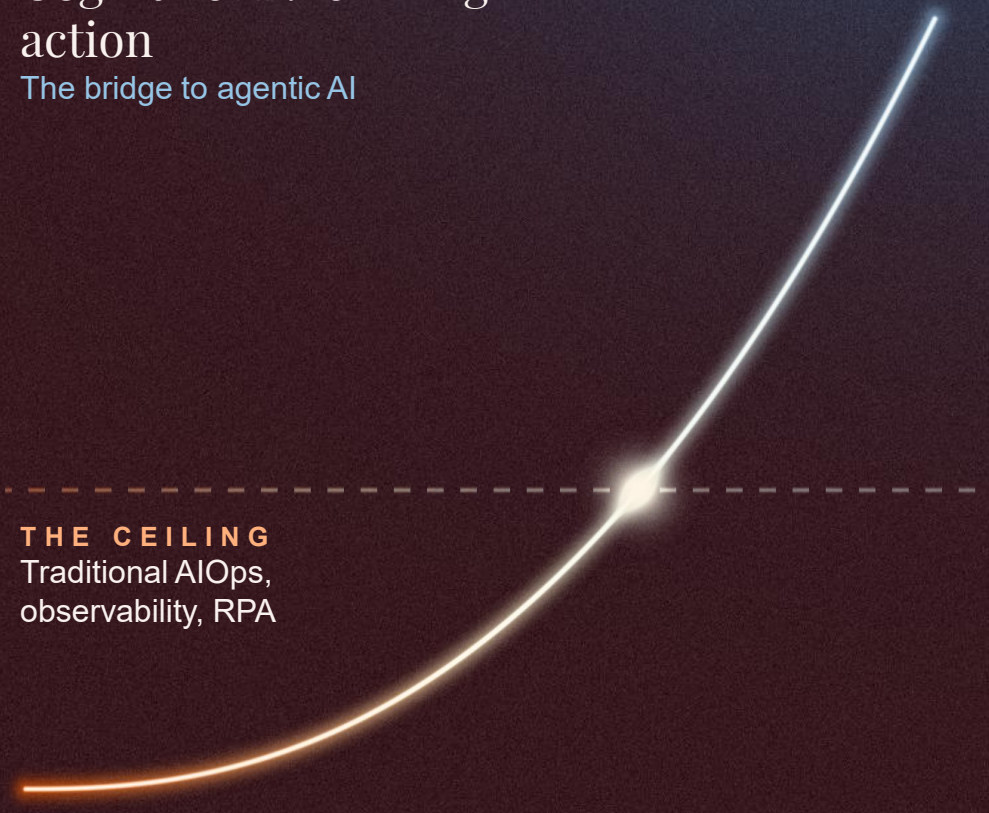

WHAT TRUSTED DELEGATION REQUIRES

The Architecture Behind Trusted Delegation

Persistent Memory

- Episodic memory and a memory graph link events, causes, actions and results
- Context carries across incidents and agents

Adaptive Learning

- 100% ML that adapts as environments change
- Improves from operator feedback and outcomes

Explainable Reasoning

- Causal root cause with traceable evidence
- Confidence behind every recommendation

Governance

- Guardrails, RBAC and tool grants per agent
- Human approval where risk or uncertainty is high

BUILT ON

Dynamic Data Fusion
One source of truth across tools

Predictive AI
Forecasts before impact

Causal AI
Explains why, not just what

Generative AI
Summaries, guidance, dialogue

ENTERPRISE AI ISN'T JUST PRODUCTIVE TECHNOLOGY – IT'S A COGNITIVE SHIFT

The Operational Metacortex

A persistent intelligence layer that learns, reasons, acts and improves from outcomes.

STAGE 1

Assistant

Answers questions, summarizes tickets, drafts updates

More productive — same operating model

STAGE 2

Investigator

Correlates signals against past incidents and proposes explanations

Less context-gathering; human still interprets

STAGE 3

Recommender

Proposes next actions, flags risk, suggests remediation

From manual diagnosis to supervised decisions

STAGE 4

Governed Actor

Executes approved actions within policy and change control

Low-risk automates; high-risk needs approval

STAGE 5

Organizational Metacortex

Persistent intelligence across perception, memory, reasoning, action and governance

The enterprise operates through AI-augmented cognition

THE FIRST ROLE-BASED AGENTIC OPERATIONS PLATFORM

Role-Based Agents That Work as a Team

ROLE-BASED AGENTS

- **L1 Operator**
- **L2 Operator**
- **SRE / Platform Owner**
- **Problem Manager**

SPECIALIZED AGENTS

- Detection Context
- Root Cause Analysis
- Knowledge Retrieval
- Next-Best Actions
- Automation Advisor
- Action Validation

SHARED COGNITIVE LAYER

Episodic memory · Memory graph · Continuous learning · Predictive, causal & generative AI

DYNAMIC DATA FUSION Events · Alarms · Changes · Incidents · Logs · Metrics

Agent identity

Every agent has its own identity, permissions and versioned reasoning

Open by design

MCP and open APIs let Droma consume context from other agents and share its intelligence back

AS AI SHIFTS FROM ASSISTING TO ACTING, SAFETY CAN'T BE BOLTED ON

Governed Autonomy, Engineered In

Prompt manipulation

System-level guardrails, intent filtering and instruction-hierarchy enforcement

Unsafe actions & tool calls

Role- and risk-based execution controls; human approval for high-impact tasks

Data leakage

Per-agent data boundaries, RBAC from the ground up, memory firewalls between agents

False confidence

Confidence gates, auto-rollback and human review for low-confidence cases

Memory corruption

Versioned writes, checkpoints and reflect-before-commit validation

Cascading failures

Dependency and impact modeling, simulation and global state awareness

GUARDRAILS AT THREE LEVELS

System · Agent · Action

TOOL GRANTS PER AGENT

Full | User-mediated | Blocked

NOT ONLY HUMAN-TO-AGENT: EARNED ACROSS EVERY BOUNDARY

Trust Is Organizational

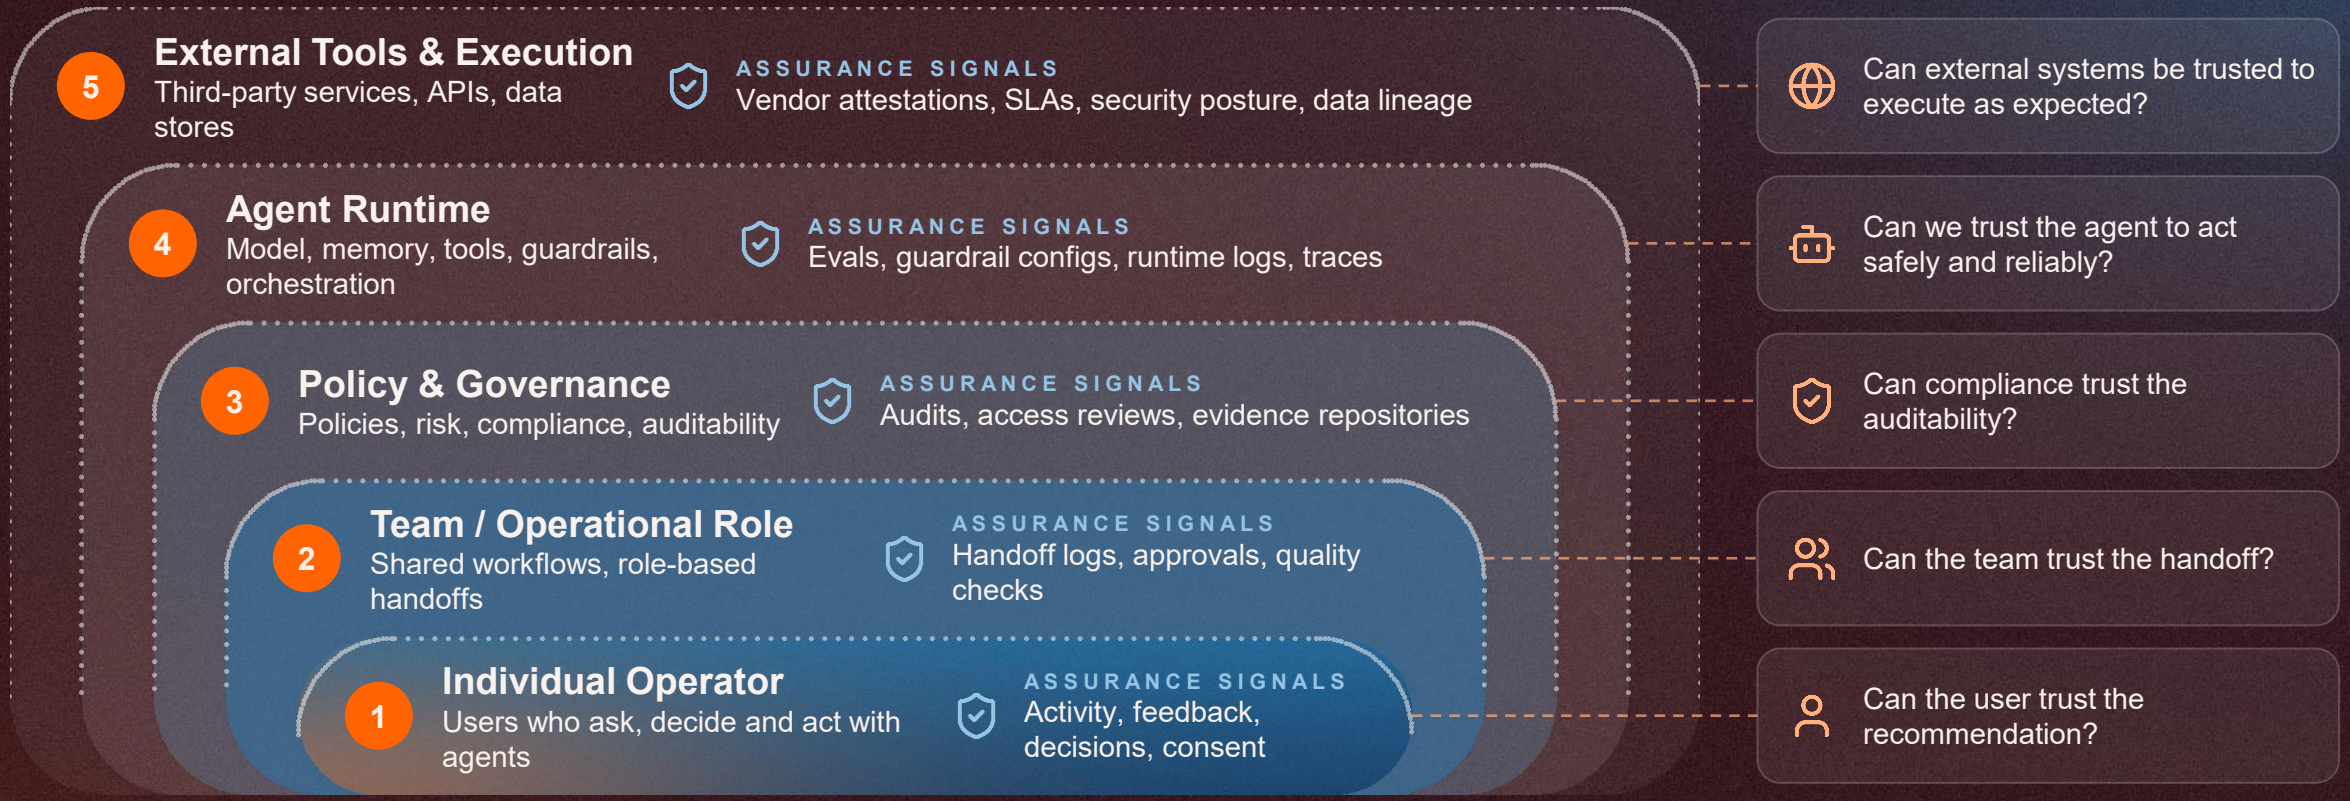

END-TO-END TRUST
Continuous and earned

Observe Collect signals › **Evaluate** Assess risk & controls › **Enforce** Apply guardrails › **Improve** Learn and adapt

BUILD TRUST, THEN EXPAND AUTONOMY

Delegate at Your Own Pace

DELEGATE NOW

Triage & investigation

Agents gather context, correlate signals and summarize. High ROI, low risk.

Humans approve

DELEGATE NEXT

Diagnosis & decisions

Agents validate root cause and choose next-best actions within guardrails.

Humans supervise

DELEGATE OVER TIME

Remediation

Agents execute and verify fixes on the way to self-healing operations.

Humans oversee and set strategy

INVESTIGATION — DELEGATION

START WHERE COGNITIVE TOIL EXISTS

L1 Operator

- Ticketing & routing
- Triage → remediation
- Root cause analysis

L2 Operator

- Cross-domain interpretation
- Hypothesis validation

SRE / Platform Owner

- Predict degradation
- Model feedback learning

Problem Manager

- Incident prediction
- Problem prevention

Droma

DELEGATE OPERATIONS. ELEVATE PEOPLE.

Droma.ai

VISIT DROMA EXHIBITION TABLE FOR DEMO

See Agents Work as a Team

Multi-agent orchestration

Cross-domain outcomes

Progressive autonomy