



ELIMINATE L1 TRIAGE WITH AGENTIC INCIDENT RESPONSE: **DROMA L1 AGENT**

AI THAT REASONS, DECIDES, AND ACTS

Every wave of AIOps investment promised to fix L1 operations. Rules compressed noise. Copilots summarized it. Automation executed what it was told. None had the whole picture. Signals stayed fragmented. Data was never trusted enough to act on alone. Lessons from the last incident rarely made it to the next. So humans kept doing the stitching, gathering context, validating root cause, deciding what came next, under pressure, every time. Droma (formerly Grok) changes that. As part of the broader Role-Based Agentic Operations Platform, powered by a cognitive AI engine that continuously learns before, during, and after every incident, Droma's L1 Agent orchestrates specialized operational agents to gather evidence, diagnose likely causes, retrieve relevant knowledge, and act, enabling governed autonomous operations across IT Operations and ITSM.

Overview

Observability solved visibility. AIOps solved noise. Neither closed the decision gap — the moment between detecting a problem and knowing what to do about it. That gap still belongs to the operator, every time, under pressure.

Droma L1 Agent and the broader Role-Based Agentic Operations Platform closes this gap, enabling organizations to progressively adopt governed autonomous operations. It evaluates every detection, investigates causally, and acts — gathering evidence, diagnosing likely causes, and retrieving relevant knowledge without waiting for a human to validate what the system already knows. Validated, previously resolved issues resolve autonomously. Everything else escalates with full context and a recommended next step already attached.

This isn't prediction bolted onto reasoning, or reasoning bolted onto automation. Classical ML and generative AI operate as one capability — investigation, diagnosis, and action under a single orchestration layer. Specialized agents handle root cause analysis, prioritization, knowledge retrieval, and remediation guidance in concert, making the L1 Agent the operational foundation for a broader, progressively autonomous system.

Interested in learning more? Contact us at info@droma.ai

HOW THE DROMA L1 AGENT WORKS


Most automation fails not because it executes poorly, but because it was never built to reason. Rules depend on clean data and predictable conditions — neither exists consistently in a live enterprise environment. The result is a hard ceiling: automation handles the known and repeatable. Everything dynamic, contextual, or previously unseen still falls back to the operator.

Droma L1 Agent is built for for what automation cannot anticipate..

Dynamic Data Fusion combines preparation and intelligence at the source, turning raw, fragmented telemetry into a clean, enriched, model-ready feed. Guru, Droma's generative AI, turns that feed into understanding: what's happening, what caused it, and what resolved it before. The role-based agent draws on both to determine the next best action — no rule, no operator required for anything it already knows how to resolve.

When an incident fires, the sequence runs without intervention:

- **Ingest** — The agent pulls enriched context from Dynamic Data Fusion layer: historical patterns, recent changes, service dependencies, and cross-domain signals already correlated.
- **Reason** — Guru analyzes the enriched incident, identifies probable root cause, and determines whether the pattern is validated, novel, or ambiguous, grounded in your SOPs, runbooks, KEDB articles, and escalation procedures.
- **Act** — Recommends or autonomously executes governed actions, updates workflows, generates scripts, and orchestrates specialized operational agents and enterprise workflows in accordance with organizational policies and risk thresholds. Validated, previously resolved issues resolve autonomously. Novel or ambiguous ones route to the right team with causal context and recommended next steps already attached.
- **Learn** — Every outcome, automated or escalated, feeds back into the model — improving accuracy and expanding autonomous coverage over time.



The screenshot displays the Droma L1 Agent interface, which is divided into several sections:

- Guru Chat:** A sidebar on the left containing navigation options like Detection Context, Root Cause Analysis, Automation Advisor, Knowledge Retrieval, Guided Triage, Conversations, and a search icon.
- L1 Agent Session:** The main workspace showing a detailed investigation of a network issue. It includes a "Guided Investigation Workflows" section with a blue callout box stating: "Work through common investigation steps — or jump anywhere. The agent supports both structured and exploratory workflows." Below this, the "Agent Working Memory" section (also with a blue callout box) states: "Access the agent's working context — so you can validate its reasoning and build trust." The "Reasoned Actions" section (with a blue callout box) states: "These actions are generated dynamically based on context — helping you move the investigation forward without guesswork." The session text describes a recurring pattern of degrading QSPF28 optic and contaminated MPO jumper on the atl-1f07 device's Eth1/49 port, leading to intermittent bit errors and performance issues. It lists recommended actions such as confirming the optic and patch lead, scheduling replacement, monitoring device and network for error rates, adding work notes, and associating the detection with a relevant incident.
- Detection Info:** A panel on the right providing metadata for the current detection, including Time (10 August 2026 at 3:36:20 PM IST), Events (16), Duration (19m 6s), Level (Cat2), Hostname (atl-1f07, atl-sp02, atl-esx11, atl-dbo3, atl-app17, atl-fw02), and Detection ID (6ab62ec9-e090-41e5-a3db-1433-e76a2f2d). It also includes expandable sections for Related Detections, Related Incidents, Change History, Configuration Items, Incident History, and Agent Internals.
- SUGGESTED:** A bottom bar with buttons for "Add a work note", "Check optic status", "List recent activities", "Next troubleshooting steps", and "Summarize session for ticket".

WHAT MAKES DROMA L1 AGENT DIFFERENT

- **Eliminate L1 Triage as a Primary Function.** Droma L1 Agent shifts incident response from manual triage to AI-led execution — investigating detections, resolving known issues, and escalating only when needed. L1 teams move from triage to exception handling.
- **Cognitive Learning Foundation.** Classical ML drives signal compression, classification, and prediction. Generative AI drives reasoning, explanation, and interaction. Both run on the same enriched signal Dynamic Data Fusion prepares, and both get sharper with every incident, learning before, during, and after detection.
- **Specialized Agents, Orchestrated for Depth.** The L1 Agent orchestrates specialized agents for root cause analysis, prioritization, automation selection, and knowledge retrieval — each contributing focused expertise to the incident lifecycle. Depth without slowing down execution.
- **Governed Autonomy, Built In.** Droma governs agent behavior across three levels. System level sets whether human approval is required. Agent level defines what agents can do and which systems they reach. Action level assigns risk, approval rules, and eligible personas to every action, including operational limits like rate limits and timeouts.
- **Foundation for Governed Autonomous Operations.** Rather than positioning autonomy as a future state, Droma delivers it incrementally. Detection, reasoning, action, and learning run as a closed loop — every incident improves the next, reducing repeat issues and advancing toward self-healing IT operations.



GOVERNED AUTONOMY BY DESIGN

Droma is built to anticipate the specific risks that come with agentic reasoning, memory, and action — not patch them after the fact.

- **Cross-Domain Cascading Failures** — Federated oversight and simulation, dependency and impact modeling, and global state awareness prevent isolated actions from triggering unintended downstream effects.
- **Memory & Learning Corruption** — Versioned writes and checkpoints, reflect-before-commit validation, and belief graph validation protect the integrity of what the agent learns over time.
- **Hallucination & False Confidence** — Confidence gates, auto-rollback and fallback plans, and human-in-the-loop review for low-confidence cases ensure the agent never acts on certainty it hasn't earned.
- **Data Leakage & Boundary Violations** — Per-agent data boundaries, role-based access control from the ground up, and memory firewalls between agents keep context properly isolated.
- **Unsafe Actions & Tool Invocations** — Role- and risk-based execution controls and human approval for high-impact tasks govern what the agent can act on and when. Agent readiness state validation ensures correct configuration and grants are in place before any agent is permitted to act.
- **Prompt Manipulation** — A governance reasoner invoked on every turn assesses unsafe elements in the pre-generation prompt, failing closed or sanitizing before injection. System-level guardrails, intent filtering, and instruction hierarchy enforcement protect the integrity of the agent's reasoning.
- **Agent Identity & Role Management** — Defines each operational agent's identity, responsibilities, permissions, governance policies, and memory boundaries.



HOW DROMA WORKS

Droma operates as a cognitive AI layer across your environment, ingesting signals, learning continuously, and orchestrating agent-driven actions across IT Operations and ITSM.

- **Ingest & Unify Signals** - Droma ingests events, alarms, logs, metrics, incidents, and changes through Dynamic Data Fusion, creating a real-time, unified operational context across systems.
- **Learn & Reason Continuously** - Using predictive, causal, and generative AI, Droma learns before, during, and after incidents, improving detection accuracy, understanding relationships, and building a continuously evolving model of your environment.
- **Detect, Predict, and Decide** - Droma compresses noise into actionable detections, classifies issues, and identifies risks hours in advance, delivering context-rich insights, root cause analysis, and recommended actions.
- **Orchestrate Agent-Driven Workflows** - Role-based agents (L1, L2, SRE, and Problem Manager) leverage specialized agent services to triage, investigate, and resolve issues, coordinating actions across tools and workflows.
- **Execute with Governed Autonomy** - Actions are executed through automation and orchestration, with human-in-the-loop approvals, policy guardrails, and full explainability, progressing toward self-healing IT operations.

Predictive & Agentic AI for IT Operations



Role-based Agents
and Workflow

L1 Ops

L2 Specialist

Platform/SRE

Problem Mgr



Specialized Agent
Services

Knowledge
Retrieval
Root Cause
Analysis

Full Context
Summaries
Troubleshooting
Orchestration

Intelligent
Recommendations
Analytics

Automation Advisor
Action Validation

Agent
Interoperability
Layer (MCP /
APIs)



Droma

Cognitive AI Intelligence Powered by
Predictive, Causal & Generative AI

Learns Before,
During and After
Incidents

Noise
Reduction

Classification

Incident
Prediction

Automation

Continuous Cognitive Learning

Dynamic Data Fusion

Unified Multi-
Domain Context

Events

Alarms

Changes

Incidents

Logs

Metrics

Observability, EMS, ITSM, CMDB



Governed Autonomy

RBAC / Security - Policy Guardrails -
Explainability & Audit - Human-in-the-
Loop Approvals & Enrichment

ABOUT DROMA

Droma delivers predictive, agentic, and self-healing IT operations through a cognitive AI platform that continuously learns from signals, incidents, actions, and outcomes. Built for rapid deployment and seamless integration, Droma enables organizations to move from reactive operations to autonomous, outcome-driven performance. We partner with high-growth and large enterprises across oil and gas, data center, fintech, and other industries, as well as CSPs and MSPs, supporting mission-critical, real-time environments where reliability, scale, and customer trust are non-negotiable.

Contact us at info@Droma.ai to learn how to modernize your IT Operations.