




HIGHLIGHTS

- Role-Based AI Agents and Workflows
- 100% AI Learning Architecture — No Rules, No Topology, No CMDB
- 3x Greater Event and Incident Compression than Other AIOps and Observability Solutions
- Actionable Predictions up to 48 Hours in Advance
- Accelerated Root Cause Analysis with Causal Intelligence
- Governed Autonomy — Three-Tier Control for Every Automated Action
- Vendor-Neutral Overlay — Deploys Over Any Existing Stack
- Intelligent Incident Response with 92%+ MTTR Reduction

FROM REACTIVE TO AUTONOMOUS: AI-LED OPERATIONS, POWERED BY DROMA

IT operations teams have spent years layering tools on top of tools — monitoring platforms, observability stacks, ITSM systems — and the result is more signal, not more control. Alert volumes grow with every new service. Resolution still depends on the right engineer being available at the right moment. And the automation that was supposed to help has become its own maintenance burden.


AI-led operations is not another layer. It is a structural shift — from teams that react to events to systems that detect, reason, and resolve before impact reaches the business. Droma (formerly Grok) is purpose-built for that shift, with a cognitive AI learning architecture that learns how environments behave, predicts what is about to go wrong, and acts through governed autonomy that keeps humans in control at every stage.

Interested in Learning More? Contact info@Droma.ai

UNMATCHED NOISE REDUCTION

Employing a combination of classical machine learning and generative AI, Droma analyzes vast amounts of telemetry to identify patterns, compress noise, predict issues, and support AI-led investigation. Unlike traditional AIOps platforms that rely heavily on static rules or topology, Droma’s AI learning architecture continuously adapts as environments evolve. There are two key ways that Droma reduces the volume of tickets:

- **Incident Prevention** – By analyzing historical data such as tickets, changes, and human actions, Droma’s AI algorithms can predict potential failures or vulnerabilities. These predictions are validated, prioritized, probability-scored, and used to initiate AI-led investigations before service impact occurs. Furthermore, Droma uniquely processes anomalies as early warning signals, continuously evaluating them to reduce false positives and improve operational trust.
- **Detections** – Droma goes beyond simple de-duplication to identify related signals, group them into meaningful detections, and surface the issues most likely to require action. By combining associative clustering, reinforced classification, and causal intelligence, Droma transforms fragmented operational data into actionable understanding for investigation and resolution.




“NTT is one of the leading managed services providers driving innovation through the latest advancement in AI technology to deliver the huge reliability and service levels to their customers. Our relationship with Droma is part of an ongoing set of investments that differentiate NTT’s service and provide networks that deliver for our clients’ businesses.”

– Amit Dhingra, Executive Vice President, Network Services, NTT Ltd.

COGNITIVE AI LEARNING ARCHITECTURE

Droma's learning architecture mirrors the cognitive layers of the human brain, excelling in causal inference and decision-making. By employing a combination of causal, predictive, and generative AI, Droma unifies signals across generation assets, transmission and distribution networks, IT infrastructure, compute, and application layers into a single operational context – bridging physical and digital infrastructure in a way that siloed tools structurally cannot. This allows Droma to associate raw data with underlying causes across domains, delivering unparalleled insight into how infrastructure actually behaves.

It analyzes vast telemetry data to identify patterns in related alerts and uses unsupervised learning to group these alerts into meaningful detections. Through reinforced (supervised) learning, Droma prioritizes alerts and adds critical context. These continuously learned insights become the foundation for specialized agent services and role-based agents, enabling agent-driven operations built on a shared operational understanding. Specialized agent services then provide intelligent recommendations, root cause analysis, contextual summaries, and guided decision support, enabling rapid and informed decision-making.



“Droma has fundamentally changed how we run IT operations. Droma is now closing more tickets than our operations team, creating the foundation for us to scale with confidence.”

— Head of Service Assurance, Zayo.

DROMA L1 AGENT


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, and Guided Triage. Below this is a list of conversations.
- 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's reasoning is shown, identifying a root cause related to a degrading QSFP28 optic and a contaminated MPO jumper. A "Reasoned Actions" section at the bottom provides dynamically generated steps like "Add a work note", "Check optic status", and "List recent activities".
- Agent Working Memory:** A blue callout box in the center states: "Access the agent's working context — so you can validate its reasoning and build trust." This refers to the detailed diagnostic information visible in the background, such as error rates and performance issues.
- 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), and Hostnames (atl-if07, atl-sp02, atl-esx11, atl-app17, atl-db03, atl-fw02).

DROMA PROBLEM MANAGER AGENT


Problem management breaks down when every recurrence becomes a new investigation. The evidence exists, but it is fragmented across incidents, detections, changes, dependencies, and past remediation. Teams spend valuable time reconstructing what has already been learned instead of using that knowledge to prevent the next failure.

Droma Problem Manager Agent is built to break that cycle.

Dynamic Data Fusion connects operational evidence across domains, bringing recurring detections, incident history, changes, service dependencies, and prior interventions into a unified context. Guru, Droma's generative AI, turns that context into operational understanding — what is recurring, why it is happening, how far it reaches, and what has already been tried. The Problem Manager agent orchestrates specialized operational agents to determine the highest-value path forward, from repair guidance to escalation and ownership.

When a recurring pattern emerges, the sequence runs continuously:

- **Detect** — The agent identifies patterns across detections and incidents, correlating frequency, similarity, severity, and history to distinguish systemic problems from isolated events.
- **Investigate** — Specialized agents analyze the Problem across its history, identify the underlying root cause, and determine why it persists, its impact, and its blast radius, grounded in correlated incidents, detections, changes, service dependencies, operational knowledge, and prior remediation outcomes.
- **Recommend** — The agent translates that understanding into the next best course of action, evaluating available remediation paths by relevance, expected impact, and operational fit — from suppression and soak periods to notification or ticketing changes, automation, or escalation.
- **Learn** — Every intervention and outcome becomes operational memory, strengthening pattern recognition, sharpening recommendations, and reducing the effort required to understand and resolve future recurrences.



Guided Investigation Workflows
Work through common investigation steps — or jump anywhere. The agent supports both structured and exploratory workflows.

Agent Working Memory
Access the agent's working context — so you can validate its reasoning and build trust.

Reasoned Actions
These actions are generated dynamically based on context — helping you move the investigation forward without guesswork.

Problem Overview

- Root Cause Analysis
- Impact & Blast Radius
- Repair Advisor
- Escalation & Ownership

Conversations 41

Problem Agent Session

Summarize this Problem and how I can help.

The root cause has been identified as a poison message in the settlement input queue combined with a long-running blocking transaction on the mktstl schema. This combination leads to a deadlock in the worker pool, which cannot self-recover, causing prolonged application alarms.

This Problem has resulted in 14 linked incidents, indicating it is a significant and persistent issue. It has never been auto-cleared historically and has been resolved previously.

How you can help:

- Understand that the core issue is the deadlock caused by a long-running blocking transaction on the mktstl schema.
- Coordinate with application and database teams to identify and terminate the blocking transaction.
- Monitor the affected hosts and services for early detection of this condition.
- Consider setting up automation or notifications to respond quickly to future occurrences.
- Investigate related incidents for additional context or patterns.
- Drive the problem toward a durable fix to prevent recurrence.

If you want, I can also help you explore related problems, analyze impact more deeply, or assist in setting up notifications or automations for this Problem.

To fix the blocking transaction, the following steps are recommended:

- Identify the Blocking Transaction:
 - Use database monitoring tools or queries to identify the exact transaction(s) on the mktstl schema that are long-running.

SUGGESTED

- Help me draft an action plan
- Explain poison message handling
- How to monitor blocking transactions?
- Suggest automation options
- Summarize session for ticket

Ask about this problem... (/ for work intents)

Problem Info

Open recognized detection

Prolonged Application Alarm Persistence

Metric	Value
OCCURRENCES	16
INCIDENTS	14
AUTO CLEARED	0

ROOT CAUSE Root cause was a poison message in the settlement input queue combined with a long-running blocking transaction on the mktstl schema; together they deadlocked the worker pool, which could not self-recover.

PROBLEM ID 3367268d-982c-424a-9056-5e45db1da045

Related Detections

Related Incidents


Related Problems

Problem Impact


Agent Internals

INCIDENT PREDICTION AND PREVENTION


Proactive Problem Identification - Droma's proactive problem identification uncovers and permanently resolves recurring issues, not just their symptoms. By consolidating alarms into meaningful groups, applying AI models trained on existing ticketing data, and learning from operational feedback, Droma empowers platform owners and problem management teams to eliminate recurring tickets, automate fixes, and drive continuous improvement across power and utility operations.



Major Incident Response - Droma's advanced anomaly detection swiftly identifies previously unobserved major incidents — such as transmission network faults, substation communication failures, or IT infrastructure outages — by analyzing event patterns, logs, and configuration changes in real time. Operations teams receive actionable intelligence the moment an incident emerges, significantly reducing time to resolution and enabling confident, streamlined response across power and utility environments.




Incident Prediction and Automation - Droma enables operations teams to predict and prevent incidents up to 48 hours in advance. By dynamically linking recommended automations to detections and prioritizing remediations by severity and frequency, Droma moves beyond single alert triggers, applying pattern-based automation across IT, network, generation, and distribution domains for smarter, more efficient remediation workflows.



AGENT INTERACTION CHANNELS

Access predictive, causal, and generative intelligence directly through Microsoft Teams and Slack — without leaving the workflows your operators already rely on.

Retrieve contextual knowledge, surface causal root cause summaries, trigger automations, and act on AI-prioritized triage guidance through two-way agentic interaction — with human-in-the-loop controls built in.



The left screenshot shows the GrokGuru interface. It displays a 'Platform error messages have been logged on some devices.' and a 'Past Incident Handling' section. A user query 'are there any past incidents related to this issue?' is shown, along with a response from GrokGuru. The response includes a summary of the issue and a list of sources.

The right screenshot shows a Slack channel named '# grok-guru'. It displays a message from GrokGuru that provides a detailed analysis of the issue, including a 'Yesterday' section, an 'Overall Assessment', and a 'Conclusion'. The message concludes with a recommendation to verify the data collection.

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.