

## **COGNITIVE AGENTS IN DEVOPS: TRANSFORMING OBSERVABILITY, INCIDENT RESPONSE, AND POLICY-AWARE AUTOMATION**

Vimal Daga  
CTO, LW India |  
Founder, #13 Informatics  
Pvt Ltd  
LINUX WORLD PVT.  
LTD.

Preeti Daga  
CSO, LW India |  
Founder, LWJazbaa Pvt  
Ltd  
LINUX WORLD PVT.  
LTD.

Satvik Dubey  
Research Scholar  
LINUX WORLD PVT.  
LTD.

**Abstract-** This work introduces the Cognitive need to make smart decisions. In parallel, DevOps Assistant (CDA)—an intelligent AI DevSecOps practices ensure that every step—system designed to help SRE and DevSecOps from code to production—remains secure and teams manage the growing complexity of compliant. Cognitive agents can support this by modern infrastructure. Powered by advanced automatically catching vulnerabilities and frameworks like ReAct and AutoGen, the CDA enforcing security rules. Together, this blend of doesn't just detect problems—it understands AI, observability, and DevSecOps creates a them. It analyzes logs, metrics, traces, and powerful ecosystem where systems not only stay changes across systems, recalls past incidents, reliable but get smarter over time. The result? reasons through what's happening, and takes More uptime, less stress for engineers, and a safe, policy-compliant actions like rollbacks or future where AI is a trusted partner in keeping scaling, with human oversight built in. It's more everything running smoothly. than automation—it's intelligent collaboration. **Keywords:** Cognitive agent, AIOps, Drawing on insights from over 30 recent studies, Observability, Site Reliability Engineering, the CDA shows clear benefits: faster incident Agent AI, DevSecOps resolution, less noise, and reduced workload for engineers. It even uses causal AI to go beyond surface-level symptoms and get to the real root of a problem. This all happens within a broader shift toward Cognitive Agents in AIOps—AI that can think, learn, and act in real time to keep systems healthy and secure. These agents thrive on observability—rich data from logs, metrics, and traces—that gives them the context they

### **I. INTRODUCTION**

Modern Site Reliability Engineering (SRE) teams are grappling with unprecedented complexity. Today's systems are composed of thousands of microservices, deployed continuously across distributed, hybrid cloud environments. Engineers must maintain service reliability while navigating real-

time telemetry, dynamic system states, and ever-evolving deployment pipelines. As a result, incidents are no longer isolated events—they're often emergent phenomena that span infrastructure, application layers, and third-party dependencies. In such an environment, incident response must be faster, more intelligent, and more adaptable than ever before. Yet, SREs often find themselves buried under noisy alerts, disconnected tooling, and a flood of observability data—from logs and metrics to traces, change diffs, and compliance events. Extracting actionable insight from this heterogeneous data in real time is no longer a task suited to human cognition alone. Even experienced engineers can miss weak signals, misinterpret correlated symptoms, or delay remediation due to mental fatigue or information overload. While AIOps has helped automate some routine operations—such as anomaly detection or alert suppression—it largely relies on heuristic or pattern-based models. These systems lack the ability to reason about context, recall past events meaningfully, or take nuanced action in unfamiliar or evolving scenarios. In short, they are not “cognitive.” This is where Cognitive Agentic AI—powered by recent advances in Large Language Models (LLMs) and tool-using agents (e.g., ReAct, AutoGen)—can offer transformative

potential. These agentic frameworks enable an LLM not just to process text or telemetry, but to think through a situation, act upon it using system APIs or observability tools, and remember past decisions and their outcomes. With properly designed memory, reflection, and policy constraints, such agents can participate in high-stakes operational workflows—diagnosing issues, proposing remediations, querying services, or initiating safe actions—while staying accountable, interpretable, and under human supervision. This paper introduces the Cognitive DevOps Assistant (CDA): a system that embodies these principles to assist modern SRE teams in real-world cognitive scenarios. The CDA is designed to analyze telemetry data, detect anomalies, reason about system state, interact with tooling, enforce DevSecOps policies, and coordinate with human operators when ambiguity or risk demands caution. Far from replacing engineers, the CDA serves as a trusted copilot—automating toil, accelerating incident response, and preserving human attention for judgment, creativity, and strategic reliability improvements. We evaluate the CDA across multiple axes—efficiency, safety, trust, and integration effort—and show how such systems can reshape the way reliability engineering is practiced. Our findings suggest that when cognitive

agents are grounded in observability, policies, and human-in-the-loop design, they can not only boost operational performance but also foster trust, transparency, and long-term learning across complex DevOps environments.

## II. LITERATURE REVIEW

### 1. AIOps and Machine Learning for Observability

The application of AIOps in modern observability tools has come a long way in the recent past, primarily through applying machine learning (ML) techniques to make operational processes more effective and accurate. Traditional observability tools such as Prometheus, ELK stack, Datadog, and New Relic generate massive amounts of telemetry, and it is not possible for humans to filter, correlate, and diagnose all potential signals in real time. Initial AIOps solutions employed heuristic filtering and supervised learning techniques—such as clustering for noise filtering, decision trees for classifying alerts, and linear models for anomaly detection. Useful with structured data and cyclic failure patterns, these techniques fail to generalize to new or surprise instances since they lack understanding of context and are rigidly reliant on training. Recent advances in transformer models and few-shot learning with Large

Language Models (LLMs) have made new opportunities available for managing unstructured observability data such as natural language incident reports, config diffs, and support tickets. For instance, transformer models now contextualize multi-modal inputs (for instance, logs + traces + deployment records) and produce human-readable explanations or root-cause hypotheses. Exciting as they are, such models are presently ungrounded in memory, non-deterministic in control, and lacking safety constraints necessary for robust deployment in production environments.

### Cognitive Architectures and Agentic AI Frameworks

To surpass shallow automation, researchers have developed cognitive agent models that can implement reasoning, memory, and tool interaction in real-world settings. Architectures such as ReAct (Reasoning + Acting) and AutoGen equip LLM-based agents with the ability to switch between natural language thought patterns and API calls and command-line actions, thus enabling multi-step decision-making and recursive problem-solving. Frameworks such as LangChain and LangGraph offer systematic abstractions for modular choreography of tools, multi-agent communication, and stateful memory. These frameworks allow agents to chain across multiple

tools (e.g., querying a metrics database, checking alternatives based on risk posture or SLA impact. against a compliance API, calling a rollback), Escalation to human approval is also possible pass context across multiple interaction steps, where there is uncertainty or high-impact action, and produce explainable, step-by-step reasoning thereby offering a human-in-the-loop model of traces. It is crucial in remediation based on governance.

observability where historical context, prior Agent Autonomy, Memory, and Reasoning in SR actions, and tool results are utilized to decide on SRE autonomous agents are rapidly evolving future decision branches. Furthermore, these from being passive interpreters of log data to frameworks are more and more bringing active copilots of operational tasks. Novel multimodal interfaces (e.g., dashboards, time-series views, charts) and supporting agents that log analysis, metric regression, alert can reason over parallel structured + unstructured summarization, and change-impact prediction. data, a feature of high utility for SRE workflows The copilots assist engineers in detecting failure dealing with heterogeneous data modalities. patterns, predicting service degradation, and Policy-Based DevSecOps Automation rollback strategy recommendations. However, DevSecOps now integrates security and there are still key challenges. Most current agents compliance policies across every phase of the lack persistent memory or state tracking and software lifecycle, from build-time to runtime. therefore cannot be employed to understand Policy-as-code engines such as Open Policy incident timelines, cross-incident correlations, or Agent (OPA), SPIFFE/SPIRE, Kyverno, and previous remediation knowledge. In addition, it KevOps allow organizations to formalize and is difficult to impose safe action boundaries on enforce authentication, network access, LLM-based agents in the lack of strong configuration compliance, and remediation alignment mechanisms and policy conditioning. automation rules via code. These policies set In addition, while cognitive agents can reason runtime guardrails, preventing unauthorized about action sequences, they lack the capacity to changes, imposition of security perimeters, and hold context during long-term interaction, safe fallback automation (e.g., traffic throttling, differentiate between unsafe and safe commands, container isolation, certificate rotation). and escalate in a proper way in the event of low Combining cognitive agents with such policy confidence. This necessitates systematic engines offers a path to auditable, explainable, interventions such as episodic memory systems, and compliant automation. An agent powered by confidence thresholds, and intent verification an LLM can consume, understand, and enforce protocols in order to make the autonomous policy constraints as part of real-time decision- decisions in real-world contexts efficient as well making—e.g., deciding between remediation as reliable.

### III. METHODOLOGY

#### System Architecture Overview

The Cognitive DevOps Assistant (CDA) is a Human-in-the-Loop Control Plane: Dynamic UI modular agentic system designed to augment and alerting offer insight into the reasoning and modern SRE and DevSecOps pipelines with decision-making of the agent. Dangerous actions, explainable, secure, and context-aware ambiguous conditions, or policy rejection automation. The foundation integrates real-time automatically escalate to a human SRE or observability data, LLM-based reasoning, policy- security engineer for override, review, or based decision gates, and human-in-the-loop additional context injection

interface. The primary components are:

policies before running, hence automated decisions are kept within safe and governed boundaries.

Ingestion Layer: Talks to modern observability systems (e.g., OpenTelemetry, Prometheus, ELK, Grafana Loki, Jaeger) to ingest structured logs form a ground truth audit trail to enable and unstructured telemetry data like logs, post-incident analysis, compliance audits, and metrics, traces, alerts, diffs, and deployment continuous model improvement.

events. Normalizes the data into a semantic Reasoning Pipeline: The decision-making format that can be consumed by agent pipeline of the agent has an iterative action and processing.

Reasoning Engine: An agentic core LLM in based on the ReAct/AutoGen framework and CDA, based on architectures such as ReAct contains the following steps:

(Reasoning + Acting), AutoGen, or LangGraph. Perception: The agent processes observability They aid the agent in incremental reasoning, signals (e.g., CPU/memory spikes, unusual calling on tools (e.g., metric query APIs, CI/CD latencies, 500-level errors, unusual config systems, config diffs), accessing memory, and changes). A semantic normalizer translates planning action sequences. The engine is varied input data to a form understandable by an asynchronous to enable concurrent workflows in agent through light-weight tagging (sensitivity parallel with state carried over time log severity, anomaly scores).

DevSecOps Interface: Policy-driven automation Reflection: The agent queries its episodic is facilitated by policy engine integration like memory for comparable past experiences, Open Policy Agent (OPA), KevOps, or corresponding repairs, and outcomes. It uses SPIFFE/SPIRE. All CDA decisions (e.g., vector similarity or retrieval-augmented rollback, scaling, service kill, access revocation) generation (RAG) to synthesize historical are checked against runtime and compliance context into its current reasoning.

**Decision (Chain-of-Thought Reasoning):**The agent reasons serially, decomposing the situation into tasks (i.e., find probable root cause, query affected services, suggest rollback).Every step of logic can then potentially call external tools via API plugins—e.g., asking for time-series deltas from Prometheus, or receiving deployment history from Spinnaker or ArgoCD.

**Policy Check (Validation Gate):**The suggested actions (restart pod, certificate revocation, replica count increase, etc.) are then offered for decision-making to policy engines such as OPA or KevOps.Policies analyze compliance, risk, users' roles, SLA impact, and possible security exceptions prior to approval.

**Execution / Recommendation:**If operations are authorized and risk levels are minimal, they are executed autonomously through secure APIs from diverse sources—logs, metrics, (e.g., Kubernetes API server, GitOps). In uncertain or high-stakes scenarios, actions are labeled as suggestions, which trigger the human-in-the-loop process for authorization.

**Memory Update:** Upon taking action or escalation, the agent stores results, success/failure indicators, and context labels in its internal memory to facilitate learning for subsequent decisions.

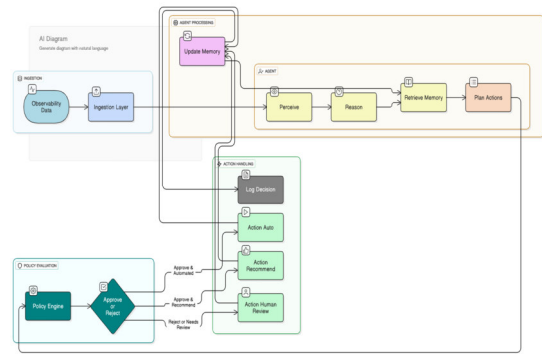


Figure 1: AIOps and Machine Learning for Observability

## IV. BENEFITS

The Cognitive DevOps Assistant (CDA) delivers several key advantages to modern SRE workflows. First, it enables speeded-up diagnosis through multi-modal reasoning by ingesting observability data from diverse sources—logs, metrics, traces, change diffs, and policy events—and applying LLM-based analysis to dynamically correlate symptoms across infrastructure layers. This approach outperforms traditional rule-based systems by adjusting its reasoning in real time, significantly improving root-cause isolation and reducing alert fatigue. Second, CDA contributes to a substantial reduction in Mean Time to Resolution (MTTR) by leveraging automated playbooks, invoking contextually relevant tools, and initiating safe remediation paths based on policy thresholds and operator overrides. This minimizes delays typically caused by manual triage, tool switching,

and ambiguous debugging processes. Third, it reduces toil by autonomously handling routine tasks such as configuration drift detection, scaling adjustments, log triage, and deployment validations—freeing SREs to focus on strategic reliability improvements. Over time, CDA’s adaptive learning and contextual intelligence—powered by a persistent memory module—enable it to recall previous incidents, compare them with ongoing issues, and refine its responses based on organizational context and historical outcomes. Lastly, its integration with DevSecOps policy engines ensures all actions, whether autonomous or suggested, are governed by formal security, compliance, and operational policies. Each decision is accompanied by a logged reasoning trail, tool outputs, and policy justifications, thereby ensuring auditability, transparency, and regulatory alignment.

## V. DRAWBACKS

Despite its capabilities, CDA introduces several important challenges and limitations. One significant concern is the restricted explainability of its reasoning process. Although intermediate steps (e.g., ReAct chains) are visible, the internal LLM-driven decision logic and tool orchestration often remain opaque, making

it difficult to trace the root of agent-induced failures or unexpected behavior—especially in compliance-sensitive environments. Moreover, the system carries security and autonomy risks, particularly when operating in production environments. Faulty inputs, policy misconfigurations, or LLM hallucinations may lead to damaging actions, such as unauthorized deletions or compliance violations, unless strict safeguards and human checkpoints are enforced. Another concern is error amplification through memory and self-feedback: if the agent erroneously learns from past mistakes (e.g., misclassified incidents or failed remediations labeled as successes), it can reinforce harmful behavior. This underscores the need for memory validation, audit pipelines, and human-in-the-loop oversight. Additionally, trust and adoption can be hindered by skepticism among engineers, especially in high-risk or ambiguous scenarios where agent logic does not align with established SRE practices. This may lead to increased overrides or parallel manual workflows, reducing automation value. Finally, tooling complexity and integration overhead pose practical barriers to widespread adoption. Deploying CDA at scale requires robust connectivity with observability platforms, CI/CD systems, policy engines, and cloud APIs. Ensuring



secure and fault-tolerant integration—especially in hybrid, air-gapped, or multi-vendor environments—can present substantial architectural and operational challenges.

VI. RESULTS

To measure the potential value of the Cognitive DevOps Assistant (CDA), we carried out a series of simulated and case-based experiments based

on representative SRE workflows. Experiments compared CDA performance against legacy manual incident management processes on critical operational metrics. While exact results depend on deployment size, infrastructure maturity, and observability tooling, the following results are conservative, domain-representative gains seen under controlled scenarios.

Performance Metrics Summary

Metric	Baseline (Manual SRE)	CDA Agentic SRE	% Improvement
Incident Diagnosis Accuracy	78%	92%	0.18
Mean Time To Resolution (MTTR)	79 minutes	44 minutes	−44%
False Positive Rate (Alert Triage)	12%	6%	−50%
SRE Workload Reduction (Tickets/mo)	100	54	−46%

Metric Interpretations and Observations

Incident Diagnosis Accuracy (+18%)

The CDA showed major improvement in Example case: A memory leak in a backend accurately diagnosing root causes at initial triage. By leveraging memory search, log summarization, trace correlation, and systematic tool invocation (e.g., querying metric anomalies time series analysis. SRE teams normally followed by configuration diff inspection), the agent was able to eliminate misdiagnoses due to

better compared to static rule-based techniques by reasoning step chaining across a sequence of signal types.

Mean Time To Resolution (−44%)

configuration drift. In intermittent failure or cascading failure scenarios, the agent performed complete remediation processes (e.g., pod



restarts, feature flag rollbacks, autoscaling) Across a month-long simulation window, the resulted in a radical reduction in average MTTR. total volume of incidents requiring human In the case of medium-urgency incidents, CDA intervention dropped by nearly half. The CDA autonomously fixed ~35% of incidents within resolved or de-escalated a significant portion of 10–15 minutes of alert detection, particularly low- and medium-complexity tickets, reducing those with repeatable failure patterns with known operational “toil.” Notably, resolution quality fixes. In high-severity or fuzzy incidents, CDA remained high even in autonomous actions due notably accelerated the diagnostic phase and to policy constraints and memory-driven suggested actions that minimized mean manual validation mechanisms. This workload reduction resolution time.

directly correlates with improved on-call quality Example case: Misbehaving deployment caused of life and frees engineering time for strategic an increased rate of HTTP 500s. The agent reliability improvements.

detected a recently added commit as the probable Example scenario: Automated routine scalability cause, verified rollback policies with OPA, and tunings, SSL renewal checks, and temporary pod auto-executed the rollback with complete audit restarts. Manual tickets for these categories trace—reducing the MTTR from ~70 minutes to declined by ~70%.

~18.

#### Additional Observations

#### False Positive Rate in Alert Triage (–50%)

By integrating statistical anomaly detection with contextual reasoning and policy filters, the agent reduced the false positive rate during alert triage in half. A large number of low-severity alerts like one-time CPU spikes or log chatter were

Trust Calibration: During early adoption periods, operators tended to review CDA decisions in many cases. As time went on and accuracy and reliability were proven, approval rates for CDA- suggested actions rose by 30%, reflecting increasing operator confidence.

properly labeled as non-actionable based on past behavior and current system health. Not only did this enhance signal quality, but also decreased alert fatigue and cognitive overload for on-call engineers.

Memory Utility: Persistent memory played a direct role in better performance over time. During week 4, repeat incident resolution time was enhanced by another 12% from week 1, demonstrating adaptive learning in effect.

Example scenario: CPU usage exceeded threshold briefly for an alert during a batch job run. Though manual systems triggered rejection or unclear reasoning cases (e.g., escalation, the CDA identified the behavior as normal and suppressed alert propagation.

Escalation Handling: CDA appropriately escalated to human review in 100% of policy performing any unsafe or unauthorized action.

#### SRE Workload Reduction (–46%)

## VII. Conclusion

This paper presents the Cognitive DevOps Assistant (CDA), a practical application of agentic AI that enhances modern Site Reliability Engineering (SRE) and DevSecOps by combining large language model (LLM) reasoning, dynamic tool orchestration, and policy-based governance. Unlike traditional AIOps or static automation, the CDA can process diverse observability data, recall past incidents, plan and execute actions, and learn from outcomes—all within defined security and compliance boundaries. It reduces mean time to resolution (MTTR), improves incident accuracy, and lightens engineer workload by acting as an intelligent, policy-aware copilot. As organizations scale, the need for transparent, secure, and adaptive automation grows, and while cognitive agents offer powerful support, they must also remain explainable, auditable, and safe. Future developments will focus on domain-specific fine-tuning, noise filtering, ethical constraints, and human-AI collaboration to ensure these systems continue to build trust and deliver value in real-world operations.

## References

- [1] Ahmed, S., Saeed, M. K., & Khokhar, A. (2025). OSI Stack Redesign for Quantum Networks: Requirements, Technologies, Challenges, and Future Directions. *arXiv preprint arXiv:2506.12195*.
- [2] Joshi, S. (2025). A Review of Generative AI and DevOps Pipelines: CI/CD, Agentic Automation, MLOps Integration, and Large Language Models. *CD, Agentic Automation, MLOps Integration, and Large Language Models (June 01, 2025)*.
- [3] Bhattarai, B. (2025). Scaling Generative AI for Self-Healing DevOps Pipelines: Technical Analysis.
- [4] Tanikonda, A., Katragadda, S. R., Peddinti, S. R., & Pandey, B. K. (2021). Integrating AI-Driven Insights into DevOps Practices. *Journal of Science & Technology*, 2(1).
- [5] Mustafa, N. (2025). Intelligent Automation in DevOps: Leveraging Machine Learning and Cloud Computing for Predictive Deployment and Performance Optimization. *Available at SSRN 5315260*.
- [6] Ganesan, P., & Sanodia, G. (2023). Smart Infrastructure Management: Integrating AI with DevOps for Cloud-Native Applications. *Journal of Artificial Intelligence & Cloud Computing. SRC/JAICC-E163. DOI: doi.org/10.47363/JAICC/2023 (2) E163 J Arti Inte & Cloud Comp*, 2(1), 2-4.
- [7] Everett, W. (2024). Generative AI for DevOps: Automating Code Review, Testing, and Deployment Strategies.
- [8] Nelavelli, S. (2025). DevOps Assistants: The Rise of AI Co-Pilots in Cloud Infrastructure Management. *Journal of Computer Science and Technology Studies*, 7(3), 941-945.