International Journal on Science and Technology (IJSAT)
E-ISSN: 2229-7677 ● Website: www.ijsat.org ● Email: editor@ijsat.org

Simple Agentic AI workflow for AIOPS (Agentic AIOps)

Surendar Raj

Student – M. Tech-DSML, PES University, Hosur Rd, Konappana Agrahara, Bengaluru, 560100, Karnataka, India.

Abstract

This paper showcases a simple Agentic AI framework aimed at improving AIOps through the deployment of autonomous, goal-oriented agents. Artificial Intelligence for IT Operations (AIOps) refers to the application of AI and machine learning to make IT operations run more efficiently and automatically. This work demonstrates AI agents for detecting problems, automated baselining deduplication of alerts and remediation. By removing unnecessary alerts, the AI agent for deduplication makes sure that only distinct and actionable alerts move to the next phase. Using anomaly detection algorithms, the operational system identifies unexpected things in operational data streams, such as logs, metrics, and traces. Agents employ ARIMA time-series modeling, and past events to baseline metrics and start resolving automatically in case of anomaly, such as by adding more resources or restarting services. Agents may make incident resolution work better by constantly improving their answers with feedback. This agentic workflow uses smart policy engines, automation frameworks, and observability tools to build a scalable base for proactive and self-healing AIOps settings. Using AI agents to accomplish different tasks in AIOPS makes things operate more smoothly and needs less interference from people.

Keywords: Artificial Intelligence for IT Operations (AIOps), AI Agents, Operations Management (ITOM), Automated deduplication, Automated baselining, Automated Remediation, Agentic AI AIOps

1. Introduction

Operations management has become more complex than ever before due to the exponential growth of modern IT infrastructures (ITOM) [1], which are fueled by cloud-native architectures, microservices, container orchestration platforms, and large-scale distributed systems. Today's businesses produce enormous amounts of diverse telemetry data at a rapid pace, such as logs, metrics, events, and distributed traces. It has become more difficult to manage such an environment using conventional rule-based monitoring techniques. In dynamic environments where workloads, resource utilization, and system behavior are constantly changing, static thresholds, manual triaging, and operator-driven root cause analysis frequently fall short. Organizations suffer from misdiagnosis of failures, delayed incident resolution, and a sharp rise in operational overhead as a result.

The introduction of Artificial Intelligence for IT Operations (AIOps), a term first used by Gartner in 2017, was a big step forward. [2]. Artificial Intelligence for IT Operations (AIOps) has become a new way of thinking about how to solve these problems. AI and machine learning are used by AIOps systems to find problems, automate root cause analysis, connect events, and fix problems before they happen. However, many current AIOps implementations are constrained because they rely on separate automation scripts, don't coordinate well between agents, and don't have a good sense of context. They might be able to find anomalies or remove duplicate alerts on their own, but they rarely show adaptive, goal-driven behaviors that can change with the situation.

Businesses that use agent-based AIOps will save a lot of money by reducing manual interventions, improve system resilience, and decrease downtime. Organizations can create a truly autonomous IT ecosystem where incidents are identified, examined, and resolved quickly by integrating intelligent policy engines, automation frameworks, and observability layers. The incorporation of autonomous, goal-driven software agents into IT operations pipelines, or Agentic AIOps, has gained attention because of this gap. Intelligent, modular agents with the ability to cooperate, share memory, and react dynamically to changing operational states are introduced by an agentic framework.

2. Research Methodology

Daniel Zota et al. [3] presents a new framework for an Agentic AIOps system, including its essential elements, features, and architectural specifications. The framework emphasizes the employment of intelligent, context-aware, self-learning agents to manage issues and incidents, fulfill requests, and prevent problems before they become problems. While staying flexible enough to accommodate the constantly changing technology present in actual organizational settings, it is made to conform to well-known service management best practices like ITIL [4] and the Unified Process Framework for IT (UPF-IT). Additionally, the framework facilitates scalable deployment for middleware, data, applications, infrastructure, and other levels of the IT ecosystem.

After filtering 140 publications based on credibility, relevance, and recentness, 42 high-impact studies were chosen as the final group. Key topics essential to the development of AIOps are reflected in these publications are:

- Anomaly detection
- Event correlation
- Automated remediation
- Intelligent decision-making in AIOps workflows

Sabharwal, N et al. [4] in his book, illustrates the three core IT operations services—automation, IT service management, and enterprise monitoring—form the foundation of AIOps systems. By fusing cutting-edge technologies and processes, the AIOps architecture creates a comprehensive AIOps platform that facilitates smooth integration across different services. The platform exhibits its applicability across many processes and functions inside the Gartner-defined IT operations value chain,

as shown in Figure 1. Examining the underlying AIOps architecture in greater detail is crucial to comprehending its function.

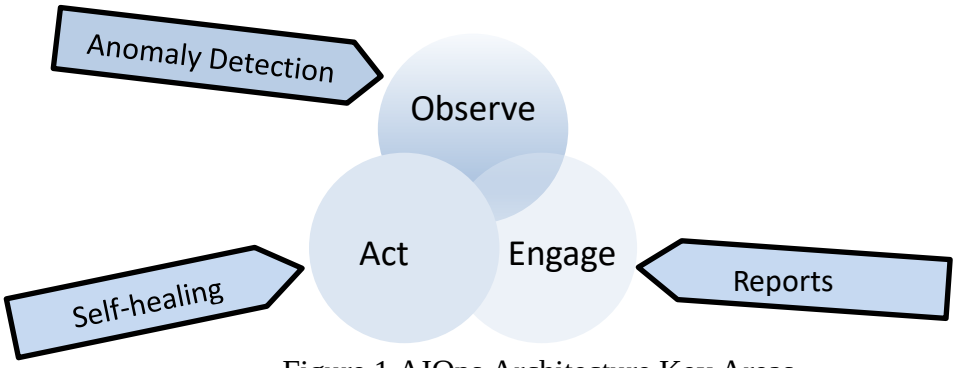

Figure 1 AIOps Architecture Key Areas

The major use cases of AIOps include deduplication, automated baselining, anomaly detection, and remediation.

Anomaly Detection

The technique of automatically spotting odd or unexpected patterns in IT systems that can point to possible issues is known as anomaly detection in AIOps. These irregularities frequently show up as unusual increases in CPU, memory, or storage utilization, such as when a disk reaches its maximum capacity. Using statistical models and machine learning, anomaly detection systems can differentiate between real anomalies that need to be addressed and typical oscillations.

Automated Baselining

The method by which IT systems dynamically learn typical behavior by adjusting to patterns, long-term trends, and seasonality is known as automated baselining. Instead of depending on fixed, manually set criteria, it automatically modifies baselines over time to take seasonal fluctuations, weekly cycles, and daily peaks into consideration. For spotting abnormalities or departures from the learnt patterns, this makes it incredibly effective.

Automated Deduplication

The process of identifying and eliminating duplicate events or alerts that originate from the same underlying issue is known as automated deduplication in AIOps. A single problem, like a server failure or network outage, may result in multiple redundant alerts from different monitoring tools in large-scale IT infrastructures. These duplicates create alert storms that overload operators and delay root cause analysis if they are not deduplicated.

Automated Remediation

The process of automatically resolving IT incidents, anomalies, or alerts after they have been identified and examined is known as automated remediation. Predefined workflows or scripts are triggered to handle the problem—such as resetting services, scaling resources, or freeing up disk space—instead of necessitating human intervention. This guarantees consistent handling of recurring operational issues, decreases downtime, and speeds up incident response. In the end, automated remediation improves dependability while freeing up IT teams to work on more difficult assignments.

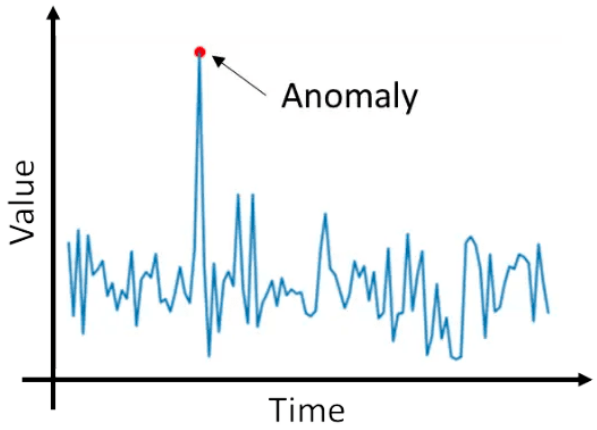

Anomaly Detection in time-based graph of a system

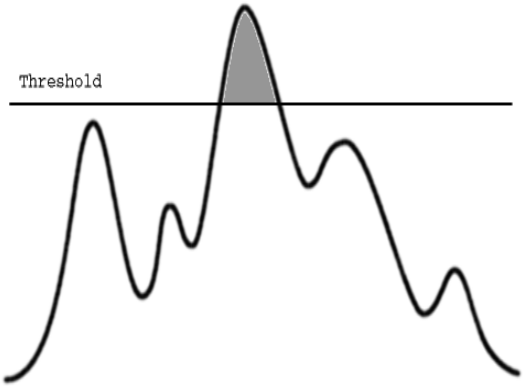

Graph of a system illustrating threshold limit exceeded over a seasonality

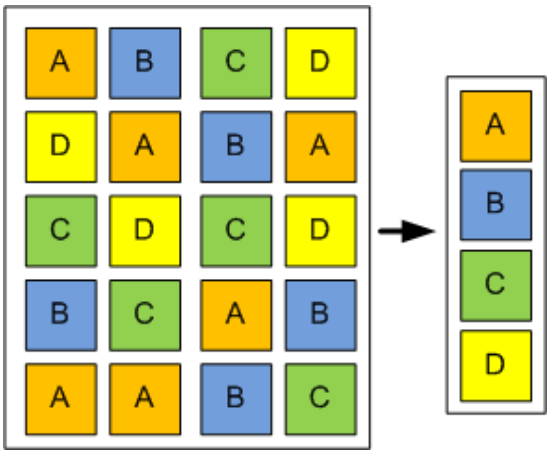

Illustration of alert deduplication

Self-Healing Script

Figure 2 AIOPS Major Cases

Yinfang Chen et. al [5] proposed an assessment system for AI agents for self-governing cloud operations. It facilitates end-to-end lifecycle management by giving agents the ability to identify, locate, diagnose, and fix issues in microservice contexts. A telemetry system, a unified Agent-Cloud Interface (ACI), a cloud-agent orchestrator, and fault and workload generators. The paper uses the word "**AgentOps**" too loosely to denote "**agent for operations,"** which is not accurate. AgentOps is a new field that focuses on managing the entire lifecycle of AI agents.

Any Microservice deployed in a broader sense is represented to have the following stages as discussed above i.e., anomaly detection, baselining, deduplication, healing which are same as discussed in other two papers.

3. Materials and Method

Preprocessing system and Environment

Within the pre-processing subsystem, monitoring tools like AWS CloudWatch and Prometheus Alertmanager can turn alerts into standard JSON payloads. As for AWS CloudWatch, the system can use the boto3 API to get the alert data and change it into JSON format. For Prometheus Alertmanager, a setup template based on YAML can be used to organize and map the alerts before they are converted. Once these JSON payloads are consistent, they are sent to the system's API or webhook destination. At the API layer, the system applies authorization mechanisms to keep this ingestion process safe. In the request header of every HTTP request, there must be an Authorization: Bearer token.

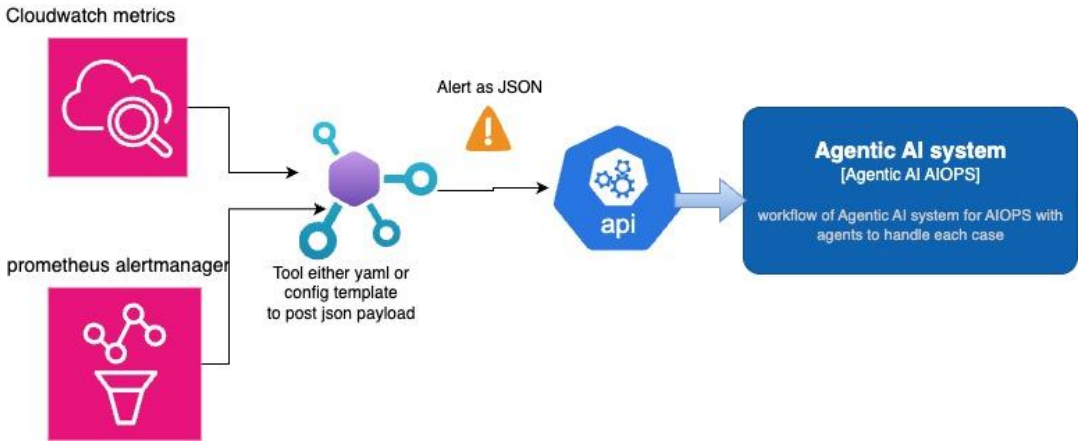

Figure 3 Preprocessing System

System architecture comprises a lightweight, modular multi-agent workflow for AIOps that makes operational administration easier and more automatic. The workflow includes several specialized agents that are designed to work together in a coordinated fashion to do things like deduplication, baseline modelling, anomaly detection, and remediation. The implementation is designed to be cloud-native and easily deployable on AWS EKS, while integrating directly with Prometheus and AWS CloudWatch for real-time data ingestion through a web hook.

System consists of metrics relevant to Prometheus Alertmanager and AWS CloudWatch. These alerts are automated based on specific defined rules and levels of thresholds which means that valuable events are picked up in almost real-time. Latency is reduced and external systems are not used when these alerts are received directly over a webhook.

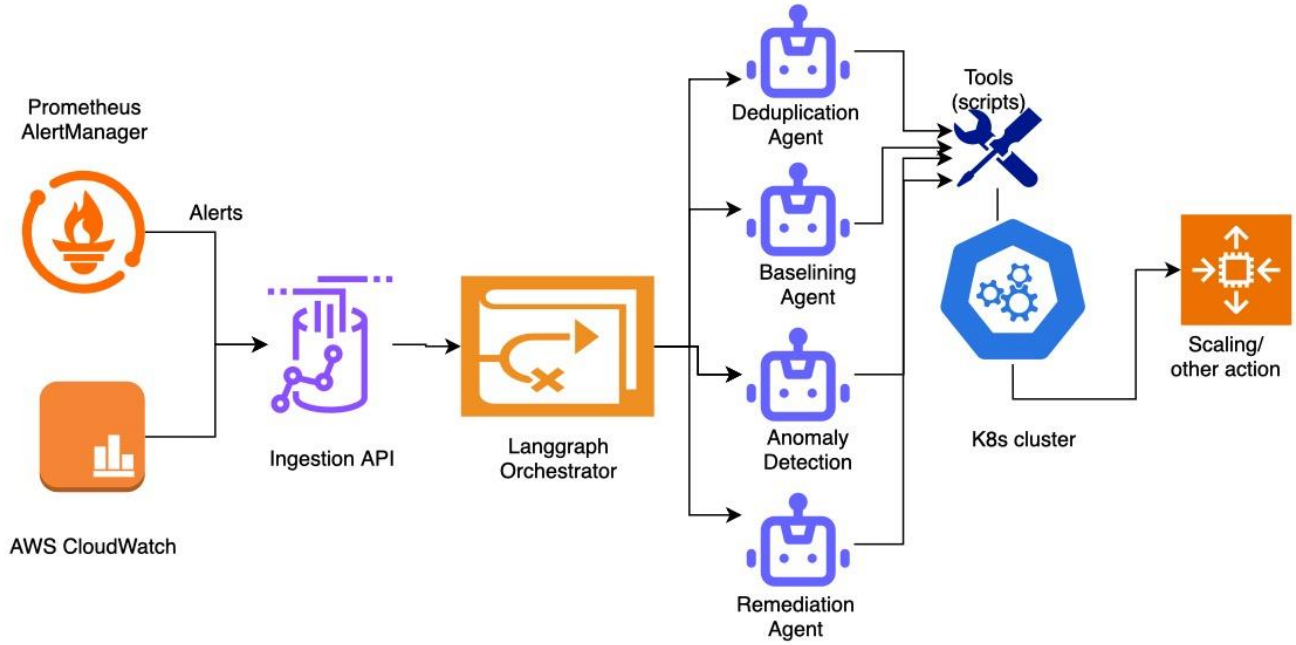

Figure 4 Agentic AIOps System

The FastAPI webhook (Ingestion API) upon which the ingestion layer is based receives notifications from external monitoring systems. The alert payloads are sent to the orchestrator for additional processing. Multiple agent systems, including Deduplication, Baseline, Policy, LLM Approval, and Remediation, are housed at the agent layer. Every agent contributes to the decision pipeline while functioning independently. As workload increases, this modular design enables flexible scaling of individual agents and enhances maintainability.

LangGraph's StateGraph is used in the orchestrator's implementation to control agent interactions. It guarantees that alerts proceed smoothly from preprocessing to evaluation and remediation by enforcing the logical flow among agents. In distributed workflows, this central coordination keeps everything consistent and avoids deadlocks. The Kubernetes client is used to carry out remediation; it has the ability to restart pods and scale deployments. Safeguards include maximum replica limitations and dry-run checks stop dangerous or erratic behaviour. This guarantees that corrective measures adhere to operational policies and are safe and reversible.

Docker is used to containerize the system, and deployment on Amazon EKS is made possible via Kubernetes manifests (k8s.yaml). Scalability, rolling updates, and resistance to pod failures are all made possible by this configuration. Reproducibility and simplicity of redeployment across environments are guaranteed by the infrastructure as code concepts.

4. Results

Mean Time to Detect (MTTD) and Mean Time to Resolution (MTTR) [6] are the two critical performance metrics that are being pursued. MTTD shows how long it takes to find out about problems or failures after they happen, while MTTR shows the average time it takes to completely investigate and fix an event from the time it is found to the time it is closed.

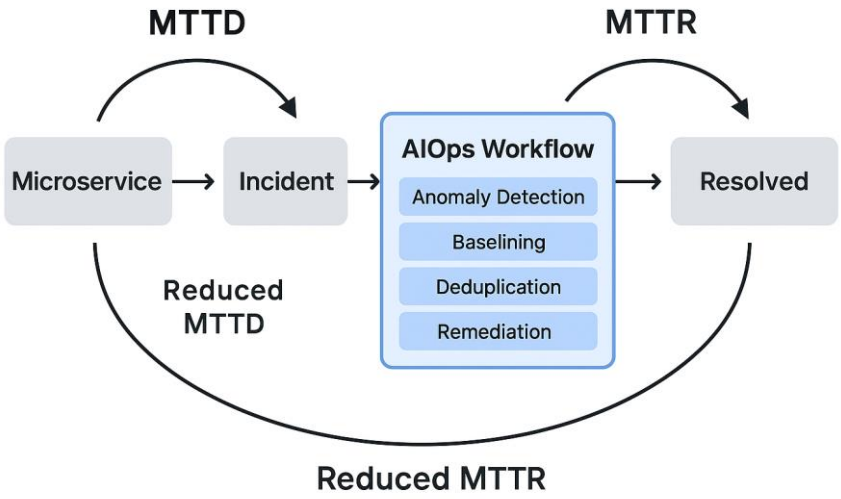

Figure 5 MTTD and MTTR

The System can cut MTTD by up to 70–80% by proactive anomaly detection and MTTR by 50–60% through automated remediation, resulting in a much faster and more dependable defect resolution process.

5. CONCLUSION

By integrating anomaly detection, deduplication, baselining, and remediation agents in a Kubernetes-native environment, the Agentic AIOps framework has effectively illustrated the potential of multi-agent orchestration. The system has demonstrated that intelligent policy evaluation, explainability, and hybrid anomaly detection (Z-score and Isolation Forest) can minimize false positives and safely and carefully automate remediation. These outcomes provide a solid foundation for enterprise-scale adoption and verify the efficacy of an agent-driven approach for contemporary IT operations.

6. Acknowledgement

For his ongoing direction, support, and encouragement during the development process, I would like to thank Dr. Milan Joshi, Faculty, Department of Data Science and Machine Learning, Great Learning. I would like to acknowledge Anurag S. as the Project Coordinator for consistent guidance throughout the development.

7. Authors' Biography

The author has worked as a Technical Specialist for top IT companies for more than 15 years. He has gained great domain experience in communications BSS systems, which have been the main focus of his work. He is currently pursuing a fourth semester of the Master of Technology in Data Science and Machine Learning at PES University, Bangalore further expanding his skills in advanced analytics and AI.

References

1. IBM 2022, What is IT operations management (ITOM)?, IBM Think, 12 July, viewed 13 July 2025, https://www.ibm.com/think/topics/itom.
2. Andrew Lerner.AIOps Platforms, Everything you need to know about AIOps https://blogs.gartner.com/andrewlerner/2017/08/09/AIOps-platforms/, August, 2017
3. Zota, R.D., Bărbulescu, C. and Constantinescu, R., 2025. A practical approach to defining a framework for developing an agentic AIOps system. Electronics, 14(9). https://doi.org/10.3390/electronics14091775.
4. Sabharwal, N. and Bhardwaj, G., 2022. “Hands-on AIOps: Best Practices Guide to Implementing AIOps. “United States: Apress.
5. Chen, Y., Shetty, M., Somashekar, G., Ma, M., Simmhan, Y., Mace, J., Bansal, C., Wang, R. and Rajmohan, S., 2025. Aiopslab: A holistic framework to evaluate ai agents for enabling autonomous clouds. arXiv preprint arXiv:2501.06706.