

Original Article

An Intelligent Agent for Root-Cause Detection and Remediation

Yuga Sai Kiran Reddy Sudini

Director Of Engineering and operations at Public Partnerships LLC, USA.

Received: 04-08-2022

Revised: 04-28-2022

Accepted: 05-06-2022

Published: 05-08-2022

ABSTRACT

Modern cloud environments generate massive volumes of operational data from infrastructure metrics, application logs, distributed traces, events, and monitoring alerts, making rapid identification of the underlying causes of failures increasingly challenging. Traditional incident-management approaches often depend on manual investigation and predefined rules, which can increase mean time to detect and resolve incidents. This paper presents an intelligent agent-based framework for automated root-cause detection and remediation across cloud infrastructure, applications, and distributed services. The proposed approach integrates artificial intelligence and machine learning techniques to continuously analyze heterogeneous operational signals, correlate abnormal behaviors, identify probable root causes, and recommend or execute appropriate remediation actions. The intelligent agent combines observability data, dependency relationships, anomaly detection, predictive analysis, and automated decision-making to support end-to-end incident resolution. By reducing dependence on manual troubleshooting, the framework aims to improve incident detection speed, diagnostic accuracy, service reliability, and operational efficiency. The approach further supports continuous learning from historical incidents and remediation outcomes, enabling the system to adapt to evolving cloud workloads and failure patterns. The proposed framework demonstrates the potential of intelligent automation to transform conventional cloud operations toward proactive, self-healing, and resilient infrastructure management.

KEYWORDS

Intelligent Agent, Root-Cause Detection, Automated Remediation, Artificial Intelligence, Machine Learning, Cloud Infrastructure, AIOps, Observability, Anomaly Detection, Self-Healing Systems, Incident Management, Cloud Operations.

1. THE PROBLEM: FROM DETECTION TO UNDERSTANDING

Traditional monitoring is good at detection and poor at explanation. It tells you that a metric crossed a threshold or that an error rate spiked, but it leaves the crucial work determining the underlying cause to human engineers who must correlate signals across many systems under time pressure. In a complex, distributed environment, a single root cause often manifests as a cascade of symptoms across many components, and disentangling the cause from its consequences is genuinely difficult. This is where operations teams spend their most stressful hours, and where mean time to resolution is won or lost.

The core insight motivating the agent was that this diagnostic work follows patterns. Experienced engineers develop intuitions about which symptoms tend to share a common cause, which signals are leading indicators, and which remediation actions resolve which failure modes. If those patterns can be encoded and augmented with the tireless, instantaneous correlation that software excels at, much of the diagnostic burden can be shifted from humans to an intelligent system.

2. DESIGNING FOR AUTOMATED DIAGNOSIS

The agent's first responsibility is diagnosis: taking the flood of signals during an incident and determining the likely root cause. Doing this well requires bringing together information that is normally siloed. Metrics reveal the shape of a problem, logs reveal its details, traces reveal how a request flows across services, and topology reveals how components depend on one another. The agent correlates across these dimensions to distinguish cause from symptom recognizing, for example, that a spike in errors across many services traces back to a single degraded dependency rather than treating each symptom as an independent problem.

This correlation across infrastructure, applications, and services is what makes the agent powerful. A problem in the infrastructure layer often surfaces as application-level symptoms, and an application defect can manifest as infrastructure stress. By reasoning across all these layers together rather than in isolation, the agent can pinpoint causes that no single-layer view would reveal. It compresses into moments the correlation work that would otherwise consume an engineer's most valuable minutes during an incident.

3. FROM DIAGNOSIS TO REMEDIATION

Detection and diagnosis are valuable on their own, but the agent's more ambitious capability is remediation: taking action to resolve issues, not merely to report them. For well-understood failure modes with known, safe corrective actions, the agent can execute remediation automatically restarting an unhealthy component, scaling a resource under pressure, or applying a known fix resolving the issue before it escalates and often before users are affected.

Automated remediation must be approached with care, and I designed the system around that principle. Not every problem should be resolved autonomously; some require human judgment, and acting incorrectly can make a situation worse. The agent therefore operates within clearly

defined boundaries: it takes autonomous action only for failure modes where the remediation is well understood and safe, and it escalates to human engineers with a diagnosis and recommended actions for anything outside those boundaries. This graduated approach captures the speed of automation where it is safe while preserving human oversight where it is needed.

The goal is not to remove humans from operations, but to free them from the mechanical work of correlation so they can focus on the judgment that genuinely requires them.

4. THE ROLE OF AI

The intelligence in the agent lies in its ability to recognize patterns and reason about relationships that would be impractical to enumerate exhaustively by hand. Rather than relying solely on rigid, predefined rules which are brittle and cannot anticipate every scenario an intelligent approach can generalize from patterns, adapt as the environment evolves, and handle the novel combinations of symptoms that real incidents so often present. This is where AI moves operations beyond static alerting toward genuine, adaptive diagnosis.

I am also deliberate about applying AI responsibly in this context. An agent that takes action in production must be trustworthy, its decisions must be explainable so that engineers can understand and verify them, and its boundaries must be well defined so that it never acts beyond its competence. Responsible design transparency, clear limits, and human oversight for consequential actions is what makes autonomous operations safe rather than reckless.

5. WHY THIS MATTERS FOR RELIABILITY

The impact of an intelligent root-cause and remediation capability is measured most directly in reduced mean time to resolution. Every minute shaved off diagnosis is a minute of reduced impact on users. By compressing the diagnostic phase and automating safe remediation, the agent enables faster incident resolution and improved system reliability. Just as importantly, it reduces the operational burden and stress on engineering teams, who no longer spend their nights manually correlating signals to find a cause the system can surface in seconds.

There is a compounding benefit as well: an agent that captures diagnostic patterns becomes an institutional memory of how the system fails and how those failures are resolved. Knowledge that would otherwise live only in the heads of experienced engineers becomes encoded, shared, and consistently applied reducing the organization's dependence on individual heroics during incidents.

6. THE FUTURE OF OPERATIONS

I regard this kind of intelligent automation as the direction operations is heading. The volume and complexity of modern systems have outpaced the capacity of purely manual operations, and the answer is not more people staring at more dashboards but smarter systems that handle the mechanical work and elevate humans to the judgment-intensive decisions where they add the most value. Building the Root-Cause Detection and Remediation Agent was an original application of AI

to a real operational problem, and it reflects a conviction I hold strongly: the future of reliable operations is human judgment amplified by intelligent automation, not replaced by it.

7. IN PRACTICE: REASONING ACROSS A REAL ESTATE

The agent's diagnostic surface is the platform described throughout this series: a global entry and web application firewall layer, regional network hubs with next-generation firewalls, a Kubernetes estate running domain microservices, a multi-tier data layer with a zone-redundant primary and purpose-segregated replicas, and an observability fabric spanning cloud-native monitoring and commercial APM tooling. Correlating across those layers is where the agent earns its keep: an error surge in one service, connection saturation at the pooling layer, and pressure on a specific replica are three symptoms the agent's job is to recognize them as one cause.

REFERENCES

- [1] Panyala, Venkatramana Reddy, and Haritha Pappu. "Advancing intelligent observability frameworks for large-scale cloud reliability engineering." *International Journal of Engineering & Extended Technologies Research (IJEETR)* 3.5 (2021): 3709-3713.
- [2] Becker, Soeren, et al. "Towards aiops in edge computing environments." *2020 IEEE International Conference on Big Data (Big Data)*. IEEE, 2020.
- [3] Samir, Areeg, and Claus Pahl. "A controller architecture for anomaly detection, root cause analysis and self-adaptation for cluster architectures." *Intl Conf Adaptive and Self-Adaptive Systems and Applications*. Vol. 63. 2019.
- [4] Aaseng, Gordon, Kevin Cavanaugh, and Somnath Deb. "An intelligent remote monitoring solution for the international space station." *Proceedings of the IEEE Aero Conference*. 2002.
- [5] Zhu, Qin. *Adaptive root cause analysis and diagnosis*. Diss. 2010.
- [6] Nedelkoski, Sasho, et al. "Multi-source distributed system data for ai-powered analytics." *European conference on service-oriented and cloud computing*. Cham: Springer International Publishing, 2020.
- [7] Iqbal, H. M., and Muhammad Bilal. "Time to automate the microbial detection and identification: the status quo." *J Pure Appl Microbiol* 14.1 (2020): 1-3.
- [8] Bogatinovski, Jasmin, et al. "Artificial intelligence for it operations (aiops) workshop white paper." *arXiv preprint arXiv:2101.06054* (2021).
- [9] Vermaak, Herman, and Johnson Kinyua. "Multi-agent systems based intelligent maintenance management for a component-handling platform." *2007 IEEE International Conference on Automation Science and Engineering*. IEEE, 2007.
- [10] Kim, Myunghwan, Roshan Sumbaly, and Sam Shah. "Root cause detection in a service-oriented architecture." *ACM SIGMETRICS Performance Evaluation Review* 41.1 (2013): 93-104.
- [11] Yuan, Yannan, et al. "Anomaly detection and root cause analysis enabled by artificial intelligence." *2020 IEEE Globecom Workshops (GC Wkshps)*. IEEE, 2020.
- [12] Solé, Marc, et al. "Survey on models and techniques for root-cause analysis." *arXiv preprint arXiv:1701.08546* (2017).
- [13] Chigurupati, Asha, and Noah Lassar. "Root cause analysis using artificial intelligence." *2017 Annual reliability and maintainability symposium (RAMS)*. IEEE, 2017.
- [14] Yue, Wenwei, et al. "What is the root cause of congestion in urban traffic networks: Road infrastructure or signal control?." *IEEE Transactions on Intelligent Transportation Systems* 23.7 (2021): 8662-8679.
- [15] Zidan, Aboelsood, et al. "Fault detection, isolation, and service restoration in distribution systems: State-of-the-art and future trends." *IEEE Transactions on Smart Grid* 8.5 (2016): 2170-2185.
- [16] Taluri, R. (2021). Cloud-Native Architectures for Enterprise Financial Data Management, Analytics, and Regulatory Reporting Compliance. *International Journal of Emerging Trends in Computer Science and Information Technology*, 2(2), 101-111. <https://doi.org/10.63282/3050-9246.IJETCSIT-V2I2P112>