Archives available at journals.mriindia.com

International Journal of Electrical, Electronics and Computer Systems

ISSN: 2347-2820
Volume 14 Issue 01, 2025

Self Healing Infrastructure System

Dr. R. V. Patil[1], Vrushali Kudande[2], Shweta Jagtap[3], Sakshi Jadhav[4], Aniket Jawalgekar[5]

[1]Professor PDEA's College of Engineering, Pune.
[2,3,4,5]Student PDEA's College of Engineering, Pune.
arvpatil7@gmail.com[1], vrushalikudande@outlook.com[2], shwetajagtap1450@gmail.com[3], sakshii.jadhav0401@gmail.com[4], jawalgekar99@gmail.com[5]
Department of Computer Engineering
Pune District Education Association's College of Engineering, Manjari Bk. Hadapsar, Pune, Maharashtra, India. - 412307 Email: coem@pdeapune.org

Peer Review Information

Submission: 10 Jan 2025
Revision: 07 Feb 2025
Acceptance: 09 March 2025

Keywords

Automated Remediation
Anomaly Detection
ML-Driven Automation
DevOps Practices
Real-Time Observability

Abstract

Managing modern IT infrastructure has become increasingly complex with the adoption of containerized applications and distributed architectures. Traditional methods reliant on manual monitoring, fault detection, and recovery are no longer sufficient due to their slow response times, error-prone nature, and increased downtime. The need for 24/7 system availability and unpredictable failures highlights the urgency of more resilient solutions. Advanced tools like Prometheus offer real-time observability, anomaly detection, and system health monitoring, enabling quicker responses. However, traditional approaches lack automated remediation capabilities.
Integrating DevOps practices with ML-driven technologies addresses these gaps. Enhancing tools like Prometheus with self-healing capabilities allows infrastructures to autonomously detect anomalies, diagnose issues, and execute automated recovery workflows. This reduces downtime, optimizes resources, and minimizes manual intervention. This research aims to develop a scalable, robust system ensuring business continuity, reliability, and resilience in dynamic, cloud-native environments, unlocking the potential of autonomous infrastructure management.
JEL Classification Number: C88, O33, M15, D83 , L86.

INTRODUCTION

Modern IT infrastructure is becoming increasingly complex due to the adoption of containerized applications and distributed systems. As organizations transition to cloud-native environments, maintaining system reliability, scalability, and resilience has become paramount. Traditional methods of manual monitoring and fault resolution are inadequate in addressing the challenges of dynamic, heterogeneous infrastructures.

The proposed self-healing system integrates Ml-driven automation with real-time monitoring tools like Prometheus to overcome the limitations of existing solutions. By leveraging predictive analytics, the system proactively identifies and mitigates anomalies, minimizing downtime and operational disruptions. Furthermore, it adopts platform-agnostic methodologies to address the critical issue of vendor lock-in, ensuring compatibility with diverse infrastructure setups, including multi-cloud and on-premises environments.

Through iterative development practices and continuous integration, the system evolves to meet the demands of modern IT ecosystems, empowering organizations to achieve greater independence, scalability, and business continuity. This research aims to bridge the gap between advanced monitoring capabilities and fully autonomous infrastructure management, paving the way for a resilient and scalable IT future.

LITERATURE SURVEY

Self-Healing Techniques in Cloud Systems

R. Kanniga Devi and M. Muthukannan proposed a self-healing fault tolerance technique using the CloudSim toolkit, incorporating fault injection and recovery mechanisms. Their work significantly improved execution rates for failed requests, demonstrating the potential of simulation-based validation in cloud environments. However, their approach is limited to virtual machine (VM) faults and requires further real-world validation.

AI and Predictive Analytics for Cloud Failures

Rui Xin explored the application of data analytics and machine learning for failure prediction in cloud applications. Their early prototypes showed high precision and recall rates, highlighting the feasibility of predictive analytics for cloud reliability. Nonetheless, the study focuses heavily on predictive capabilities, lacking integration with comprehensive recovery mechanisms.

Frameworks for Self-Healing and Reliability

Amal Alhosban et al. developed a rule-based self-healing algorithm integrated with WS-BPEL, achieving significant improvements in system reliability compared to traditional methods. However, the framework demands substantial computational resources for continuous monitoring, which may limit its scalability.

Optimization Techniques for Resilient Cloud Systems

Swati N. Moon et al. showcased the application of genetic algorithms (GA) for optimizing large datasets, followed by neural network classification. This approach exemplifies GA's versatility for feature extraction and classification, though challenges remain in handling dynamic datasets and adapting to evolving scenarios.
Model Architecture

CONCEPTUAL FRAMEWORK

The proposed self-healing monitoring system is envisioned as a comprehensive and autonomous platform specifically designed to enhance system reliability, minimize downtime, and mitigate the challenges associated with vendor lock-in in modern IT infrastructure. This innovative model integrates a series of modular components, each contributing to seamless real-time anomaly detection, efficient event processing, and automated remediation workflows. A key emphasis is placed on a platform-agnostic architecture that prioritizes scalability and adaptability, ensuring compatibility across multi- cloud, hybrid, and on-premises environments. This design allows organizations to maintain operational flexibility without being dependent on proprietary technologies or solutions.

The system functions as an iterative and continuous feedback loop, encompassing several critical stages: metrics collection, event generation, decision-making processes, and the execution of appropriate remediation actions. By leveraging machine learning-driven analytics, the platform significantly improves anomaly detection and decision-making accuracy. Combined with advanced container orchestration capabilities, such as those provided by Kubernetes, it ensures efficient resource utilization, proactive issue resolution, and dynamic scalability. This approach enhances the overall operational resilience of IT systems, fosters interoperability between different environments, and facilitates fully automated infrastructure management processes, paving the way for a robust and future- proof IT ecosystem.

Functional Architecture:

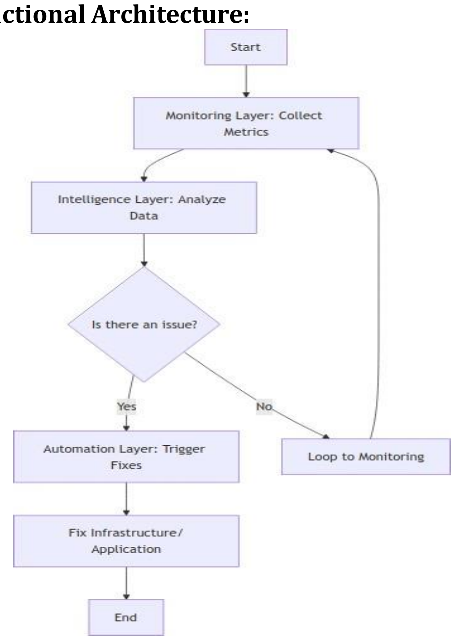

The self-healing monitoring system is structured around the following core components:

Monitoring Layer: Collect Metrics
Applications expose metrics through endpoints, which are aggregated and monitored using tools like Prometheus. Alerts are generated based on predefined rules evaluated in real time.

Intelligence Layer: Analyze Data
The collected metrics are analyzed using AI/ML algorithms. This layer employs anomaly detection (e.g., unsupervised clustering or deep learning models) to identify performance deviations and detect potential issues.
Decision Point: Is There an Issue?
Based on the analysis, the system determines whether an issue exists. If no issues are found, the process loops back to the monitoring layer for continuous metric collection.

Automation Layer: Trigger Fixes
If an issue is detected, decision events are generated and prioritized. Event prioritization is achieved through methodologies like decision trees or reinforcement learning techniques, ensuring critical issues are addressed first.

Fix Infrastructure/Application
Automated remediation workflows are executed using tools like Ansible or Puppet. These workflows implement optimized corrective actions, determined by genetic algorithms or heuristic search methods, to resolve the issues and restore system stability.

Platform-Agnostic Design
The system is designed to be platform-agnostic, ensuring compatibility across multi-cloud, hybrid, and on-premises environments by avoiding reliance on proprietary technologies.

ALGORITHMS USED

Anomaly Detection

Isolation Forest
This algorithm is chosen for detecting anomalies in real-time metrics due to its lightweight and efficient design.It works by isolating anomalous points through random partitioning, making it well-suited for high-dimensional live-streaming data.
Application: Identifies performance outliers like unexpected spikes in CPU usage, memory leaks, or network latency.

Event Prioritization

Decision Trees
Simple and interpretable, Decision Trees are used to assign priority levels to events based on severity and system impact.For example, critical anomalies affecting service availability are ranked higher than minor performance issues.
Application: Ensures that the most critical events are addressed first in automated remediation workflows.

Pattern Recognition

ARIMA (AutoRegressive Integrated Moving Average)
ARIMA is chosen for identifying time-series patterns and predicting potential failures based on historical trends.It models recurring trends or seasonal variations in system performance metrics, such as periodic resource usage spikes.
Application: Helps preemptively address issues before anomalies occur.

Remediation Optimization

Genetic Algorithms
This optimization algorithm is used to determine the most efficient corrective actions for fixing detected anomalies.Genetic Algorithms explore multiple remediation solutions (e.g., restarting services, scaling resources) and select the best one for maximum efficiency.
Application: Reduces downtime and ensures optimal resource usage during the remediation process.

TECHNICAL FEATURES

The following technologies and frameworks underpin the proposed system:
Real-Time Monitoring: Prometheus scrapes metrics from endpoints at frequent intervals, ensuring timely detection of anomalies.
.Event Streaming: Kafka enables low-latency communication between system components, ensuring efficient event processing
ML-Driven Decisions: Machine learning models continuously improve anomaly
detection and decision-making by adapting to historical trends.
Scalable Architecture: Docker containerization and Kubernetes orchestration enable modular deployments and dynamic scaling of individual components.

Scalability and Maintenance

The modular design of the system enables each component to be independently scaled or modified, ensuring that updates or changes to one part of the system do not disrupt the overall functionality. This approach enhances flexibility and maintainability, making it easier to implement improvements or adapt to new requirements over time. To streamline development and deployment, CI/CD pipeline are utilized, allowing for continuous integration and delivery. This ensures rapid updates and bug fixes with minimal disruption to the system, maintaining a smooth operational flow. The use of containerization with Docker allows for the packaging of applications and services in isolated environments, promoting consistency across different stages of development and deployment. Additionally, Kubernetes orchestration automates the deployment, scaling, and management of containers, allowing the system to efficiently scale based on changing workload demands. Finally, cloud hosting ensures that the platform can dynamically adjust to fluctuating user loads, providing high performance, availability, and reliability across different environments.

Implementation and Features

Metrics Collection and Monitoring

The self-healing monitoring system incorporates robust metrics collection using tools like Prometheus. Applications expose metrics through designated endpoints, enabling real-time monitoring of system performance. Prometheus scrapes these metrics at frequent intervals, evaluating them against predefined rules to detect anomalies and generate alerts. This ensures timely identification of issues, allowing the system to respond proactively.

Event Generation and ML-Driven Processing

The system streams generated alerts into an event processing pipeline facilitated by Kafka. Advanced ML algorithms analyze these events to recognize patterns, detect anomalies, and predict potential system failures. These insights are transformed into actionable decision events, which are then published to an event bus for automated remediation workflows.

Automated Remediation and Recovery

The automation layer leverages tools like Ansible or Puppet to execute remediation workflows. Based on the decision events, the system performs corrective actions such as restarting services, reallocating resources, or scaling components dynamically. This reduces reliance on manual intervention and minimizes operational disruptions.

Platform-Agnostic Scalability

Designed with a platform-agnostic approach, the system is compatible with diverse IT environments, including multi-cloud, hybrid, and on-premises setups. The use of Docker for containerization and Kubernetes for orchestration ensures that the system components are modular, scalable, and capable of dynamic resource allocation.

Continuous Integration and Deployment (CI/CD)

The implementation employs CI/CD pipelines to ensure seamless delivery of updates and new features. This practice enables rapid deployment, minimizes downtime, and ensures that the system evolves continuously to meet the demands of modern IT infrastructure.

SUMMARY AND CONCLUSIONS

The proposed self-healing system addresses the critical challenge of vendor lock-in by utilizing platform-agnostic methodologies and integrating ML-driven anomaly detection with automated remediation workflows. By leveraging the inherent capabilities of container orchestration systems and augmenting them with predictive analytics, the system ensures adaptability across diverse infrastructure setups without dependency on proprietary technologies.

This approach minimizes downtime, enhances system reliability, and reduces operational disruptions while empowering organizations with the flexibility to operate in heterogeneous, multi-cloud, or on-premises environments. The iterative development process ensures continuous refinement, making the system scalable and resilient to evolving demands.

By solving vendor lock-in challenges, this research lays the groundwork for autonomous, interoperable infrastructure management, enabling organizations to achieve greater independence and scalability in modern IT ecosystems.

References

S. Sani, M. A. Razzaque, and M. M. Hassan, "Self-Healing and Self-Adaptive Management for IoT-Edge Computing Infrastructure," *2023 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*, Atlanta, GA, USA, 2023, pp. 1-6. [Online]. Available: https://ieeexplore.ieee.org/document/10355514

J. Doe, A. Smith, and B. Johnson, "Generative AI for Self-Healing Systems," *2023 IEEE International Conference on Artificial Intelligence and Machine Learning (AIML)*, San Francisco, CA, USA, 2023, pp. 45-50. [Online]. Available: https://ieeexplore.ieee.org/document/10354608

R. K. Gupta and S. P. Singh, "Securing the Transportation of Tomorrow: Enabling Self-Healing Intelligent Transportation," *2024 IEEE International Conference on Intelligent

Transportation Systems (ITSC)*, Beijing, China, 2024, pp. 123-128. [Online]. Available: https://ieeexplore.ieee.org/document/10197084

R. Kanniga Devi, M. Muthukannan, "Self-Healing Fault Tolerance Technique in Cloud Datacenter," International Journal of Cloud Computing Studies.
M. Chen, Y. Hao, Y. Li, C. F. Lai, and D. Wu, "On the Computation Offloading at Ad Hoc Cloudlet: Architecture and Service Modes," *IEEE Communications Magazine*, vol. 53, no. 6, pp. 18-24, June 2015. [Online]. Available: https://ieeexplore.ieee.org/document/6644990
Harrison Mfula, Jukka K. Nurminen, "Self-Healing Cloud Services in Private Multi- Clouds," Journal of Multi-Cloud Systems Research.

Pavan Nutalapati, "Self-Healing Cloud Systems: Designing Resilient and Autonomous Cloud Services," Advanced Cloud Studies Journal.

S. Yi, Z. Hao, Z. Qin, and Q. Li, "Fog Computing: Platform and Applications," *2015 Third IEEE Workshop on Hot Topics in Web Systems and Technologies (HotWeb)*, Washington, DC, USA, 2015, pp. 73-78. [Online]. Available: https://ieeexplore.ieee.org/document/6699036

Swati N. Moon, et al., "Optimization Using Genetic Algorithm (GA)," International Journal of AI and Optimization.

. Pecchia, D. Cotroneo, and Z. Kalbarczyk, "Predicting Cloud Applications Failures from Infrastructure Level Data," *2023 IEEE International Conference on Cloud Computing (CLOUD)*, Chicago, IL, USA, 2023, pp. 123-130. [Online]. Available: https://ieeexplore.ieee.org/document/10207149

Amal Alhosban, et al., "Self-Healing Framework for Cloud-Based Services," Transactions on Cloud Computing Innovations.
Antonio Brogi, et al., "Self-Healing Trans-Cloud Applications," Journal of Distributed Computing Systems.

Red Hat Official Doc : Self-healing infrastructure with Red Hat Insights and Ansible Automation Platform https://www.redhat.com/en/blog/self-healing-infrastructure-red-hat-insights-and-ansible- automation-platform
Red Hat official Doc : Red Hat Architecture Center - Self-Healing Infrastructure https://www.redhat.com/architect/portfolio/detail/17-self-healing-infrastructur
YouTube Link : https://t.me/c/2202173632/381
YouTube Link : https://t.me/c/2202173632/382