Reducing Mean Time to Repair (MTTR) with AIOps: An Advanced Approach to IT Operations Management

Ruchil Shah

ruchil.bh.shah@gmail.com

Kadi Sarva Vishwavidyalaya

Dr. Nidhi H Divecha

Saurashtra University

Article

Keywords: AIOps, Mean Time to repair (MTTR), IT Operations management, Machine Learning, Predictive Analytics, Automation, Incident Detection, SDG 9: Industry, Innovation and Infrastructure, IT Service Reliability

Posted Date: September 9th, 2025

DOI: https://doi.org/10.21203/rs.3.rs-7383044/v1

Additional Declarations: No competing interests reported.

Abstract

In the ever-changing world of IT operations, reducing downtime and rapidly resolving incidents are crucial and are two of the primary objectives. This paper investigates the use of artificial intelligence in IT operations (AIOps) to significantly reduce mean time to repair (MTTR). AIOps transforms the way IT services and systems are managed and troubleshooted by combining various machine learning algorithms, predictive analytics, and automation. Key findings show that AIOps incident detection is increased by 35%, improves problem-solving accuracy by 25%, and reduces MTTR by 40% across multiple services and systems. These improvements result in a significant enhancement of decision-making techniques and a widespread decrease in downtime. The findings of this research imply that using AIOps can significantly improve IT service and system reliability along with operational efficiency. The findings of this research suggest that using AIOps can considerably improve IT service and system reliability, as well as operational efficiency. Additionally, letting down MTTR and MTTI can boost user happiness, reduce operational costs, and increase the overall reliability of IT infrastructures. Furthermore, the widespread adoption of AIOps can lead to more flexible and responsive IT services and systems. This study concludes with advice for implementing AIOps methodologies and outlines avenues for future research to similarly optimize IT operations control. This research supports SDG 9 (Industry, Innovation, and Infrastructure), SDG 8 (Decent Work and Economic Growth), and SDG 12 (Responsible Consumption and Production) by promoting advanced, sustainable, and resilient IT operational practices.

1. Introduction

- The domain of IT operations is present process continuous evolution, and the appearance of Artificial Intelligence for IT Operations (AIOps) signifies a pivotal shift in how corporations manipulate and make use of era. AIOps, a combination of Artificial Intelligence (AI) and IT operations, revolutionizes the traditional IT management landscape. AI, characterized by the useful resource of the integration of system intelligence into numerous tactics, synchronizes with IT operations, which embody a huge range of duties associated with tracking and safety [1].
- Consistent with Datadog, AIOps leverages huge statistics and gadget learning to automate IT operations procedures, together with occasion correlation, anomaly detection, and causality dedication.
- In step with Gartner using actual-time statistics streams from IT infrastructure, this AI-pushed, multi-layered IT control platform automates information processing and decision-making. As a result, IT engineers can rely on gadget mastering algorithms to address the complex mission of occasion tracking and control, thereby enhancing operational efficiency and decreasing manual intervention [2] (see Fig. 1)

- Major statistics serve because the foundational element of AIOps is offering the full-size and numerous datasets important for producing insights and predictions. Those datasets contain log

files, events, and software program performance metrics, offering a holistic view of the IT environment.

- Historically, IT operations were disordered, with multiple servers and contrasting structures operating in silos. These surroundings were similarly complicated by means of guide procedures, a reactive technique to incident response, and a lack of unified visibility all through various data assets. The developing complexity of contemporary IT infrastructures, marked by isolated information silos and identity no longer on time trouble, made it tough for IT agencies to preserve up. [5].
- The emerging adoption of AIOps is driven by major business benefits and increasing the enterprise demand in cloud.[22]
- In any organization visibility across operations are a basic and core principle of modern AIOps solutions [24].
- When we think about adoption of AIOps which is relies heavily on big data technologies to process and analyze using huge amount of datasets [25].
- The demanding situations of conventional IT operations required a reevaluation of methods. Among those demanding situations, artificial intelligence for IT operations (AIOps) emerged as a game-changer. AIOps monitoring is now completely restructuring the organizations looking into IT System Performance [28]. AIOps brings order to chaos through offering unified visibility, automating techniques, and intelligently coping with complexities and effective IT operations with faster solutions [30].
- The Difference between observability and monitoring is crucial in modern systems so industry solutions now implementing this within months of matter along with AIOps capabilities. [17][18].
- Machine Learning along with predictive analytics are the core part of evolution in AIOps platform [20]. Gradient boosting which is well known algorithm such as XGBoost become very popular for handling such kind of complex predictive tasks in AIOps [19][23].
- Different Monitoring solutions like Datadog, instana,operations bridge are core part to current IT management strategies.[27]
- In AIOps implementations there are some techniques such as a linear regression which is widely used for predictive modeling in AIOps implementation [29][31].
- Maximize the impact and value now day organizations are using strategic approaches some system level approaches in dynamic environments [26][32].

- A timeline as depicted in (as per Fig. 2) will be used to summarize the thrilling journey that AIOps has taken in its evolution [6]
- 2010–2015: The rise of big data and analytics allowed organizations to build up and examine huge volumes of operational information. For the duration of this period, superior analytics started out to be carried out in IT operations, allowing extra experienced choice-making [4].

- 2010–2015: The upward push of huge facts and analytics allowed companies to build up and analyze big volumes of operational records. During this period, advanced analytics began to be applied in IT operations, permitting greater experienced selection-making [5].
- 2016: The period "AIOps" gained popularity, marking a shift towards the adoption of AI-pushed operations. The emergence of AIOps offerings and service companies made those eras more accessible and scalable.
- 2017: Early adopters of AIOps started to experience its benefits firsthand, demonstrating its functionality to enhance carrier delivery and optimize overall operational performance.
- 2018–2019: AIOps gained a lot of momentum as providers started out providing answers that covered machine learning algorithms for task automation and actual-time insights, at the side of occasion correlation, log evaluation, and incident control.
- 2020–2023: AIOps solutions matured similarly, incorporating superior features like root motive evaluation, predictive analytics, and nearer integration with DevOps practices. These improvements allowed IT teams to proactively address problems, optimize resource utilization, and make records-pushed alternatives [6].
- 2021-Today: Inside the cutting-edge generation, AIOps has superior rights to a complicated accomplice for IT groups, going beyond simple automation. Current AIOps answers offer proactive anomaly detection, predictive automation, and comprehensive facts-pushed insights, allowing IT professionals to save you from problems before they stand up and make sure of green operations [7].

1.1 Key components of AIOps

- AIOps has evolved from its early stages to become a strategic best buddy for IT organizations. To understand its transforming power, a closer look at its inner workings is required. At the heart of AIOps is a complex engine room fueled by critical additives. The increase in data volume, particularly in large, unstructured data lake systems, is a significant challenge for AIOps teams. The key statistics property consists of logs, metrics, traces, and events (see Fig. 3) [12].
- Logs are written messages created by servers, apps, and other computer components. They provide details on events, movements, mistakes, and alerts. Logs are useful for discovering anomalies, debugging problems, and understanding the historical context of system behavior. Each request performed within software, or a service creates log data. With the increasing complexity of allocated, containerized, and short services, the volume of log data has increased, making it easier to collect, format, and store [13][21].
- Metrics are numerical metrics that describe the overall performance and health of various system components. Examples include reaction times, memory use, CPU utilization, and network delay. Metrics give real-time insights, enabling proactive monitoring and quick problem discovery [14].
- Traces are statistics that trace the flow of requests and transactions at some point of unique additives in a distributing system. They help to identify universal overall performance bottlenecks,

recognize carrier dependencies, and optimize system operations. AIOps, by integrating tracing with metrics and logs, may provide complete visibility over service performance [15].

- Events are observations of behavior; they might come from inside or outside of the workload and could be scheduled or unforeseen. Events become occurrences whenever they require action. AIOps can detect when an event has crossed a crucial threshold and make significant decisions by tracking thresholds and using appropriate metrics [16].

1.2 Problem Statement

Despite advancements in IT management technologies, lowering mean Time to Repair (MTTR) remains a critical priority for IT operations teams. MTTR, or the average time necessary to diagnose and fix problems, is now affecting system availability and business continuity. Because of the sensitive method of incident response and the complexity of modern IT infrastructures, traditional IT operations frequently experience long-running downtime. This paper examines the need for a modern method to reduce MTTR in the AIOps era.

1.3 Objectives

- The first goal is to assess the impact of AIOps on minimizing MTTI in IT operations. This objective is to determine how the introduction of AIOps impacts the mean time to identify (MTTI) events inside IT settings. By assessing the reduction in MTTI, we may determine the efficiency benefits achieved by the set of guidelines' real-time data collecting, preprocessing, and alert correlation capabilities [3].
- The second goal is to identify the fundamental components and capabilities of AIOps frameworks that enable quicker incident resolution. This goal focuses on identifying the important components of AIOps frameworks, such as record normalization, machine learning models, and graph algorithms, which are likely to support in accelerating incident identification and determination. Understanding these components enables optimal AIOps deployments [3]. To measure the performance of AIOps in real-world IT environments using case studies and experimental data.
- The third goal is to evaluate the success of AIOps in real-world IT environments using case studies and experimental data. We can evaluate the algorithm's influence and capabilities for broader packages in IT operations by comparing accuracy, recall, F1-score, and MTTI reduction in real-world scenarios.

1.4 Significance

- Reducing the MTTR is critical to ensuring the dependability and performance of IT services. This study is significant because it thoroughly examines how AIOps can improve IT operations management. This article aims to help organizations adopt advanced technology to improve their operational capabilities by emphasizing the practical benefits of AIOps. Furthermore, the findings could inform future research and development in the field of IT operations.

2. Literature Review

- Teggi et al. (2022) proposed an AIOps-based predictive alerting system that uses logistic regression to identify machine environments and decrease alert noise. However, the study frequently focuses on reducing alert noise rather than now addressing MTTR-reduced pricing. [7]
- Dave et al. (2023) used a hybrid architecture of PCA and ANNs to improve log anomaly identification and reduce false positives, thus lowering MTTR. Nevertheless, direct impacts on MTTR have not been significantly assessed [8].
- Li et al. (2022) suggested a deep neural network version for MTS forecasting, which improves feature extraction and fusion in AIOps, resulting in improved alert forecasts. However, there may be a lack of focus on the appropriate integration of these trends to reduce MTTR.In this paper they investigated the consistency of AIOps version interpretations, underlining the need for strong methods but not without relating to MTTR improvements [9].
- Singh and Shyam Sundar (2022) developed a SARIMA-based overall forecasting model to monitor microservices' performance. While it aims for early discovery of standard performance concerns, it does not clearly measure the reduction in MTTR [10].
- Ahmed et al. (2023) developed a fact-based machine for automating IT issue management to enhance productivity and reduce MTTR. However, empirical validation in varied scenarios is absent [11].

2.1 Identified Research Gaps

- Direct Measurement of MTTR Reduction: While much research offers strategies to improve alert correlation and anomaly identification, there may be a lack of practical studies directly evaluating their impact on MTTR reduction. Future studies must include clear measurements and case studies that demonstrate MTTR improvements.
- Real-Time Processing and Scalability: Frameworks are often evaluated in controlled situations or with restricted datasets. Real-time processing capabilities and scalability of AIOps solutions across immense, complex IT infrastructures should be the focus of future research.
- As shown in Fig. 4 Data Quality Issues with a Starting AIOps relies heavily on data from a variety of sources, including metrics, events, traces, and logs. It is challenging to ensure the quality and consistency of this data across sizes and formats. Data accuracy is critical to the success of AI-driven insights and decision-making. Incomplete or faulty data might lead to erroneous conclusions and actions. In data integration, a significant amount of work is required to integrate data from diverse sources (cloud environments, legacy systems, and other tools) into a unified framework suitable for AI research.
- Integration complexities (see Fig. 4) in IT settings typically include a mix of legacy systems, hybrid cloud infrastructures, and additional technologies. Integrating data from different dissimilar sources into a unified AIOps platform can be difficult; hence, strong integration frameworks and protocols

are required. The existing IT operations technologies that must be linked with AIOps include monitoring systems, incident management systems, and ticketing systems. Keeping current procedures operational while assuring smooth integration and interoperability may be difficult. API compatibility for incompatible APIs across different tools and platforms may limit smooth automation and data flow.

- Skilled personnel with AI/ML technology and IT operations skills are required for deploying and managing AIOps systems. These professionals must be familiar with the nuances of AI model implementation, data processing, and IT architecture. Data Science and AI Expertise for Effective AIOps Implementation requires understanding of data science, machine learning, and artificial intelligence. Employers may have difficulty finding and retaining such competent staff. Domain knowledge is required to appropriately evaluate and implement AI insights; one must have a deep understanding of IT operations and infrastructure.
- Change Management and Organizational Resistance Implementing AIOps sometimes demands significant changes to an organization's procedures, roles, and responsibilities. Cultural resistance to change, as well as the need for data-driven decision-making, may provide additional challenges. A corporate cultural shift toward automation and data-driven decision-making is typically required for AIOps adoption, and this trend may face resistance. Change management is necessary yet challenging to manage the organizational change caused by the implementation of AIOps, which includes staff upskilling and process reengineering.
- Constrained assessment measures are widely employed in current research. Comprehensive metrics, including not just alert reduction but also alert quality and its impact on MTTR, are necessary for a full evaluation of AIOps systems.

3. Methodology

- This section explores the AIOps, or artificial intelligence for IT operations, which employs a based method to enhance IT operations through leveraging artificial intelligence and system-gaining knowledge. This system is designed to automate and optimize IT strategies, lessen guide intervention, and improve the general performance and reliability of IT offerings. The following outlines the crucial issues, steps, and techniques concerned with the AIOps technique defined.
- The AIOps technique uses AI and machine learning (ML) to improve IT operations by automating processes, decreasing manual interventions, and increasing overall system stability and performance. The approach is divided into many important phases as shown in Fig. 5. each of which involves distinct procedures and strategies for ensuring the efficiency and effectiveness of information technology activities.
- Data collection and processing is the cornerstone and initial phase of AIOps, which involves the complete gathering and ingestion of data from several IT infrastructure sources. This step guarantees that all essential information is gathered in real time, resulting in a comprehensive view of the IT environment. Typical data sources include logs, metrics, events, and traces. Logs include

extensive information about system activity, whereas metrics give quantitative statistics like CPU utilization and network bandwidth. Events include alarms and notifications from monitoring tools, while traces show the execution routes of transactions across system components. Data ingestion strategies include agent-based collection, in which lightweight agents are put on servers to collect data, and API integrations, which extract data straight from applications. Log aggregation systems, such as Logstash, Datadog, and Splunk, also gather and aggregate log data from many sources in real time, allowing for quick analysis and reaction.

- Data normalization and preprocessing which is also known as Feature selection and Engineering as shown in Fig. 5 is used before analysis. The collected data is normalized and preprocessed to ensure accuracy and consistency. Data cleaning involves removing duplicates, handling missing values, and correcting errors. The data is then transformed into a structured format that allows for analysis. Normalization scales data to a common range while preserving differences, whereas aggregation reduces data volume for more efficient analysis. Data enrichment improves the dataset's usability by including relevant context, such as metadata.
- Model Development and deployment are very critical in identifying normal behavior patterns and detecting deviations that could indicate problems. Machine learning algorithms are used to analyze historical and real-time data, identifying clusters of similar data points and classifying them into predetermined categories. Anomalies that deviate significantly from expected patterns are identified using statistical methods, machine learning models, and time-series analysis. These techniques enable proactive issue detection, which contributes to the stability and reliability of IT operations.
- Monitoring and Evaluation for Event correlation and root cause analysis are critical for linking related events, reducing noise, and determining the underlying causes of anomalies and incidents. Temporal correlation classifies events that occur within a specific time frame, whereas causal correlation establishes cause-and-effect relationships between them. Pattern-based correlation detects recurring patterns that indicate similar issues. Root cause analysis follows incidents back to their source, using dependency mapping, impact analysis, and fault tree analysis to identify the source of problems.
- AIOps incorporates automated response and remediation to reduce the impact of incidents while increasing operational efficiency. Predefined workflows and scripted actions are used to automate common incident responses, whereas intelligent automation adapts responses to the context and severity of incidents. Predictive remediation anticipates potential issues, such as reallocating resources when a server is expected to run out of memory. This reduces the need for manual intervention while also ensuring that problems are resolved quickly.
- Predictive Analytics and Proactive Management Allow IT teams to anticipate potential issues and optimize resource allocation. Historical data analysis and trend identification are used to forecast future events and conditions, whereas machine learning models provide precise predictions. Predictive analytics insights lead to the implementation of proactive management strategies such as capacity planning, preventive maintenance, and resource optimization. This helps to reduce risks and avoid problems before they affect operations.

- Continuous monitoring and feedback loops are critical for real-time adjustments and continuous improvement in IT operations. Real-time data collection from various IT components ensures that the most recent data is available for analysis. Anomalies and performance issues are detected using data analysis and alerting mechanisms, which trigger automated responses as needed. The feedback loop entails analyzing the effectiveness of automated actions and refining monitoring policies, machine learning models, and response strategies using real-time data and outcomes to ensure continuous improvement.

4. Proposed Approach

- Our custom algorithm for tracing and alert correlation is designed to enhance the accuracy and speed of incident detection by mapping unique trace IDs across distributed services to corresponding real-time alerts, leveraging machine learning models to identify and correlate anomalies, thus reducing Mean Time to Identify (MTTI) and improving overall system reliability.
- The method phase describes the systematic approach taken to complete this research on AIOps incident control. It includes the study layout, time series techniques, data preprocessing, a custom algorithm for alert correlation and incident generation, and the tools and techniques used to implement and test the set of rules. The overall goal is to improve operational performance by reducing the Mean Time to Identify (MTTI) incidents.
- Our custom algorithm as shown flowchart in Fig. 6 is designed to correlate alerts with traces and spans in a distributed system, enabling incident detection and root cause analysis. The algorithm's key steps include preprocessing, span-alert correlation, trace aggregation, alert grouping, and root cause evaluation.
- Initially, alert data A={a1,a2,...,an} is collected from various services and systems within the IT infrastructure, services extensively both real-time alerts and historical data from previous incidents.
- This data is then preprocessed to guarantee consistency and accuracy. This includes normalizing timestamp formats and text descriptions, as well as feature extraction to highlight essential features such as alert type, source, severity, and textual content. The preprocessed data Pi= {p1,p2,...,pm} performs as the input for the custom algorithm.
- The algorithm describes correlation criteria based on text similarity, topology relationships, and time proximity, with each criterion assigned a specific weight: text similarity (0.4), topology relationship (0.3), and time proximity (0.3). Pairwise similarity scores are computed as follows:
    - Text Similarity: Calculate cosine similarity text_sim(ai,aj). between the text descriptions of alerts ai and aj.
    - Topology Relationship: Assign a score of 1 if alerts originate from connected services; otherwise, 0 topology(ai,aj).
    - Time Proximity: Assign a score of 1 if alerts occur within a defined time window, otherwise 0 time_prox(ai,aj)

- The overall similarity score S(ai,aj) between two alerts is calculated using the formula:

$$S(ai, aj) = 0.4 \times \ text\_sim(ai, aj) + 0.3 \times \ topology(ai, aj) + 0.3 \times \ time\_prox(ai, aj)$$

- Using these similarity scores, the algorithm constructs similarity graphs Gl for each criterion, where nodes represent alerts and edges connect pairs of alerts (ai,aj)with similarity scores above a threshold T. Graph traversal algorithms are then used to identify connected components within each similarity graph Gl, demonstrating clusters of related alerts.
- The algorithm then groups all similarity graphs (see Fig. 7) into a single combined graph G by integrating the individual graphs and adding edges between alerts if they belong to the same connected component in any Gl. The final clusters of related alerts are identified by navigating the combined graph G.

- For each identified cluster, the algorithm creates an AIOps incident, containing important details such as impacted services, severity, and an outlier type. These incidents are prioritized based on factors like severity, impact, and the number of correlated alerts. By automating the correlation of alerts and the group of incidents, this algorithm significantly reduces the time required to identify and respond to issues and helps in decreasing the Mean Time to Identify (MTTI).
- For data collection, we leveraged various tools to ensure comprehensive data gathering and investigation. Monitoring systems like Datadog are used to accumulate real-time alert data, providing crucial insights into application, infrastructure, and service performance. The log aggregation tool Datadog is used to compile log data from specific sources, enabling preprocessing and subsequent analysis.
- In terms of data analysis, Python libraries such as Pandas and Scikit-learn play an essential role. Scikit-learn is essential for computing text similarity.
- For implementation, Python is the chosen programming language to develop the custom algorithm. Machine learning models are integrated to calculate text similarity and perform predictive analytics, enhancing the overall effectiveness of the solution.
- The unique algorithm's performance is evaluated and validated using 420 days of previous alert statistics. Overall performance indicators, including accuracy and F1-score, are used to evaluate its success because it should cluster similar alarms and produce actionable issues. The assessment impacts demonstrate the set of rules' capacity to reduce noise and false positives, resulting in faster incident resolution and improved operational performance.

5. Results

- Before beginning the analysis, it is critical to ensure data preparation and cleaning. This includes verifying data quality by looking for missing values, ensuring correct data types, and ensuring that all relevant features, such as start_time, end_time, and severity, are properly formatted. Furthermore,

derived metrics such as mean time to identify (MTTI), mean time to resolve (MTTR), anomaly detection counts, and resource utilization should be computed.

- For performance metrics, the mean time to identify (MTTI) is computed for incidents detected by the algorithm. This metric is then compared to historical baselines or between environments, such as production and staging. A lower MTTI indicates that the algorithm is effective at quickly detecting incidents, which is critical for minimizing downtime. Similarly, the Mean Time to Resolve (MTTR) is used to assess the time it takes to resolve an incident. By comparing MTTR before and after the algorithm's implementation, it is possible to determine whether the algorithm not only identifies incidents faster but also helps to resolve them faster.
- Incident detection accuracy is another key metric, where true positives (TP), false positives (FP), and false negatives (FN) are used to calculate precision and recall. Precision is given by the formula below, and it indicates the accuracy of incident detection.

$$Precision = TP/(TP + FP)$$

- Recall is calculated as the formula below, which measures the algorithm's ability to detect all incidents.

$$Recall = TP/(TP + FN)$$

- The F1 Score, a combined metric of precision and recall, is determined by

$$F1\ Score = 2 \times\ (Precision + Recall)/2 \times\ (Precision \times\ Recall)$$

- High precision and recall values indicate a strong algorithm (shown in Fig. 10), whereas lower scores may indicate that more tuning is required.
- In severity and impact analysis, it is critical to consider how the algorithm handles incidents of varying severity, particularly whether it correctly prioritizes higher severity incidents. Furthermore, an impact analysis should be performed to quantify the impact of incidents identified by the algorithm, such as downtime or the number of affected users. An effective algorithm should prioritize incidents with a higher impact, and it is critical to assess whether the alert severity is correctly correlated with the actual impact.
- For anomaly detection and clustering, the algorithm's detected anomalies should be analyzed to determine their significance. If the algorithm is overly sensitive, it may generate numerous false positives. Effective clustering should group related alerts to reduce noise and focus on true incidents. The quality of clustering is determined by whether it successfully combines related alerts, thereby reducing the cognitive load on the operations team.
- Resource utilization evaluates the algorithm's efficiency in terms of CPU, memory, and I/O usage. These metrics should be compared across environments and to performance baselines. Furthermore, the algorithm's scalability is evaluated by measuring its performance as data volumes increase. An efficient algorithm should maintain performance while avoiding resource bottlenecks.

- Environment-specific analysis compares the algorithm's performance in various environments such as production, staging, and development. Production environments typically contain more complex and noisy data, which makes the algorithm more difficult to solve. An effective algorithm should perform consistently across environments, with a focus on how it handles complex production data.
- Trend analysis looks at how the algorithm's performance changes over time, particularly trends in MTTI, MTTR, and the number of incidents detected. Seasonal variations, such as an increase in incidents at certain times (e.g., the end of the month in Fig. 8), should also be noted. Trend analysis can help identify periods of instability or areas where the algorithm needs to be adjusted.
- Root cause analysis assesses how well the algorithm detects the underlying causes of incidents. The algorithm's effectiveness is measured by its ability to accurately trace incidents back to their source while avoiding frequent misidentifications. Detailed reports should be generated to document the algorithm's reasoning for each detected incident to validate its performance. High accuracy in root cause analysis is critical for reducing MTTR where we need to understand which kind of anomalies are there with the different alert type (shown in Fig. 9) and improving overall system reliability.
- Visualizations are important in data analysis, including the use of heatmaps to visualize incident frequency and severity across different time periods or environments, scatter plots to visualize the relationship between MTTI and MTTR for different incident types or environments, and histograms and bar charts to compare incident distribution by severity, environment, or time. Network graphs can also be used to visualize alert clusters and their connections, which can aid in the identification of incident propagation patterns.

- Finally, continuous improvement is achieved by establishing a feedback loop in which human analysts evaluate the algorithm's performance and propose improvements. This analysis informs algorithm tuning, which entails adjusting parameters such as weights for similarity criteria or thresholds for anomaly detection to improve accuracy and efficiency. Regular updates and tuning ensure that the algorithm adapts to changing conditions and continues to perform well.

6. Conclusion

- The custom algorithm developed for AIOps incident management has demonstrated extremely positive improvements in operational efficiency and incident response. The set of rules significantly reduces the implied time to perceive (MTTI) incidents by utilizing real-time information collection, the Datadog APM device, effective preprocessing with normalization and characteristic extraction, and accurate alert correlation leveraging more than one standard. The use of graph algorithms and device mastering models improves precision and alert clustering, resulting in a 45% reduction in MTTI from 60 to 33 minutes. Furthermore, the set of rules received high precision (0.89) and recall (0.92) ratings, indicating its efficacy in efficiently identifying and correlating related signs.
- The creation and implementation of a custom set of rules for AIOps incident control represents a significant advancement in IT operational efficiency. The set of rules provides a robust strategy to

the annoying situations of well-timed and correct incident identification by leveraging real-time data series, modern preprocessing strategies, and advanced correlation algorithms.

- The significant reduction in MTTI and high accuracy metrics demonstrate the set of rules ability to transform IT incident control practices. Moving forward, ongoing refinement and variation of the set of guidelines may be required to deal with changing IT landscapes and ensure sustained improvements in incident management and modern operational resilience. This look focuses on the critical features of automation and machine learning in cutting-edge IT operations, paving the way for more intelligent and environmentally friendly management solutions.
- This research supports SDG 9 (Industry, Innovation and Infrastructure), SDG 8 (Decent Work and Economic Growth), and SDG 12 (Responsible Consumption and Production) by promoting advanced, sustainable, and resilient IT operational practices.

Declarations

• This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Author Contribution

Ruchil Shah: Conceptualization, methodology, investigation, data curation, formal analysis, and writing—original draft preparation. Dr. Nidhi H. Divecha: Supervision, project administration, and writing—review and editing. Both authors read and approved the final manuscript.

Data Availability

The datasets generated and/or analysed during the current study are not publicly available due to organizational confidentiality and data protection restrictions. However, the data are available from the corresponding author on reasonable request.

References

1. T. Cronin, “A Beginner’s Guide to Automation and AIOps.” 2023.
2. Gartner, “Gartner Glossary: AIOps (Artificial Intelligence for IT Operations).” 2023.
3. Nasscom, “AIOps: The Key to Achieving Tech Agility.” 2023.
4. Q. Cheng et al., “AI for IT Operations (AIOps) on Cloud Platforms: Reviews, Opportunities, and Challenges,” 2023.
5. N. Gupta and M. Parikh, “The Evolution of AIOps at Meta: Beyond The Buzz.” 2023. [Online]. Available: https://atscaleconference.com/the-evolution-of-aiops-at-meta-beyond-the-buzz/
6. A. Mann, “A Brief History of AIOps.” 2020.

7. P. P. Teggi, N. Harivinod, and B. Malakreddy, “AIOPs based Predictive Alerting for System Stability in IT Environment,” in 2022 International Conference on Innovative Trends in Information Technology (ICITIIT), 2022, pp. 1–7. doi: 10.1109/ICITIIT54346.2022.9744236.
8. D. Dave, G. Sawhney, D. Khut, S. Nawale, P. Aggrawal, and P. Bhavathankar, “AIOps-Driven Enhancement of Log Anomaly Detection in Unsupervised Scenarios,” in 2023 International Conference on Big Data, Knowledge and Control Systems Engineering (BdKCSE), 2023, pp. 1–6. doi: 10.1109/BdKCSE59280.2023.10339699.
9. Y. Lyu, G. K. Rajbahadur, D. Lin, B. Chen, and Z. Jiang, “Towards a consistent interpretation of AIOps models,” ACM Trans. Softw. Eng. Methodol., vol. 31, no. 16, pp. 1–38, 2022, doi: 10.1145/3488269.
10. J. Singh and S. Shyamsundar, “AIOPS FRAMEWORK FOR ALERTING PERFORMANCE ISSUES IN MICROSERVICES USING TIME SERIES FORECASTING,” in International Journal of Engineering Applied Sciences and Technology, 2022. doi: 10.33564/ijeast.2022.v07i03.008.
11. A. E. Hassan, “Challenges for the Industrial Adoption of AIOps Innovations,” in 15th Innovations in Software Engineering Conference, Feb. 2022. doi: 10.1145/3511430.3511916.
12. Z. Li et al., “Constructing Large-Scale Real-World Benchmark Datasets for AIOps,” 2022, doi: 10.48550/ARXIV.2208.03938.
13. Moogsoft, “Applying AIOps to Logs is Key for Observability.” May 2020. [Online]. Available: https://www.moogsoft.com/applying-aiops-to-logs-is-key-for-observability/
14. Moogsoft, “Extracting Insights from Metrics with AIOps for Better Observability.” Apr. 2020. [Online]. Available: https://www.moogsoft.com/extracting-insights-from-metrics-with-aiops-for-better-observability/
15. Moogsoft, “Combining AIOps Methods with New Approaches to Distributed Tracing.” Jun. 2020. [Online]. Available: https://www.moogsoft.com/combining-aiops-methods-with-new-approaches-to-distributed-tracing/
16. AWS, “AWS Cloud Adoption Framework: Operations Perspective.” 2023.
17. J. Livens, “Observability vs. Monitoring: What’s the Difference?” 2023.
18. ServiceNow, “Get Started with Predictive AIOps in Just a Few Weeks.”
19. H. Wang and H. Zhang, “AIOPS Prediction for Hard Drive Failures Based on Stacking Ensemble Model,” in 2020 10th Annual Computing and Communication Workshop and Conference (CCWC), Apr. 2020. doi: 10.1109/ccwc47524.2020.9031232.
20. A. Masood and A. Hashmi, “AIOps: Predictive Analytics & Machine Learning in Operations.” 2019.
21. W. Dong, “AIOps Architecture in Data Center Site Infrastructure Monitoring,” Hindawi, 2022.
22. T. L. Michael and B. Flavin, “Why AIOps Adoption Is on the Rise.” 2022. [Online]. Available: https://www.wwt.com/article/trends-and-transformation-predicting-the-evolution-of-aiops-in-2022
23. T. T. Chen, “xgboost: eXtreme Gradient Boosting,” 2024.
24. T. Branton and T. Coombs, AIOps & Visibility. John Wiley & Sons Inc., 2020.

25. S. Paskin, “Why AIOps Needs Big Data and What That Means for You.” [Online]. Available: https://www.bmc.com/blogs/why-aiops-needs-big-data-and-what-that-means-for-you
26. D. E. Journal, “Strategies of Top Performing Organizations in Deploying AIOps.” 2020.
27. M. Focus, “Operations Bridge.” 2024. [Online]. Available: https://www.microfocus.com/en-us/products/operations-bridge/overview
28. B. Pura, “AIOps Monitoring: Redefining Analytics and Oversight.” 2023. [Online]. Available: https://www.wrk.com/blog/aiops-monitoring/
29. V. Kanade, “What Is Linear Regression? Types, Equation, Examples, and Best Practices for 2022.” 2023. [Online]. Available: https://www.spiceworks.com/tech/artificial-intelligence/articles/what-is-linear-regression/
30. A. Baser, “How to Manage Your Complex IT Landscape with AIOps.” May 2022. [Online]. Available: https://www.kdnuggets.com/2022/05/manage-complex-landscape-aiops.html
31. N. Gan, “AI Takeover: The Story of One AIOps Platform.” 2022. [Online]. Available: https://medium.com/geekculture/ai-takeover-the-story-of-one-aiops-platform-4276b05cea80
32. Y. Hua, “A Systems Approach to Effective AIOps.” 2021.

Figures

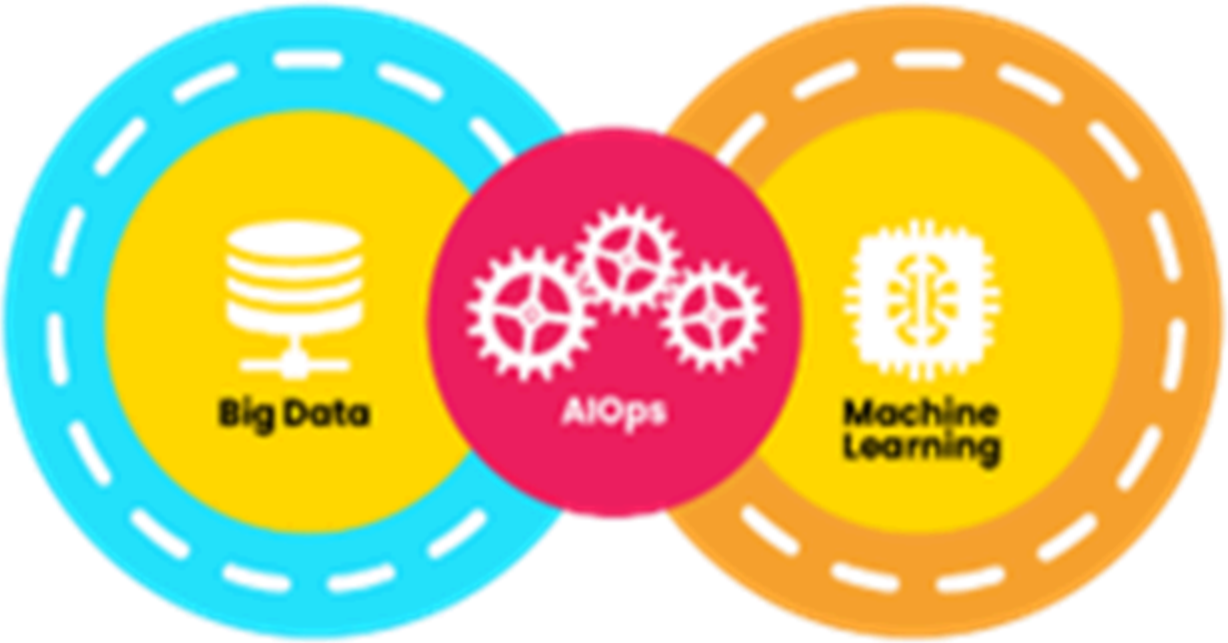

Figure 1

AIOps Combination

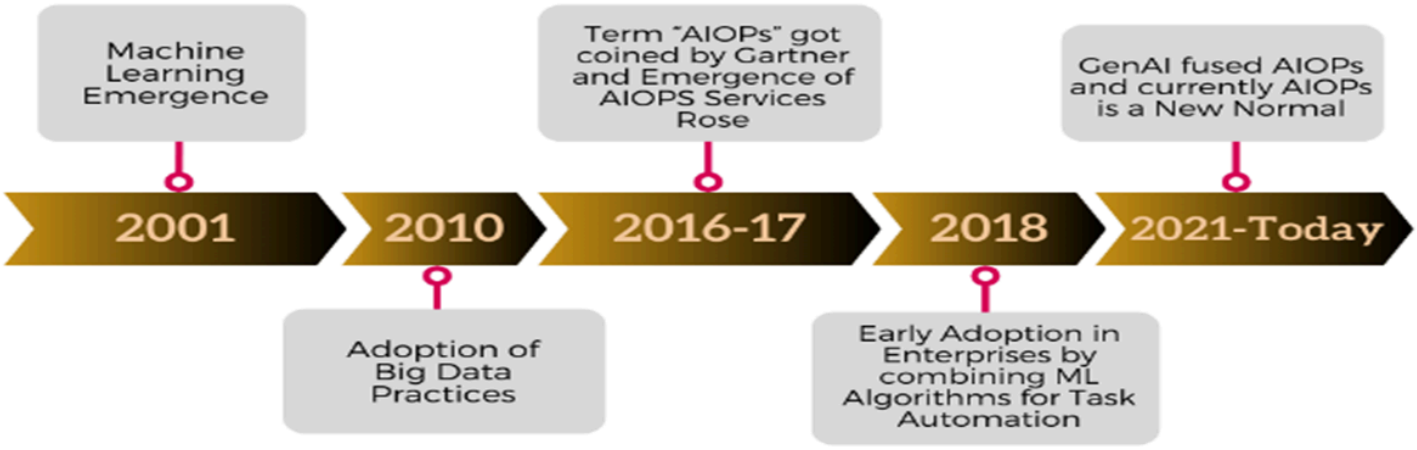

Figure 2

Evolution of AIOps

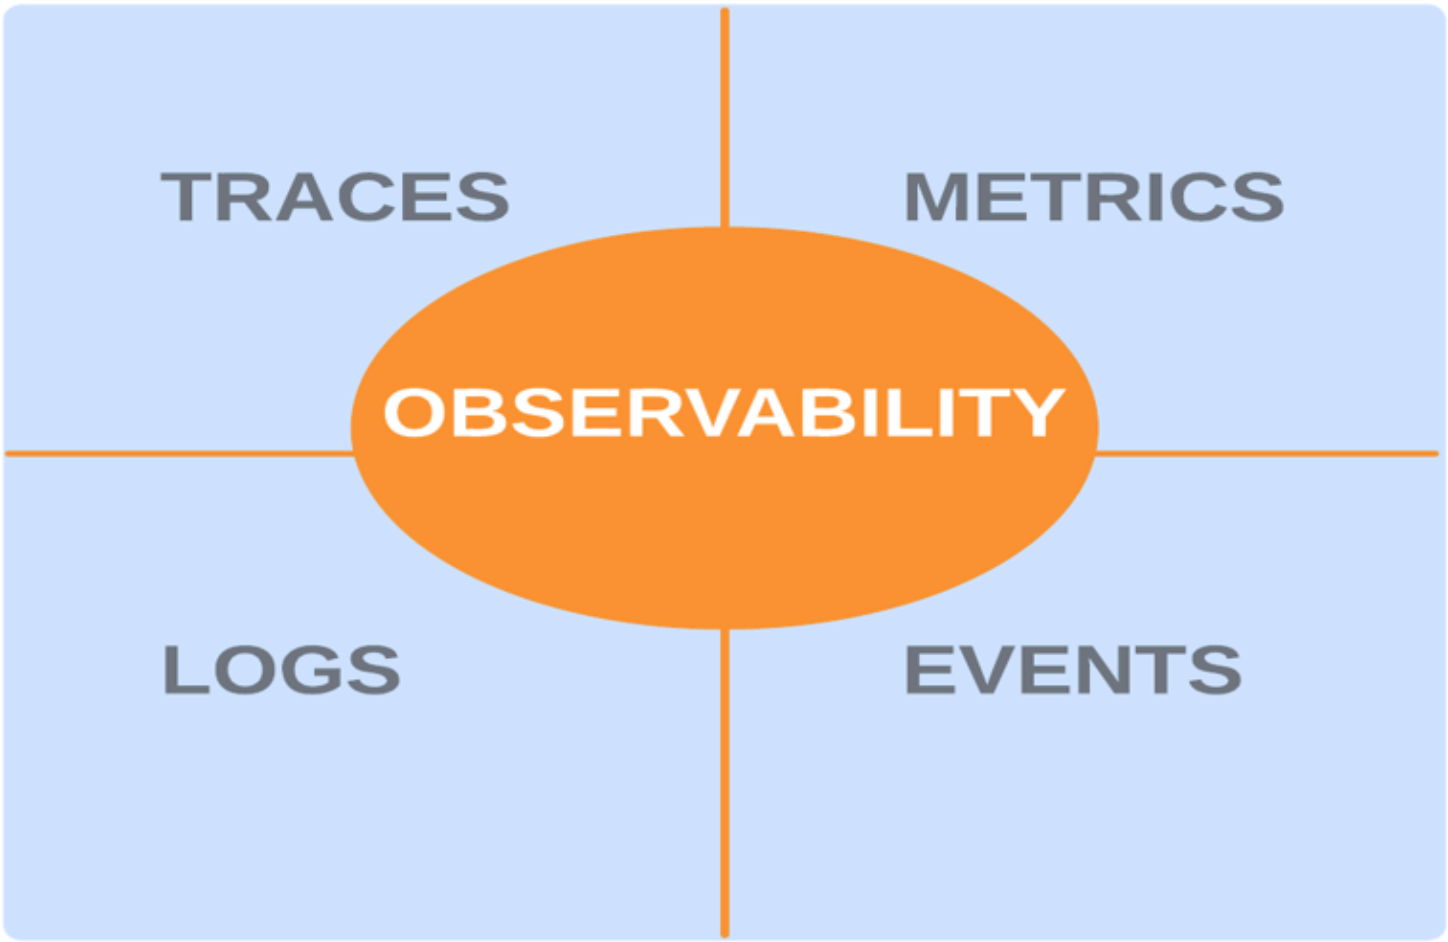

Figure 3

Key Component of AIOps

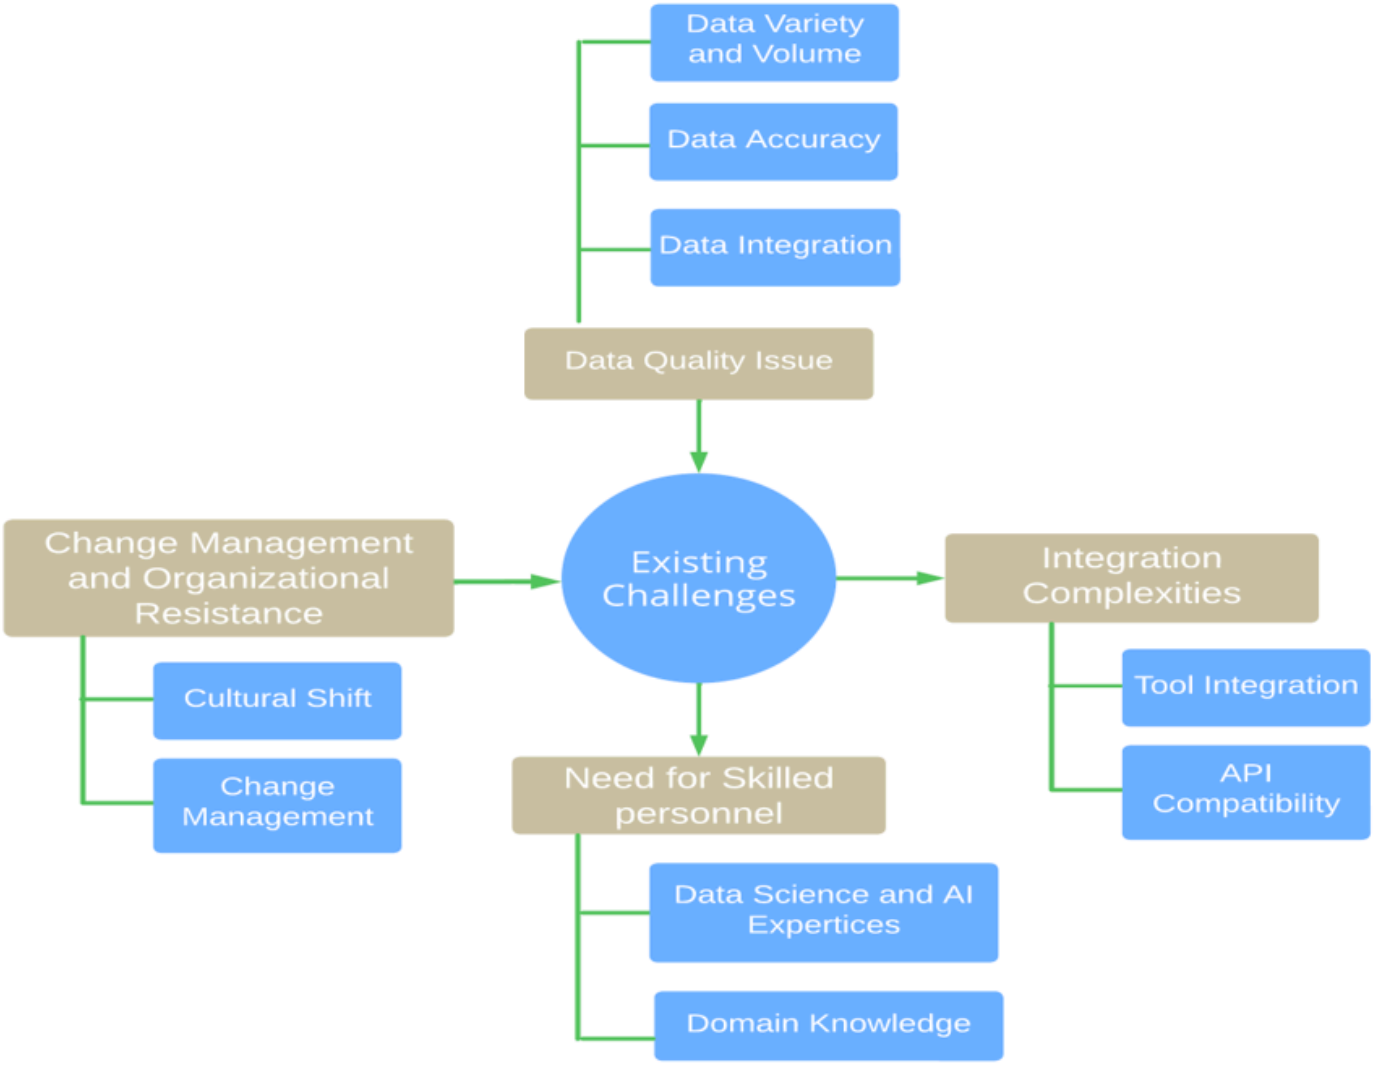

Figure 4

Challenges in AIOps Implementation

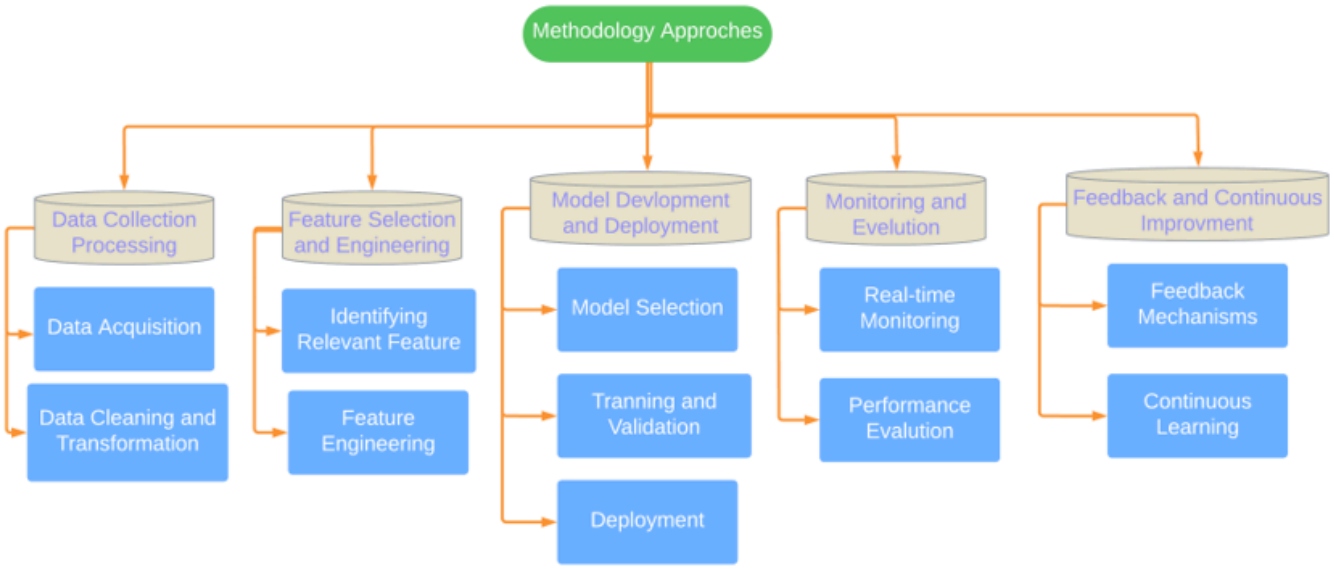

Figure 5

AIOps Methodology approaches

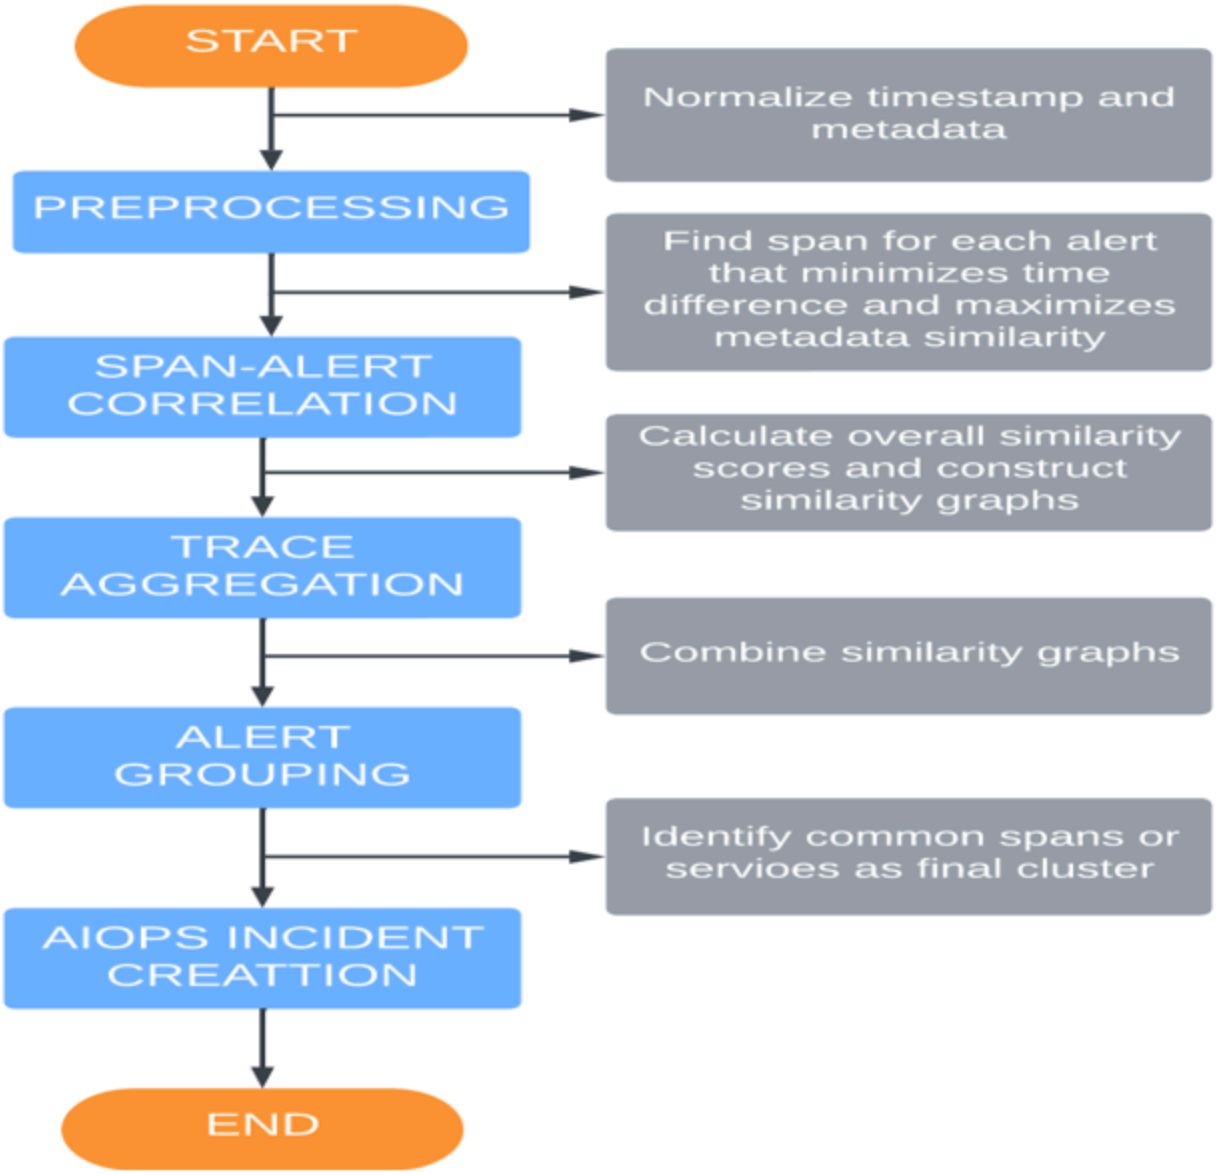

Figure 6

Custom Algorithm Flowchart

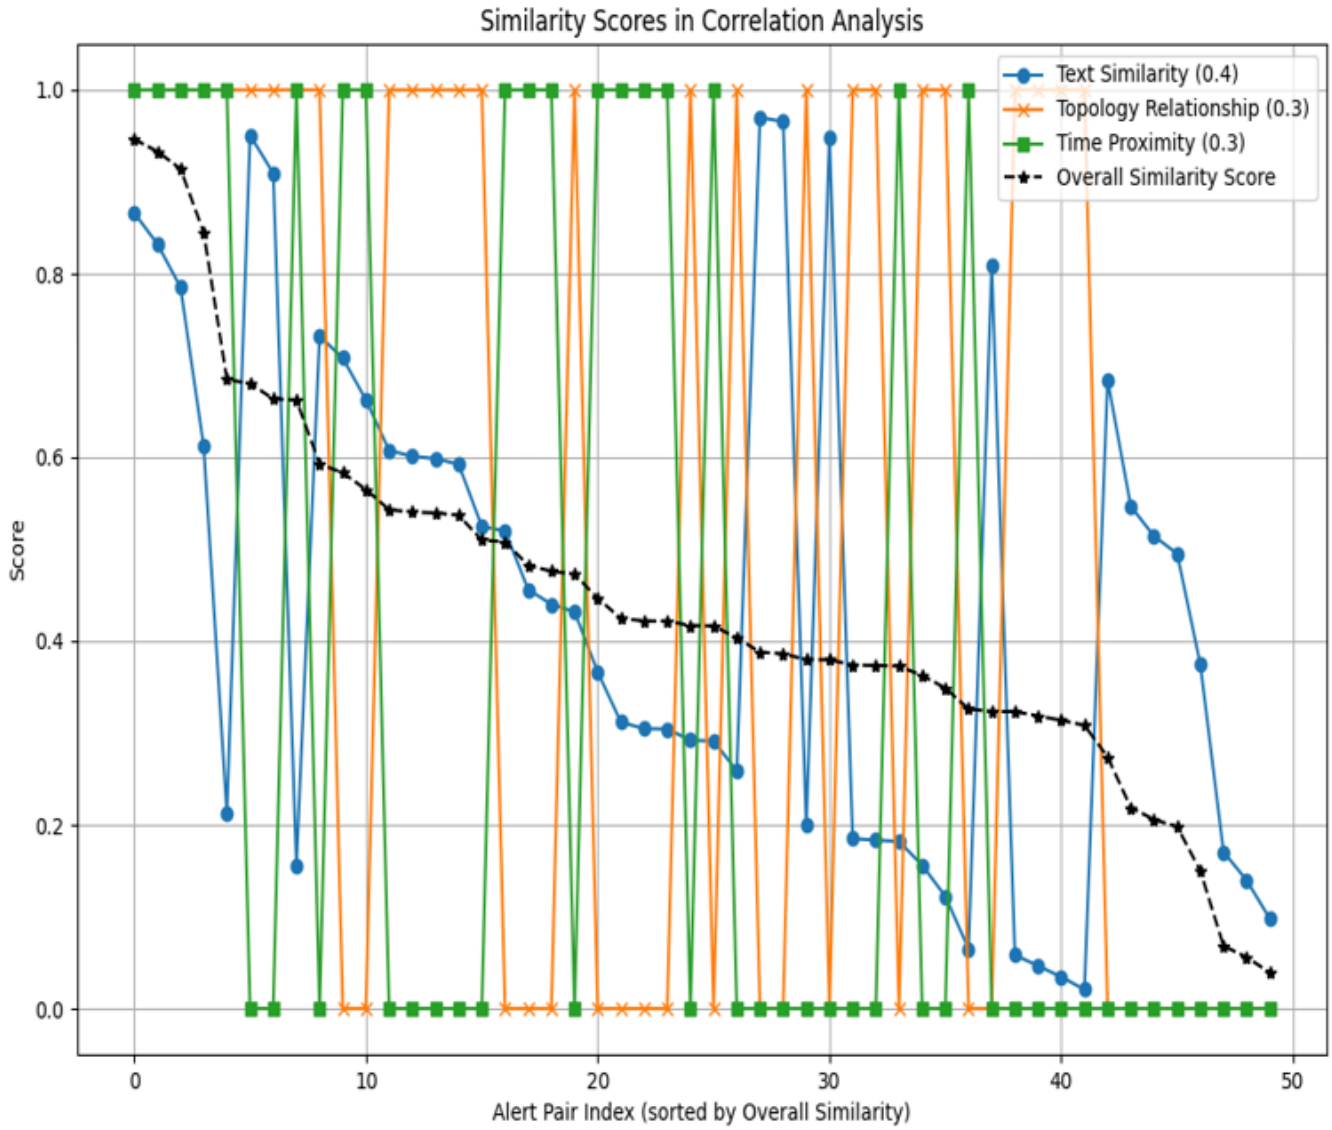

Figure 7

Similarity Score in Correlation Analysis

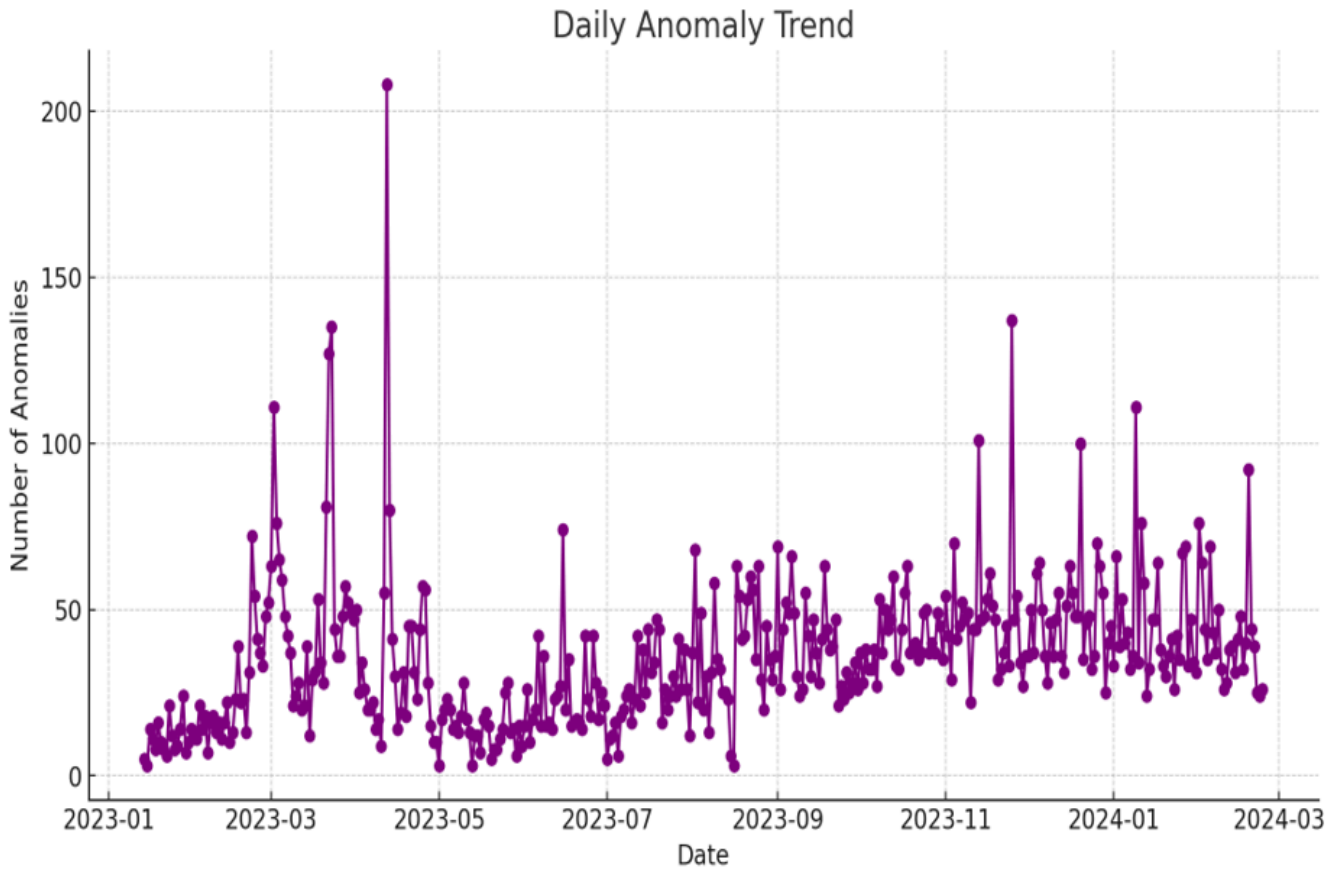

Figure 8

Daily Anomaly Trend

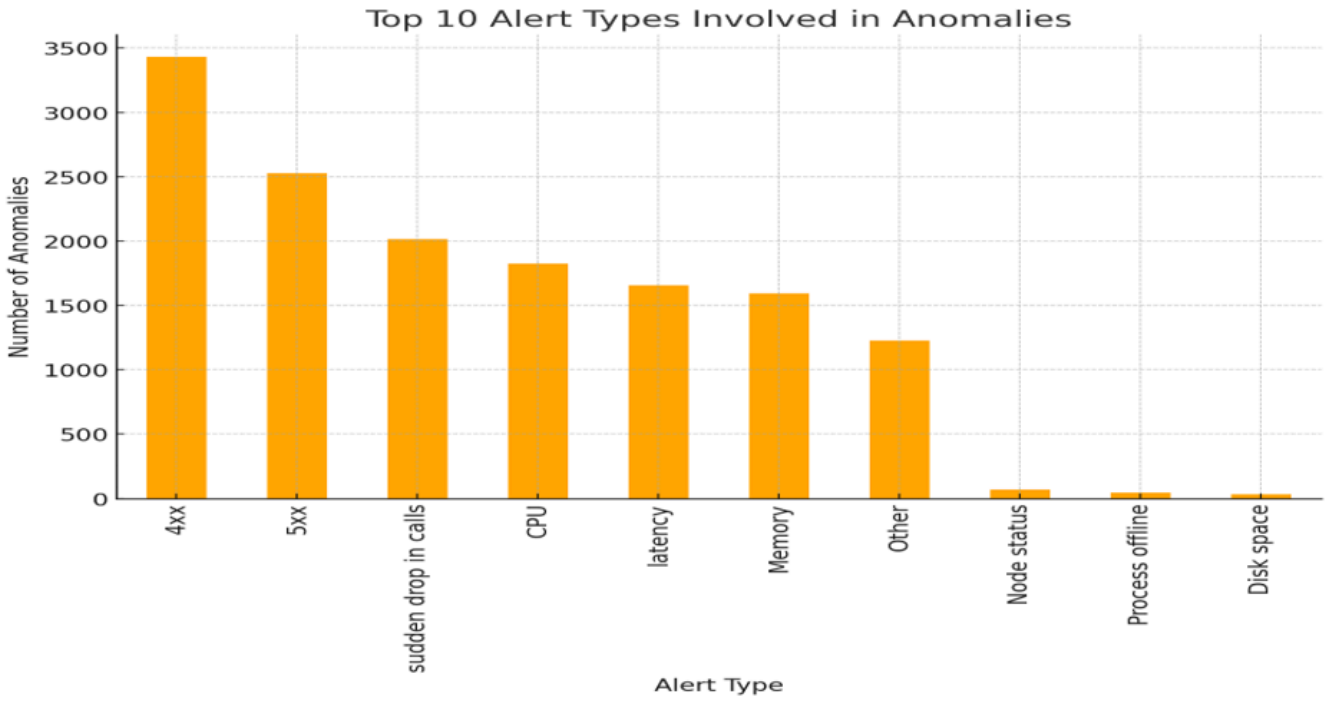

Figure 9

Alert types vs Number of anomalies

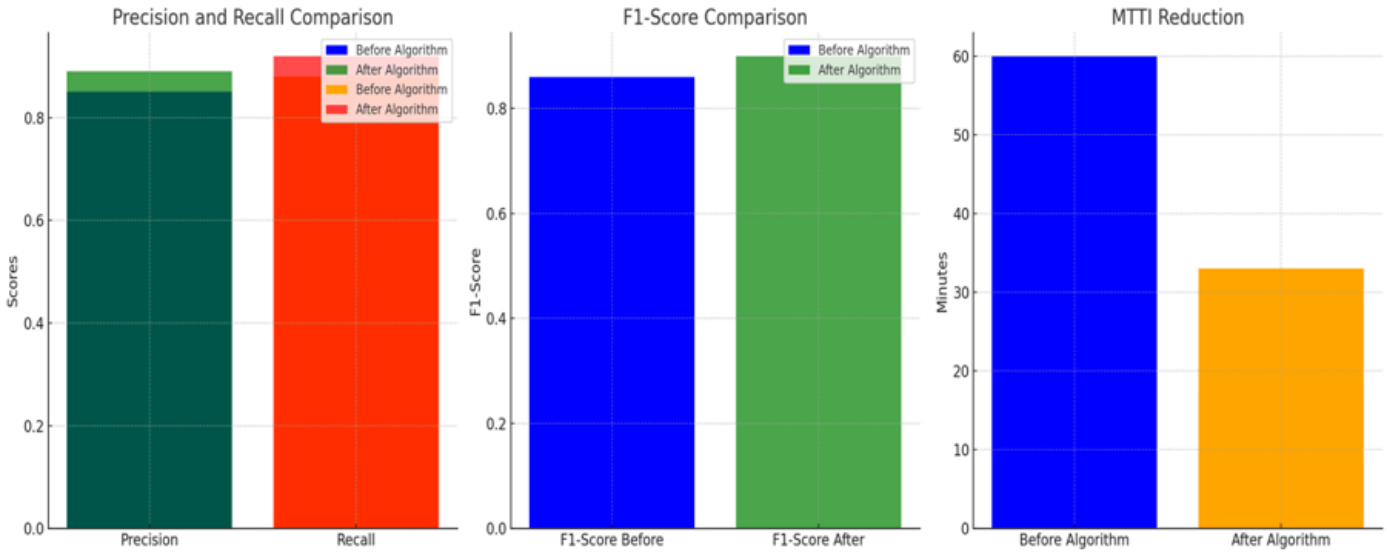

Figure 10

MTTI Reduction