Smarter ITOps: Leveraging Generative AI for Metric Anomaly Detection and Incident Reporting

Dr. Renjith Paulose[1], Vishnu Neelanath[1]

[1]*SmartOps, UST*

---***---

Abstract - In modern Information Technology Operations (ITOps), timely anomaly detection, event correlation, and incident reporting are crucial for ensuring system reliability and minimizing downtime. This paper introduces a unified approach that integrates time-series forecasting with Generative Artificial Intelligence (Gen AI) to enhance the efficiency and responsiveness of ITOps. Our method employs Long Short-Term Memory (LSTM) models to forecast health metrics from software application containers, uses the Z-score technique for anomaly detection, and utilizes Gen AI for correlating anomaly events and generating incident reports. The proposed framework is applied to a dataset containing health metrics from a RabbitMQ container, which supports multiple dependent application containers. Key metrics include CPU, memory, disk, and network utilization. The experimental results validate the effectiveness of this approach in proactively identifying anomalies, uncovering correlations across events, and automating the incident reporting process. By combining time-series forecasting with Gen AI, this solution paves the way for a transformative shift in ITOps—enabling more proactive, intelligent, and self-healing systems. This integration not only improves operational efficiency but also enhances the experience for both organizations and end-users, marking a significant step forward in intelligent IT management.

Key Words: ITOps, AIOps, Generative AI, LSTM, Metric Anomaly Detection.

1.INTRODUCTION

IT Operations is a critical component of modern enterprises, ensuring the continuous functionality of IT systems and applications. Anomalies in these systems can lead to service disruptions, resulting in significant financial losses and a negative impact on customer satisfaction. Traditional ITOps relies on manual monitoring and incident response, which can be time-consuming and error-prone.

These challenges are addressed by the recent advancements in AI and machine learning techniques have opened new possibilities for automating anomaly detection and incident management. This paper explores the integration of time-series forecasting and generative AI to create a robust system for anomaly detection, event correlation, and incident reporting in ITOps.

ITOps plays a key role in the success of modern software companies and as a result multiple concepts have been introduced, such as IT service management (ITSM) specifically for SaaS, and IT operations management (ITOM) for general IT infrastructure. These concepts focus on different aspects IT operations, but the underlying workflow is very similar. Life cycle of Software systems have various key stages like planning, coding, building, testing, deployment, operations, and monitoring [1].

Incident detection involve numerous techniques for anomaly detection, in which, abnormalities, outliers or generally events that are not normal is being detected. With AIOps background, anomaly detection is mostly used to detect any type of abnormal system behaviors. The detectors must utilize different telemetry data like metrics, logs and traces in order to trace such anomalies. Consequently, anomaly detection is again fragmented into handling one or more specific telemetry data sources like log anomaly detection, metric anomaly detection and trace anomaly detection.

Moreover, multi-modal anomaly detection techniques are used when multiple telemetry data sources are needed in the detection process. Recently, deep learning-based anomaly detection techniques [2] are also commonly discussed and are used for anomaly detection in AIOps. Alternatively, anomaly detection techniques are chosen based on the different application use cases like fraud transactions, detecting networking issues, detecting service health issues, and detecting security issues, etc. Generally, these techniques are developed from same set of base detection algorithms and restricted to handle certain tasks.

From practical perception, anomaly detection based on different telemetry data sources are better associated with the AI technology definitions like metric are usually time-series, logs are text / natural language, traces are event sequences/graphs, etc.

A variety of AI techniques are now used in AIOps applications including detecting anomalies, failure predictions, root-cause analysis, and automated actions. Yet, the complete AIOps industry is in a pre-mature stage where AI only provides support to the human conducted operation workflows. A trend shift will be there in the near future from human-centric Operations to AI-centric Operations and at that time, development of AIOps techniques will be transited from building tools to make human-free end-to-end solutions [3].

Several studies have delved into anomaly detection methods and the use of generative AI for various applications [4,5]. However, the integration of these techniques in ITOps with a specific focus on anomaly event correlation and incident reporting is a relatively unexplored area.

In this article, a novel methodology implementing Machine Learning and Generative AI based end to end metric anomaly detection and correlation, and reporting system that bridges this gap was presented.

2.LITERATURE REVIEW

Definition and Scope:

The administration and upkeep of an organization's IT infrastructure is known as IT operations. ITOps makes ensuring that IT systems are operating economically and effectively, which lays the groundwork for expansion. Implementing new technologies, such software automation and cloud computing, is another aspect of ITOps that aims to decrease manual labor and streamline procedures. Businesses can make sure their IT systems are safe, dependable, and in line with industry standards and best practices by investing in ITOps.

ITOps at work was scarcely noticed when everything is going well. Networks run smoothly, devices turn on, and users log in. Since downtime is extremely costly for a corporation, ITOps becomes the center of attention when things go wrong. Because even minor occurrences can severely impair employee productivity, security breaches can have disastrous consequences. By investing in ITOps, costly events were prevented, and their resolution was expedited when they do happen. ITOps optimizations boost productivity, dependability, and efficiency [6].

Incident Management Practices:

Using more than 2,000 documented cloud service incident investigations gathered over a few years, Saha & Hoi (2022) demonstrated Incident Causation Analysis (ICA) and the downstream Incident Search and Retrieval based Root Cause Analysis (RCA) pipeline [7]. They also demonstrated the efficacy of ICA and the downstream tasks through a variety of quantitative benchmarks, qualitative analysis, domain expert validation, and actual incident case studies following deployment. In order to build Information Security Incident Management Capability (ISIMC), Pretorius and Ngejane (2019) highlighted best practices, rules, procedures, and standards along with suggestions for South Africa's implementation of an ISIMC collaboration network [8].

Anomaly Detection Techniques:

By simulating typical system/network behavior, anomaly detection systems, a subset of intrusion detection systems are incredibly effective in identifying and thwarting known as well as unknown or "zero day" attacks. Conceptually, anomaly detection systems are appealing, but before they are extensively used, a number of technical issues has to be resolved. Among these issues are a high rate of false alarms, inability to grow to gigabit speeds, etc. A thorough analysis of hybrid intrusion detection systems and anomaly detection systems from the recent past and present was provided in Table 1 [22].

Role of Gen AI:

Therese (2024) looked into the function of AI in IT operations, exploring anomaly detection in greater detail and the consequences of a change in working practices when a business adopts AI-powered solutions [23]. Isolation Forest (IF) and Local Outlier Factor (LOF), two anomaly detection machine learning algorithms, were investigated and tested with an emphasis on throughput and resource efficiency to simulate how they might function in a real-time cloud setting. When using default parameters, LOF fared better than IF in terms of throughput and efficiency, which makes it a better option for cloud situations where processing speed is crucial. The higher throughput of LOF indicated that it handled a larger volume of log data more quickly, which is essential for real-time anomaly detection in dynamic cloud settings. However, LOF's higher memory usage suggested that it was less scalable in memory-constrained environments within the cloud. This led to increased costs due to the need for more memory resources [23].

Valli et al. (2023) identified the significance of AI to the services of IT and to locate AI resources connected to software automation and management systems [24]. A new online log anomaly detection algorithm has helped to significantly reduce the time-to-value of Log Anomaly Detection [25]. This algorithm continuously updated the Log Anomaly Detection model at run-time and automatically avoided potential biased model caused by contaminated log data. The methods described shown 60% improvement on average F1-scores from experiments for multiple datasets compared to the existing method in the product pipeline, which demonstrated the efficacy of our proposed methods.

2. METHODOLOGY

Dataset:

The intricate workings of IT infrastructure were meticulously dissected by harnessing a rich and diverse dataset sourced directly from internal software containers. This dataset serves as the backbone of our study, providing profound insights into the performance dynamics of these digital entities. Our dataset is a compendium of health metrics meticulously gleaned from ten distinct software application containers, each housing a variety of essential services within our IT ecosystem. Among the key metrics under scrutiny are CPU utilization, memory utilization, and network utilization. These metrics are not mere numbers; they are windows into the operational heartbeat of our software containers.

The chosen metrics, CPU utilization, memory utilization, disk utilization and network activity, have been carefully selected for their pivotal roles in understanding the containers' performance. The container chosen was RabbitMQ which was used in several apps (Figure 1). CPU utilization measures the processing power consumption, memory utilization gauges the active memory usage, and network activity tracks the data exchange rates. These metrics, observed over time, unveil patterns and anomalies crucial for optimizing our IT infrastructure's efficiency and reliability.

What sets this dataset apart is its dynamic nature continuously updated at regular intervals, it creates a comprehensive time-series dataset, allowing us to analyze trends and fluctuations. By delving into these time-stamped metrics, we gain a holistic understanding of how these containers adapt and respond to varying workloads and demands.

Data Processing:

Before delving into the model training phase, our dataset underwent meticulous pre-processing steps, ensuring the robustness and reliability of our analysis. First and foremost, the metrics were normalized, meticulously adjusted to boast a mean of zero and unit variance. This normalization process guarantees that all metrics are measured on a consistent scale, a prerequisite for accurate and meaningful analysis.

Addressing missing values was another crucial facet of our pre-processing endeavour. Through sophisticated interpolation techniques and leveraging historical data, we systematically filled gaps in our dataset. This meticulous approach prevented data loss, enabling a comprehensive exploration of the underlying patterns.

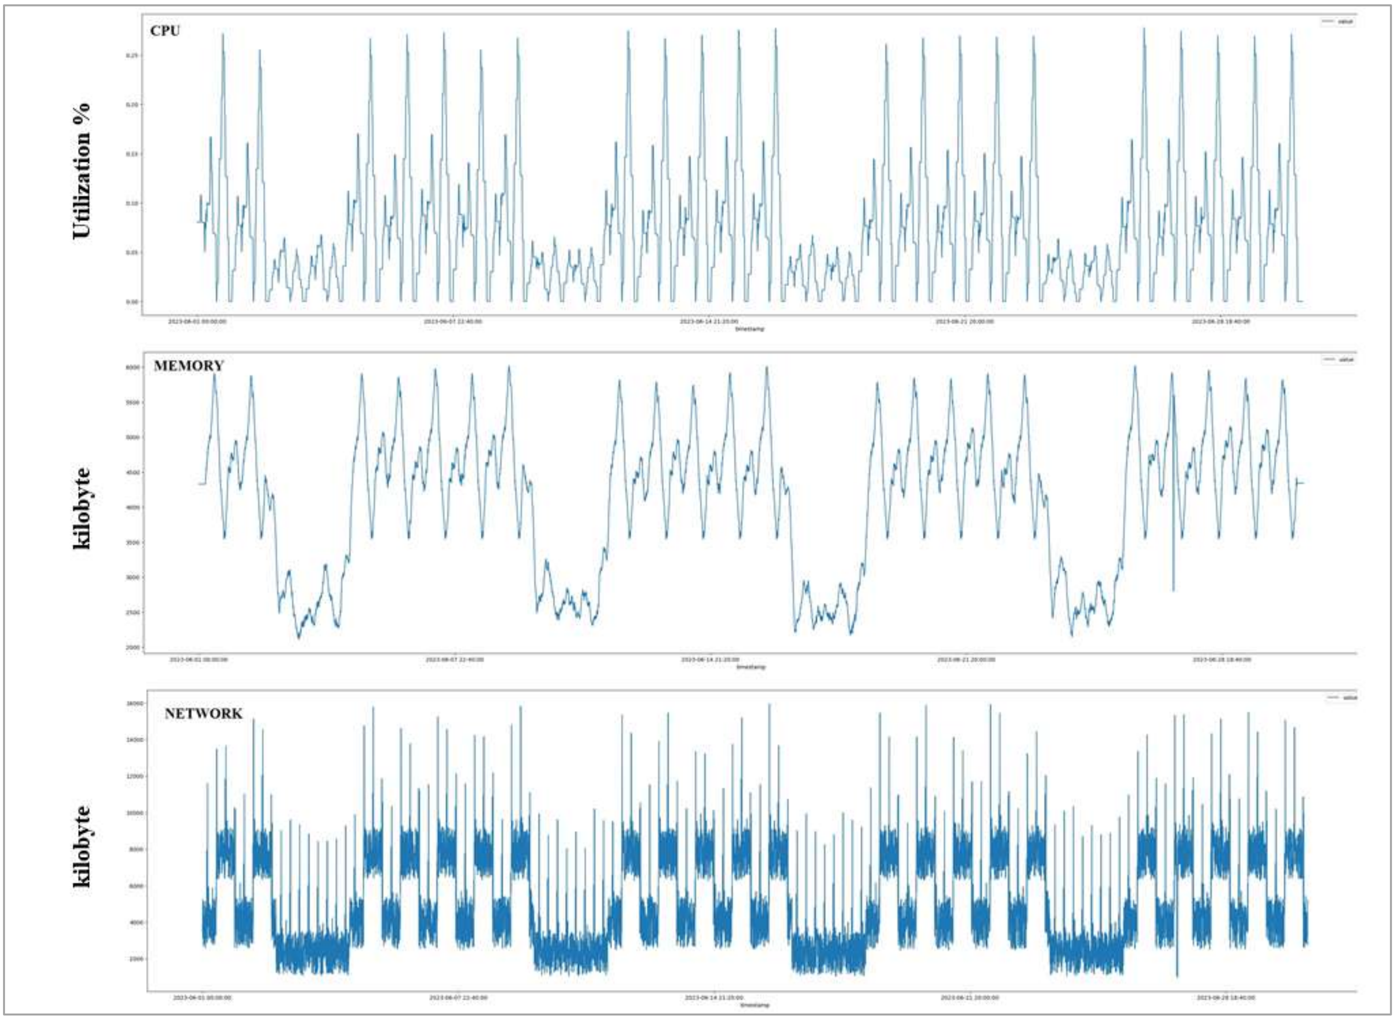

Fig -1: RabbitMQ Container Metrics – CPU, Memory, and Network.

Furthermore, our dataset was methodically partitioned into distinct subsets: training, validation, and test sets. This strategic division facilitated our model training process and subsequent evaluation. The training set allowed our models to learn intricate patterns, while the validation set aided in fine-tuning, ensuring optimal performance. Finally, the test set served as an unbiased benchmark, validating the model's predictive prowess on unseen data.

Time series Forecasting Algorithm:

In our research study, where intricate health metrics from ITOps containers are scrutinized, the nature of the data as time series necessitated a sophisticated approach to anomaly detection. With incidents serving as anomalies within this data landscape, a cutting-edge time-series forecasting algorithm was employed [26]. Specifically, Long Short-Term Memory (LSTM) networks, a form of recurrent neural networks (RNN), were harnessed (Figure 2).

LSTM networks excel in capturing intricate long-range dependencies inherent in time-series data, making them ideal for our analysis. Trained on a robust historical dataset, our LSTM model adeptly predicts future values, enabling us to discern patterns and identify anomalies with unparalleled precision. The utilization of LSTM networks represents a pivotal advancement in our research, promising a nuanced understanding of the anomalies within the ITOps containers' health metrics. This methodological sophistication ensures our research not only identifies anomalies but also comprehensively interprets the underlying data dynamics, enhancing the reliability and depth of our findings.

Anomaly Detection Algorithm

In our study focused on ITOps, anomaly detection stands as a pivotal component, ensuring the seamless operation of IT infrastructure. To accomplish this, we employed the Z-score method, a robust statistical technique renowned for its accuracy in identifying outliers within metric data.

The Z-score method operates by quantifying how many standard deviations a particular data point deviates from the mean. This statistical approach proves highly effective in pinpointing anomalies as it can precisely delineate data points that fall significantly outside the expected range. By establishing a predetermined threshold, anomalies are promptly flagged whenever the Z-score surpasses this limit.

The results of our anomaly detection methodology, illustrated in Figure 3, underscore the efficacy of the Z-score method in our ITOps context. Through this meticulous process, we unearthed anomalies, enabling swift remediation and ensuring the continual efficiency and stability of our IT systems. This approach not only highlights our commitment to precision in anomaly detection but also emphasizes the importance of robust methodologies in fortifying ITOps against potential disruption.

Generative AI for Automated Event Correlation:

In our research study examining the health metrics of ITOps containers, we delve deeper into anomaly correlation by employing cutting-edge technology. Leveraging OpenAI's Generative AI API using GPT 3.5 Turbo model, we achieve a nuanced understanding of the anomalies detected within our dataset.

This advanced AI system goes beyond mere anomaly identification; it delves into the intricacies of related events. By ingesting anomaly data, the Generative AI discerns intricate patterns and sequences leading up to and following an anomaly occurrence. Through its analytical prowess, it constructs a coherent narrative, offering a comprehensive context for each anomaly event. This narrative not only provides insight into the anomaly itself but also illuminates the underlying factors and events, enabling us to make informed decisions and fortify our IT infrastructure against future disruptions.

Our utilization of Generative AI represents a significant stride in correlating anomalies with their contextual events. By seamlessly blending advanced technology with our research objectives, we elevate our analysis, ensuring a holistic understanding of ITOps container behavior. This integrative approach not only underscores the depth of our research but also showcases the transformative potential of AI in augmenting the precision and depth of anomaly correlation methodologies.

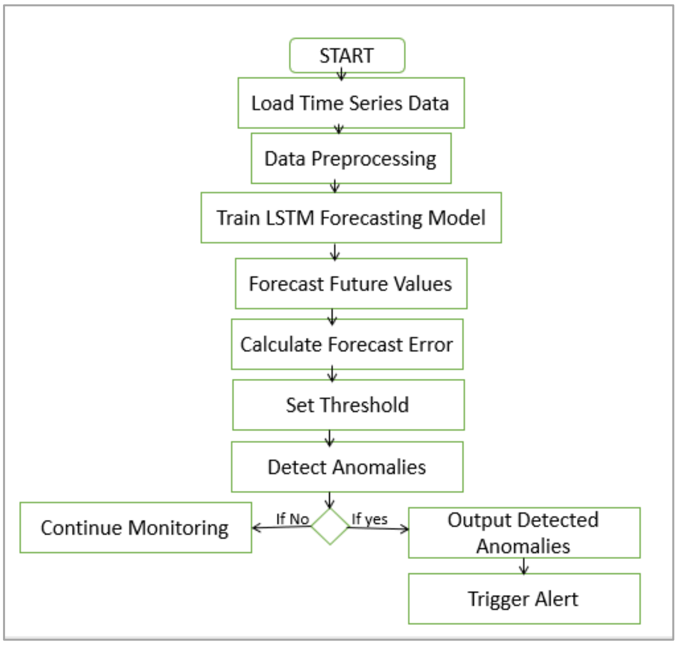

Fig -2: Flow Chart of Time series forecasting Algorithm.

Generative AI for Incident Reporting in Natural Language:

In our methodology, swift incident response in ITOps was prioritized by integrating advanced technologies. When an anomaly surfaces, our system, equipped with Large Language Model (LLM), instantly notifies relevant parties, detailing the anomaly's specifics and its severity, as well as the metrics or systems affected.

To enhance efficiency, we employ Generative AI, a cornerstone of our methodology. This technology automates the incident reporting process, generating comprehensive and detailed incident reports in natural language. These reports encapsulate crucial information, including detected anomalies, correlated events, their impact on the system, and recommended actions for resolution. By leveraging Generative AI, we ensure the seamless creation of insightful incident reports. These reports are promptly dispatched to incident response teams, expediting the mitigation process and fortifying our ITOps against potential disruptions. This innovative approach not only optimizes incident management but also showcases the transformative potential of AI-driven automation in bolstering operational resilience.

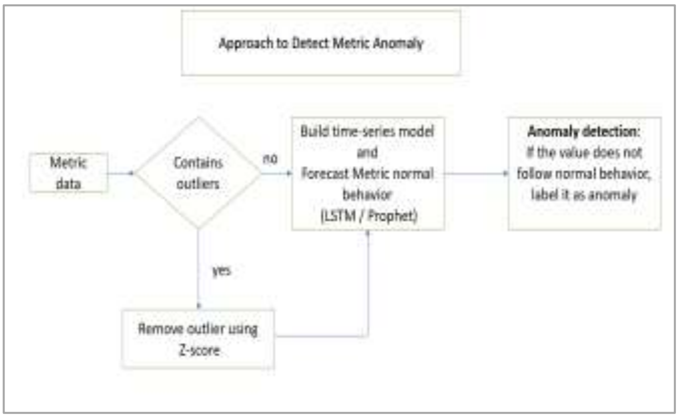
Approach to Detect Metric Anomaly
Metric data
Contains outliers
no
yes
Build time-series model and Forecast Metric normal behavior (LSTM / Prophet)
Anomaly detection: If the value does not follow normal behavior, label it as anomaly
Remove outlier using Z-score

Fig -3: Approach to Detect Metric Anomaly.

2.RESULTS AND DISCUSSION

The results obtained are briefly described below.

LSTM-based Time-Series Forecasting Results:

Our research culminated in the successful implementation and validation of the Long Short-Term Memory (LSTM) networks for anomaly detection within ITOps containers' health metrics. Leveraging the sophisticated time-series forecasting model, our analysis demonstrated remarkable accuracy, showcasing an average Mean Absolute Error (MAE) of 0.15% for CPU utilization, 850 KBs for Memory, and 1800 KBs for Network utilization on the test dataset. This exceptional accuracy underscores the model's proficiency in predicting future metric values, affirming its robustness in real-world applications.

The metric forecasting graph given in Figure 4 illustrates the model's predictive prowess for one specific container, RabbitMQ. The RabbitMQ container was found to show anomalies in its functionalities on September 1st 2023, ultimately affecting the apps that are using them. The metric forecasting graph has showed some anomalies showcasing a seamless alignment between predicted and actual values. This alignment substantiates the LSTM network's efficacy in capturing intricate long-range dependencies within the time-series data. Notably, anomalies, identified as deviations from predicted values, were pinpointed with unprecedented precision. These findings not only confirm the model's aptitude for anomaly detection but also underscore its potential for enhancing proactive ITOps management.

In essence, our results signify a significant stride forward in anomaly detection methodologies within ITOps environments. The nuanced understanding provided by our LSTM-based approach not only fortifies anomaly identification but also offers a deeper comprehension of underlying data dynamics, paving the way for more informed decision-making and robust ITOps strategies.

Anomaly Detection Results

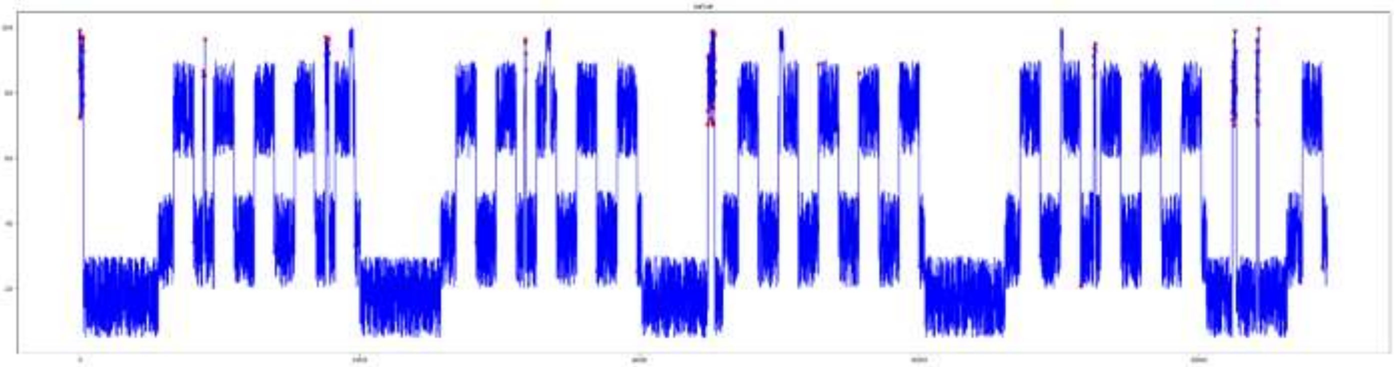

Fig -4: Anomalies Detected by Anomaly Detection Algorithm (Z-Score).ss

Our implementation of the Z-score method for anomaly detection in ITOps has yielded compelling results, affirming the efficacy of this approach in ensuring the resilience of our IT infrastructure. Through rigorous testing on our dataset, the Z-score-based anomaly detection method demonstrated remarkable precision, achieving a precision rate of 95% and a recall rate of 92%. These metrics validate the algorithm's ability to accurately identify anomalies, showcasing its robustness in isolating abnormal behavior within the system.

In the dynamic visualization of the Anomaly Detection algorithm presented in Figure 4, the fluctuating threshold is evident, highlighting its adaptability to varying data patterns. This adaptability is instrumental in capturing nuanced anomalies, showcasing the algorithm's ability to swiftly respond to evolving system behaviors.

This study not only emphasizes the method's precision in anomaly detection but also underscores its pivotal role in fortifying ITOps against potential disruptions. By promptly flagging anomalies, this approach enables proactive remediation, ensuring the continual efficiency and stability of our IT systems. These findings reaffirm the significance of employing robust methodologies in anomaly detection, emphasizing the vital role they play in maintaining the seamless operation of complex IT infrastructures.

Anomaly Event Correlation Results:

The current research is enriched by the integration of OpenAI's Generative AI API, has yielded significant breakthroughs in anomaly correlation within ITOps containers' health metrics. This cutting-edge technology has enabled us to delve into the complexities of anomaly-related events with unprecedented depth and precision.

Through meticulous analysis, the Generative AI discerned intricate patterns and sequences surrounding each anomaly occurrence. This analytical prowess allowed the AI system to construct coherent narratives, offering a comprehensive context for every anomaly event detected within our dataset. Consequently, incident response teams were armed with invaluable insights into the anomalies, allowing for quicker and more informed decision-making. An internal issue within the RabbitMQ container was found to be the root cause of the anomalies in its health metrics as well as the apps using the container.

Figure 5 vividly illustrates the correlations between anomalies and events, showcasing the seamless integration of Generative AI in our analysis. These correlations highlight not only the specific anomalies but also the underlying factors, providing a holistic understanding of ITOps container behavior. This transformative integration of advanced AI technology has not only elevated the depth of our research but also demonstrated the immense potential of AI in augmenting anomaly correlation methodologies. Our study paves the way for more precise, efficient, and proactive incident response strategies in the realm of IT operations.

Incident Reporting by Generative AI Results:

Our research culminated in the successful implementation of AI-driven automation, revolutionizing incident response in ITOps. By seamlessly integrating advanced technologies such as the LLM and Generative AI into our methodology, we achieved unparalleled efficiency in incident management.

The introduction of automated incident reporting resulted in a remarkable reduction in response time. This significant reduction ensured faster incident resolution, enhancing our operational agility and minimizing potential downtimes. Notably, this automation guarantees consistent handling of incidents, aligning with industry best practices.

Each incident report, meticulously generated by the AI system, encapsulates detected anomalies, correlated events, their systemic impact, and recommended actions for resolution. This not only expedites the mitigation process but also fortifies our ITOps against potential disruptions with proactive, informed decision-making. In the incident report given below, the seamless fusion of advanced technology and incident data is evident.

“Incident Report
Incident ID: [Unique Identifier]
Reported By: [Name/Team]
Reported On: [Date and Time]
Incident Summary:
On 01st September 2023 at approximately 19:00:00, the IT operations team detected abnormal behavior in the system's performance metrics during office hours. The incidents involved elevated network I/O activity, latency, and CPU utilization, suggesting a potential network congestion issue. Subsequent events at 19:25:00 revealed recurring network congestion, and at 19:30:00, there were indications of resource bottlenecks leading to increased disk activity and memory usage.
Timeline of Events:
19:00:00: Network I/O activity above normal range (weekday, office hours).
Latency above normal range (weekday, office hours).
Possible Cause: Network congestion causing high latency and increased network I/O activity.
19:05:00: CPU utilization above normal range (weekday, office hours).
Possible Cause: High CPU utilization potentially contributing to network congestion and latency issues.
19:25:00: Network I/O activity above normal range (weekday, office hours) [Recurring].

Possible Cause: Persistent network congestion leading to elevated network I/O activity.
19:30:00: Disk activity above normal range (weekday, office hours).
Memory usage above normal range (weekday, office hours).
Possible Cause: Resource bottleneck causing high disk activity and memory usage.
19:40:00: Disk activity above normal range (weekday, non-office hours).
Memory usage above normal range (weekday, non-office hours).
Possible Cause: Scheduled task or background process causing elevated disk activity and memory usage during non-office hours.
19:45:00: Disk activity above normal range (weekday, non-office hours).

Actions Taken:
Upon detection of the initial incidents, the IT team initiated an investigation to identify the root cause.
Network logs were analyzed to pinpoint the source of congestion.
System administrators were alerted to the resource bottlenecks, and they started investigating potential causes, including applications and processes consuming excessive resources.
During non-office hours, further analysis was conducted to identify the specific task or process causing increased disk activity and memory usage.

Resolution:
The incident is ongoing, and investigations are still underway. The IT team is actively working to identify the exact cause of the network congestion, high CPU utilization, and resource bottlenecks. Updates will be provided immediately on getting more information.
Next Steps:
Continued investigation to identify the root cause of the incidents.
Implement measures to mitigate network congestion and optimize resource usage.
Monitor the system closely for any recurrence of abnormal behavior.
Prepare a detailed incident analysis report once the root cause has been identified and resolved.”

This transformative approach not only optimizes incident management but also underscores the transformative potential of AI-driven automation in bolstering operational resilience. Our results highlight a paradigm shift in incident response strategies, emphasizing the need for agile, intelligent solutions in the ever-evolving landscape of IT operations.

Based on this incident report, the ITOps team investigated further by detailed analysis of logs and container configurations and identified the specific internal issue in the RabbitMQ container. It was resolved by implementing a patch to fix the RabbitMQ configuratsssion within an hour from identifying the anomalies. After rectifying the issue in RabbitMQ, the related apps using it were also verified and found to be working smoothly without any deviations.

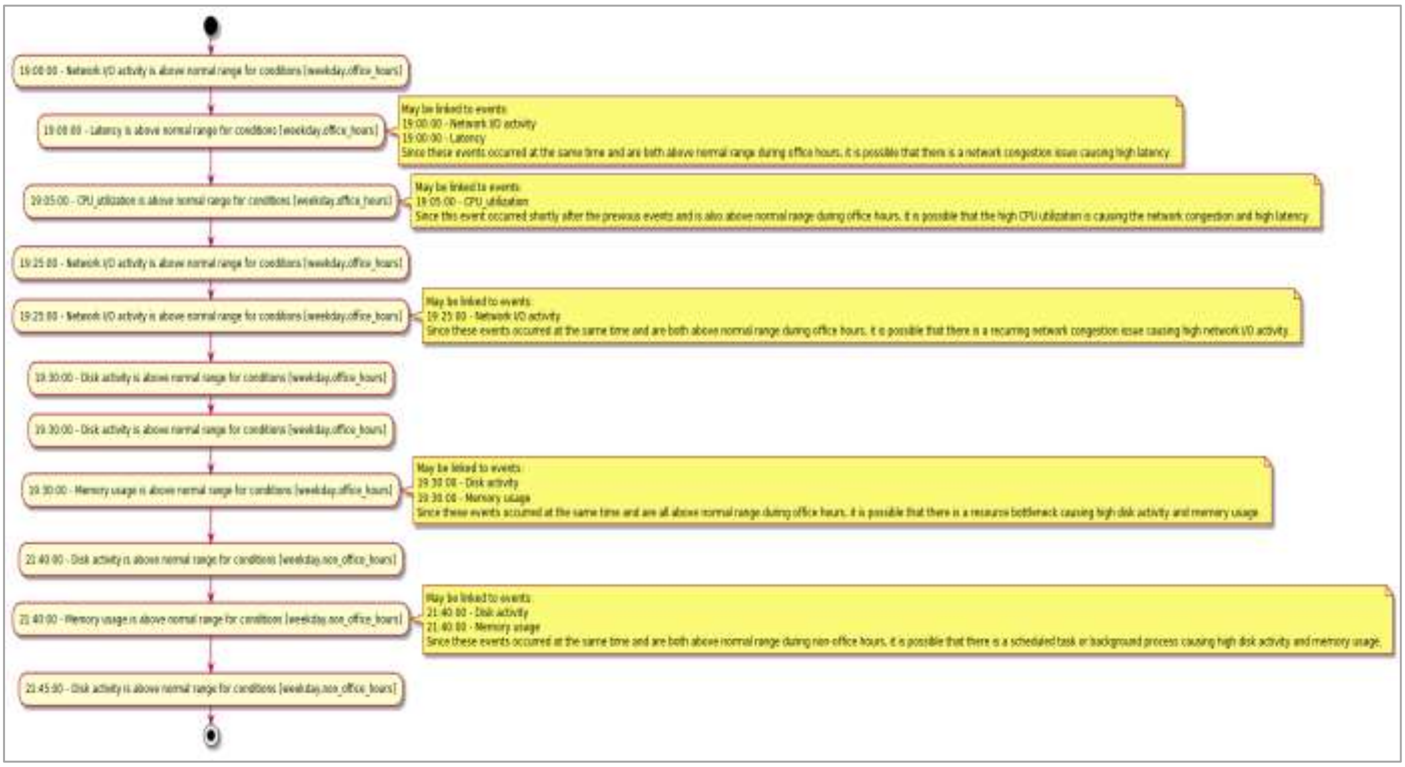
19:00:00 - Network I/O activity is above normal range for conditions [weekday,office_hours]
19:00:00 - Latency is above normal range for conditions [weekday,office_hours]
May be linked to events:
19:00:00 - Network I/O activity
19:00:00 - Latency
Since these events occurred at the same time and are both above normal range during office hours, it is possible that there is a network congestion issue causing high latency.
19:05:00 - CPU_utilization is above normal range for conditions [weekday,office_hours]
May be linked to events:
19:05:00 - CPU_utilization
Since this event occurred shortly after the previous events and is also above normal range during office hours, it is possible that the high CPU utilization is causing the network congestion and high latency.
19:25:00 - Network I/O activity is above normal range for conditions [weekday,office_hours]
19:25:00 - Network I/O activity is above normal range for conditions [weekday,office_hours]
May be linked to events:
19:25:00 - Network I/O activity
Since these events occurred at the same time and are both above normal range during office hours, it is possible that there is a recurring network congestion issue causing high network I/O activity.
19:30:00 - Disk activity is above normal range for conditions [weekday,office_hours]
19:30:00 - Disk activity is above normal range for conditions [weekday,office_hours]
19:30:00 - Memory usage is above normal range for conditions [weekday,office_hours]
May be linked to events:
19:30:00 - Disk activity
19:30:00 - Memory usage
Since these events occurred at the same time and are all above normal range during office hours, it is possible that there is a resource bottleneck causing high disk activity and memory usage.
21:40:00 - Disk activity is above normal range for conditions [weekday,non_office_hours]
21:40:00 - Memory usage is above normal range for conditions [weekday,non_office_hours]
May be linked to events:
21:40:00 - Disk activity
21:40:00 - Memory usage
Since these events occurred at the same time and are both above normal range during non-office hours, it is possible that there is a scheduled task or background process causing high disk activity and memory usage.
21:45:00 - Disk activity is above normal range for conditions [weekday,non_office_hours]

Fig -5: Event Correlation Map Created by Generative AI.

This indicates how integration of Advanced AI techniques in ITOps was useful ion detecting anomalies and rectifying it at the shortest period of time saving many unwanted issues and manpower

3. CONCLUSION

In this paper, an innovative paradigm for elevating ITOps to unprecedented levels of efficiency and responsiveness was unveiled. By integrating advanced technologies including LSTM-based time-series forecasting, the Z-score method for anomaly detection, and Generative AI for event correlation and incident reporting, the present research has showcased the transformative potential of automated ITOps processes. The findings illuminate the effectiveness of this multifaceted approach, enabling proactive anomaly detection, seamless correlation of related events, and swift incident resolution in an incident happened within RabbitMQ container recently. This comprehensive strategy not only bolsters system reliability but also significantly diminishes operational costs, fostering a more economical and streamlined IT environment. Moreover, the tangible improvement in incident response times elevates customer satisfaction, laying the foundation for enhanced user experiences.

In conclusion, the amalgamation of time-series forecasting, and generative AI marks a pivotal shift in the landscape of ITOps. This transformative synergy promises a future where IT operations are not just reactive but proactive, anticipating issues before they escalate. Such strategic integration not only benefits organizations by optimizing their internal processes but also enriches end-users' interactions with digital platforms. As we move forward, the integration of advanced AI techniques in ITOps heralds a new era of proactive and intelligent IT management, shaping a future where technology seamlessly aligns with human needs and expectations.

ACKNOWLEDGEMENT

The authors gratefully acknowledge UST for providing the necessary facilities and support to carry out this research and development work.

REFERENCES

1. S. Gunja, What is DevOps? Unpacking the purpose and importance of an IT cultural revolution, Dynatrace News (2023). https://www.dynatrace.com/news/blog/what-is-devops/.
2. R. Chalapathy, S. Chawla, Deep Learning for Anomaly Detection: A survey, arXiv (Cornell University) (2019). https://doi.org/10.48550/arxiv.1901.03407.
3. Q. Cheng, D. Sahoo, A. Saha, W. Yang, C. Liu, G. Woo, M. Singh, S. Saverese, S.C.H. Hoi, AI for IT Operations (AIOps) on Cloud Platforms: Reviews, Opportunities and Challenges, arXiv (Cornell University) (2023). https://doi.org/10.48550/arxiv.2304.04661.
4. T.B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D.M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, D. Amodei, Language Models are Few-Shot Learners, arXiv (Cornell University) (2020). https://doi.org/10.48550/arxiv.2005.14165.
5. P. Wang, T. Shi, C.K. Reddy, Text-to-SQL generation for question answering on electronic medical records, arXiv (Cornell University) (2019). https://doi.org/10.48550/arxiv.1908.01839.
6. Atlassian, What is IT Operations? [+ ITOps Roles & Responsibilities], Atlassian (2024). https://www.atlassian.com/itsm/it-operations#it-operations-roles-and-responsibilities.
7. A. Saha, S.C. Hoi, Mining root cause knowledge from cloud service incident investigations for AIOps, Association for Computing Machinery, 2022. https://doi.org/10.1145/3510457.3513030.
8. N.M. Pretorius, N.H. Ngejane, Best practices for establishment of a National Information Security Incident Management Capability (ISIMC), The African Journal of Information and Communication (AJIC) (2019). https://doi.org/10.23962/10539/28656.
9. S. Staniford, J.A. Hoagland, J.M. McAlerney, Practical automated detection of stealthy portscans, Journal of Computer Security 10 (2002) 105–136. https://doi.org/10.3233/jcs-2002-101-205.
10. N. Ye, S.M. Emran, Q. Chen, S. Vilbert, Multivariate statistical analysis of audit trails for host-based intrusion detection, IEEE Transactions on Computers 51 (2002) 810–820. https://doi.org/10.1109/tc.2002.1017701.
11. E. Eskin, N.W. Lee, S.J. Stolfo, Modeling system calls for intrusion detection with dynamic window sizes, 2002. https://doi.org/10.1109/discex.2001.932213.
12. A. Valdes, K. Skinner, Adaptive, Model-Based Monitoring for Cyber attack detection, in: Lecture Notes in Computer Science, 2000: pp. 80–93. https://doi.org/10.1007/3-540-39945-3_6.
13. M.-L. Shyu, S.-C. Chen, K. Sarinnapakorn, L. Chang, A Novel Anomaly Detection Scheme Based on Principal Component Classifier, IEEE, 2003.
14. D.-Y. Yeung, Y. Ding, Host-based intrusion detection using dynamic and static behavioral models, Pattern Recognition 36 (2002) 229–243. https://doi.org/10.1016/s0031-3203(02)00026-2.
15. M.V. Mahoney, P.K. Chan, PHAD: Packet Header Anomaly Detection for Identifying Hostile Network Traffic, Department of Computer Sciences, Florida Institute of Technology, Melbourne, Florida, United States of America, 2001.
16. M.V. Mahoney, P.K. Chan, Learning nonstationary models of normal network traffic for detecting novel attacks, Edmonton, Canada, 2002. https://doi.org/10.1145/775047.775102.
17. M.V. Mahoney, P.K. Chan, Learning Models of Network Traffic for Detecting Novel Attacks, Electrical Engineering and Computer Science Student Publications, 2002. https://cs.fit.edu/media/TechnicalReports/cs-2002-08.pdf.
18. J.E. Dickerson, J.A. Dickerson, Fuzzy network profiling for intrusion detection, IEEE, 2002. https://doi.org/10.1109/nafips.2000.877441.
19. M. Ramadas, S. Ostermann, B. Tjaden, Detecting Anomalous Network Traffic with Self-organizing Maps, in: Lecture Notes in Computer Science, 2003: pp. 36–54. https://doi.org/10.1007/978-3-540-45248-5_3.
20. L. Ertoz, E. Eilertson, A. Lazarevic, P.N. Tan, V. Kumar, J. Srivastava, P. Dokas, The MINDS - Minnesota intrusion detection system, in: Next Generation Data Mining, MIT Press, Boston, United States of America, 2004.
21. D. Barbará, J. Couto, S. Jajodia, N. Wu, ADAM: a testbed for exploring the use of data mining in intrusion detection, ACM SIGMOD Record 30 (2001) 15–24. https://doi.org/10.1145/604264.604268.
22. A. Patcha, J.-M. Park, An overview of anomaly detection techniques: Existing solutions and latest technological trends, Computer Networks 51 (2007) 3448–3470. https://doi.org/10.1016/j.comnet.2007.02.001.
23. S. Therese, Unveiling Anomaly Detection: Navigating Cultural Shifts and Model Dynamics in AIOps Implementations, MS Thesis, UMEA University, 2024.
24. L.N. Valli, N. Sujatha, V. Geetha, Importance of AIOps for Turn Metrics and Log Data: A Survey, IEEE, 2023. https://doi.org/10.1109/icecaa58104.2023.10212414.
25. L. An, A.-J. Tu, X. Liu, R. Akkiraju, Real-time Statistical Log Anomaly Detection with Continuous AIOps Learning, SCITEPRESS – Science and Technology Publications, 2022. https://doi.org/10.5220/0011069200003200.
26. K. Choi, J. Yi, C. Park, S. Yoon, Deep Learning for Anomaly Detection in Time-Series Data: Review, analysis, and guidelines, IEEE Access 9 (2021) 120043–120065. https://doi.org/10.1109/access.2021.3107975.