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

AI and ML for Automated Incident Resolution: Enhancing Speed and Accuracy

Lakshmi Narasimha Rohith Samudrala

Abstract
Modern IT environments are extremely complex in nature. Traditional incident resolution methods which rely on manual troubleshooting, static alerts, and reactive monitoring are no longer sufficient to handle the complex IT Environment. Artificial Intelligence (AI) and Machine Learning (ML) provide abilities that can significantly improve the traditional incident management process.
This paper explores how AI-driven technologies leverage historical data to proactively identify issues, automatically detect root cause of the incident, and remediate them. This allows organizations to shift from reactive incident response to proactive and autonomous resolution strategies.
As organizations continue to grow, AI-powered incident management will play a crucial role in ensuring business continuity, operational efficiency, and a seamless digital experience for users.

Keywords: Artificial Intelligence (AI), AI-Driven Incident Management, Machine Learning (ML), Automated Root Cause Analysis, Anomaly Detection, Self-Healing, Predictive Analytics, Intelligent Alert Correlation, Mean Time To Resolve (MTTR), Mean Time To Detect (MTTD), Event Correlation, IT Service Management (ITSM)

INTRODUCTION

Modern IT environments are becoming increasingly complex and distributed, with cloud-native architectures, microservices, and hybrid infrastructures generating vast amounts of data. Traditional incident management approaches are struggling to keep up with these changes. Traditional methods often rely on manual troubleshooting, static threshold-based monitoring, and rule-based alerting, which result in slow response times, prolonged downtime, and high operational costs.
With the growth in the IT landscape, organizations face several challenges such as alert fatigue, siloed monitoring tools, and inefficient root cause analysis (RCA). To resolve incidents, IT teams would have to spend hours manually correlating logs, traces, and metrics, delaying resolution times, and increasing the impact of incidents on business operations [4]. As traditional incident management is reactive in nature, these issues are typically addressed after they have caused disruption to the service.
The introduction of Artificial Intelligence (AI) and Machine Learning (ML) in Incident Management is transforming how organizations identify, analyze, and resolve incidents. AI-driven solutions provide automated anomaly detection, intelligent alert correlation, and predictive capabilities [7][8]. These features allow organizations to identify problems before they start impacting the users. Self-healing IT systems powered by AI and ML can take actions such as restarting services, reallocating resources, rolling back changes, or apply patches without human intervention [1]. These advancements significantly reduce Mean Time to Detect (MTTD) and Mean Time to Resolve (MTTR).

AI-Driven Anomaly Detection and Incident Identification

Incident detection is the first and most important step in incident resolution. Traditional monitoring tools rely on static threshold-based predefined alerts. This approach is not compatible with the dynamic nature of modern IT landscape. Threshold-based alerts generate excessive false positives or miss subtle anomalies, leading to delayed incident detection.

AI-based systems learn the "normal" behavior of the application by analyzing logs, metrics, and traces over time. AI leverages unsupervised learning algorithms to identify outliers or deviations from the normal behavior of the application, indicating potential problem. Below figure 1 shows automated baseline identification by AI.

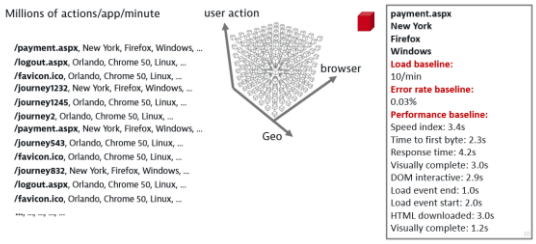

Figure 1 – Automated Baseline Identification by AI

Traditional monitoring tools generate isolated alerts for each metric, making it difficult to determine if alerts are related. AI-driven systems use event correlation techniques to connect related anomalies and group them into a single incident. This significantly resolves alert fatigue allowing IT teams to focus on problems that have an impact on the business.

Automated Root Cause Analysis (RCA)

The next step in incident management after incident identification is root cause analysis (RCA). Traditionally RCA involves manual log analysis, manual correlation of multiple alerts from siloed monitoring tools, and cross team collaboration [6]. This approach significantly increased Mean Time to Resolve (MTTR).

AI and ML significantly change the RCA process by automating the identification of potential root cause. AI-based monitoring tools unify different aspect of monitoring in one place. This removes data silos and provides a single pane of glass for the IT teams to use.

AI is capable to understand the complexity of the modern IT landscape, allowing it to create a complex system dependency graph [5]. This graph enabled the AI to trace the failure upstream and downstream to identify where the issue originated [5]. Below figure 2 shows an example of multi-dimensional graph generated by AI.

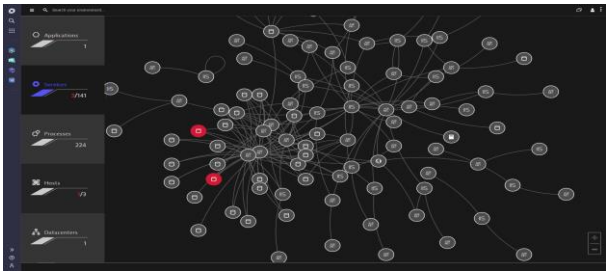

Figure 2 – Dynatrace Smartscape (multi-dimensional graph generated by DAVIS AI)

Also, as system dependencies are understood, the AI can intelligently correlate related alerts and identify recurring patterns in system behavior, helping IT teams to focus on the root cause rather than chasing symptoms. Below figures 3 shows example of alert correlation by AI.

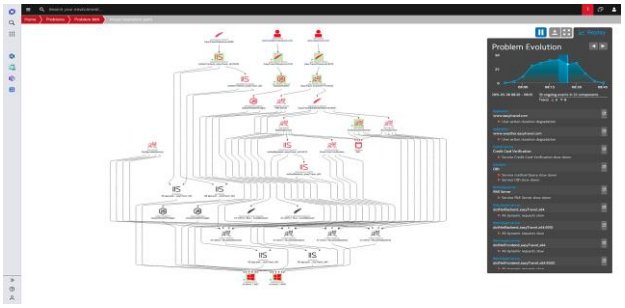

Figure 3 – Example of Alert Correlation by Dynatrace's DAVIS AI

Natural Language Processing (NLP) enables AI to analyze logs, extract meaningful information, and automatically bring that in-context to the incident [6].

AUTOMATED INCIDENT RESOLUTION AND SELF-HEALING IT SYSTEMS

AI and ML have proven to be very useful in proactively identifying the incidents and automatically get to its root cause. With the advancement of AI and ML, organizations can now leverage automated incident resolution and self-healing IT systems to minimize human involvement, accelerate issue resolution, and even prevent failures before they occur [2].

AI-based monitoring tools, monitor the health of the applications continuously. These tools automatically detect performance anomalies or failures in real time. The AI correlates the incident with historical incident data, it then selects the optimal remediation action. The remediation action is based on AI-driven runbooks that dynamically adapt based on real-time system conditions. The selection remediation action is executed, fixing the incident automatically [2].

Finally, the AI models learn from past incidents, improving remediation strategies over time [1]. Below figure 4 explains the auto-remediation workflow.

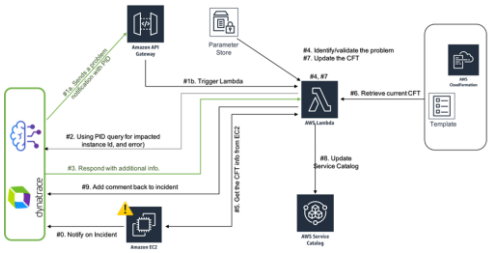

Figure 4 – Self-Healing Architecture [3]

BUSINESS BENEFITS OF AI-DRIVEN INCIDENT RESOLUTION

The adoption of AI and ML in automated incident resolution delivers significant business advantages for organizations. AI-driven incident management streamlines IT operations, increases the speed of incident resolution, and transforms incident management from a reactive to proactive process. Some key benefits of AI and ML in incident management are:

- **Faster Incident Detection and Resolution:** Traditionally incident management relies on users to report issues and manual analysis of data. This leads to high Mean Time to Detect (MTTD) and Mean Time to Resolve (MTTR). AI leverages historical data to identify incidents before they impact users. It is then able to correlates system logs, metrics, and traces to pinpoint the failure point instantly. Finally, AI is also able to execute automated fixes, reducing resolution time from hours to minutes. This significantly improves MTTD and MTTR. Thus, enhancing customer trust and experience.
- **Lower Operational Cost:** Manual incident resolution requires large IT teams and results in high labor costs. The automated incident management eliminates the need for human intervention in routine incidents, reducing IT labor costs.
- **Increased IT Productivity:** AI automates incident detection, RCA, and resolution. It frees up IT teams to focus on high-value projects, significantly benefiting the business.
- **Proactive Incident Prevention:** AI can leverage predictive algorithms to identify potential failures in advance, allowing preemptive remediation [7].

CONCLUSION

Traditional incident management process is slow, resource-intensive, and prone to human error. This could lead to extended downtime, increased operational costs, and reduced customer satisfaction. By leveraging AI-driven anomaly detection, automated root cause analysis (RCA), and intelligent self-healing workflows, organizations can significantly reduce Mean Time to Detect (MTTD) and Mean Time to Resolve (MTTR) [1][8]. With predictive analytics, intelligent alert correlation, and AI-driven automated runbooks, organizations can free up the time of the IT teams and allow them to focus on innovation and strategic growth of the organization [7].

As AI adoption continues to evolve, the future of incident management will be defined by fully autonomous IT operations. The organizations that adopt AI-powered automated incident management, gain a competitive edge over their peers.

REFERENCES

1. W. Beer, “Self-healing: Ansible Tower fixes Dynatrace-detected problems in real time,” Dynatrace News, May 08, 2017. [Online]. Available: https://www.dynatrace.com/news/blog/self-healing-ansible-tower-fixes-dynatrace-detected-problems-in-real-time/
2. A. Grabner, “Auto-Mitigation with Dynatrace AI – or shall we call it Self-Healing?,” Dynatrace News, Nov. 07, 2015. [Online]. Available: https://www.dynatrace.com/news/blog/auto-mitigation-with-dynatrace-ai-or-shall-we-call-it-self-healing/
3. “Building Self-Healing Infrastructure-as-Code with Dynatrace, AWS Lambda, and AWS Service Catalog | Amazon Web Services,” Amazon Web Services, Mar. 18, 2019. https://aws.amazon.com/blogs/apn/building-self-healing-infrastructure-as-code-with-dynatrace-aws-lambda-and-aws-service-catalog/
4. S. Kempter, “Incident Management | IT Process Wiki,” IT Process Wiki - the ITIL® Wiki, Dec. 31, 2023. https://wiki.en.it-processmaps.com/index.php/Incident_Management
5. W. Beer, “Dynatrace innovates again with the release of topology-driven auto-adaptive metric baselines,” Dynatrace News, Mar. 11, 2021. [Online]. Available: https://www.dynatrace.com/news/blog/dynatrace-innovates-again-with-the-release-of-topology-driven-auto-adaptive-metric-baselines/

6. A. Gardner, “Enhanced root cause analysis using events,” Dynatrace News, Nov. 11, 2022. [Online]. Available: https://www.dynatrace.com/news/blog/enhanced-root-cause-analysis-using-events/
7. R. Mora, “Predictive monitoring,” IBM, Aug. 13, 2020. https://www.ibm.com/think/topics/predictive-monitoring
8. Tuychlev, “A Comprehensive Introduction to Anomaly Detection,” Datacamp, Nov. 28, 2023. https://www.datacamp.com/tutorial/introduction-to-anomaly-detection