JOURNAL OF INTERNATIONAL CRISIS AND RISK COMMUNICATION RESEARCH
ISSN: 2576-0017
2025, VOL 8, NO S12

# AI-Enhanced Cloud Reliability Engineering: Integrating Incident Automation And Aiops For National Resilience

**Saravanan Raj**

*Independent Researcher, USA*

**Abstract**
The growth of cloud infrastructures has raised challenges in achieving stable digital services, as conventional manual methods of monitoring have not been sufficient to meet the size and complexity of the current distributed systems. IT operations, artificial intelligence, and automated incident response technology help in resolving these constraints through the use of machine learning algorithms to identify anomalies, alert discrepancies, anticipated failures, and perform self-remedies without human intervention. Experience in various industry areas indicates that organizations that bring these capabilities into site reliability engineering systems realize significant improvements in the count of alerts, quicker incident detection and reaction periods, automated management of common infrastructure problems, and reduced change failure rates. These technologies have been effectively implemented by financial institutions, e-commerce, healthcare, telecommunications providers, and small business consortia with higher levels of availability, and at the same time, lowered operational expenses and increased the productivity of engineers. In addition to the short-term benefits in reliability, AI-enhanced reliability practices enhance the resilience of critical digital infrastructure, facilitate environmental sustainability due to efficient resource exploitation, and aid workforce satisfaction, removing tedious toil. The change involves the initial investments in data quality, cultural embracement of blame-free post-mortem operations and error budget models, and incremental change beginning with low-stakes systems, and later to services to the customer. The challenge to policymakers and industry leaders presents the prospects of accelerating adoption by investing in open interoperability standards, workforce development initiatives, and investing in explainable artificial intelligence capabilities that provide automated systems to operate transparently and ethically and protect the interests of the nation in preserving reliable digital infrastructure underpinning economic activity, providing healthcare services, government services, and national security. While the results demonstrate consistent operational and societal benefits across sectors, this study is limited by its reliance on secondary case studies and industry reports, highlighting the need for future work on longitudinal evaluations, standardized benchmarks, and controlled empirical validation of AIOps-driven reliability outcomes.

**Keywords**: AIOps, Site Reliability Engineering, Incident Automation, Cloud Infrastructure Resilience, Machine Learning Operations.

## 1. Introduction

The pace of cloud infrastructure has grown exponentially, constructing sophisticated distributed solutions with billions of users of banking, healthcare, and retail services. Maintaining the reliability of these operations becomes more challenging as the systems continue to grow, and failure often costs the company thousands of dollars an hour as the customers demand 24/7 access. Conventional IT departments are very much dependent on manual monitoring and a reactive level of troubleshooting, which by all means cannot cope with the enormous volumes of operational data traversing the cloud environments of today.

It is important to distinguish between Site Reliability Engineering (SRE) and AIOps, as the two are often discussed interchangeably. SRE is an organizational and engineering discipline that applies software engineering principles—such as service level objectives, error budgets, and blameless postmortems—to system reliability. AIOps, by contrast, refers to the application of machine learning and data analytics to operational telemetry for tasks such as anomaly detection, event correlation, and automated remediation. In practice, AIOps acts as an enabling capability within modern SRE implementations, augmenting human decision-making rather than replacing the foundational SRE principles.

Technical documentation from AWS describes how AIOps fundamentally changes operations by combining big data analytics with machine learning to automatically spot and resolve routine IT problems, helping teams address slowdowns and outages far quicker than manual approaches allow [1]. These technologies solve a critical limitation: people cannot process endless telemetry streams, link scattered events across distributed systems, or execute repairs within tight service agreement deadlines.

Academic work in the World Journal of Advanced Engineering and Technology Sciences shows how site reliability engineering has revolutionized operations through automated monitoring, intelligent incident response mechanisms, and constant improvement frameworks [2]. The old reactive approach—where engineers burn most hours fighting fires and answering alerts—has become unworkable for cloud-native configurations running hundreds of microservices with tangled dependencies. Financial stakes run high, too: downtime bleeds thousands per minute from company budgets, and customers demand always-on access, cranking competitive pressure higher.

This examination explores how SRE methods merge with AI-powered automation, gathering evidence from research and real deployments showing concrete gains in operational performance, cost reductions, and system durability. The review covers technical designs, deployment tactics, measured outcomes across industries, and implications for maintaining a strong national digital infrastructure during times when digital services support the economy, defense, and everyday life.

## 2. Financial Services Sector Implementation

The banking and investment companies work under especially severe conditions where reliability meets the regulations and customer trust, as well as economic stability; thus, these types of organizations are natural adopters of advanced reliability methods. A study published by the Asian Journal of Research in Computer Science investigated how large financial institutions integrated AIOps systems into the existing reliability models to consider the individual banking infrastructure issues [3]. The technical designs of major banks include machine learning engines that will continuously scan the operational data system metrics, application logs, transaction alerts, and user trends, building dynamic profiles of normal behavior and intercepting minor clues before deterioration into service failures or security violations.

AWS technical documents describe how these platforms provide advanced functionality that groups related alerts and identifies trends in apparently unrelated events, and autonomously initiates the repair processes so that operations departments can respond to incidents quickly and with precision tasks that cannot be done manually [1]. In the case of banks, this would be automated response systems to detect unusual patterns of transactions indicating fraud, spot bottlenecks in the infrastructure before customers start complaining about delays, and failover usage when regions suffer outages without the need to wait till a human responds. Major institutions formed cross-functional teams that consisted of reliability engineers, data scientists, security experts, and compliance officers who went together designing AI models, validating predictions against past incidents, and setting up governance such that automated actions did not go outside the regulatory lines, such as Sarbanes-Oxley guidelines and PCI-DSS standards.

The strategy of deployment was initiated carefully, first to internal systems where automated repair was less risky, and was later extended to customer-facing platforms as more and more quarters passed and proved successful. Studies of effective SRE programs reported by ResearchGate have identified that effective change needs more than the deployment of new tools; it needs cultural changes in which engineering teams learn to make errors, conduct reviews of incidents without blame, and experiment continuously, becoming the norm [4]. The trip was not without opposition: the teams of the operations that were used to manual processes were initially resisting, the unsystematic or disorganized logging complicated model training, and the changing systems also required that the models be retrained regularly as new failure patterns emerged.

Findings of financial deployments demonstrated significant returns in various aspects. Problems facing the customers were reduced significantly because predictive systems identified problems before the customers were affected. When automated correlation was adopted and used in place of manual investigation (looking through thousands of alerts and log entries to find root causes), the recovery time was reduced significantly. Embarking on automated systems to eliminate the regular infrastructure headaches, such as full disks, memory leaks, and network congestion, allowed engineers to do strategic work such as system redesign, capacity planning, and security hardening. The savings were realized in the form of more efficient use of resources and fewer overtime, as well as less dependence on expensive emergency support with vendors, without sacrificing the availability, which is above the industry standards and regulatory demands.

**Table 1:** Financial Services Implementation Outcomes [3, 4]

| Metric Category | Before Implementation | After Implementation | Key Improvement |
|---|---|---|---|
| Customer-Facing Incidents | High frequency | Substantially reduced | Predictive intervention enabled |
| Mean Time to Recovery | Extended duration | Dramatically shortened | Automated correlation replaced manual investigation |
| Common Infrastructure Issues | Manual resolution required | Automated handling | Freed engineering capacity for strategic work |
| System Availability | Industry baseline | Exceeded benchmarks | Maintained regulatory compliance |
| Operational Costs | Standard enterprise-level | Significant reduction | Optimized resource utilization |

## 3. E-Commerce and Healthcare Applications

### 3.1 E-Commerce Platform

The unique pressure of online retailers is that the interruption of order processing in the most active shopping hours will directly translate into the loss of profits, the tarnished image, and the abandonment of carts to other competitors, which will make their availability a business need and not a technical issue. The rapid value analysis of AIOps by BigPanda in the retail setting revealed how organizations use smart alert correlation and incident enrichment to convert the chaotic alert floods into actionable intelligence that expedites the resolution of issues [5]. Major retailers have technical architectures with several AI-powered layers: anomaly detection code that learns normal behavior by individual microservice and customer group, alert correlation engines to identify cause-and-effect relationships between infrastructure events and application symptoms, and predictive analytics to predict capacity needs by promotion calendar, seasonal trends, and real-time demand indicators. Work from Inspyr Solutions about harnessing artificial intelligence for operational efficiency describes how machine learning models trained on historical incident data automatically categorize and prioritize incoming alerts by business impact, customer exposure, and probable causes, letting support teams concentrate on critical issues while safely deferring or auto-resolving

minor events [6]. Retailers linked these capabilities to continuous deployment pipelines, forming closed-loop automation in which anomaly detection in test rollouts can roll back and so the problematic code does not reach live customers' traffic. Embedded conversational AI assistants in incident response processes can offer to on-call engineers plain-language summaries of the state of a system, common suggested troubleshooting suggestions given based on similar past incidents, and verified repair workflows that can be executed with a single approval click rather than a manual command sequence prone to human error during crises.

The retailer transformation experience showed intriguing trends about the development of AI-driven processes. Initial adoptions have occasionally had a short-term peak of identified issues as more vigilant monitoring revealed formerly obscured problems. However, as models were honed by trial and error and engineering teams' optimization work under the newfound visibility was both reduced in the frequency of incidents, and the time it took to fix them. The effects of the business were not limited to the operational figures but also to the real customer experience, as the faster page loading and a reduction in errors contributed to the increase in conversion rates and customer lifetime value in a measurable way. Co-locating developer and operations processes with common platforms encouraged joint problem solving as opposed to finger-pointing, and faster development of automated repair methods that embodied knowledge in expert experience that had been previously hidden in the experience of individual engineers.

### 3.2 Healthcare Provider

The special reliability concerns of healthcare organizations include health service interruptions, which may negatively impact patient safety and clinical outcomes, along with technical excellence and compliance with the regulations of medical equipment, privacy of patient data, and clinical decision support options. The studies published under the CEUR Workshop Proceedings investigated predictive autoscaling in healthcare cloud systems and showed that machine learning models can predict resource usage based on appointment schedules, seasonal diseases, and external variables such as weather conditions related to the number of visits to the emergency room [7]. DeepAR forecasting models were designed into the healthcare network designs, taking the historical usage data to predict future capacity requirements with enough warning to allocate serverless computing infrastructure to avoid the service slowdowns caused by under-provisioning, as well as the resource waste caused by excessive over-provisioning.

Published works in the ACM Digital Library on AI-based incident management in healthcare IT infrastructure explain that the combination of AIOps features with electronic health records and telemedicine systems assists providers in maintaining their services during peak hours and ensures automated responses are subjected to a sufficient safety check before implementation [8]. Its rollout practice was based on defense-in-depth principles whereby it had many layers of oversight to prevent potentially harmful automated responses: simulation environments whereby proposed fixes were tested against system models before being deployed live, human-in-the-loop approval of automation to services with direct patient interaction, and audit trails of all automated decisions, showing confidence scores and factors explaining why the automation met regulatory requirements and continued to improve.

The transition of the healthcare provider to AI-enhanced reliability also involved significant spending on the underlying capabilities: extensive instrumentation that records the detailed telemetry of every system component, information governance mechanisms that ensure that patient information is not disclosed and allow analytics to be performed anonymously, and security mechanisms to prevent the manipulation of AI models by adversarial manipulators of critical services. It was observed that predictive autoscaling implementation had quantifiable responsiveness improvements by reducing cold start delays that would worsen user experience with surging demand in the past, and also decreasing unprocessed requests that might imply lost appointments or delayed clinical communication. False positive alert notifications had to be filtered out by alert correlation, which allowed IT staff to be vigilant enough to notice real incidents rather than to get used to noise.

**Table 2:** E-Commerce and Healthcare Performance Metrics [5, 6]

| Sector | Technology Deployed | Primary Challenge Addressed | Measurable Outcome | Business Impact |
|---|---|---|---|---|
| E-Commerce | Alert correlation engine with conversational AI | Alert storms during peak shopping | Alert noise has been dramatically reduced | Higher conversion rates and customer satisfaction |
| E-Commerce | Automated rollback integration | Deployment errors reaching production | Error rates declined substantially | Faster page loads and fewer customer-facing issues |
| Healthcare | DeepAR predictive autoscaling | Resource demand forecasting | Cold start delays have been reduced | Improved telemedicine responsiveness |
| Healthcare | AIOps with human approval loops | Regulatory compliance and patient safety | False positives trimmed significantly | Maintained critical service availability |

**4 SME and Telecommunications Implementations.**

Telecommunications infrastructure poses special reliability problems since the distributed network components include a wide variety of technologies, radio access networks, optical transport systems, and software-defined networking controllers, all forming distinct telemetry patterns and with varying failure modes. The studies of the AIOps architecture to support higher-order IT operations underline that telecommunication providers require advanced root cause analysis software that traverses complicated dependency graphs comprising thousands of network components that identify the underlying infrastructure failures that result in reported customer service issues [9]. At large operators, technical implementations typically apply graph neural networks, which model network topology as dynamic graphs with nodes modeling physical connection network elements and edges modeling logical connections and measured patterns of correlations between events that happen at various locations in infrastructure.

AWS documentation on AIOps functionality describes how event correlation is of particular importance in telecom settings where a fiber cut or router failure can create an alert storm in dozens of dependent services, making it difficult to see actual root causes that require remediation through the traditional monitoring scheme [1]. The AIOps platform of the telecom provider deals with this by performing multi-stage correlation algorithms that first cluster temporally and spatially related alerts, which is followed by causal inference methods that identify symptoms and root causes, and finally prioritization of repair solutions based on customer impact and not just on technical severity. Chat-based conversational interfaces allow the personnel at the network operations center to inquire about the state of the system in a common language, demand diagnostic reports of a specific circuit in a customer, and have the system make an automated suggestion of the steps to carry out during the same incident, as it has been identified to work before.

The predictive maintenance features of telecom implementations are important evolutionary advances from reactive incident response to proactive reliability engineering, where machine learning models can predict component failures days or weeks before service failures happen. Models take into account various signal error rates, temperature change, power consumption trend, manufacturer recall notice, creating a maintenance work order that allows replacement to be planned during a routine maintenance window rather than emergency repairs that break service, and cost premium rates. The change involved a significant change management effort in the organization because field technicians and network planners had to adjust their workflows to include AI-generated predictions, with a skeptical attitude toward them, but slowly gaining trust as prediction quality proved to be worthwhile.

**4.2 Small and Medium Enterprise (SMEs).**

SMBs face specific hurdles in implementing more sophisticated reliability practices, as most of them do not have the size to justify a specific SRE department or the investment of capital in an enterprise-level monitoring system, but are nonetheless expected to maintain high availability and fast responsiveness to issues by their customers. ResearchGate recorded case studies concerning organizational change that

analyzed how organizations can use groups of SMEs to achieve economies of scale through the sharing of resources that deploy shared AIOps infrastructure that would not be economically viable to individual organizations [4]. The cooperative model involves the companies involved in sharing platform costs with individualized dashboards, alert routing, and runbook execution on individual application architectures and operational requirements.

Articles in the Asian Journal of Research in Computer Science on the topic of improving site reliability engineering using AIOps frameworks on the role of open-source tools and cloud-native monitoring solutions in making advanced observability features affordable to small and medium enterprises with limited budgets [3]. Implementations of SME consortium used open-source software such as Prometheus to collect metrics, Elasticsearch to aggregate logs, Grafana to visualize data, and custom Python-based correlation engines using machine learning algorithms to detect anomalies and recommend fixes. Lightweight runbook automation targeted more popular infrastructure problems, such as container restarts when memory limits are reached, database failover when primary instances are not responding, and autoscale when request queues reach quotas.

There was a shared platform model that had to be carefully governed to guarantee equitable distribution of resources, safeguard proprietary information amongst competing organizations, and set up service level agreements as to the anticipated platform availability and responsiveness. Cost-sharing plans saved personal costs by much more than standalone implementations, allowing SMEs access to features such as twenty-four-hour monitoring coverage, AI-based triage, and automatic remediation that had to be staffed would otherwise impose a financial strain on budgets. The change produced quantifiable value, such as less manual labor to deal with tickets due to automation of routine processing errors, a lower rate of false positives due to smart correlation, and quicker response to an incident due to recommended diagnostic tests.

**Table 3:** Telecommunications and SME Deployment Results [7, 8]

| Implementation Type | Core Technology | Primary Benefit | Cost Impact | Operational Transformation |
|---|---|---|---|---|
| Telecommunications Root Cause Analysis | Graph neural networks | Identified failures in complex dependencies | Emergency dispatch costs decreased | Network operations center efficiency improved |
| Telecommunications Predictive Maintenance | Machine learning forecasting | Component failure prediction | Premium emergency repair costs avoided | Planned maintenance replaced reactive repairs |
| SME Consortium Platform | Open-source AIOps tools | Shared infrastructure economies of scale | Individual expenses have been reduced substantially | Access to enterprise-grade capabilities enabled |
| SME Automated Runbooks | Lightweight automation scripts | Routine infrastructure issue resolution | Avoided additional hiring costs | After-hours disruptions minimized |

**5 Comparative Analysis and Broader Implications.**

Academic research and real-world implementations of AI-based site reliability engineering demonstrate that the same trends can be observed in the implementation of operational results in different industries, scopes, and technical structures, where the principles of change can be considered instead of industry

factors. The analysis of the World Journal of Advanced Engineering and Technology Sciences of SRE practices indicates that organizations that have managed to incorporate automation and machine learning into the process of reliability have observed radical changes in the distribution of engineering capacity, with most of the companies that are currently involved in reliability workflows switching to emphasis on proactive system improvement, capacity planning, and architectural development [2].

The comparison of applications in the financial services, e-commerce, healthcare, telecommunications, and SME implementations shows that although each metric may vary depending on the maturity of the baseline and the context of the organization, directional improvement along the way is remarkably similar. Organizations have reported high percentages of reduction in alert noise due to intelligent correlation suppressing duplicate and low-value notifications, reduction in mean detection time as anomaly detection is applied before customer-impacting outages, and reduction in mean resolution time as automated remediation suppresses manual investigation state.

Infosys' views of site reliability engineering best practices observe that successful realizations involve regular updates to configuration based on the changes in system architecture, periodic review of service level indicators to ensure that they correspond to the business goals, and consistent investment in automation that captures operational knowledge into runbooks that can be executed [10]. Greater environmental sustainability implications arise in the area of resource utilization, which lessens unnecessary server provisioning and the ensuing energy use, and predictive autoscaling allows organizations to run near their actual demand as compared to significant buffers of capacity.

The Long-term perspective implies that as AI models keep improving, the ratio between human judgment and machine automation will change, with habitual work-specific operational choices being transferred to AI, as the work of engineers will be to identify new failure modes, architectural design, and continuous improvement projects. The future way forward involves investing in explainable AI features that allow engineers to comprehend and verify automated suggestions, bias identification systems that ensure the AI systems do not reinforce injustices, and a unified telemetry format that allows AIOps platforms to function in a heterogeneous multi-cloud setup.

**Table 4:** Cross-Sector Comparative Outcomes [9, 10]

| Performance Dimension | Traditional Operations Baseline | AI-Enhanced Implementation | Transformation Achieved |
|---|---|---|---|
| Engineer Time on Reactive Work | The majority of working hours | Minimal fraction of capacity | Shift toward proactive improvement enabled |
| Alert Volume Management | Overwhelming notification floods | Intelligently filtered actionable items | Operations teams maintain focus on critical issues |
| Incident Detection Speed | Delayed recognition after customer impact | Predictive identification before service degradation | Customer experience protected |
| Resolution Execution | Manual investigation and remediation | Automated correlation and repair workflows | Dramatic time savings realized |
| Service Availability | Declining trend under increasing load | Maintained or improved despite growth | Scalability without proportional resource increases |

## Conclusion

This study synthesizes evidence across multiple industries to show that AI-enhanced reliability engineering consistently shifts organizations from reactive incident response toward predictive, risk-aware system operations, regardless of sector-specific architectures or regulatory constraints. It has been demonstrated by financial services institutions that machine learning platforms can reduce customer-facing incidents by a dramatic factor and remain regulation-compliant, and save costs through efficient resource allocation and less overtime. E-commerce services can use intelligent alert correlation as a means of converting the flood of notifications into actionable intelligence to resolve the incidents faster, and directly into better customer conversion rates and lifetime value. To satisfy two conflicting priorities of healthcare provision, predictive autoscaling models and human-in-the-loop approval mechanisms allow healthcare providers to provide services and protect patients simultaneously by making sure that automated operations are properly validated and approved before being implemented. Telecommunications operators use graph neural networks to navigate complex dependencies of networks within network dependency structures to find root causes hidden in cascading alert storms and predictive maintenance plans, which lower the cost of emergency repairs. Small business consortia get the benefits of economies of scale because they can use AIOps platforms on a shared basis, thus democratizing access to more advanced observability features that were previously only accessible to large enterprises with large technology budgets. Quantity of real results under these various implementations show repeated trends: alert noise is reduced significantly by correlation engines that remove duplicates and notifications of low value, detection times are reduced by anomaly algorithms that detect problems before their effects on customers, resolution times are reduced by automated remediation devices that by-pass the manual investigation phase, and engineering capacity is redirected towards strategic efforts such as system redesign and capacity planning. Environmental benefits come in the form of optimized usage of resources, which leads to reduced provisioning of unused servers along with the energy consumption; social benefits are seen in enhanced reliability of digital services in facilitating access to healthcare, financial transactions, education, and communication. The way ahead requires investments in explainable AI functionality that can ensure that engineers can confirm automated recommendations, bias propagation frameworks that eliminate the continuation of inequities, and consistent telemetry forms that allow the interoperability of platforms operating in heterogeneous multi-cloud settings. The policy makers must focus on the development of open standards, workforce training that will result in multidisciplinary skills in the form of reliability engineering and machine learning skills, and investments in safety frameworks that will enable automated systems to work with reliability and ethics. The collegiate framework between academia, industry, and government will be critical in the actualization of the full potential of AI-driven reliability engineering, as well as in protecting against algorithmic bias, vulnerabilities, and over-reliance on automation, which is likely to undermine human experience. With the digital infrastructure taking on a central role in the economic prosperity, social welfare, and national security, incorporating artificial intelligence in reliability practices would not only be the technical evolution but strategic necessity to ensuring the continuity, sustainability, and security of digitally based services to support societal well being in an increasingly interconnected world where downtime liability rate runs in the thousands of dollars per minute and customer demands of 24-7 service in all sectors of the economy places competitive pressures on firms that require strategic choices to stay sustainable and viable. Despite demonstrated benefits, several challenges remain unresolved. These include the absence of standardized benchmarks for evaluating AIOps effectiveness across industries, limited transparency in proprietary machine learning models, and the operational risks associated with over-reliance on automation in safety-critical systems. Addressing these gaps will require deeper collaboration between academia, industry, and regulators, along with investments in explainable AI, shared telemetry standards, and formal validation frameworks for autonomous remediation.

## References

[1] Amazon Web Services, "What is AIOps?" [Online]. Available: https://aws.amazon.com/what-is/aiops/

[2] Saravanakumar Baskaran, "Evaluating the Impact of Site Reliability Engineering on Cloud Services Availability," World Journal of Advanced Engineering and Technology Sciences, 2020. [Online]. Available: https://wjaets.com/sites/default/files/WJAETS-2020-0016.pdf
[3] Mahender Singh, "Enhancing Site Reliability Engineering Through AIOps: A Framework for Next-Generation IT Operations," Asian Journal of Research in Computer Science, 2025. [Online]. Available: https://journalajrcos.com/index.php/AJRCOS/article/view/619
[4] Nitin Mukhi, "E-ISSN:2073-607X Case Studies on Successful SRE Implementations: An In-depth Analysis of Organizational Transformation and Career Pathways," ResearchGate, 2025. [Online]. Available: https://www.researchgate.net/publication/398004184
[5] Nathan Bao, "Achieving quick time to value with AIOps," BigPanda, 2024. [Online]. Available: https://www.bigpanda.io/blog/aiops-time-to-value/
[6] Inspyr Solutions, "Harnessing AI and Machine Learning for Enhanced IT Operational Efficiency". [Online]. Available: https://www.inspyrsolutions.com/harnessing-ai-and-machine-learning-for-efficiency/
[7] Oleksandr Kyrychenko et al., "Predictive autoscaling in AWS Serverless by means of machine learning and SQS metrics⋆," Intelitsis'25: The 6th International Workshop on Intelligent Information Technologies & Systems of Information Security, 2025. [Online]. Available: http://ceur-ws.org/Vol-3963/paper7.pdf
[8] Paolo Notaro et al., "A Survey of AIOps Methods for Failure Management," ACM Transactions on Intelligent Systems and Technology (TIST), 2021. doi: 10.1145/3483424. [Online]. Available: https://dl.acm.org/doi/10.1145/3483424
[9] UST "AI-powered root cause analysis: From RAN to Core," Asian Journal of Research in Computer Science, vol. 17, no. 11, pp. 52-73, 2024. [Online]. Available: https://www.ust.com/en/insights/ai-powered-root-cause-analysis-from-ran-to-core#:~:text=The%20advantages%20of%20AI%2Dpowered%20RCA&text=Faster%20resolution%20times%3A%20AI%20narrows,to%20focus%20on%20strategic%20projects.
[10] Udaykumar Gupta and Vanishree Mahesh, "A strategic roadmap for implementing site reliability engineering practices," Infosys, 2025. [Online]. Available: https://www.infosys.com/iki/perspectives/site-reliability-engineering-practices.html
[11] P. Notaro, J. Cardoso, and M. Gerndt, "A survey of AIOps methods for failure management," ACM Transactions on Intelligent Systems and Technology, vol. 12, no. 5, pp. 1–31, Sep. 2021, doi: 10.1145/3483424. https://dl.acm.org/doi/10.1145/3483424
[12] Z. Li, J. Chen, Y. Huang, and J. Wang, "InterFusion: Multivariate time series anomaly detection with interpretation," in Proc. 27th ACM SIGKDD Conf. on Knowledge Discovery and Data Mining (KDD), Singapore, Aug. 2021, pp. 983–994, doi: 10.1145/3447548.3467075. https://dl.acm.org/doi/10.1145/3447548.3467075
[13] S. Blázquez-García, A. Conde, U. Mori, and J. A. Lozano, "A review on outlier/anomaly detection in time series data," ACM Computing Surveys, vol. 54, no. 3, pp. 1–33, Apr. 2021, doi: 10.1145/3444690. https://dl.acm.org/doi/10.1145/3444690