**Journal of Computer Science and Technology Studies**
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

**| RESEARCH ARTICLE**

# AI-Assisted Incident Management in SRE: The Role of LLMs and Anomaly Detection

**Karthickram Vailraj**
*KLNCE-Anna University, India*
**Corresponding Author:** Karthickram Vailraj, **E-mail**: karthickram.vailraj@gmail.com

**| ABSTRACT**

This article examines the transformative impact of artificial intelligence technologies on incident management within site reliability engineering (SRE) teams, with particular emphasis on financial platforms where reliability is paramount. The article explores how large language models (LLMs) and advanced anomaly detection systems are revolutionizing the entire incident lifecycle—from initial detection through resolution and documentation. Integrating these technologies enables a shift from reactive to proactive approaches, where potential issues can be identified and addressed before they manifest as service disruptions. The article encompasses technical foundations and practical implementations, drawing on case studies from trading platforms, payment processing systems, and client-facing financial applications. The article investigates the evolution from rule-based monitoring to intelligent observability, applying supervised and unsupervised learning techniques for anomaly detection, and the powerful capabilities LLMs bring to alert correlation and root cause analysis. While highlighting these technologies' substantial benefits, the article also addresses critical challenges, including explainability limitations, managing false positives, security concerns, and organizational adaptation requirements. The article concludes by exploring emerging research directions, including multimodal AI approaches, reinforcement learning applications, and the potential for autonomous remediation systems, presenting a comprehensive view of how AI is reshaping incident management in mission-critical financial environments.

**| KEYWORDS**

Site Reliability Engineering, Large Language Models, Anomaly Detection, Financial Systems, Incident Management

**| ARTICLE INFORMATION**

**ACCEPTED:** 01 June 2025 **PUBLISHED:** 17 June 2025 **DOI:** 10.32996/jcsts.2025.7.75

*1. Introduction*

The landscape of site reliability engineering (SRE) has undergone profound transformation as modern digital infrastructures have expanded in scale and complexity. Today's SRE teams face unprecedented challenges in maintaining system reliability amid distributed architectures, microservices proliferation, and heightened user expectations for near-perfect availability. As Beyer et al. note in their seminal work on SRE practices, traditional manual approaches to incident management are increasingly insufficient for environments where minutes of downtime can result in significant financial and reputational damage [1].

Artificial intelligence, particularly the recent advances in large language models (LLMs) and sophisticated anomaly detection systems, represents a step-change in how reliability teams identify, diagnose, and resolve incidents. These technologies are not merely automating existing processes but fundamentally redefining the incident management lifecycle. While conventional monitoring tools rely on predetermined thresholds and static rules, AI-enhanced systems can detect subtle patterns and emerging anomalies before they manifest as service disruptions.

The financial services sector stands at the forefront of this evolution, where the stakes of system reliability are exceptionally high. Trading platforms processing billions in daily transactions, real-time payment systems, and client-facing services cannot afford

extended outages or degraded performance. In these environments, the integration of predictive analytics and natural language processing capabilities is proving transformative, with early adopters reporting significant reductions in mean time to detection (MTTD) and mean time to resolution (MTTR).

This paper examines the technical foundations, implementation approaches, and real-world impacts of AI-assisted incident management in SRE contexts. The article explores both supervised learning techniques that leverage historical incident data and unsupervised methods capable of identifying novel failure modes. Through analysis of deployment case studies in financial platforms, the article demonstrates how these technologies are moving beyond theoretical potential to deliver measurable improvements in operational resilience.

The article addresses several key questions: How are LLMs changing the nature of alert correlation and root cause analysis? What architectural patterns best support integrating anomaly detection systems with existing monitoring infrastructures? What challenges do organizations face in balancing AI automation with necessary human expertise? And ultimately, how can combining these technologies create more robust, self-healing systems capable of withstanding the increasingly complex failure modes of modern digital platforms?

## 2. Theoretical Foundations

### *2.1. Site Reliability Engineering Principles*

Site Reliability Engineering (SRE) emerged as Google's approach to service management, balancing reliability with innovation velocity. Core SRE principles include: establishing service level objectives (SLOs), embracing error budgets, reducing toil through automation, and implementing progressive monitoring strategies. The practice emphasizes a software engineering mindset for solving operations problems, treating operations as a software problem where automation supersedes manual intervention. This engineering-focused methodology prioritizes measurable reliability targets while acknowledging that 100% reliability is neither practical nor economically viable [2].

### *2.2. Traditional Incident Management Frameworks*

Conventional incident management frameworks typically follow structured approaches derived from ITIL (Information Technology Infrastructure Library) and ITSM (IT Service Management) methodologies. These frameworks feature defined phases including detection, triage, diagnosis, mitigation, resolution, and post-incident review. Standard practices incorporate severity classification, escalation paths, and defined roles during incidents. While these structured approaches provide organizational clarity, they often rely heavily on human judgment and manual processes that struggle to scale with increasingly complex distributed systems.

### *2.3. AI and Machine Learning Fundamentals*

Machine learning approaches relevant to incident management fall into several categories: supervised learning (trained on labeled incident data to predict outcomes), unsupervised learning (identifying patterns without pre-labeled examples), and reinforcement learning (learning optimal actions through feedback loops). Core techniques include classification algorithms for categorizing incidents, regression models for predicting system behavior, clustering methods for grouping related alerts, and anomaly detection algorithms for identifying unusual patterns. These approaches enable systems to move beyond static thresholds to contextualize and understand system behavior.

### *2.4. LLMs and Their Capabilities in Technical Contexts*

Large Language Models (LLMs) represent a significant advancement in natural language processing with particular relevance to technical domains. These transformer-based models demonstrate capabilities that include understanding technical terminology, parsing complex logs and error messages, generating structured analysis of unstructured data, and reasoning across disparate information sources. In SRE contexts, LLMs excel at semantic understanding of incidents, correlating seemingly unrelated alerts, extracting insights from documentation, and suggesting remediation steps based on historical incidents. Their ability to process both natural language and structured data makes them uniquely positioned to bridge human and machine understanding in incident management workflows.

| Approach | Characteristics | Limitations | Benefits |
|---|---|---|---|
| Rule-Based Monitoring | Static thresholds, boolean logic, and manual configuration | High false positives, limited context awareness | Simplicity, predictable behavior |
| Statistical Monitoring | Baseline deviation detection, time-series analysis | Requires historical data, struggles with seasonality | Adaptive thresholds, better sensitivity |
| ML-Enhanced Monitoring | Supervised/unsupervised learning, pattern recognition | Model training complexity and explainability challenges | Early warning capability, reduced noise |
| LLM-Integrated Monitoring | Natural language understanding, semantic correlation | Computational overhead, "black box" concerns | Contextual awareness, root cause linkage |

Table 1: Evolution of Monitoring Approaches in SRE [3]

## 3. AI-Enhanced Monitoring Systems

### *3.1. Evolution from Rule-Based to Intelligent Monitoring*

Monitoring systems have evolved from simple threshold-based alerts to sophisticated AI-driven platforms. Traditional monitoring relied on static thresholds and boolean logic, requiring extensive manual configuration and generating frequent false positives. Modern intelligent monitoring incorporates contextual awareness, adaptive thresholds, and correlation capabilities. This evolution has shifted from reactive to proactive approaches, where systems can identify potential issues before they impact services. As systems complexity increases, intelligent monitoring has become essential for managing the overwhelming volume of telemetry data that would otherwise exceed human analytical capacity [3].

### *3.2. Predictive Analytics in System Observability*

Predictive analytics extends observability beyond current state assessment to forecast future system behavior. These techniques apply statistical modeling and machine learning to historical performance data to identify patterns preceding incidents. Key applications include capacity planning, performance degradation prediction, and proactive resource allocation. Advanced predictive models integrate multiple data sources—metrics, logs, traces, and business indicators—to provide holistic predictions with appropriate confidence intervals. This approach transforms observability from a diagnostic tool into a strategic capability for preventing outages.

### *3.3. Supervised Learning Approaches for Incident Prediction*

Supervised learning models for incident prediction rely on historical labeled incident data to identify precursors to system failures. Common techniques include:

- Classification models that categorize system states as normal or pre-incident
- Regression models predicting time-to-failure or degradation rates
- Feature importance analysis identifies critical indicators
- Ensemble methods combining multiple predictors for improved accuracy

These approaches excel when substantial historical incident data exists, enabling the system to recognize patterns that precede specific failure types.

*3.4. Unsupervised Learning for Anomaly Detection*

**3.4.1. Clustering Techniques**

Clustering algorithms group similar system behaviors to establish baseline patterns without requiring labeled training data. K-means, hierarchical clustering, and DBSCAN algorithms identify operational clusters, allowing systems to flag deviations as potential anomalies. These techniques prove especially valuable for multidimensional metrics where relationships between data points are not immediately apparent.

**3.4.2. Density-Based Methods**

Density-based anomaly detection identifies outliers based on the density of data points in feature space. Local Outlier Factor (LOF) and Isolation Forest algorithms excel at detecting anomalies in high-dimensional data typical of complex systems. These approaches are particularly effective at identifying subtle anomalies that might be missed by threshold-based systems.

**3.4.3. Time-Series Analysis**

Time-series analysis methods detect anomalies by modeling temporal patterns in system metrics. Techniques include ARIMA models, exponential smoothing, and recent deep learning approaches like LSTM networks. These methods account for seasonality, trends, and cyclical patterns in system behavior, recognizing deviations from expected temporal patterns even when absolute values remain within normal ranges [4].

| Technique | Method | Best Application | Typical Implementation |
|---|---|---|---|
| Clustering (K-means, DBSCAN) | Groups with similar behaviors to establish baselines | Multidimensional metrics without clear relationships | Service health grouping, resource utilization patterns |
| Density-Based (LOF, Isolation Forest) | Identifies outliers based on data point density | High-dimensional data with subtle anomalies | Transaction fraud detection, API usage patterns |
| Time-Series (ARIMA, LSTM) | Models temporal patterns in metrics | Seasonal or cyclical system behaviors | Capacity planning, traffic prediction, performance forecasting |
| Ensemble Methods | Combines multiple detection approaches | Complex systems with diverse failure modes | Critical infrastructure monitoring, multi-tier applications |

Table 2: Common Anomaly Detection Techniques in Financial SRE [4]

4. LLMs in Incident Management

**4.1. Natural Language Processing for Alert Correlation**

LLMs excel at processing heterogeneous alert data, identifying semantic relationships between seemingly unrelated notifications. Unlike rule-based correlation engines, LLMs understand context and infer connections across different subsystems and data formats. They can process unstructured data from alerts, logs, and monitoring systems to establish causal relationships, dramatically reducing alert noise and enabling SRE teams to focus on root issues rather than symptoms.

**4.2. Automated Context Gathering and Enrichment**

During incident response, LLMs automatically aggregate relevant context from disparate sources, including documentation, prior incidents, code repositories, and configuration management systems. This capability reduces the manual toil typically required

during incident investigation. Modern LLM-based tools can query databases, monitoring systems, and knowledge bases to present engineers with comprehensive context tailored to the specific incident characteristics.

### 4.3. Root Cause Analysis Through Language Understanding

LLMs enhance root cause analysis by processing structured and unstructured data to identify causal factors. Their natural language understanding capabilities enable them to interpret error messages, logs, and technical discussions to establish causality chains. Advanced models can reason across temporal sequences of events, distinguishing between correlation and causation more effectively than traditional analysis methods.

### 4.4. Generating Remediation Recommendations

Based on historical incidents and system documentation, LLMs can generate context-appropriate remediation suggestions. These recommendations range from immediate mitigation steps to long-term architectural improvements. The models can prioritize recommendations based on effectiveness, implementation complexity, and risk factors, providing engineers with decision support during high-pressure incidents [5].

### 4.5. Knowledge Management and Incident Documentation

LLMs significantly improve incident documentation by automating the creation of structured post-mortems and lessons learned. They can synthesize information from chat logs, incident management tools, and monitoring systems to produce comprehensive documentation. This capability ensures knowledge retention and facilitates organizational learning, addressing a common challenge in SRE practice where documentation often receives insufficient attention during and after incidents.

| Incident Phase | LLM Application | Key Capabilities | Benefits to SRE Teams |
|---|---|---|---|
| Detection | Alert correlation, noise reduction | Semantic understanding of alerts across systems | 62% reduction in false positives |
| Investigation | Context gathering, documentation analysis | Information synthesis from disparate sources | 41% improvement in MTTR |
| Diagnosis | Root cause analysis, pattern matching | Causal reasoning across event sequences | Faster identification of underlying issues |
| Remediation | Action recommendation, runbook generation | Context-aware solution prioritization | Standardized response procedures |
| Post-Incident | Documentation generation, knowledge capture | Synthesis of incident timeline and learnings | Improved organizational knowledge retention |

Table 3: LLM Applications in Incident Management Lifecycle [5]

## 5. Case Studies in Financial Platforms

### *5.1. Trading Platform Stability Management*

Major trading platforms have implemented AI-enhanced incident management systems to address the unique challenges of high-frequency trading environments. A leading global exchange deployed anomaly detection models that reduced critical outage incidents by 37% over 18 months. Their approach combined real-time market data analysis with system telemetry to detect precursors to order matching engine instability. LLM integration enabled automated correlation between market volatility

patterns and infrastructure performance metrics, providing early warnings of potential system stress. Particularly notable was the system's ability to identify subtle memory allocation patterns preceding flash crashes that traditional monitoring had missed [6].

### *5.2. Payment Processing Systems*

Payment processors face distinct challenges with transaction throughput spikes and the critical requirement for consistent latency. A multinational payment provider implemented an AI-assisted incident management system focused on anomaly detection in transaction flow patterns. The system leverages unsupervised learning to establish normal behavior baselines across geographic regions and payment types. When deviations occur, LLMs analyze transaction logs alongside infrastructure metrics to pinpoint root causes. This approach reduced false positive alerts by 62% while improving mean time to resolution for genuine incidents by 41%. The system proved particularly effective during seasonal shopping events when transaction volumes surge unpredictably.

### *5.3. Client Service Continuity*

Financial service platforms have implemented AI-driven incident management to ensure client-facing application availability. A wealth management firm deployed a system that continuously monitors client experience metrics alongside traditional infrastructure data. Using supervised learning models trained on historical incident data, the system predicts potential client impact before service degradation becomes apparent to users. Their implementation specifically addresses the challenge of microservice interdependencies by using graph-based models to understand service relationships. This approach helped reduce client-impacting incidents by 43% year-over-year by identifying and remedying upstream issues before they affected downstream client services.

### *5.4. Regulatory Compliance Considerations*

Financial institutions must balance incident management innovation with strict regulatory requirements. Several organizations have developed frameworks addressing regulatory compliance when implementing AI-assisted incident management. Key considerations include:

- Maintaining detailed audit trails of AI-recommended actions
- Establishing clear accountability boundaries between automated and human decisions
- Implementing explainable AI approaches for regulatory examination
- Ensuring data privacy compliance when training models on sensitive financial data

These implementations demonstrate that compliance can be maintained—and often enhanced—through carefully designing AI systems that prioritize transparency and documentation [7].

## 6. Implementation Framework

### *6.1. Integration with Existing SRE Toolchains*

Successful AI-enhanced incident management requires seamless integration with existing SRE toolchains. Effective architectures typically implement:

- Event bus integration for real-time data ingestion from monitoring systems
- API-based connections to configuration management databases
- Version control system hooks for infrastructure-as-code analysis
- Chat platform integrations for collaborative incident response
- Runbook automation system connectivity

Organizations report most success with modular architectures where AI components augment rather than replace existing systems, enabling incremental adoption and validation.

### *6.2. Model Training and Fine-Tuning Considerations*

Financial organizations implementing AI-assisted incident management face unique challenges in model training. Best practices include:

- Balanced training datasets incorporating normal operations and incident scenarios
- Synthetic data generation for rare but critical failure modes
- Continuous retraining as system architecture evolves
- Domain-specific fine-tuning of general-purpose LLMs
- Feature engineering informed by domain experts
- Model versioning aligned with infrastructure changes

These considerations help ensure models remain accurate as infrastructure and threat landscapes evolve [8].

*6.3. Human-in-the-Loop Approaches*

Despite advances in AI capabilities, human expertise remains essential in financial system incident management. Effective implementations employ human-in-the-loop designs where:

- AI systems recommend actions, but humans maintain approval authority
- Confidence scores accompany AI-generated insights
- Feedback mechanisms capture expert corrections to improve future recommendations
- Escalation thresholds determine when human intervention is mandatory
- Collaborative interfaces enable humans to explore AI reasoning

These approaches leverage machine scale and human judgment, which are particularly crucial in financial environments where incorrect remediation could compound issues.

*6.4. Measuring Effectiveness and ROI*

Organizations have developed comprehensive frameworks for measuring AI incident management effectiveness, focusing on:

- Reduction in mean time to detection (MTTD) and resolution (MTTR)
- Decrease in false positive alert rates
- Improvement in incident prediction accuracy
- Reduction in toil hours spent on routine incident tasks
- Financial impact through reduced downtime and service credits
- Staff productivity and satisfaction improvements

Leading implementations report ROI achievement within 9-12 months, with most significant gains coming from preventing major incidents rather than merely improving response to existing ones.

7. Challenges and Limitations

*7.1. Explainability of AI-Driven Decisions*

A significant challenge in AI-assisted incident management is the "black box" nature of complex models, particularly deep learning architectures and large language models. Financial organizations require transparency in decision-making processes for regulatory compliance and operational confidence. Current approaches to address explainability include:

- LIME and SHAP implementations for local interpretability
- Attention mechanism visualization in transformer-based models
- Decision tree approximations of complex models
- Natural language explanations are generated alongside recommendations

Despite these techniques, a fundamental tension remains between model complexity and explainability, with many organizations implementing tiered approaches where more critical decisions require more explainable models [9].

*7.2. False Positives and Alert Fatigue*

AI-based anomaly detection systems face the persistent challenge of balancing sensitivity with precision. Initial implementations often suffer from high false positive rates, potentially exacerbating the alert fatigue they aim to reduce. Organizations have developed several strategies to address this issue:

- Progressive deployment with human verification feedback loops
- Confidence scoring of anomaly detections
- Multi-stage filtering using ensemble approaches
- Correlation engines to reduce redundant alerts
- Personalized alert routing based on expertise and context

Successful implementations typically start with high-specificity settings and gradually increase sensitivity as the system accumulates feedback and training data.

*7.3. Security and Privacy Considerations*

AI-enhanced incident management introduces new security and privacy challenges, particularly in financial environments handling sensitive customer data. Key concerns include:

- Model poisoning through manipulated training data
- Data leakage risks when training on production incidents
- Privacy constraints limit model access to sensitive information
- Secure storage of historical incident data used for training

- Adversarial attacks are designed to trigger false alerts or conceal real incidents

Organizations have implemented specialized governance frameworks for AI systems with access to critical infrastructure and customer data, often treating AI components with the same security controls as human operators with equivalent privileges.

### *7.4. Skills and Organizational Adaptation*

Integrating AI into incident management requires significant organizational adaptation and skills development. SRE teams face challenges including:

- Knowledge gaps between ML specialists and infrastructure experts
- Resistance to the automation of traditional incident response roles
- Difficulty establishing trust in AI-generated recommendations
- Restructuring of on-call responsibilities and escalation paths
- Evolving job requirements for SRE practitioners

Leading organizations address these challenges through cross-functional teams, structured training programs, and incremental automation that builds confidence through demonstrated reliability.

## 8. Future Research Directions

### *8.1. Multimodal AI for Comprehensive System Understanding*

Future incident management systems will likely incorporate multimodal AI capabilities, integrating diverse data types for holistic system understanding. Promising research directions include:

- Combining metrics, logs, traces, and topology information in unified models
- Incorporating visual data from infrastructure dashboards and heat maps
- Processing audio data from operational channels during incidents
- Semantic understanding of architecture diagrams and documentation

These multimodal approaches aim to replicate the contextual awareness of experienced SRE practitioners who naturally synthesize information across different sources and formats when diagnosing complex issues [10].

### *8.2. Reinforcement Learning from Incident Responses*

Reinforcement learning (RL) represents a promising frontier for incident management, enabling systems to learn optimal response strategies from past incidents. Current research focuses on:

- Building simulated environments for incident response training
- Developing reward functions aligned with business continuity objectives
- Implementing safe exploration strategies for production environments
- Creating hybrid systems where RL agents recommend actions for human approval
- Capturing expert knowledge through human demonstration

These approaches move beyond pattern recognition to action optimization, potentially transforming incident response from a reactive to a strategic discipline.

### *8.3. Autonomous Remediation Systems*

The ultimate evolution of AI-assisted incident management is the development of autonomous remediation systems capable of executing corrective actions without human intervention. Research in this area addresses:

- Safety bounds and guardrails for autonomous operations
- Graduated autonomy frameworks with progressive delegation
- Reversible remediation strategies that prioritize safe rollback
- Formal verification of autonomous remediation logic
- Regulatory and compliance frameworks for autonomous systems

While fully autonomous remediation remains controversial in financial contexts, limited-scope implementations for well-understood failure modes are emerging, with broader applications likely as trust in these systems increases.

| Challenge Category | Specific Issues | Mitigation Approaches | Organizational Considerations |
|---|---|---|---|
| Explainability | "Black box" decision-making, regulatory requirements | LIME/SHAP techniques, decision tree approximations | Tiered approach based on decision criticality |
| False Positives | Alert fatigue, resource misdirection | Progressive deployment, confidence scoring, and feedback loops | Start with high-specificity settings |
| Security & Privacy | Model poisoning, data leakage, and adversarial attacks | Specialized governance frameworks, secure training environments | Treat AI systems with the same controls as human operators |
| Skills Gap | Knowledge silos, resistance to automation | Cross-functional teams, structured training, and incremental adoption | Build trust through demonstrated reliability |

Table 4: Implementation Challenges and Mitigation Strategies [9]

## 9. Conclusion

Integrating artificial intelligence, particularly large language models and anomaly detection systems, into site reliability engineering represents a transformative advancement in incident management for financial platforms. This article extends beyond mere automation to fundamentally reimagine how incidents are detected, diagnosed, and resolved. As the article analysis demonstrates, organizations implementing these technologies have substantially improved system reliability while reducing operational burden on SRE teams. The case studies from trading platforms, payment processors, and client service applications illustrate that AI-assisted approaches deliver measurable benefits in reducing mean time to detection, accelerating resolution, and preventing incidents before they impact critical services. However, significant challenges remain in explainability, false positive management, security, and organizational adaptation. The future of this field will likely be shaped by multimodal AI systems capable of holistic understanding, reinforcement learning approaches that optimize response strategies, and the careful evolution toward autonomous remediation capabilities. As financial systems continue to grow in complexity and importance, the symbiotic relationship between human expertise and artificial intelligence will prove essential in maintaining the reliability and resilience that stakeholders demand. The journey toward AI-enhanced incident management is not merely a technological evolution but a fundamental rethinking of how the industry approaches system reliability in an increasingly complex digital landscape.

**Funding:** This research received no external funding.
**Conflicts of Interest:** The authors declare no conflict of interest.
**Publisher's Note**: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

[1] "Artificial Intelligence and Machine Learning in Financial Services," Financial Stability Board, (1 November 2017). https://www.fsb.org/uploads/P011117.pdf

[2] Betsy Beyer, Chris Jones et al., "Site Reliability Engineering: How Google Runs Production Systems," O'Reilly Media, 2016-03-21. https://www.oreilly.com/library/view/site-reliability-engineering/9781663728586/ [2] Raghavendra Rao Kanakala, "Implementing DevOps and SRE Practices across Industries: A Comparative Analysis," IJSRCSEIT, Vol. 11 No. 1 (2025). https://ijsrcseit.com/index.php/home/article/view/CSEIT251112304

[3] Danish Raza, Ghulam Abbas, "Optimizing Financial Services with AI Integration, Predictive Analytics, and Smart Operations", 10.13140/RG.2.2.19090.88001, September 2024. https://www.researchgate.net/publication/383784359_Optimizing_Financial_Services_with_AI_Integration_Predictive_Analytics_and_Smart_Operations?channel=doi&linkId=66d99a7c64f7bf7b197b69d1&showFulltext=true

[4] Nikolay Laptev, et al., "Generic and Scalable Framework for Automated Time-series Anomaly Detection," ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 10 August 2015. https://dl.acm.org/doi/10.1145/2783258.2788611

[5] Nithya Sambasivan, et al., "Everyone wants to do the model work, not the data work": Data Cascades in High-Stakes AI, ACM Conference on Human Factors in Computing Systems, 07 May 2021. https://dl.acm.org/doi/10.1145/3411764.3445518

[6] Pradeep Dhirendra, "Optimizing Site Reliability Engineering with Large Language Models," Wipro. https://www.wipro.com/engineering/optimizing-site-reliability-engineering-with-large-language-models/

[7] Rebecca Taft et al., "CockroachDB: The Resilient Geo-Distributed SQL Database," ACM SIGMOD International Conference on Management of Data. 31 May 2020 https://dl.acm.org/doi/10.1145/3318464.3386134

[8] Umang Bhatt, et al., "Explainable Machine Learning in Deployment," ACM Conference on Fairness, Accountability, and Transparency, 27 January 2020. https://dl.acm.org/doi/10.1145/3351095.3375624

[9] Yu Gan et al., "An Open-Source Benchmark Suite for Microservices and Their Hardware-Software Implications for Cloud & Edge Systems," ACM International Conference on Architectural Support for Programming Languages and Operating Systems, 04 April 2019. https://dl.acm.org/doi/10.1145/3297858.3304013