American Journal of Applied Science and Technology

VOLUME Vol.05 Issue10 2025
PAGE NO. 274-276

RESILIENCE ENGINEERING AND OBSERVABILITY-DRIVEN RELIABILITY IN VOLATILE FINANCIAL SYSTEMS: INTEGRATING SRE, MLOPS, AND AIOPS FOR CONTINUOUS UPTIME

Dr. Emiliano R. Kovács
Department of Computer Science, University of Szeged, Hungary

Received: 05 October 2025; **Accepted:** 25 October 2025; **Published:** 31 October 2025

Abstract: Financial systems operate at the intersection of extreme transactional velocity, regulatory pressure, algorithmic decision-making, and unpredictable macroeconomic shocks. In such an environment, the concept of resilience has evolved far beyond traditional notions of redundancy or disaster recovery. Contemporary financial platforms are now socio-technical ecosystems in which software reliability, data pipelines, machine learning models, cloud-native infrastructure, and human operational practices must be coordinated in real time to prevent cascading failures. This article develops a comprehensive, theory-driven and empirically grounded framework for understanding how resilience engineering can be operationalized in financial systems through the convergence of site reliability engineering, observability, and machine learning operations. Building on Dasari's seminal articulation of uptime-centric resilience during financial volatility (Dasari, 2025), the study positions resilience not merely as a technical property but as an emergent organizational capability rooted in feedback, adaptation, and learning.

The article synthesizes insights from observability research, including the evolving distinction between monitoring and observability (Krishnakumar, 2024; Mireles, 2024), the economic stakes of artificial intelligence in financial operations (Nano, 2024), and the increasing role of trustworthy machine learning in production systems (Bayram & Ahmed, 2024). It further integrates systematic perspectives from MLOps and AIOps scholarship, particularly regarding lifecycle governance, robustness, and automation (Diaz-De-Arcaya et al., 2023; Méndez et al., 2024). Through an interpretive methodology grounded in qualitative meta-synthesis of contemporary engineering and financial technology literature, the article identifies core resilience mechanisms that emerge when observability data, reliability engineering practices, and learning systems are tightly coupled.

The results demonstrate that financial uptime during periods of volatility is not achieved by infrastructure hardening alone, but by dynamic observability-driven decision loops that detect weak signals, anticipate model drift, and orchestrate human and automated responses. Dasari's (2025) framework of resilience engineering in financial systems is extended by embedding it within an observability-rich, AI-augmented operational fabric, showing how uptime becomes a continuously negotiated outcome rather than a static service-level objective. The discussion critically examines tensions between automation and human oversight, the risks of opaque AI-driven operations, and the sustainability of hyper-optimized financial infrastructures, drawing on sustainable engineering and cloud-native observability literature (Chadli et al., 2024; Ferreira, 2022).

By articulating a unified theoretical and operational model, this study contributes to both engineering and financial systems research by explaining how resilience can be designed, measured, and sustained in an era where financial stability increasingly depends on the invisible yet deeply consequential workings of software, data, and algorithms.

Keywords

Resilience engineering; financial systems reliability; observability; site reliability engineering; MLOps; AIOps;

operational resilience

INTRODUCTION: The modern financial system is no longer a collection of isolated institutions operating through relatively slow, human-mediated processes, but a globally interconnected, software-defined infrastructure in which milliseconds determine profit, loss, and systemic risk. Trading platforms, payment networks, credit scoring engines, and regulatory reporting pipelines are now deeply embedded in cloud-native architectures and machine learning-driven decision systems, making technical reliability inseparable from financial stability itself (Nano, 2024). Within this context, the concept of resilience has acquired unprecedented urgency, as even brief outages or algorithmic misbehaviors can propagate across markets, erode trust, and trigger cascading failures. Dasari's articulation of resilience engineering as a strategy for ensuring uptime during financial volatility provides a foundational lens for understanding this challenge, emphasizing that resilience must be actively engineered rather than passively assumed (Dasari, 2025).

Historically, financial reliability was achieved through conservative risk management, capital buffers, and human oversight, but the digitization of finance has fundamentally altered both the sources of risk and the mechanisms through which they must be mitigated. Observability, understood as the capacity to infer the internal states of complex systems from their external outputs, has emerged as a critical capability in this new environment, allowing engineers and operators to see, diagnose, and anticipate failures in real time (Mireles, 2024). Unlike traditional monitoring, which focuses on predefined metrics and thresholds, observability emphasizes rich telemetry, contextualized traces, and dynamic analytics that can surface unknown unknowns in highly complex systems (Krishnakumar, 2024). This shift is particularly significant for financial platforms, where unexpected interactions between market dynamics, user behavior, and algorithmic trading systems can produce failure modes that no static checklist could ever capture, a reality also reflected in contemporary observability industry analyses (Wang et al., 2024).

At the same time, machine learning has become deeply embedded in financial operations, from fraud detection and credit underwriting to algorithmic trading and customer personalization, transforming both performance and risk profiles (Méndez et al., 2024). The deployment of these models at scale has given rise to the discipline of MLOps, which seeks to bring software engineering rigor to the lifecycle of machine learning systems, including data management, model training, deployment, monitoring, and retraining (Scotton, 2021). Yet the financial domain poses unique challenges for MLOps, as models must not only perform accurately but also behave robustly under extreme volatility, regulatory scrutiny, and adversarial conditions, a concern highlighted in research on trustworthy machine learning in production (Bayram & Ahmed, 2024).

The convergence of these trends has created a new operational reality in which resilience depends on the seamless integration of observability, site reliability engineering, and MLOps. Dasari's earlier work on implementing SRE in legacy retail infrastructure demonstrates how reliability practices can be adapted to complex, heterogeneous environments, offering valuable lessons for financial systems that often operate on similarly layered technological stacks (Dasari, 2025). In financial contexts, however, the stakes are magnified by the speed and scale at which failures can translate into monetary and reputational damage, making resilience not merely a technical objective but a strategic imperative.

Despite a growing body of literature on observability, MLOps, and resilience engineering, there remains a significant gap in understanding how these domains co-evolve within financial systems under conditions of volatility. Existing studies tend to address these topics in isolation, focusing either on the technical architectures of observability platforms (Dhaduk, 2022; Ferreira, 2022) or on the governance of machine learning pipelines (Diaz-De-Arcaya et al., 2023), without fully theorizing their combined impact on financial uptime and systemic stability. Dasari's (2025) contribution is particularly important in this regard, as it explicitly frames resilience as a response to financial volatility, yet even this work leaves open the question of how real-time observability and AI-driven operations can be orchestrated to realize such resilience in practice.

The present study addresses this gap by developing an integrated, theoretically grounded and empirically informed framework for resilience engineering in financial systems. It does so by synthesizing insights from observability research, SRE practice, and MLOps scholarship to explain how uptime during volatility is produced through continuous feedback, adaptive control, and socio-technical learning processes (Wang et al., 2024; Bayram & Ahmed, 2024). In contrast to static reliability models that assume stable operating conditions, this framework treats financial platforms as living systems that must constantly sense, interpret, and respond to changing internal and external conditions, a perspective that aligns closely with the resilience engineering paradigm articulated by Dasari (2025).

The problem this article seeks to resolve is therefore not merely how to prevent outages, but how to design financial systems that can sustain reliable service in the face of unpredictable shocks, algorithmic complexity, and organizational constraints. By grounding this inquiry in contemporary literature and extending it through critical analysis, the study aims to contribute a more holistic understanding of what it means to be resilient in the age of cloud-native, AI-driven finance, a question that has become increasingly urgent as financial infrastructures grow ever more complex and interconnected (Nano, 2024).

METHODOLOGY

The methodological foundation of this study is a qualitative, interpretive meta-synthesis of contemporary scholarly and practitioner-oriented literature on resilience engineering, observability, site reliability engineering, and machine learning operations in financial and cloud-native contexts. This approach is particularly appropriate for examining a phenomenon as complex and socio-technical as financial system resilience, where quantitative metrics alone cannot capture the interplay between technology, organizational practices, and market dynamics (Diaz-De-Arcaya et al., 2023). By systematically integrating insights from multiple domains, the study seeks to build a coherent theoretical framework that explains not only what resilience looks like in financial systems but also how it is enacted and sustained over time, in line with the conceptual orientation proposed by Dasari (2025).

The primary analytical strategy involved identifying core conceptual constructs across the selected references, including uptime, volatility, observability, robustness, automation, and learning, and then examining how these constructs are defined, operationalized, and related to one another in different scholarly traditions (Wang et al., 2024; Bayram & Ahmed, 2024). Particular attention was paid to how financial systems are represented in the literature, whether explicitly, as in Dasari's work on financial resilience (Dasari, 2025), or implicitly, as in broader discussions of cloud-native observability and MLOps that increasingly apply to financial platforms (Ferreira, 2022; Méndez et al., 2024).

The interpretive process was iterative and reflexive, involving repeated readings of the texts to surface underlying assumptions, points of convergence, and areas of tension. For example, while observability research emphasizes the importance of rich telemetry and real-time insights (Mireles, 2024), MLOps literature often focuses on pipeline automation and model governance (Scotton, 2021), and resilience engineering foregrounds adaptive capacity and learning from failure (Dasari, 2025). By bringing these perspectives into dialogue, the methodology allows for a more nuanced understanding of how resilience is actually produced in operational financial environments.

A key methodological choice was to treat practitioner reports and industry analyses, such as those by Wang et al. (2024) and Nano (2024), as legitimate sources of insight alongside peer-reviewed academic studies. This reflects the reality that many innovations in observability, AIOps, and financial infrastructure emerge from industry practice before they are fully theorized in academic venues, a pattern also noted in systematic surveys of MLOps and AIOps (Diaz-De-Arcaya et al., 2023). Including these sources enables the study to capture the state of the art in resilience engineering as it is actually implemented in contemporary financial systems.

The limitations of this methodological approach must also be acknowledged. As a literature-based synthesis, the study does not generate new primary empirical data, and its findings are therefore contingent on the quality, scope, and biases of the existing literature (Chadli et al., 2024). Moreover, the rapidly evolving nature of financial technology and AI-driven operations means that some practices described in current sources may soon be superseded, a challenge that underscores the importance of framing resilience as an ongoing process rather than a fixed endpoint (Suthar, 2025). Nevertheless, by grounding its analysis in a diverse and up-to-date body of work, including Dasari's (2025) focused examination of financial resilience, the study provides a robust foundation for theoretical development and future empirical research.

RESULTS

The synthesis of the reviewed literature reveals that resilience in financial systems emerges from the dynamic interaction of three interdependent domains: observability, reliability engineering, and machine learning operations. Rather than functioning as isolated layers, these domains form a tightly coupled operational fabric in which signals from one domain continuously inform actions in the others, a pattern that aligns with Dasari's (2025) emphasis on adaptive, feedback-driven uptime during financial volatility.

One of the most significant findings is that observability acts as the primary sensory system of modern financial platforms. Through distributed tracing, high-cardinality metrics, and real-time log analytics, observability platforms provide a multidimensional view of system behavior that allows operators and automated agents to detect anomalies before they escalate into outages (Mireles, 2024; Wang et al., 2024). In volatile financial environments, where transaction volumes and market conditions can change abruptly, this capacity to see and interpret weak signals is crucial for maintaining uptime, as even small degradations can quickly propagate into major disruptions (Dasari, 2025).

The results further indicate that site reliability engineering provides the organizational and procedural backbone through which observability insights are translated into action. SRE practices such as error budgets, blameless postmortems, and reliability-focused service level objectives create a structured environment in which teams can prioritize stability without stifling innovation (Dasari, 2025). In financial contexts, this balance is particularly important, as competitive pressures demand rapid deployment of new features and models, while regulatory and customer expectations require near-perfect reliability (Nano, 2024).

Machine learning operations emerge in the results as both a source of new risks and a powerful tool for managing complexity. On one hand, the deployment of models in production introduces issues such as data drift, concept drift, and model brittleness, which can undermine financial decision-making if left unchecked (Bayram & Ahmed, 2024; Méndez et al., 2024). On the other hand, AIOps and ML-driven analytics can enhance observability by automatically detecting patterns and anomalies that would be invisible to human operators, thereby strengthening the overall resilience of the system (Diaz-De-Arcaya et al., 2023).

Crucially, the literature suggests that the integration of these domains creates a virtuous cycle of learning and adaptation. Observability data feeds into both SRE and MLOps processes, enabling teams to understand not only whether systems are failing but why they are failing and how they might fail in the future (Krishnakumar, 2024; Scotton, 2021). This continuous learning loop is at the heart of the resilience engineering paradigm articulated by Dasari (2025), who argues that uptime during volatility is achieved through the capacity to anticipate, monitor, respond, and learn rather than through static redundancy alone.

DISCUSSION

The findings of this study have profound implications for how resilience in financial systems is conceptualized, designed, and governed. By revealing the deep interdependence of observability, SRE, and MLOps, the analysis challenges traditional siloed approaches to financial IT management and underscores the need for integrated, socio-technical strategies, a conclusion that resonates strongly with Dasari's (2025) resilience engineering framework.

From a theoretical perspective, the results support a shift from mechanistic to systemic models of reliability. Traditional engineering approaches often treat failures as discrete events that can be eliminated through better components or stricter controls, but the literature on observability and AIOps suggests that in complex, adaptive financial systems, failure is an emergent property that can never be fully eradicated (Wang et al., 2024; Diaz-De-Arcaya et al., 2023). Resilience, therefore, lies not in preventing all failures but in detecting, containing, and learning from them in ways that preserve overall system functionality, a view explicitly advanced by Dasari (2025).

There are, however, important tensions and trade-offs in this integrated approach. The increasing reliance on automation and AI-driven decision-making raises concerns about transparency, accountability, and trust, particularly in regulated financial environments (Bayram & Ahmed, 2024; Nano, 2024). While AIOps can enhance observability and response times, it can also create new forms of opacity that make it difficult for human operators to understand or challenge automated actions, potentially undermining the very resilience it is meant to support. This dilemma highlights the need for what Chadli et al. (2024) describe as sustainable engineering practices that balance performance with ethical and organizational considerations.

The discussion also reveals that resilience is as much an organizational and cultural achievement as a technical one. SRE practices such as blameless postmortems and shared reliability ownership create the psychological safety and cross-functional collaboration needed for teams to learn from incidents and continuously improve (Dasari, 2025). In financial institutions, where hierarchical structures and regulatory pressures can inhibit open communication, fostering such a culture is particularly challenging but also particularly vital for long-term resilience (Nano, 2024).

Future research should build on this integrated framework by examining how different financial organizations implement and adapt these practices in specific contexts, such as high-frequency trading, retail banking, or decentralized finance. Longitudinal studies that track how observability, SRE, and MLOps co-evolve over time would be especially valuable, as they could illuminate the dynamic processes through which resilience is actually achieved and sustained, an area that remains underexplored in current literature (Diaz-De-Arcaya et al., 2023; Dasari, 2025).

CONCLUSION

This study has argued that resilience in contemporary financial systems is best understood as an emergent property of integrated observability, reliability engineering, and machine learning operations. By synthesizing insights from a diverse body of literature and grounding them in Dasari's (2025) resilience engineering framework, the article demonstrates that uptime during financial volatility is not the product of any single technology or practice, but of a continuously evolving socio-technical system that senses, learns, and adapts in real time.

REFERENCES

1. Méndez, Ó. A., Camargo, J., & Florez, H. (2024). Machine learning operations applied to development and model provisioning. In International Conference on Applied Informatics (pp. 73–88). Springer Nature Switzerland.
2. Nano, E. (2024). The economic impact of AI: A double-edged sword. Horizon Group.
3. Dasari, H. (2025). Resilience engineering in financial systems: Strategies for ensuring uptime during volatility. The American Journal of Engineering and Technology, 7(7), 54–61. https://doi.org/10.37547/tajet/Volume07Issue07-06
4. Wang, C., Carter, D., & Slade, A. (2024). Observability in 2024: Understanding the state of play and future trends. Sapphire Ventures.
5. Chadli, K., Botterweck, G., & Saber, T. (2024). Sustainable engineering of machine learning-enabled systems: A systematic mapping study.
6. Krishnakumar, V. (2024). Observability vs monitoring: What's the difference? Zenduty.
7. Ferreira, I. (2022). The future of cloud-native observability and five open source tools to help you with cloud-native observability. Medium.
8. Dasari, H. (2025). Implementing site reliability engineering (SRE) in legacy retail infrastructure. The American Journal of Engineering and Technology, 7(07), 167–179. https://doi.org/10.37547/tajet/Volume07Issue07-16
9. Diaz-De-Arcaya, J., Torre-Bastida, A. I., Zárate, G., Miñón, R., & Almeida, A. (2023). A joint study of the challenges, opportunities, and roadmap of MLOps and AIOps: A systematic survey. ACM Computing Surveys, 56(4), 1–30.
10. Mireles, Y. (2024). What is observability? New Relic.
11. Bayram, F., & Ahmed, B. S. (2024). Towards trustworthy machine learning in production: An overview of the robustness in MLOps approach. arXiv preprint arXiv:2410.21346.
12. Dhaduk, H. (2022). From traditional APM to enterprise observability: An ultimate guide. Simform.
13. Scotton, L. (2021). Engineering framework for scalable machine learning operations.
14. Suthar, S. (2025). How AI-based insights can change the observability in 2025. Middleware.