International Journal of AI, BigData, Computational and Management Studies
Noble Scholar Research Group | Volume 7, Issue 1, 123-125, 2026
ISSN: 3050-9416 | https://doi.org/10.63282/3050-9416.IJAIBDCMS-V7I1P120

Original Article

Self-Healing Infrastructure - Predictive Automation for High-Availability Banking Systems

Tripatjeet Singh
Senior Cloud Engineer, Dallas-Fort Worth, USA.

Received On: 26/12/2025 ***Revised On: 30/01/2026*** ***Accepted On: 06/02/2026*** ***Published On: 18/02/2026***

Abstract - *Modern banking platforms rely even more on intricate, cloud-native architectures to provide real-time, always-on financial services. Reliability of infrastructure has increased significantly, but operational downtime from unsafe changes, configuration drift and human error remain a leading cause of service disruption. This paper presents a prevention-first, self-healing infrastructure paradigm based on predictive automation, built-in governance and analytics-driven use of AI. Instead of concentrating on the reactive remediation, this approach is interested in proactive aversion of high-risk system states before they actualize and affect customers.*

Keywords - *Self-Healing Infrastructure, Predictive Automation, Artificial Intelligence, Banking and Financial Systems, Cloud Governance, Operational Resilience, Security, Compliance.*

1. Introduction

Digital adoption, cloud computing, and increased reliance on platforms that are always available, have had a significant impact on banking and financial services. Payments, lending, and account management are core banking functions that are now performed by highly distributed, cloud-native systems that operate continuously and on large scale. Customers want uninterrupted access to their financial data and services, and regulators demand strong controls, traceability, resilience, and accountability at every level of the technology stack in use. In this setting, even short outages or problems with operations can cause customers to lose trust, draw the attention of regulators, create reputational damage, and have direct financial impact [1][2].

Redundancy, disaster recovery planning, and reactive incident management have traditionally been used to make banking systems highly available. These mechanisms are still important, but they are not enough in today's world, where failures are more often caused by unsafe changes, configuration drift, or operational complexity than by problems with the underlying infrastructure. Given the rapid evolution of technology, systems are increasingly expanding their capabilities to automate usability. This growth and pace lead to an increase in the chances of human error, which makes the reactive approach both expensive and ineffective.

To the best of the author's knowledge, no previous research has established self-healing in regulated banking systems as a prevention-oriented, governance-based framework enhanced by cost-aware artificial intelligence. This paper contends that real self-healing must surpass post-failure remediation and concentrate on proactively averting unsafe system states, prior to customer impact. Banking platforms can become more resilient, improve their compliance posture, and keep customer trust without adding unnecessary complexity or cost by adding predictive automation, governance controls, and selective AI-driven risk assessment directly into operational workflows.

2. Challenges in Modern Banking Operations

Banking systems operate under the very strict rule of regulation, security, and compliance constraints. Auditing and repeating every change made to the infrastructure and application systems is required for both internal governance and external compliance needs. These regulations make it very difficult to adopt any fixes or experiments hence, operational discipline is mandatory with these systems.

As banking platforms scale across multiple environments and services, the volume and frequency of change increases significantly which in turn introduces associated risk in the form of configuration drift, inconsistent deployments, or unintended system interactions. Even a minor misconfiguration can cascade into larger service disruptions, particularly in systems supporting real-time financial transactions.

Conventionally used incident response processes heavily depend on human intervention for resolving incidents successfully and reducing the risk compared to an automated system. Though these processes are intended to reduce risk, they often increase Mean Time to Resolution (MTTR) [4] and result in corrective actions only after customers have experienced service degradation. This reactive approach negatively impacts customer trust, increases operational cost, and exposes institutions to heightened regulatory scrutiny, indicating the importance of a more proactive and resilient operational model.

3. Limitations of Reactive Self-Healing

The existing self-healing solutions are heavily concentrated on runtime remediation mechanisms such as automated application restarts, dynamic scaling of resources, or traffic rerouting after a failure has already occurred[7][8]. Though this might function in a consumer environment or an environment that is lightly governed, they are often insufficient in banking systems where post-failure behavior still must comply, coexist with audits, and operate in a maximally controlled environment.

Not just that, but even after responding to the issue in a post-customer incidence situation in a bank and a more governed environment, the response to the failure may still not comply with regulations and may still pose a risk of exposing sensitive information. Additionally, post-failure responses from self-healing do not really address the underlying issues by just trying to fix problems that originally created the failure. These limitations highlight the need for prevention-first approaches that reduce the possibility of failures before they impact customers or regulatory stance.

4. Predictive Automation Model

Predictive automation has shifted the paradigm of resiliency from reactive response to proactive prevention by injecting risk awareness into operational workflows. Rather than waiting for failures to occur, this model promotes the concept of pre-validated changes, standardized architectural patterns, and automated policy enforcement. It analyzes the changes even before they are promoted into production to minimize the risk of bad misconfiguration, policy violations, and the overall consistency.

Standard patterns ensure consistency across environments, minimizing configuration drift and reducing human error. Automated validation and policy enforcement act as guardrails to prevent overly operationally risky changes to the environment from being successfully propelled through the delivery pipeline without appropriate controls or approvals. As a result, reliability becomes embedded into the system, and not a retrofitted one.

By integrating predictive automation into delivery and operational workflows[4][5], it significantly reduces dependency on manual intervention, reduces resolution timelines, and enables banking systems to achieve higher resilience while maintaining compliance, auditability, and operational efficiency. While predictive automation lays the foundational controls for prevention, its effectiveness can be further improved through selective use of artificial intelligence to identify complex risk patterns and detect subtle anomalies that are difficult to capture otherwise.

5. AI-Assisted Self-Healing Loop

Artificial intelligence helps improve predictive automation by identifying risk patterns that are difficult to catch using fixed rules or manual reviews. Past deployment data, incident history, configuration changes, and operational signals are reviewed to assess the risk of proposed changes before they are promoted into production. At the same time, anomaly detection is used to monitor system behavior and identify unusual patterns that may point to configuration drift, access issues, or early signs of failure.

AI is not used to make automatic repair decisions on its own[7] [9]. Instead, it supports existing controls and human judgment within defined governance boundaries. As shown in Figure 1, this loop connects change review, risk assessment, policy checks, controlled execution, and continuous feedback to help prevent issues before they affect customers. To ensure that cost is under control, AI analysis runs mainly in batch or low-frequency mode, using existing data rather than expensive real-time processing.

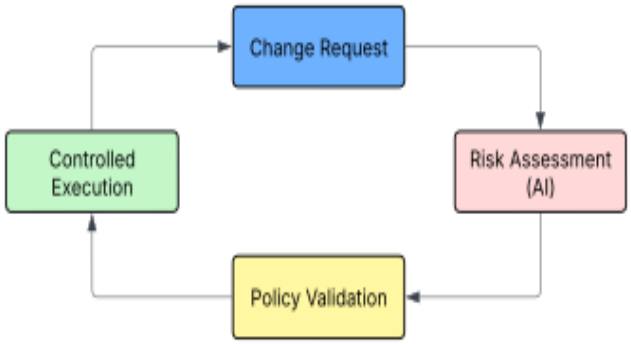

Fig 1: AI-Assisted Self-Healing Loop for Prevention-First Banking Systems

6. Governance and Compliance as Reliability Enablers

In regulated financial institutions, governance and compliance are often viewed as barriers that slow down innovation and delivery. When correctly applied, in practice, they can enhance reliability and confidence in the system. In the proposed model, governance is baked directly into everyday workflows rather than enforced as a separate, manual process[1].

By collecting evidence through automation, having checks on policy, and using standardized ways of delivering, one meets with continuous compliance, rather than just during periodic audits. This diminishes last-minute fixes, lessens audit stress, and decreases operational risk. Due to governance being predictable and repeatable, the teams could move faster with fewer errors, while regulators get a sharper look at how the systems are built, changed, and operated. Over time, this approach strengthens both system stability and organizational trust.

7. Architectural View of Self-Healing Systems

Figure 2 Illustrates the high-level conceptual architecture of a predictive self-healing banking platform showing how operational signals, predictive analytics, governance controls, and safe execution mechanisms interact to prevent unsafe system states [2] [4].

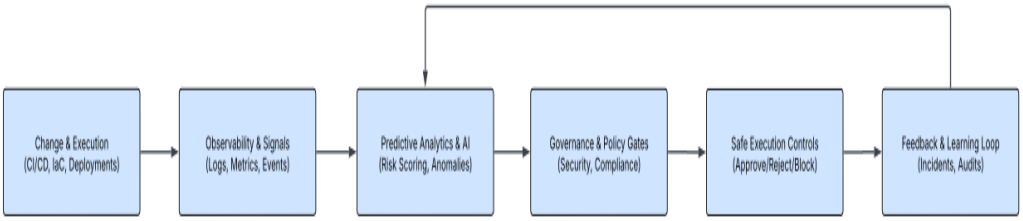

Fig 2: Conceptual Architecture of AI-Assisted Predictive Self-Healing Infrastructure

8. Operational and Business Impact

When implemented within a large banking environment, the self-healing model was successful in reducing the number of incidents in the production system, as well as reducing the amount of time needed to identify and address system operational challenges. The system became stable since the introduction of the model prevented unsafe operational modifications from entering the production system. From a business perspective, the system was successful in reducing operational costs while generating trust within customers. Most importantly, the model demonstrated that reliability, governance, and efficiency can be improved together through thoughtful architectural design rather than reactive operational fixes[2].

9. Cost-Aware Design Considerations

To ensure a practically feasible approach, a measure of cost awareness has also been incorporated into the design. Rather than adopting a continuous real-time analysis approach, batch analysis and evaluation based on meaningful events[5] [9], such as a deployments or configuration changes are encouraged. In this way, a better understanding of the operation can be achieved, all without increasing infrastructure costs. Ultimately, artificial intelligence should only used when it can be demonstrated to be adding real value, thus improving reliability and stability while reducing the overall operation and infrastructure costing.

10. Future Direction

Future improvements to the model can focus on refining existing capabilities rather than introducing major complexity. Risk thresholds can be adjusted gradually based on past incidents and system behavior to improve decision accuracy over time. Rollback strategies may be enhanced by using predefined recovery paths for common failure scenarios, allowing faster response with less manual effort. Integrating basic business indicators, such as service criticality or transaction volumes, can help prioritize operational decisions more effectively. Additionally, extending the model to support hybrid and multi-cloud setups would allow organizations to apply the same reliability and governance practices across environments they already operate today [2][5].

11. Conclusion

This paper introduced a prevention-first approach to self-healing infrastructure designed specifically for regulated banking environments [1]. The model focuses on preventing risky system states through predictive and proactive automation, integrated governance, and strategically used artificial intelligence rather than depending on reactive fixes. By embedding reliability and compliance directly into everyday workflows, the approach supports stable, high-availability platforms without adding unnecessary complexity or cost. The model demonstrates that operational resilience, regulatory alignment, and efficiency can be achieved together, offering financial institutions a practical path toward more reliable and trustworthy digital banking systems.

References

1. Smallstep. SSH Certificate Login Tutorial. [Online]. Available: https://smallstep.com/docs/tutorials/ssh-certificate-login/
2. Basel Committee on Banking Supervision, Principles for Operational Resilience, Bank for International Settlements, 2021. https://www.bis.org/bcbs/publ/d516.pdf
3. Amazon Web Services, Resilience on AWS, AWS Whitepaper (Online), 2024. https://pages.awscloud.com/rs/112-TZM-766/images/01%20Resilience%20on%20AWS%20-%20Final.pdf
4. Amazon Web Services, Amazon Web Services' Approach to Operational Resilience in the Financial Sector & Beyond, AWS Whitepaper (Online), 2023. https://docs.aws.amazon.com/pdfs/whitepapers/latest/aws-operational-resilience/aws-operational-resilience.pdf
5. Google, Site Reliability Engineering (SRE) Book – Table of Contents, sre.google. https://sre.google/sre-book/table-of-contents/
6. Microsoft, Azure Well-Architected Framework – Reliability, Microsoft Learn (Online), 2023. https://learn.microsoft.com/en-us/azure/well-architected/reliability/
7. Microsoft, Reliability Design Principles, Microsoft Learn (Online), 2023. https://learn.microsoft.com/en-us/azure/well-architected/reliability/principles
8. Q. Cheng et al., AI for IT Operations (AIOps) on Cloud Platforms (Online), 2023. https://arxiv.org/pdf/2304.04661
9. Z. Yazdanparast et al., A Survey on Self-healing Software System (Online), 2024. https://arxiv.org/pdf/2403.00455
10. Z. Zhong et al., A Survey of Time Series Anomaly Detection Methods in the Context of AIOps (Online), 2023. https://arxiv.org/abs/2308.00393.