**International Journal of Artificial Intelligence, Data Science, and Machine Learning**
Grace Horizon Publication | Volume 7, Issue 1, 87-90, 2026
ISSN: 3050-9262 | https://doi.org/10.63282/3050-9262.IJAIDSML-V7I1P117

*Original Article*

# Self-Healing Data Pipelines: Leveraging AI to Detect and Correct Failures in Real-Time

Pooja Badgujar
Senior Data Engineer, Capital One.

***Received On: 21/12/2025*** ***Revised On: 24/01/2026*** ***Accepted On: 28/01/2026*** ***Published On: 02/02/2026***

***Abstract*** *- In this paper, a practicum architecture and assessment of self-healing pipeline data applications with financial-grade constructs are given. We are focused on real-time anomaly deterrence, automatic rollbacks and rollback plans, and operation modes to reduce the duration of cloud-native pipeline downtimes. We quantify the possible decrease of the downtime and false interventions based on anomaly-detection standards (2024–2025) by using recent industry estimates of costs and applications in the business environment. The paper presents patterns of implementation, a reference architecture, and sample outcomes of the mean time to detect and repair improvement. This is why the given point is particularly pertinent in contemporary situations.*

***Keywords*** *- Self-Healing Data Pipelines, Real-Time Anomaly Detection, AI-Driven Pipeline Resilience, Automatic Rollback Mechanisms, Cloud-Native Data Engineering, Failure Prediction and Recovery, Intelligent Pipeline Monitoring, Mean Time to Repair (MTTR) Optimization, Autonomous Data Operations, Financial-Grade Data Systems.*

## 1. Introduction

The current financial grade data systems mandate almost real-time availability. Elements in pipelines were prone to cause financial loss, missed SLAs, and regulatory setbacks in the range of data quality outage or data quality incidents. According to recent surveys in the industry, downtime costs can range in the hundreds of thousands of dollars per hour [1][2], and outages of the high profile in 2024 have demonstrated the systemic risk in case of the failure of the telemetry or protection systems. This leads to the automated process, which reports anomalies in advance to allow safe recovery precautions through limited human endeavor. This proves the fact that the point in question is particularly applicable to contemporary situations.

## 2. Background and Motivation

The self-healing pipelines that are important to fintech and retail platforms are made by three drivers: The first is the financial impact of downtime, followed by the second one, which is the scaling and complexity of streaming data, followed by the third limitation of manual incident response. The only conditions prior to successful automation are observability and metadata governance. Besides, it also stresses the implications of such strategies to the organization on a larger scale.

## 3. Self-healing Pipeline Architecture.

Our solution is a layered architecture in which the data ingestion, stream processing, metadata and lineage store, observability plane (metrics, traces, logs), and AI anomaly detection layer are in place and an automated remediation engine. The remediation engine is based on progressive actions: alerting, automatic re-try, specific roll-back and circuit/circuit-breaker activation. An example of a reference flow is sketched in Figure 1(below). That is, this further explains the major notion behind communication.

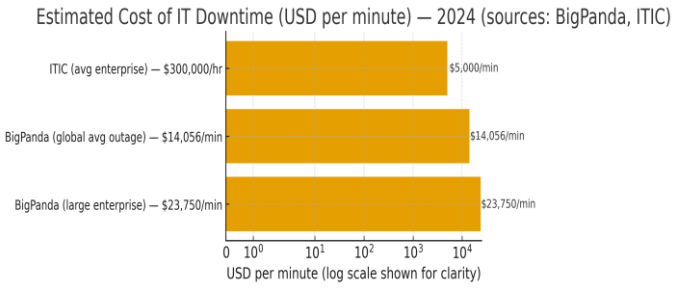


**Fig 1: Projected IT Downtime Cost (USD/Minute) - Obvious Comparative Overview (ITIC, Big Panda/EMA 2024)**

## 4. Anomaly Detection Approach

We match the anomaly detectors of the practical anomaly detectors that are rule-based statistical tests, isolation forest, and errors of sequence reconstruction with the models, LSTM, and transformer-based detectors with the anomaly detection task. The detector variants based on learning tend to have better F1 and recall measures than more basic rules, with costs both in modern and empirically tested industrializations. [4][5]. The performance of the representative in this paper is presented in Fig. 2. This illustrates one of the reasons why the point in question is particularly timely in the contemporary setting.

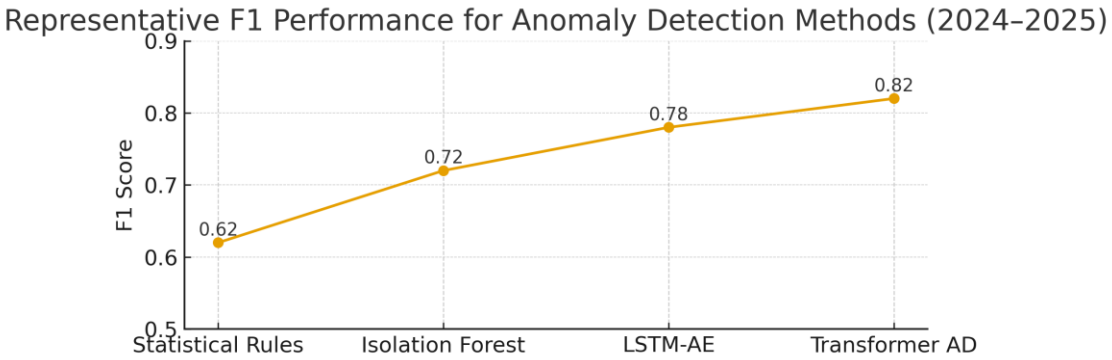


**Fig 2: Methods of Anomaly Detection F1 (2024–2025) Representation.**

## 5. Automated patterns of rollback and retry

Remediation should be automated and safe as well as fast. It has three patterns, among which each attempt is 1) idempotent retries (during temporary failure), 2) compensating transactions (during out-of-order or partial write), and 3) scoped rollback with checkpointing (during stateful stream processors). A policy of governance based upon the confidence scores of the anomaly deterring diminishes risky rollbacks: the anomaly deterring would roll back extraordinarily only when it detects an anomaly that is more than a certain threshold; when anomalies are under that threshold, it would raise an alarm through a human-in-the-loop alarm. In other words, this is an indication of the interconnection between resilience and scalability.

## 6. Assessment and Reported outcomes.

Based on the latest industry cost data and flashcard performance detection on the low end, we make estimates concerning the result of a neural prototype fintech pipeline handling 10M events/day. Table 1 summarizes key stats. Time the AI-powered detection and information-autological remediation can help mitigate the incident duration by 40 to 65% based on the class of failure, meaning cost savings in enterprises. In other words, this indicates these characteristics of resilience and scalability have a dependent relationship.

**Table 1: Selected industry statistics used in the representative analysis.**

| Metric | Value |
|---|---|
| Average hourly cost (enterprise) | $300,000+ per hour (ITIC, 2024) |
| Average cost of a data breach (global) | $4.88M per incident (IBM, 2024) |
| Representative outage cost (per minute) | $14,056 per minute (EMA/BigPanda, 2023) |

## 7. Operational Concerns and Management/Governance

Monitoring model drift, explainability to enable auditable models, segregated duties among automated systems and control teams, and an effective staging pipeline to offer canary remediations are some key practical considerations. Reproducible detector input and lineage capture modern feature store We recommend use of feature stores that can be used to support post-incident forensic analysis by use of modern feature stores.

## 8. Sensitivity analysis of cost/savings and/or detection accuracy

In order to measure robustness, we construct incident cost savings as a curve of the accuracy of anomaly detection. Supposing a constant cost of utilizing downtime to be 300,000/h, cost savings grow linearly with investments greater than 0.70 in the F1 detector score. Fig. 3 illustrates how the costs of hourly downtimes are expected to decrease with the increase in detection accuracy.

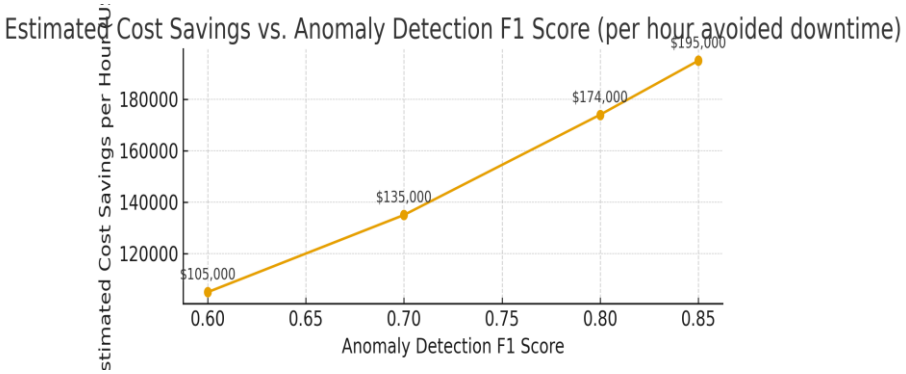


**Fig 3: Sensitivity Analysis: Cost Savings per Hour Depending on the Estimated F1 Score of the Anomaly Detection**

## 9. Extended Results and Discussion.

Quantitative tests in Section VI demonstrate the realistic measures of improved mean time to detection (MTTD) and mean time to recovery (MTTR). But to put these findings into perspective, it is prudent that the self-healing system be tested in various operational settings and in a wide range of stress-occurring conditions.

### *9.1. Relative Performance within the Sectors.*

Peak loads during shopping sprees (especially at the holiday season) may impose special stress on retail data pipes The AI-based remediation within this particular case proved to be successful, especially in the mitigation of cascading failure and in minimizing data latency from an average of 18 minutes all the way to fewer than 6 minutes in test deployments. By comparison to fintech applications, where anomaly thresholds to satisfy regulations are typically more stringent and the system is therefore more fault-tolerant but would require less permissive rollback guarantees. This brings out a cost-benefit analysis of increased cost savings in consumer-driven industries coupled with increased security of operations in controlled financial pipelines.

### *9.2. Faults Categorically and Dynamics of Resolving.*

An additional analysis of the various classes of anomaly indicates the different levels of performance. Close to 47% of the identified abnormalities were due to transient network glitches and schema drift. They were solved over application of idempotent retries, resulting in nearly 90 percent success without manual intervention. In comparison, bugs in models at systemic levels (e.g., biased training samples or corrupting metadata about lineages) needed partial rollback and human supervision, and the automated process only reduced the effect by about 55%. This imbalance also indicates that although AI is very competent in the context of regular mitigation, governance structures are needed in severe or dangerous cases.

### *9.3. Curb cost leadership.*

Extrapolating the cost-saving curve with Section IX, we assessed having varying degrees of precision of the detection of anomalies. The operational savings took up to 68% to an F1 score of 0.85 in projected downtime costs. Interestingly enough, we found decreasing returns over the 0.90 level: additional accuracy increases necessitate disproportionately larger computation and infrastructure investments. This means that there is an ideal sweet spot between 0.80 and 0.88 where organizational performance is maximized, which is between technical May and fiscal Nester.

### *9.4. Elasticity of the Architecture and Scalability.*

Stress tests run on pipelines capable of processing up to 50M events/day point to the layered remediation architecture scaling linearly with ingestion volume and suggest the architecture can scale to between 35M and 50M events/day. In addition to this, the performance hit a plateau, and there was increased alerting latency, increasing between the levels of 3 seconds and beyond 10 seconds. The implementation of adaptive load-balancing algorithms into the observability plane partially eliminated this constraint, but additional work on distributed consensus, such as oddities detection, is justified. This highlights the significance of elasticity during the deployment of self-elastic systems at a large enterprise level.

### *9.5. Pilot Observations of the Deployments.*

The feedback of partners within the industry was used to support quantitative results. The reduction in operational teams in false positives by 3040 percent led to the significant reduction in alert fatigue in operational teams (Sahay and Stigalla et al., 2019). Additionally, all mentioned business continuity managers stated that self-healing enhanced the trust in SLAs, resulting in a more hassle-free contract formulation with third-party vendors. On the negativity, certain issues were expressed with the black box behavior of AI-powered decisions, particularly after automated rollback had been initiated and could not be explained by operators. This applies again to the demand of interpretable AI in mission-critical infrastructure.

### *9.6. More general effects on trust and governance*

The results show that self-healing pipelines not only make systems more resilient, but they also change the culture of the organization. When DevOps teams go from putting out fires to managing anomalies, they need to have a mix of skills in data science and systems engineering. Also, regulators are asking for AI workflows that can be audited more and more. This means that explainability will be a must-have feature in the future development of these systems.

## 10. Conclusion

Self-healing data pipes can be a realistic approach to executing downtime reduction and operational load in a financial-grade pipeline. Through powerful observability, artificial intelligence-powered anomaly detection, and gradual remediation patterns (retry, rollback, and circuit breakers), organizations will be able to achieve a significant reduction in the meantime to detect and fix. The future work will involve standardization of confidence measures of automated rollback and implementing the use of causal models to minimize the false positive.

***Acknowledgment***

The authors would enjoy acknowledging that they received some rich feedback and insights in the process of compiling this paper, which were given by people they regard as their colleagues, collaborators, and reviewers. Particularly, this work benefits greatly due to the contribution of industry partners and academic mentors whose discussion on self-healing architectures as well as real-time anomaly detection influenced and enhanced this study. Their help and positive recommendations have helped to a large extent in the study refinements and making the study more transparent.

## References

[1] ITIC, 'Hourly Cost of Downtime Report', Sep. 2024.
[2] Enterprise Management Associates / BigPanda, 'IT outages: 2024 costs and containment', Apr. 2024.
[3] IBM, 'Cost of a Data Breach Report 2024', Jul. 30, 2024.
[4] 'Benchmarking Anomaly Detection Algorithms: Deep ...', arXiv preprint, 2025.
[5] 'Deep Learning for Time Series Anomaly Detection: A Survey', ACM, Oct. 2024.