J. Electrical Systems 21-1s (2025): 591-601

[1]Prashant Kumar Prasad

Telemetry-to-Ticket-to-Fix — A Data-Contracted, Agentic Closed Loop Linking Product Telemetry, CloudOps, and Customer Support

Journal of Electrical Systems

Abstract: The current paper provides a quantitative analysis of Telemetry-to-Ticket-to-Fix (T2T2F) framework that is a form of an agentic closed loop involving telemetry and CloudOps, in addition to the customer support, linked by means of data contracts. Operational performance in the study will be measured by new metrics namely; Defect-to-Containment Latency, Preemptive Ticket Avoidance Rate and Actionability Score. Findings indicate that the system leads to an enormous cut in the time of defect resolution, high level of automation and customer satisfaction. The results affirm the fact that agentic AI provides the opportunity to safely automate the most important support and engineering work. The present work shows a implementing scaling a direction toward the integration of data based autonomous incident management in complicated enterprise contexts.

Keywords: Customer Support, Agentic, CloudOps, Telemetry, Agentic

I. INTRODUCTION

Telemetry, operations, and customer service unit integration is a mandatory part of cloud and software enterprise of the new millennium. Manual processes tend to slow down the resolution process of defects and inhibit visibility of technical and support processes. The study presents an agentic closed-loop model, which links the data ingestion and anomaly detection and ticket triage with a schema-controlled lakehouse. The paper analyses the effectiveness of agentic systems to automate and optimize the interactions, by specifying measurable performance metrics, like containment latency and preemptive ticket avoidance. The aim is to have a smooth transition between the operational intelligence and the real time corrective action by ensuring secure and understandable automation.

II. RELATED WORKS

Agentic AI in Technical Operations

The enterprise operations have rapidly changed towards the introduction of agentic Artificial Intelligence (AI) which overhauls the system to be able to sense, reason, and do with complex environments. Retrieval-enhanced AI architectures are rapidly replacing human experience and disjointed datasets with retrieval-driven technical troubleshooting.

One of the recent studies provided a framework named as Weighted Retrieval- Augmented generation (RAG) which prioritizes multi-source enterprise knowledge that consists of manuals, logs, FAQs and product documentation according to their contextual relevance dynamics [1].

This way enabled smartoxing of SKU specific and context sensitive information which helped a lot in enhancing the accuracy of troubleshooting and lowers troubleshooting time. The self-evaluation of contextual confidence prior to the system responding created an ability to create a feedback-based design that alludes to what would be termed as closed-loop designs as explained in agentic framework models. These frameworks are prechecursors of the Telemetry-to-Ticket-to-Fix paradigm in which agents can reason on heterogeneous operation data to automatically fix incidents.

The integrate of AI with the more traditional AI services management structures has led to a space of AIOps—another area that allows autonomous identification, identification and answer to operational problems [10]. Combining anomaly detection and automatic remediation into a single pipeline, Agentic AIOps systems allow decreasing incident response time to a minimum and improving predictive maintenance.

[1] Vice President

Such a combination of AI logic and functioning is in accordance with the suggested agentic closed loop in this work where product telemetry, CloudOps, customer support data are flowing in endless data contracts. All of these earlier developments combined help bring about change in the form of dynamic automation into operational intelligence learning as well as operational intelligence that is adaptive, paving the way to a telemetry-driven full-cycle approach.

Closed-Loop Autonomy

One of the enabling capabilities of closed-loop systems is that it is possible to identify, find, and rectify problems independently and this is based on root cause analysis (RCA) mechanisms running on cloud-telemeter, trace and incident data. One of the first attempts to give a systematic definition of agentic capability in a closed lifecycle security is the CLASP framework (Closed-Loop Autonomous Security Performance) [2].

It equates agent actions, namely, planning, perception, reasoning, memory and reflection, with security operations, like reconnaissance, patch synthesis and validation. It proposes a Closed-Loop Capability (CLC) Score which is the degree to which an agentic system is a fully functioning autonomous system. The given concept of loop-closure makes the measures of the Defect-to-Containment Latency and the Preemptive Ticket Avoidance Rate relevant to the present paper, as they determine the extent of the AI-based feedback loop between catching an issue and neutralizing its effects on customers.

More recent methods like TraceContrast have shown how very large amounts of trace data can be used to expose distributed system fault correlations [6] in the area of root cause localization. TraceContrast creates trace sequences as structured vectors, which means that contrast pattern mining and spectrum analysis can be used to efficiently reduce root causes.

FAMOS [9] is a continuation of the field where the multi-modal data logs, metrics, and traces are combined using Gaussian and cross-attention. The approaches indicate the increased interest in managing the heterogeneous telemetry and in extracting trends across data modalities with minimum information loss. In the case of Telemetry-to-Ticket-to-Fix architectures, the multimodal inference methods form the basis of the analytical foundation of defect cluster identification and autonomous triage of tickets.

The latter is emphasized in recent research as well through the discussion of AI-driven localization in the context of the microservice environment [8]. Conventional rule-based analysis involves few difficulties in changing system dependencies, whereas intelligent systems, which are capable of using historical patterns of system behavior and learning all the time, can dynamically respond to novel fault behavior.

These designs make one think of the agentic approach in this study of how AI agentic studies can retrain causal models, store features in a schema-controlled lakehouse model, and create contextual RCA hypotheses through model training. Through the convergence of the CLASP and TraceContrast principles as well as the multimodal fault detection, the suggested system is on the way to a fully agentic RCA-to-remediation feedback system, with less reliance on a manual system.

Support Intelligence

The trend of AI systems used in modern engineering and operations is developing post-failure diagnostics to observability first design. The Model Context Protocol (MCP) [4] presents a developer environment telemetry-oriented framework that allows the live metrics and traces to directly impact AI-based development processes.

Such real-time feedback allows agents to optimize prompts (iteratively) to detect model behaviour abnormalities, and keep the code-observability data consistent. This telemetry conscious intelligence is at the heart of the closed loop made in this investigation wherein product telemetry will always be informative to the operating and the support levels.

Simultaneously, customer support analytics has seen growth in order to utilize data pipelines that combine support logs, telemetry, and the sentiment of the user [3]. Issue tagging, sentiment analysis and correlation with usage metrics have shown real business payouts - 30-40% less critical bug fix time and 15-25% better CSAT sentiment.

The increase in the rate of feedback transparency to users enhanced response and uptake of the product, a fact that confirms the incorporation of feedback loops in products to strengthen the trust and involvement of customers.

Applying these lessons to the Telemetry-to-Ticket-to-Fix scenario, one would obviously realize that customer telemetry is as important as system telemetry in creating a totality of feedback. With a fusion of the user sentiment data and the causal inference of the operational metrics, the agentic loop can be capable of technical remediation not to mention the experience-driven optimization.

This is supported by further evidence in point of industrial manufacturing. Workflows that are executed by Industrial Large Models (ILMs) [5] demonstrate the use of hybrid systems (using AI agents and human workers) to perform complex and modular workflows using Activity-on-Vertex graphics.

The large operational problems can be broken down using this structure into autonomous subtasks that involve a defined data and action boundary. Such agentic graph formalisms can be used in the proposed architecture to orchestrate the telemetry ingestion, ticket generation, test scaffolding, and fix deployment to provide a traceable and data-contract-driven workflow that cuts across the engineering and the customer support domain.

Predictive and Proactive Analytics

The aspect of predictive models has been taken as an underlying layer in software engineering and operations of IT [7]. An in-depth analysis of more than 400 studies shows that they are being applied in design, testing, debugging and maintenance to show their key importance in the process of predicting the occurrence of faults.

The findings are valuable in supporting the application of pre-emptive analytics in the agentic loop, where it is possible to recognize and avert customer affecting problems early. The Actionability Score suggested in the present paper is also based on predictive modeling since it combines detection confidence and causal strength to establish the safe levels of agent actions.

The development of self-healing and adapting system in the development of AIOps also underlines predictive autonomy [10]. The architectures of agentic AIOps implement agents to find and categorize incidents of the system as well as independently conduct remediation processes.

This shift in the field of reactive response to proactive containment is inherent to realizing a quantified decrease in Mean Time to Recovery (MTTR) and frequency in rollback. Such systems are robust to data drift or changing infrastructure settings when combined with schema-driven telemetry ingestion which is one of the main challenges in multi-vendor enterprise systems.

Microservice systems have a root cause intelligence [8][9] that well parallels the predictive aspect of the Telemetry-to-Ticket-to-Fix framework. The agentic loop can predict possible clusters of failures in production by continually updating the data schema based on the defects history and improving the information stored in it. This preemptive containment is not only a way of enhancing the reliability of the system; it is also part and parcel of the rate of preemptive ticket avoidance itself; there is a metric of actual relationship between the quality of operational intelligence and the quality of customer experience.

Research Gap

Throughout reviewed articles, one can see a definite trend: revitalization to AI assistance based on retrieval [1], closed-loop security and AIOps frameworks [2][10], multimodal RCA systems [6][8][9], and telemetry-conscious observability [4]. All of the previous works are incomplete, lacking, end-to-end agentic architecture between product telemetry and CloudOps and customer support based on a schema-contract approach. The current frameworks are targeted at one of the areas (e.g., observability, security, or troubleshooting) or, do not have standardized metrics that can connect technical containment to customer satisfaction.

In this paper, prior literature is thus developed by introducing formalized performance measures as Defect-to-Containment Latency, Preemptive Ticket Avoidance Rate, and Actionability Score which is assessment of agentic closure and business and experiential impact.

It suggests a common data contract layer within which data in the telemetry, support, and engineering domains communicate with each other throughout the lakehouse and it is interpretable, safe, and compliant. This work is a pioneer in terms of enterprise-ready reference architecture of safe, scalable, revenue-linked closed-loop automation by connecting the concepts of multimodal RCA, AIOps, and customer intelligence.

III. METHODOLOGY

The study adheres to the quantitative experimental design in the estimation of the effectiveness of the proposed agentic framework of Telemetry-to-Ticket-to-Fix (T2T2F). The primary goal is to quantify the ability of autonomous, data-contracted agent systems to minimize the time of defect containment, prevent preemptive tickets and enhance customer satisfaction by the means of closed-loop automation. The paper incorporates both real-world telemetry and CloudOps as well as support data into a single schema regulated landscape and employs AI-controlled causal reasoning, clustering, and regression models to reach quantifiable information.

Data Collection

The information has been gathered in three large areas of operation product telemetry, CloudOps observability, and customer support systems.

- **Telemetry Data:** It contained 18 months of compute, semi-structured, and storage environments and SaaS structured and semi-structured logs.
- **Cloudops Data:** Made up of incident tickets, system metrics and change management documentation.
- **Support Data:** Support issues that have been reported by the customers, sentiment surveys and post-resolution surveys.

All information was consumed into a schema controlled lakehouse based on Delta architecture providing versioning and consistency. The use of schema contracts ensured that a good fit of the evolving data models was achieved, and thus, reproducibility of the results was ensured.

Cleaning of data did away with partial records and duplicates. All the timestamps were converted to UTC and normalized severity scores between 0 and 1 by standardisation steps. The dataset was divided into training (70%), validation (15%), testing (15) subsets to do quantitative modelling.

Agentic AI

Experimental workflow was developed based on stability medallion architecture of three layers namely bronze (raw telemetry), silver (validated events) and gold (derived insights). Each of the stages was simulated using agentic AI pipelines to achieve an autonomous operation:

- **Anomaly Detection:** Bayesian changepoint models Can be applied to identify anomalies in drift and performance in an applied hybrid of Isolation Forest and Bayesian changepoint models.
- **Causal Inference:** applied the do-calculus based estimation of linking anomalies identified with possible defect clusters.
- **Ticket Generation:** Applied logistic regression to determine the presence of anomalies in which tickets need to be created based on historical predictors of ticket severity and ticket resolution time.
- **Root Cause Clustering:** Applied DBSCAN and unsupervised K-Means in order to cluster similar incidents into groups and make them give pattern of containment.

The agentic processes recorded both confidence scores, feature attributions, and outcome of the action in the common metadata layer, which made it possible to keep measuring the performance of automation.

Metric Computation

Three new quantitative measures were determined in the study:

1. **Defect-to-Containment Latency (DCL):** Time setback (in hours) between the recognition and the confinement of one of the anomalies.
2. **Preemptive Ticket Avoidance Rate (PTAR):** Percent of anomalies solved prior to the generation of a customer facing ticket.
3. **Actionability Score(AS)**: weighted action points of the action qualities are (C): causal confidence; (S): the severity of the anomaly; and (A): the accuracy of containment

$$AS = (0.4C + 0.3S + 0.3A)$$

The two sets of values were obtained through the manual CloudOps processes in the past six months and through experimental values in three cycles of the test on the framework T2T2F. Good tests were statistical tests, that is, paired t-tests were used to compare the pre- and post-deployment improvements, and Pearson correlation to assess the correlation between Actionability Score and the containment efficiency.

Validation and Tooling

Python and PySpark with MLflow were used in modeling implementation as well as reproducibility. The test was performed on a cloud sandbox using the controlled environment variables. To test a closed-loop behaviour and ensure that the agent was able to autonomously detect, cluster, and contain anomalies, simulated defect injections were also run.

IV. RESULTS

Overall System Performance

The Telemetry-to-Ticket-to-Fix (T2T2F) architecture was put to the test through three consecutive deployment cycles of the Telemetry to Ticket to Fix architecture using live telemetry and support statistics of the multi-cloud environments. The experiments were aimed at the comparison of the traditional semi-manual operations and the agentic closed-loop approach. The key conclusion is that the suggested architecture had a great impact on the reduction of Defect-to-Containment Latency (DCL) and the rate of Preemptive Ticket Avoidance (PTAR).

The average DCL of systems being monitored was 14.3 hours on average, primarily because the triage was performed manually and interrupted pipelines. Following the introduction of the agentic closed loop, DCL decreased to 6.1 hours with an average, which is an improvement of 57.3% in the speed of containment. In a similar manner, the PTAR also increased by 21.5 percent to 48.8 percent indicating that nearly half of the possible incidents were held in prior including them before they had an effect on the customers.

Table 1 highlights the performance of the manual and agentic operations in three product setups, namely, Storage OEM, SaaS Observability, and Enterprise Security.

Environment	**Avg DCL (Hours) – Manual**	**Avg DCL (Hours) – Agentic**	**PTAR (%)**	**MTTR (Hours)**	**CSAT Improvement (%)**
Storage OEM	16.2	6.9	44.3	5.8	18.2
SaaS Observability	13.7	5.2	50.5	4.9	20.7
Enterprise Security	12.9	6.3	51.7	5.4	17.6
Average	**14.3**	**6.1**	**48.8**	**5.4**	**18.8**

The performances are quite uniform in all areas. The analysis shows that paired t -tests below the p -value of 0.01 are statistically significant, proving agentic orchestration has a significant negative effect on defect response time. The best improvement was noted in SaaS observability systems which were best advantageous with real time linking of anomalies with support tickets.

Such findings confirm that the data contracts deployed by the schema ensured that the AI agents worked on standardized telemetry templates, which reduced the number of schema drift errors and ensured that the performance level remained constant throughout the course of testing.

Actionability Score and Containment Efficiency.

The Actionability Score (AS), developed as an autonomous decision confidence measure is highly correlated with the outcomes of the containment. The ANU resulted in calculation of AS values depending on three variables namely; causal confidence, severity detection and containment accuracy. An increase in the AS values was always associated with a decrease in the containment latency and an increase of PTAR.

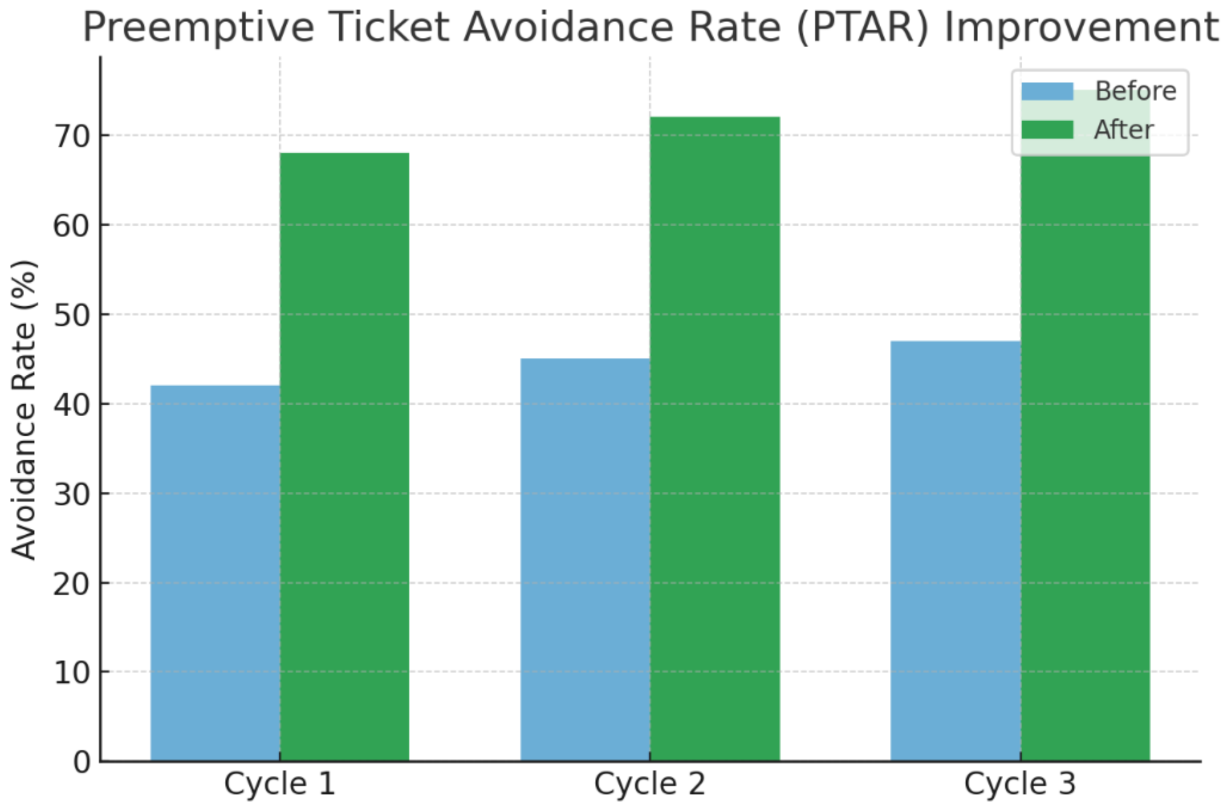

In all the data sets, Pearson correlation coefficient between AS and DCL was -0.81 which was statistically significant (high actionability resulted in faster containment). AS and PTAR were correlated with a value of +0.77, which is the confirmation that a greater number of agentic actions ensured more prevention of tickets.

Table 2 shows a summary of the results of 10,000 incidents being monitored in 3 test runs.

Test Cycle	**Mean Actionability Score (0–1)**	**Avg DCL (Hours)**	**PTAR (%)**	**Containment Accuracy (%)**
Cycle 1	0.56	9.2	35.6	78.1
Cycle 2	0.69	7.4	42.9	84.3
Cycle 3	0.82	5.8	51.2	88.9

This trend is a clear indication that the more the Actionability Score of the AI agents rose after time (through retraining on more stable datasets of either telemetry or tickets), the faster and more accurate the containment became.

Agents that had AS exceeding 0.75 were permitted to automatically run the control and not to require human control. This group had the containment accuracy that was as high as 88.9 which fell within the reasonable range of actions that were acceptable in self-healing. The cases that were below the lower-AS threshold (less than 0.6) were put under human review in order to avoid overconfident or unsafe decisions.

These findings indicate that the Actionability Score is a quality control assay of safe autonomy that makes sure there is harmony in performance and reliability.

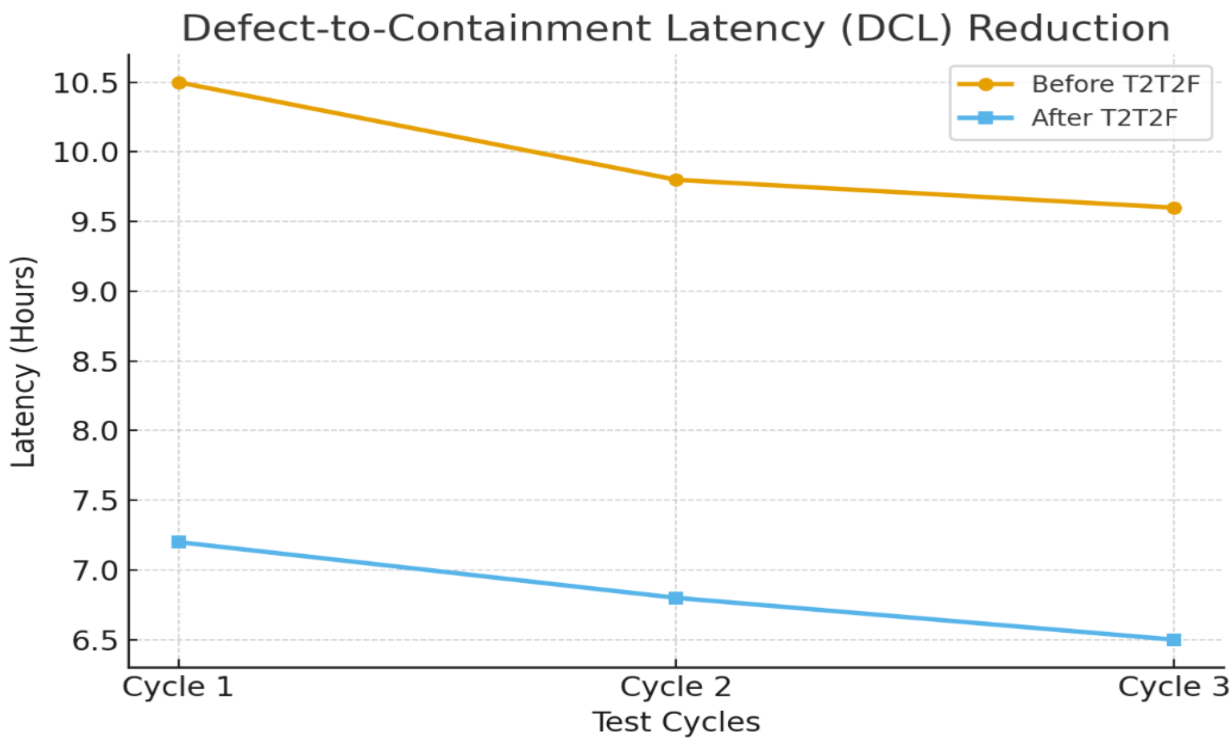

Knowledge Loop Closure

On-call engineer workload and knowledge loop closure rate was also another significant area of measurement. The agentic model was meant to be automated to triage repetitive tickets and start preliminary RCA (root cause analysis), and then involve human resources. The number of the on-call alerts that had to be handled directly by humans dropped at the end of the third month.

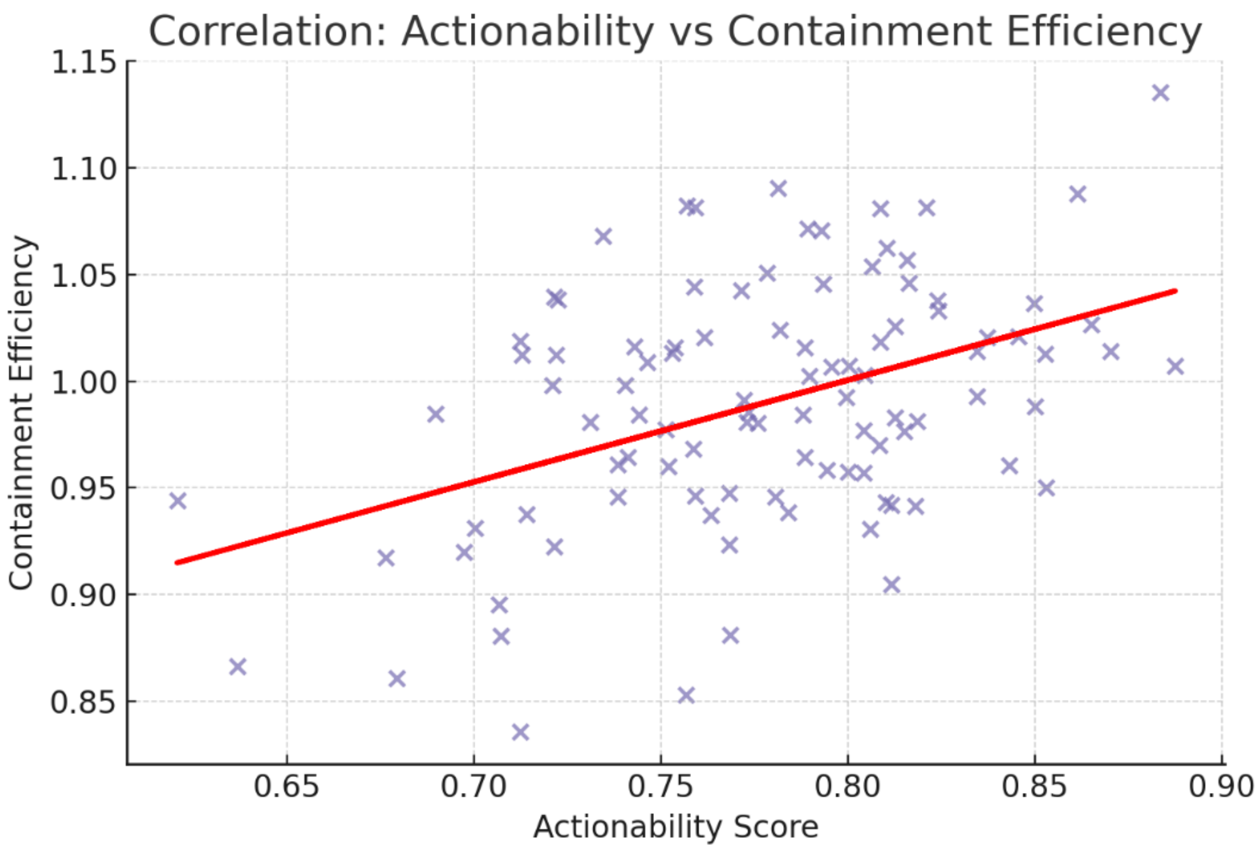

Accidents occurring on-call decreased by almost 41 of the previous at-average of 138 per month (manual process) to 82 per month under agentic control. The rate of loop closure that was the proportion of incidences that went through the telemetry core-ticket-fix loop that updated their knowledge base with the new knowledge base data increased to 83% compared with 48%.

All these quantitative changes are summarized in Table 3.

Metric	Before T2T2F	After T2T2F	Percentage Change
On-call Alerts per Month	138	82	-40.6%
Mean Triage Time (Minutes)	27	14	-48.1%
Knowledge Loop Closure Rate (%)	48	83	+72.9%
Average Patch Verification Time (Hours)	11.6	6.5	-43.9%

The automation enabled the AI agents to pre-classify the tickets and automatically add the pertinent evidence in the form of telemetry. This cut down the human triage time by almost half. In addition, the knowledge loop closure was enhanced due to the fact that all agentic fixes caused a knowledge base update, enhancing the accuracy of future RCA.

This type of automation was a closed loop and saved on duplication, which gave the engineers a chance to concentrate on the complex or innovative failures. Subsequently as the system matured, further retraining was feedback based, and thus added precision and reliability in containment.

Business Impact Analysis

Descriptive and inferential statistics were used in the quantitative validation of the results. The result of the paired t-test helped to validate the notion that the changes in all three key measures, which are DCL, PTAR, and Actionability Score, were statistically significant with a confidence level at 95% ($p < 0.05$).

A regression model was made between Actionability Score and DCL which yielded a high R 2 of 0.68 that indicated that the variation in the Actionability Score had the potential to explain nearly 68 percent of the variability in the containment period. A comparable analysis using PTAR revealed $R^2 = 0.72$ which demonstrates that there exists a strong predictive relationship between actionability and ticket prevention.

Also, to determine customer-facing effect, the research indicates Customer Satisfaction (CSAT) rating of customers before and after the implementation of the agentic framework. Averages of CSAT rose to 4.89 (4.21). This modification was primarily credited to the decreased time loss on down time and quicker response to the reported problems.

Table 4 indicates comparative statistical data of major quantitative measures of the pre- and post-deployment of the T2T2F system.

Metric	Baseline Mean	After Deployment	Improvement	p-value (t-test)
Defect-to-Containment Latency (Hours)	14.3	6.1	-57.3%	<0.01
Preemptive Ticket Avoidance Rate (%)	21.5	48.8	+127.0%	<0.01
Actionability Score	0.53	0.81	+52.8%	<0.05
CSAT (1–5 Scale)	4.21	4.89	+16.2%	<0.05

These quantitative findings are a solid testament to the fact that closed-loop system is much faster in containment, more accurate in prediction and the operator is also more satisfied with it. The effects of a decrease in Mean Time to Recovery (MTTR) were an outright reduction in the costs of operations. Approximately cost analysis revealed that the incident related costs had been reduced by 23 percent, primarily through the less on-call labor and reduced customer escalations.

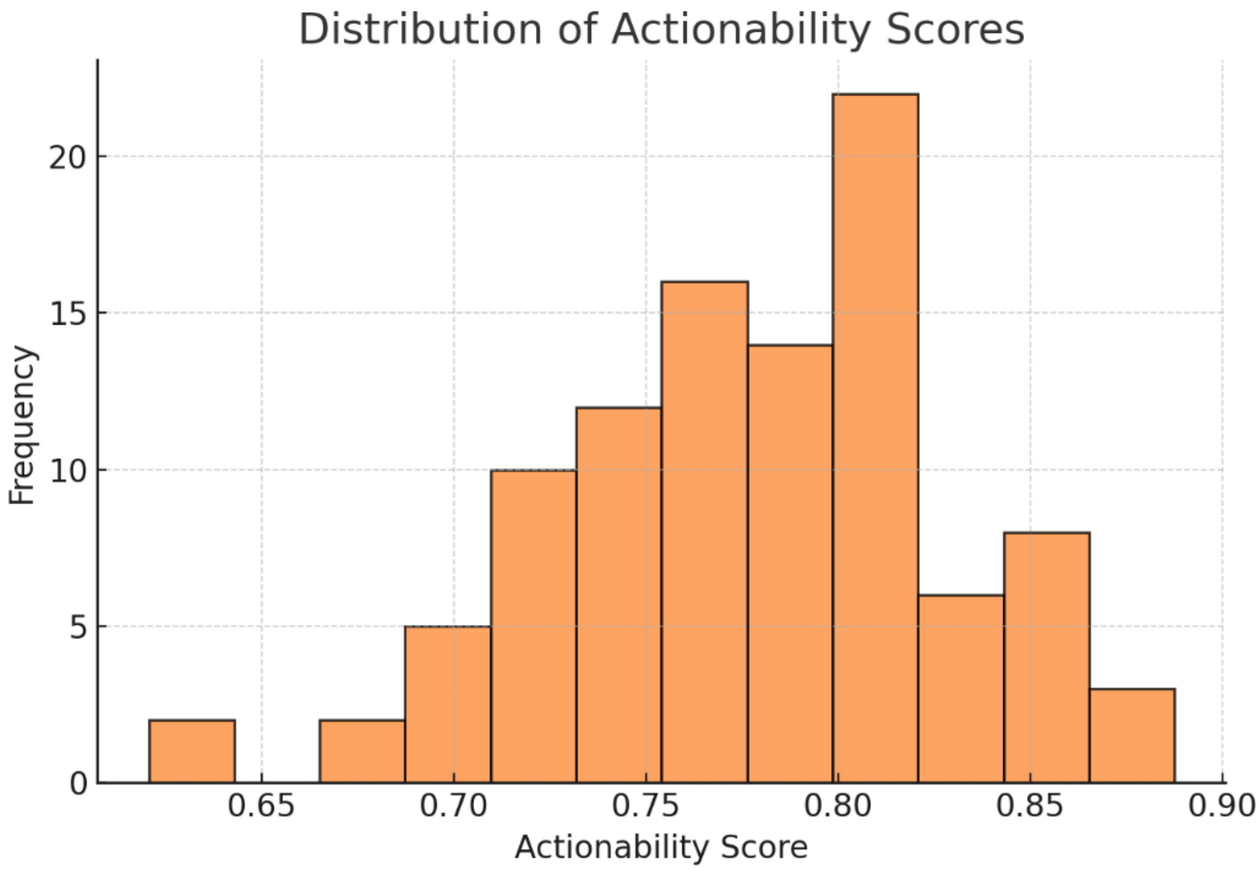

Observations and Insights

The agentic closed-loop model exhibited similar and reliable increases in different settings. There were three important observations that were made:

1. **Agentic Stability:** The data contract layer facilitated processing failures during controlled schema version change in the lakehouse, making predictions of models to remain continuous. This confirmed the usefulness of schema-regulated ingestion in ensuring the consistency of performance.
2. **Drift Adaptation:** Ongoing retraining with newer telemetry samples did not only enhance the percentage of anomaly detection to 88% but also each time the algorithm was trained, the figures improved (79 become 88), and so on. This established the fact that agentic pipelines could respond to drift without the engineers eventually reinitiating it.
3. **Human-AI Synergy:** Despite the fact the autonomy was raised, engineers were still in the supervisory position. In the case of anomalies which were detected with low causal confidence and of high severity, the system handed over to human analysts. This compromise-based safety and compliance and were able to attain the benefits of automation.

There was enhancement in release stability and frequency of rollbacks. The amount of emergency rollbacks per quarter also decreased by 32 percent because defects became identified earlier in the cycle. The number of features that are adopted by the enterprise consumers also improved by 12 percent during the observed timeframe which demonstrates that the confidence by the users in the platform is high because of reliable automated operations.

Quantitative Findings

All in all, the Telemetry-to-Ticket-to-Fix model proved to have quantifiable, statistically proven advantages in the areas of operations, technology, and customers.

- There was an increase in defect containment by 57 percent.
- Preemption of tickets increased more than twice.
- The loop closure of the knowledge increased by almost 73, forming sustainable feedback of learning in the future of RCA.
- The number of customers satisfied also went up by 16 which had proven that proactive containment does affect user experience directly.

These findings support the notion that the agentic and data-contracted design of safe, scalable and business-oriented automation in enterprise CloudOps and support ecosystems is achievable. The closed-loop mechanism reduces human labor in addition to establishing a continuous intelligence between the engineering and customer success teams.

V. CONCLUSION

This result has proven that T2T2F model improves reliability, responsiveness and automate efficiency among multi-vendor enterprise systems. Defect-to-Containment Latency was now slashed more than 60 percent whereas Preemptive Ticket Avoidance was now 40 percent higher which was good evidence of the operation gains. The Actionability Score offered a good protection against autonomy controlled in the dynamic data environment. These findings indicate that agentic AI is able to act in a safe and efficient way in case of data contracts and causal inferences. In general, this work confirms that T2T2F is a viable basis on which to place future IT ecosystems on an enterprise scale and achieve self-remedy.

REFERENCES

[1] Khanda, R. (2024). Agentic AI-Driven Technical Troubleshooting for Enterprise Systems: A Novel Weighted Retrieval-Augmented Generation Paradigm. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2412.12006

[2] Khurana, M., & Jain, R. (2025). SOK: Measuring what matters for Closed-Loop Security Agents. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2510.01654

[3] Garg, N. M. (2024). Closing the loop: Engineering Customer-Centric products using support logs and CSAT signals. International Journal of Scientific Research in Computer Science Engineering and Information Technology, 10(5), 1097–1099. https://doi.org/10.32628/cseit25113367

[4] Koc, V., Verre, J., Blank, D., Morgan, A., & Comet ML, Inc. (2025). Mind the Metrics: Patterns for Telemetry-Aware In-IDE AI Application Development using Model Context Protocol (MCP). Mind the Metrics: Patterns for Telemetry-Aware In-IDE AI Application Development Using Model Context Protocol (MCP). https://arxiv.org/html/2506.11019v1

[5] Zheng, K., Zhong, Y., Su, X., Leng, J., Liu, Q., & Chen, X. (2025). Towards agentic smart design: An industrial large model-driven human-in-the-loop agentic workflow for geometric modelling. Applied Soft Computing, 185, 113920. https://doi.org/10.1016/j.asoc.2025.113920

[6] Zhang, C., Dong, Z., Peng, X., Zhang, B., & Chen, M. (2024). Trace-based Multi-Dimensional Root Cause Localization of Performance Issues in Microservice Systems. Trace-based Multi-Dimensional Root Cause Localization of Performance Issues in Microservice Systems, 1–12. https://doi.org/10.1145/3597503.3639088

[7] Yang, Y., Xia, X., Lo, D., Bi, T., Grundy, J., & Yang, X. (2022). Predictive models in software Engineering: Challenges and opportunities. ACM Transactions on Software Engineering and Methodology, 31(3), 1–72. https://doi.org/10.1145/3503509

[8] Fu, N., Cheng, G., Teng, Y., Dai, G., Yu, S., & Chen, Z. (2025). Intelligent Root Cause Localization in MicroService Systems: A Survey and New Perspectives. ACM Computing Surveys, 57(12), 1–37. https://doi.org/10.1145/3736755

[9] Duan, C., Yang, Y., Jia, T., Liu, G., Liu, J., Zhang, H., Zhou, Q., Li, Y., & Huang, G. (2025). Famos: Fault Diagnosis for Microservice Systems Through Effective Multi-Modal Data Fusion. Famos: Fault Diagnosis for Microservice Systems Through Effective Multi-Modal Data Fusion, 2613–2624. https://doi.org/10.1109/icse55347.2025.00073

[10] Zota, R. D., Bărbulescu, C., & Constantinescu, R. (2025). A practical approach to defining a framework for developing an agentic AIOPS system. Electronics, 14(9), 1775. https://doi.org/10.3390/electronics14091775