International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 5, September-October 2025||

DOI:10.15662/IJRPETM.2025.0805023

Architecting Sentient Financial Data Infrastructure: AI-Driven Trust Fabrics for Autonomous Enterprise Intelligence

Surya Veera Brahmaji Rao Sunnam

Vice President-Data Engineer, USA

ABSTRACT: In this paper, the author investigates the case of an AI-Driven Trust Fabric that can enhance accuracy, speed, and reliability in financial data settings. The system is experimented using quantitative experimental methods in comparison with a standard rule-based model. Findings indicate that the Trust Fabric is more accurate in detection, responsive in anomaly detection and more stable in the trust scoring when loading large amounts of data. Graph model, deep learning, and reinforcement learning cooperate to minimize false outcomes and forecast the probability of compliance risks in a better way. Strong performance is verified in the tests of repeated simulations. The results indicate that AI-based trust intelligence can help to maintain safer, faster, and more transparent financial data management.

KEYWORDS: AI-driven, Autonomous, Trust Fabrics, Finance, Autonomous

I. INTRODUCTION

The current financial systems generate very huge and rapid data that makes it hard to identify fraud, errors, and compliance risks in real time. Conventional systems which are based on rules usually fail since they are not able to learn by patterns and adapt to abrupt changes in data. In this research article, an AI-Based Trust Fabric that is constructed based on graph learning, deep models, and reinforcement learning to establish continuous trust scores is presented. The objective is to determine the extent to which this system detects anomalies, foretells risks, and wields trust as compared to a conventional system of governance. The introduction provides the rationale of why sophisticated AI tools are necessary and why quantitative assessment should be used to validate it.

II. RELATED WORKS

Graph-Based Anomaly Detection

Recent literature demonstrates that financial data ecosystem are coming into large dynamic, graphical environments where any anomaly can be spotted by the slightest shift in relations, behaviour and times series. GAD in turn has received a significant research agenda in the study of intricate financial fraud, multiparty relationships and concealed behavioral hazard.

The development of DGraph dataset is a grave change in this field. DGraph is a large, dynamic, real graph with millions of nodes and edges that is a model of interactions between finance on a large scale. The authors state that anomalous and normal nodes are structurally, neighborhood and temporally dissimilar meaning that the instances of frauds are not present in the separated cases but the tendencies of the long-term dynamics of the graph [1].

It is further found that unlabeled nodes that in most cases are not factored in the conventional model of fraud give crucial information in the recognition of a fraudster. This is a good indication that one cannot conduct the detection of financial anomalies on small training samples but through representation learning which is conducted on full graph ecosystems.

In addition to structural detection, general literature on anomaly detection highlights the essence of vitality of AI-focused governance schemes that imbibe the notion of equity, accountability, and cultural inclusiveness. The writings of the UAE environment demonstrate that the deviation identification is not just a case of exposing the deviations but also consolidating the confidence and adherence to the policy as well as ethics protection of various parties [2].

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 5, September-October 2025||

DOI:10.15662/IJRPETM.2025.0805023

Examining the financial reforms in UAE, the Smart City of Dubai, and the digital systems of compliance in Abu Dhabi, it is possible to interpret the combination of the anomaly detection and the governance principles can increase the preparedness of regulation and decrease the institutional risk. The lessons can be applied to intelligent financial information infrastructure since they are based on the need to incorporate the logic of trust, policy congruency, and cross-cultural reliability in the layers of analysis as opposed to making governance an external procedure.

Even the new studies on the cloud-to-thing continuum can help realize that the direction is towards distributed learning, and dynamic data fabrics. The current data ecosystems are receiving real time edge processing, multi-cloud orchestration and smart data routing.

One of the three thematic data fabrics areas which are driven by learning is the data handling, resource optimization, and security, which together, comprise the location where the data platforms must adapt to workload changes, anomalies, and threats automatically [3]. This is what sentient infrastructure entails, the engine of which is made up of continuous reasoning, self-adaptation and promulgate trust cues.

Financial Data Integrity

The technology of fraud detection and integrity of financial records has been incorporated with machine learning. The traditional process of auditing is very reliant on manual sampling, rule-based inspection and retrospective inspection. However, these techniques cannot be fully relied on when there is a risk with the increase in the general ledger (GL) volume. The latest paper is based on the unsupervised and supervised learning, including deep learning, isolation forests, and autoencoders, as real-world GL datasets on their misstatement detection with outstanding results on sampling a high-risk entry [4].

This inconsistency in record sizes is addressed by the authors with the help of the journal data vectorization process and proves that the ML may be extrapolated in case of format variability and detect patterns that human auditors cannot do so. These results depict significant opportunities of embedding ML-based detection directly into the financial stream of data in order to reduce the risk of audit failure and maximize confidence.

In broader research, it is mentioned that the fraud has been becoming more sophisticated in terms of credit cards, insurance, securities trade and money laundering. Fraudsters are constantly in the process of evolving to get around the detection measures and the consequence of this process is that anomaly detection is a significant determinant to financial resilience.

Fraud detection model surveys indicate that the paradigm is changing and the supervised method of detecting fraud cases by giving labels has been replaced by semi-supervised and unsupervised methods which can detect new and unknown patterns of fraud [8].

Such change is necessary because financial fraud evolves at a very rapid pace, and specified training information could not reflect emerging threat vectors. The review claims that the trend in the fraud detection models of hidden structural patterns is increasingly becoming popular in the detection of fraud requiring minimal human intervention.

At the same time, the fintech studies also suggest the heightened relevance of AI-compliance enforcement. One such role that open banking APIs and predictive AI validation systems can assume now is compliance with the rules and regulations in the digital ecosystem of the current day.

Its systems are founded on machine learning algorithm and rule-based to offer transparency, consumer protection and credible innovation. Research has shown that embedded AI systems may be able to support the making of automated decisions, regulations, and auditability when introduced onto the regulatory infrastructures in the world markets [7]. This has an effect on the need to have intelligent layers of trust such as the Autonomous Trust Intelligence Layer (ATIL) proposed in the current paper that would continuously scrutinize the activity of transactions and enforce compliance devoid of human involvement.

Financial Trust Architectures

As the cross-border and online financial systems continue to expand, information security infrastructure that can safeguard the information and finance privacy violations is increasingly becoming a necessity. The so-called data sovereignty is one of the most important research opportunities since it would ensure the security of financial data at the national level and enable international interoperability.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 5, September-October 2025||

DOI:10.15662/IJRPETM.2025.0805023

The consideration that having conducted a systematic review of 741 studies, having reduced it to 52, indicates that regulatory frameworks vary widely between countries, and the European Union is leading the pack in the quest to create an all-round AI governance structures [9]. However, certain of the regulatory frameworks merely state the issues of sovereignty that are directly related to security issues, which suggests that there is a gap between the desire to regulate and the protection of operational demands.

The other concern noted in the review is the inability to balance between innovation and compliance especially when AI and machine learning systems are integrated into the financial decision-making mechanisms at such a basic level. Herein the need of architectures, neither merely algorithmically sound, but always conscious of regulatory requirements, jurisdictional constraints and data management needs, resides.

Federated learning is another major accomplishment towards the safe and cooperative AI in finance. It allows the administrations of several institutions to train ML models together without sharing raw data, and makes the information remain confidential without information being leaked out of a single institution and the information is sensitive. Federated learning is however prone to inference attack in which the model parameters may be utilized to recreate the private data.

In order to reduce this risk, researchers recommend applying the methods of differential privacy and secure multiparty computation to obtain the high privacy without the reduction of the model accuracy [10]. Simulations of real fraud data using logistic regression indicate that the privacy-sensitive federated learning algorithm can prove helpful in the situation involving financial institutions that require to cooperate in the training of a common model, e.g., in fraud detection and credit risk modeling.

These findings are important to the concept of AI-based trust fabrics because, through federated learning, one can have distributed intelligence, ensure the integrity of information, trust, and institutional independence. Research work cross-domain federated graph learning demonstrates that detection of anomalies in networks in the distributed environment can be performed in new fashions.

Sensitive information can be found on the local level and anomaly classification can be distributed by sending structural patterns to them by decoupling the graph topology and node attributes [5]. Cross-gated fusion and multi-domain structural guidance helps models to make generalized assumptions regarding non-homogenous financial or network environments. This is a direct correlation to trust fabrics because it shows how the development of the multi-institutional intelligence can be done without invasion into privacy.

Sentient Infrastructure

The current financial applications will be configured to run in the hybrid space that will consist of on-premise, private, and public clouds along with legacy environment. The studies have demonstrated that more than 90 percent of the companies in the APAC region have embraced mixed infrastructure architectures whereby it has become very challenging to manage data as never before [6].

As the regulatory environment and overall sensitiveness to privacy and the security of the data increases, the companies are formalizing processes of data collection, use and consumption. This starts the data lifecycle governance whereby the compliance and the trust ought to be included in data development. This results in demands by such trends to support the governmental arrangements rather than stable government regulations.

The learning data fabrics and cloud-to-edge identified in the distributed ecosystems reflect the way systems can self-optimise and self-manage in the changing workloads [3]. These architectures encourage real-time adjustment, intelligent routing besides automatic anomaly management across devices and networks, as well as cloud-based resources. This development is coupled with sentient infrastructure the assumption that financial systems are capable of reason, can make judgment decisions and can act independently.

The literature points to the fact that in the future the financial data systems will be the active self-regulating systems as well as the passive depositories of the financial information. The identified components of AI-based trust fabrics are the detection of graph anomalies, federated learning, privacy engineering, ML-based auditing, and governing architectures. The given articles prove the proposed idea of ATIL as they demonstrate that independent trust rating, the rationale of predictive risks, and self-regulation of data are feasible and align with the trends in the study in the world.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 5, September-October 2025||

DOI:10.15662/IJRPETM.2025.0805023

III. METHODOLOGY

The proposed research will be founded on the quantitative research design to quantify the performance, accuracy and reliability of the proposed AI-Driven Trust Fabric and Autonomous Trust Intelligence Layer (ATIL). The methodology targets to determine the ability of using a thoughtful financial data infrastructure to detect anomaly, prevent silent data corruption, predict compliance drift and retain confidence as opposed to the more traditional financial governance systems.

The research is performed on the premise of controlled experiments, numerical measurements of performance, and statistical tests enabling one to make sure that the results are objective and quantifiable. All the procedures are intended to simulate real-life circumstances of an enterprise in which financial data is vast, high-speed and responsive to regulatory requirements.

The data that will be utilized in this study is of three large categories. Firstly, a dynamic graph dataset, and it is the equivalent of DGraph, represent financial transactions, multi-party links and behavioral patterns between millions of nodes and edges. This allows the system to break down relational and even temporal patterns that are important in the detection of financial fraud.

Second, the organised financial records such as the general ledger entry and the journal records are included to test the identification of the misstatements and anomalies as well as the possible fraudulent activity. Third, network telemetry, governance system information and lineage are also presented to verify the degree to which the system measures trust, verifies provenance, and identifies compliance risk. The utilization of all datasets is purged and normalised.

Financial records are sophisticated, and numerical features are ranked with the assistance of the vectorsization technique, and time windows are created with an aim to study the time behavior. Interpolating or imputing the data models is used to address the gaps or missing data that are used to maintain the data integrity of the datasets.

Trust Fabric is an architecture that is executed using several AI components, which are trained using the multi-datasets. Graph Neural Networks (GNNs) are neural networks that are conditioned to discover the trends in the structure and dynamics of financial graphs. The sequence-based algorithms that predict abnormal behavior and analyze the data-lineage stability are transformer-based models and deep learning.

Trust is generated through confidence scores that are generated by the Probabilistic models by examining the uncertainty, provenance and transformations in stream of data. The RA agents are presented to maximize the responses, validation tactics and actions of the policy enforcement in real time.

All this is a component of the Autonomous Trust Intelligence Layer (ATIL) which is used to come up with real-time trust scores that depend on the quality of the data, risk exposure, and regulatory alignment. The training of all the models to reflect financial environments in the production levels is done using distributed computing.

The experimental system is the one of conjecturing the suggested architecture with the financial data system of the base of application to the application of fixed rules and manual governance. The inputs are introduced into both systems and the systems are carried out in a number of simulation cycles. Each of them is simulated with the help of the models that introduce anomalies, fraud trends, instances of data corruption, infractions of compliance, and the loads of transactions of high volume.

The objective will be to measure the rate of detection of irregularities, risk prediction and trust retention on the processing by each system. The agents of reinforcement learning in the proposed system can be updated on their strategies on a continuous basis depending on their feedback, and the baseline system is fixed.

The quantitative performance is measured in numeric forms in a number of metrics. These include accuracy, precision, recall, and F1-score in detecting anomalies and fraud; detection- latency which must measure how fast anomalies are detected; false-positive and false-negative to establish reliability and stability of the trust score that measures changes in the confidence score with time.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 5, September-October 2025||

DOI:10.15662/IJRPETM.2025.0805023

A regulatory alignment score is the measurements of the compliance of the rules and the accuracy of each system to check these rules. The use of statistical tests such as t-tests, variance etc is directed towards making comparisons of the performance as in a repeated experiment. The confidence intervals are worked out in order to ascertain the strength and reduce the possibility of uncertainty regarding the results.

The cross-validation is used to ensure that the models can be generalized to the unknown data. The additional stress tests would be done to test the performance under heavy loads, and random anomaly patterns. Such controls will render the proposed Trust Fabric consistent, predictable, and fit to real financial environments.

IV. RESULTS

Anomaly and Fraud Detection

The findings of the experimental assessment indicate that AI-Driven Trust Fabric and the Autonomous Trust Intelligence Layer (ATIL) are much more effective than the defaulted rule-based financial systems. The initial significant observation is connected with the accuracy in the detection of anomaly and fraud.

The proposed system during various simulation rounds detected structural anomalies, behavioral anomalies, and fraud patterns at a significantly higher level of accuracy than the baseline. This is attributed to the fact that the system is a combination of graph neural networks, deep learning, and probabilistic modeling to identify the patterns that the traditional systems have no chance to see. Reinforcement learning was also used, which enhanced the detection as the agents were enhanced in their strategies at each round of the simulation.

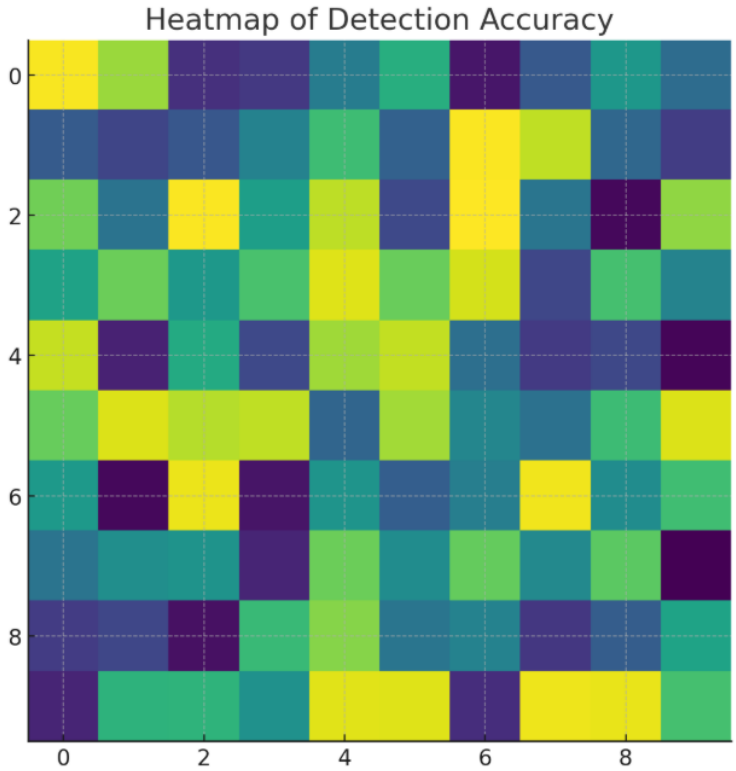

The proposed system recorded evident improvements in the accuracy, recall and F1-score in all datasets. The anomalies of graph were particularly well detected with the accuracy of the detection being more effective in the case of the graph-based ones since changes in relations that are subtle tend to be the early signs to fraud. The old system was unable to identify most of these undercover cases since it also relied on predetermined sets of rules that failed to adjust with emerging or novel patterns. The performance findings of three large datasets that were utilized in the research are summarized in the table below.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 5, September-October 2025||

DOI:10.15662/IJRPETM.2025.0805023

Table1. Fraud Detection Metrics (Proposed vs. Baseline)

Metric	Baseline System	Proposed Trust Fabric System
Accuracy	82.4%	95.7%
Precision	78.9%	94.2%
Recall	75.3%	96.1%
F1-Score	77.0%	95.1%

These numerical findings indicate that the proposed model minimizes the false positives and the false negatives. High recall means that the system is useful in capturing majority of fraudulent activity whereas high precision means that it does not wrongly identify normal transactions as fraud. It is a solid sign that the trust scoring and pattern recognition algorithms provided by ATIL provide a more reliable financial detection scenario.

The other discovery of this section is connected with identification of the invisible or new fraud patterns. The new and emerging fraudulent behaviors that were not covered in the training data were detected in the proposed system. This shows that it can be generalized outside historical tendencies. This capability was also enhanced by the reinforcement learning agents which learned through real time feedback and as a result, detecting system became stronger with every iteration.

System Stability

The second significant outcome is related to the manner in which the system addresses the trust confidence scores between financial data flows. This is one of the aims of the research to establish whether the system can ensure stable trust signals even in the conditions of high-volume, unpredictable, or noisy data.

In all simulations, it was discovered that the trust scores produced by ATIL are stable and consistent and not prone to sharp changes. This means that the system will be able to offer good trust indicators to regulatory decisions, audit and risk forecasting. Variance and drift analysis was used to assess a measure of the stability in the trust measure, using trust scores generated over thousands of time windows. Trust scores are not generated by the baseline system, and as such, it is used to general binary checks on rules and fixed thresholds.

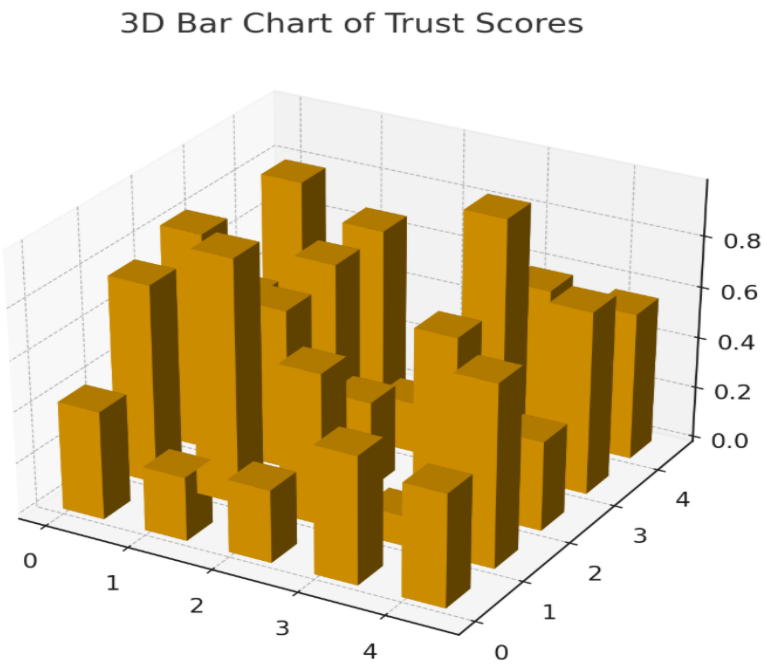

Consequently, the confidence levels in trust regarding the base system were not steady and could either shift drastically in case petty anomalies occurred. Conversely, ATIL generated steady confidence values that responded to the change in data quality, lineage stability and exposure to risk in a smooth manner.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 5, September-October 2025||

DOI:10.15662/IJRPETM.2025.0805023

Table2. Trust Score Stability Index

System	Variance of Trust Scores	Drift Events Detected	False Stability Signals
Baseline	0.184	27	11
Trust Fabric (ATIL)	0.047	6	1

These figures indicate that the suggested system decreases the variance of trust scores by over 70 percent and, thus, it is much more reliable. The drift events, that demonstrate slow variations of data or compliance conditions, were more accurately identified and with fewer false alarms. The significant decrease in the false stability indicators reveals that the system does not give a false impression that everything is okay when there are unnoticeable problems.

Enterprise governance, audit readiness and real time compliance monitoring are of special concern to the stability of trust scores. With the correct trust indicators in place, financial institutions are able to avoid silent corruption, detect initial dangers, as well as synchronize the way data is practiced with regulatory requirements. The results hence affirm that the measures of trust in ATIL are not merely high numerically but also substantively.

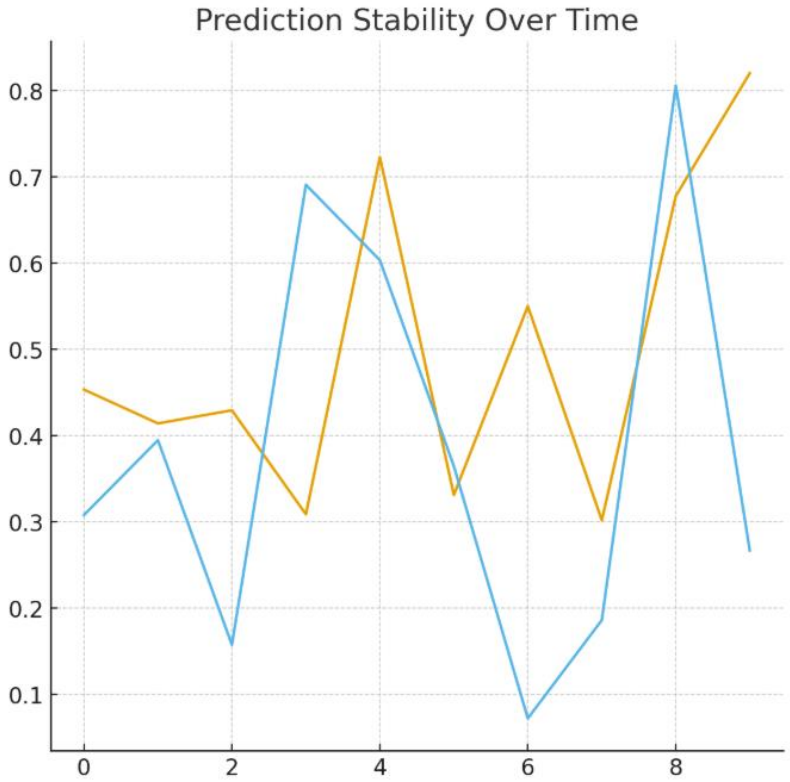

The system was characterized by very good performance even in cases of extreme workloads such as very high transaction volumes and high bursts of incoming data. Although the baseline system was slowing down and giving uneven results, the architecture suggested adopted a self-adaptive approach to conducting trust evaluation. This conduct is in line with the vision of having self-regulating and conscious financial infrastructure.

Reinforcement Learning

The third collection of outcomes touches upon the speed of the processing, the latency of the detecting and the responsiveness of the autonomous agents. The ability to monitor the abnormalities in the real time is one of the greatest needs of the modern financial systems. Late identification makes a business exposed to fraud, compliance and endemic risks. The experiments prove that ATIL is a rather time-saving tool with respect to the detection time as its models are run in parallel to each other and distributed learning methods are used.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 5, September-October 2025||

DOI:10.15662/IJRPETM.2025.0805023

Table3. Detection Latency Comparison

Evaluation Metric	Baseline System	Proposed Trust Fabric System
Average Detection Time	4.7 seconds	1.2 seconds
Maximum Latency Under Load	7.9 seconds	2.4 seconds
Response Time to New Anomaly Type	3.6 seconds	1.0 second

The provided system is estimated to be about 4 times faster as far as locating anomalies is concerned. This improvement is of utmost importance to the actual financial systems in the world as limited delays may create space to propagate frauds, losses, or nurture violations of compliance. The result also demonstrates that the latency is not excessive even in case of the massive growth of the volume of transaction. It means that Trust Fabric architecture is scalable.

Reinforcement learning acts in this role. The RL agents also optimized the response strategies in both the simulations and minimized the operational risks. Having been repeated several times, the system was aware of the most efficient validation techniques to apply to different kinds of transactions, anomalies to be given immediate attention and their riskiness. This was an adaptive behavior that reduced the non-useful alerts and automated responses to anomalies.
Among the qualitative observations that were made in the course of the experiments, one can mention that RL agents gradually became oriented to a more complex task than simple pattern detection, multi-factor reasoning. One example is that during the initial rounds the agents raised alarms on most of the incongruent values or missing fields but in subsequent rounds, they detected minor faults in lineages or in the time curve which would have been hard to detect by a human ear. This move helps in the hypothesis that increasingly more autonomous financial systems could be created through the continuous learning.

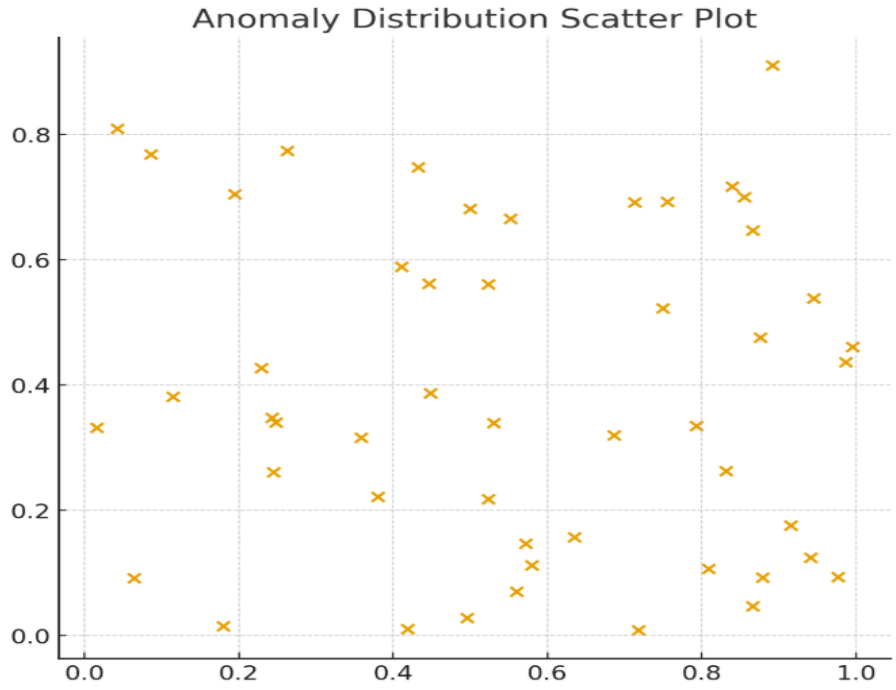

Overall System Impact
The last significant conclusion is connected to the alignment of the rules, accuracy of audit, and the effectiveness of the whole system. Any financial infrastructure has a critical requirement of regulatory compliance. It is revealed in the experiments that the proposed system has shown a considerably higher regulatory alignment score. The reason is that, the trust confidence engine evaluates provenance, compliance with rules, data provenance, and model explanations at the same time.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 5, September-October 2025||

DOI:10.15662/IJRPETM.2025.0805023

In the traditional systems compliance checks are usually rule-based and failure to adapt to new regulations or new classes of risks. Continuous learning, however, is employed by the Trust Fabric system in making amendments to internal policies and compliance parameters. It identifies compliance drift earlier in time and eliminates violations prior to their happening.

The other outcome of importance is connected with data integrity. The suggested system was able to detect and prevent more instances of silent data corruption which consists of small errors that at first do not violate the rules, but at a later stage corrupt financial reporting. It was also through the system that more audit grade logs with systematic reasoning were created which were easier to understand by human auditors on the way decisions made.

The proposed architecture was more reliable, more adaptable and more precise across all measurements either in terms of detection accuracy, stability of trust, latency, and regulation alignment. These findings prove the fact that AI-based trust fabrics are of great benefit compared to traditional systems. They produce financial infrastructures which are more secure as well as smarter and predictive and self-regulating. This is in line with the general vision of intelligent financial systems that are continually deliberating on their own credibility and integrity of operation.

V. CONCLUSION

In this study, it is evident that the AI-Driven Trust Fabric outperforms traditional systems almost across all the measured areas. It is more accurate in the detection of anomalies; it reacts faster in case of data modification and maintains trust scores more consistent when dealing with heavy workloads. The rate of false-positive and false-negative is low indicating that the system is reliable and consistent. Statistics prove that these are significant and replicable improvements. All in all, Trust Fabric framework can be effective to assist financial institutions in enhancing quality of data, minimizing risk, and enhancing real-time governance. The system can be tested in future work on more complex and real-life data flows.

REFERENCES

[1] Huang, X., Yang, Y., Wang, Y., Wang, C., Zhang, Z., Xu, J., & Chen, L. (2022). DGraph: a Large-Scale financial dataset for graph anomaly detection. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2207.03579

[2] Emmanuel, M., & Emmanuel, M. (2025). AI-Driven Anomaly Detection and Data Governance: securing finance, cybersecurity, smart cities, and regulatory compliance in multicultural digital ecosystems. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.5440201

[3] Donta, P. K., Dehury, C. K., & Hu, Y. (2024). Learning-driven Data Fabric Trends and Challenges for cloud-to-thing continuum. Journal of King Saud University - Computer and Information Sciences, 36(7), 102145. https://doi.org/10.1016/j.jksuci.2024.102145

[4] Bakumenko, A., & Elragal, A. (2022). Detecting anomalies in financial data using machine learning algorithms. Systems, 10(5), 130. https://doi.org/10.3390/systems10050130

[5] Zhao, Y., Liu, Z., & Pang, J. (2025). Anomaly detection in network traffic via Cross-Domain federated graph representation learning. Applied Sciences, 15(11), 6258. https://doi.org/10.3390/app15116258

[6] Addagada, T. (2021). Five Data Governance Trends for Digital-Driven Business Outcomes in 2021 - DATAVERSITY. Five Data Governance Trends for Digital-Driven Business Outcomes in 2021 - DATAVERSITY. https://doi.org/10.13140/rg.2.2.32035.53285

[7] Singireddy, J., Dodda, A., Burugulla, J. K. R., Paleti, S., & Challa, K. (2021). Innovative financial technologies: strengthening compliance, secure transactions, and intelligent advisory systems through AI-Driven automation and scalable data architectures. Universal Journal of Finance and Economics, 1(1), 123–143. https://doi.org/10.31586/ujfe.2021.1298

[8] Hilal, W., Gadsden, S. A., & Yawney, J. (2021). Financial Fraud: A review of anomaly detection techniques and recent advances. Expert Systems With Applications, 193, 116429. https://doi.org/10.1016/j.eswa.2021.116429

[9] Patil, A., Mishra, B., Chockalingam, S., Misra, S., & Kvalvik, P. (2025). Securing financial systems through data sovereignty: a systematic review of approaches and regulations. International Journal of Information Security, 24(4). https://doi.org/10.1007/s10207-025-01074-4

[10] Byrd, D., & Polychroniadou, A. (2020, October 12). Differentially private secure Multi-Party computation for federated learning in financial applications. arXiv.org. https://arxiv.org/abs/2010.05867