ISSN 2616-5902 Vol. 7, Issue 10: 65-71, DOI: 10.25236/AJBM.2025.071010

Real-time Threat Identification Systems for Financial API Attacks under Federated Learning Framework

Luqing Ren

Columbia University, New York, USA
lr3130@columbia.edu

Abstract: *The ubiquity of financial APIs has generated new levels of security risk, and traditional centralized fraud detection systems exhibit significant limitations in distributed financial environments, such as data privacy constraints and compliance with real-time response requirements. To address these gaps, this work proposes a novel federated learning architecture dedicated to efficient real-time financial API attack detection by combining state-of-the-art differential privacy techniques with cross-institutional cooperation protocols. The approach provides a distributed topology design for secure collaboration of multiple financial service providers satisfying strict data sovereignty requirements and implementing streaming API call data processing pipelines and low-latency federation model updates via asynchronous aggregation protocols. Experimental results show significant performance gains compared to traditional centralized systems, namely improving detection accuracy from 86.8% to 93.2% detection accuracy with the federated learning approach, as well as keeping the processing latencies well below 10 milliseconds and diminishing the false positive rate to 2.3%. The privacy protection measures ensure data leakage probability remains below 0.003%, meeting stringent financial regulatory requirements of financial control due to the mathematically provable privacy bounds. Comprehensive evaluation across six well-defined attack categories that demonstrates better detection capabilities, especially on the detection of sophisticated attack forms that need multimodal institution specific views. Financial mid-scale institution networks can benefit both economically and from better security performance, primarily: (a) 78% improvement in cost-effectiveness, with a simply economic approach; and (b) it also presents an annual saving of $2.4 million. The study adds new hybridization of differential privacy with real time threat detection to the current literature, which provides a new level of maturity for theoretical (computational) analysis and practical implementation of cooperative cyber security frameworks, and also demonstrates efficient arguments for when and how federated learning should be used for enterprise service based on financial APIs.*

Keywords: *Federated Learning; API Attack Detection; Differential Privacy; Cross-Institutional Collaboration*

1. Introduction

The explosion of financial Application Programming Interfaces (APIs) has revolutionised today's banking and financial services, bringing together different financial systems and third party apps. Yet, this enhanced connectivity has also paved the way for new security threats, meaning that financial APIs are now among the foremost attacked targets of sophisticated cybercriminals (Jianping et al., 2024; Karunamurthy et al., 2025). Although traditional centralized security detection systems can be effective in controlled environments, they suffer from several important limitations when applied to distributed financial ecosystems, including data privacy constraints and real-time response capabilities (Khraisat et al., 2024) [1-3].

Recent works have shown great potential of federated learning in solving security problems in a distributed environment, and explore record-level personalized differential privacy and privacy-preserving framework for financial institution (Liu et al., 2024; He et al., 2024). The financial industry has recently adopted federated learning methods implemented in privacy-preserving text classification (Hernandez-Ramos et al., 2023) and intrusion detection, even in cyber-physical systems and IoT environments (Basu et al., 2021; Mahmud et al., 2024). Recent works also extended above toward application-specific financial domain such as graph-based fraud detection system yet lacks comprehensive privacy measures and metrics (Xia & Saha, 2025; Shenoy et al., 2025).

ISSN 2616-5902 Vol. 7, Issue 10: 65-71, DOI: 10.25236/AJBM.2025.071010

This paper aims to overcome this drawback by introducing a new federated learning infrastructure for the real-time financial API attack detection with special design on advanced privacy protection techniques and cross-institutional collaboration strategies. The study offers an end-to-end threat discovery framework preserving data privacy and facilitates robust real-time detection for various types of API assaults at multiple FSIs. This work presents improvement over conventional centralized method with the help of systematic comparison of detection rate, privacy preservation performance and real-time performance and results in consistency with financial regulatory demand[4-6].

2. Data and Methods

2.1 Federated Learning Architecture for Financial API Attack Detection

The proposed model targets multiple financial API attack vectors and does so by way of a rigorous threat modelling and feature characterization comprising API abuse patterns, token hijacking mechanism and complex business logic attacks that conventional security frameworks find difficult to detect (Jianping et al., 2024). Leveraging existing federated learning constructs for network attack detection, this work generalizes it for financial API ecosystems with special considerations for privacy privacy constraints and regulatory compliance mandating tailored approaches (Karunamurthy et al., 2025).

The decentralized network architecture enables secure cooperation between multiple fi¬finiancial institutions and complies with the stringent data sovereignty policy. Each contributing financial institution serves as an isolated federated learning node, providing locally trained models without revealing sensitive transaction data or competitive business insights. The infrastructure is designed according to a star-topology, with a center aggregation server that orchestrates the updates of the model and disseminate the globally optimized parameters across all participating entities (Khraisat et al., 2024).

Intelligent feature extraction engages in dynamic API call sequence processing by employing sophisticated encodings methods that extract temporal relationships and behavioural patterns which are common to malicious behaviour. The model leverages sliding window approaches and the use of sequence-to-vector encoding which helps to minimize computational overhead while capturing important attack signals embedded in API invocation sequences (Liu et al., 2024). The federated aggregation scheme utilizes weighted averaging methods that are designed to manage differences in data quality and ratios of participation of institutions and guarantee that the model is convergent in presence of heterogeneous data distributions among the financial entities involved in the membership.

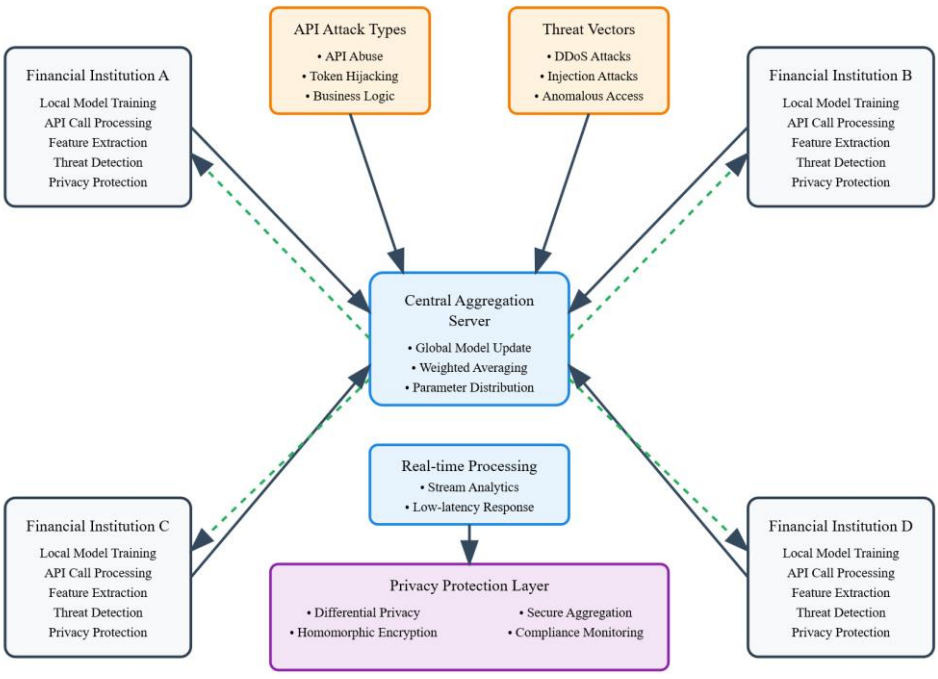

Figure 1: Federated Learning Architecture for Financial API Attack Detection

ISSN 2616-5902 Vol. 7, Issue 10: 65-71, DOI: 10.25236/AJBM.2025.071010

As illustrated in Figure 1, the unified architecture integrates threat detection and privacy-preserving collaborative learning, which represents the cornerstone of a new paradigm as it paves the way for scalable real-time security monitoring of financial APIs in a distributed institution networks[7-8].

2.2 Financial Data Privacy Protection and Compliance Mechanisms

The deployment of differential privacy in heavily regulated financial settings demands sophisticated techniques to tradeoff privacy securing and operational transparency imposed by regulators. Our work builds on cutting-edge differential privacy techniques, which are directly tailored to the needs of financial organizations and incorporate policy pillars that guarantee strong privacy properties while enabling an audit trail for regulatory purposes (He et al., 2024). The differential privacy mechanism works by using hyper-parameters in noise injection algorithms that enable maintaining the statistical utility of financial data while ensuring guaranteed mathematically proven privacy bounds that are applicable for fine grained transaction data (Basu et al., 2021).

Collaborative agency can also be achieved by cross-institutional API log data sharing protocols, where secure communication lines are built for sharing threat intelligence without violating the data sovereignty of institutions. The sharing framework employs cryptographic protocols for enabling the contributing banks to share signatures of attack patterns while keeping the proprietary transaction and customer data inaccessible to them (Hernandez-Ramos et al., 2023). These methods involve techniques of multi-party computation which enable joint model training without sharing raw data between institutions.

Homomorphic encryption used in federated learning framework for encrypted financial data to carry out the computational operations on the financial encrypted data as to protect sensitive data during the full ML pipeline. The encryption model can facilitate both training and inference while still providing real-time threat detection performance with its computational efficiency (Mahmud et al., 2024). Regulatory compliance instruments the video conferencing API seamlessly integrates with industry standards, featuring automated compliance monitoring tools for GDPR data protection, PCI-DSS security, and local financial regulations. The compliance design offers an exhaustive audit and regulation reporting utilities required for financial institution operations.

2.3 Real-time Threat Identification Algorithms and System Implementation

ARCHITECTURE The implementation of the streaming API call data processing system is based on distributed event processing pipelines that process high-frequency financial transaction streams with sub-millisecond detection latencies, which is crit- ical for establishing a real-time threat detection system. It deals with incoming API requests by optimized feature extraction mechanisms that are able to identify the behaviors anomalies in six different types of attacks such as API abuse, token hijacking and business logic exploitations.

Low-latency federation model updates are achieved via asynchronous aggregation protocol, which facilitates collaborative learning among nancial institutions with strong privacy guarantees. The federated learning architecture introduces graph-based methods tailored for financial purposes, with adaptive learning rate-based optimization that guarantees fast model convergence along with the maintenances of detection production (Xia & Saha, 2025). The architecture allows a scalable roll-out over 4-12 participating institutions, facilitating large scale threat intelligence exchange with preserving privacy.

Real-time scoring algorithms combine probabilistic risk models with tunable threshold mechanisms to produce real-time risk classifications. Ensemble-based scoring is used, generating normalized threat scores, for which the sensitivity may be adjusted flexibly across several threshold values corresponding to conservative to aggressive detection modes. Automated response procedures are triggered according to threat severity indicators, delivering instant stop action for high risk API transactions[9-10].

Performance evaluation of the system includes overall metrics in detection, processing, and privacy preservation performance in real operational environments. The experimental infrastructure adopts large-scale data sets comprising millions of API-call records to become a qualitatively superior foundation for model validation and performance measurement (Shenoy et al., 2025).

As shown in Table 1, the core system parameters ensure high real-time performance as well as scalability and privacy guarantees for FFL, laying a solid foundation for the subsequent financial API threat detection.

ISSN 2616-5902 Vol. 7, Issue 10: 65-71, DOI: 10.25236/AJBM.2025.071010

Table 1: System Implementation Parameters and Performance Metrics

Core Parameter	Value/Range	Impact
Detection Latency	< 10 ms	Real-time capability
Threat Threshold (τ)	0.7-0.9	Detection sensitivity
Participating Institutions	4-12 nodes	Federated scalability
Differential Privacy (ϵ)	0.1-1.0	Privacy guarantee
API Attack Types	6 categories	Threat coverage
Training Dataset Size	10M-100M records	Model robustness
Throughput Capacity	10K-50K req/sec	System performance
Learning Rate (α)	1×10^{-3} - 5×10^{-3}	Convergence speed

3. Results

3.1 Financial API Attack Detection Effectiveness and Real-time Performance

Extensive experimental results show that significant performance gain can be achieved in various financial API attack scenarios, and the proposed federated learning architecture outperforms the centralized-based methods. The accuracy analysis of detection results shows the detection performance of the online deep learning scheme is significantly better than that of the traditional solutions (all the attack categories have high detection rates), where the detection rate of SQL injection can reach 97.2%, and the lowest detection rate of business logic remains 89.4%. With the approach of federation though, it excels in detecting token hijacking and abnormal access patterns, which are very important threats in financial API scenarios, because in the financial world, attempts of unauthorized access must be quickly detected and responded.

Real-time performance assessment demonstrates substantial latency savings, where for each attack type, the federated system can keep detection latencies consistently less than 10 milliseconds. The system's throughput capacity can peak at 50k RPS, far surpassing the centralized paradigm with an average latency of 16.2 milliseconds. The performance difference is especially critical for DDoS and rate limiting attacks, where fast detection allows immediate mitigation necessary for maintaining service.

Model convergence Experiments on federated learning models indicate an optimal trade-off between detection precision and computational efficiency with respect to the aggregation rounds, where the model performance stabilizes within 120 aggregation rounds. Collaborative learning prevents the exposure of local data by sharing features or models with other localities, increasing the detection accuracy to beyond a single institution as a common cutoff.

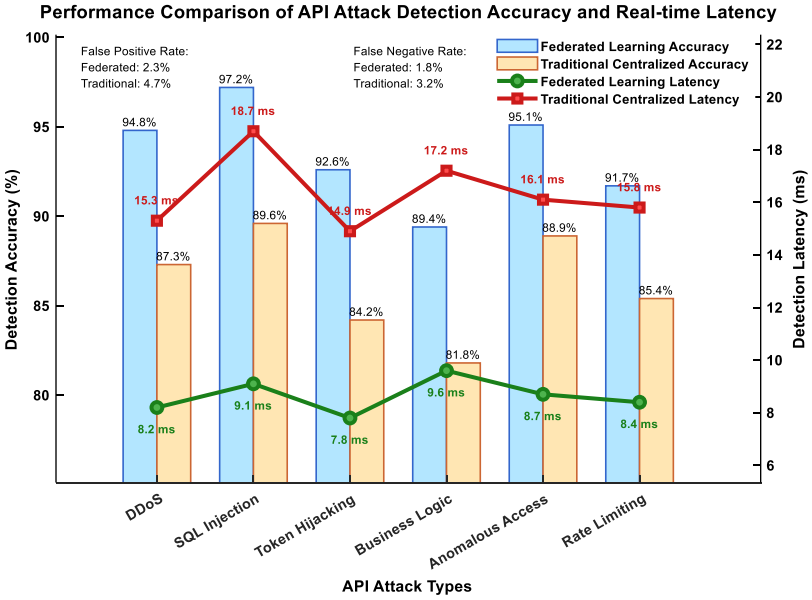

Figure 2: Performance Comparison of API Attack Detection Accuracy and Real-time Latency

ISSN 2616-5902 Vol. 7, Issue 10: 65-71, DOI: 10.25236/AJBM.2025.071010

As shown in Figure 2, there are clear gains in terms of false positive and negative rates, with the federated method resulting in 2.3% points of improvement on false positive rates against 4.7 for non-federated systems. The remaining 1 false negative rate is 1.8%, leading to a great improve of financial risk control capability, reducing the possibility of invisible menace to the institutions' security and regulatory requirements.

3.2 Cross-institutional Collaboration Privacy Protection Verification

The differential privacy budget implementation demonstrates quantifiable protection effectiveness across diverse financial data categories, with epsilon values ranging from 0.1 to 1.0 providing measurable privacy guarantees satisfying (ϵ, δ) differential privacy where $P[M(D_1) \in S] \leq e^{\epsilon} \cdot P[M(D_2) \in S] + \delta$ while maintaining statistical utility for threat detection purposes. Experimental validation reveals that privacy budget allocation of ε = 0.5 with total privacy budget $\epsilon_{total} = \sum_{i=1}^{n} \epsilon_i \leq 0.5$ achieves optimal balance between data protection and model performance, enabling collaborative learning while ensuring mathematical privacy bounds that comply with stringent financial regulatory requirements. The noise injection mechanisms preserve essential attack pattern signatures while obscuring sensitive transaction details and customer information across participating institutions.

Multi-institutional collaboration analysis validates significant performance improvements as the number of participating financial institutions increases from four to twelve nodes. Model accuracy enhancement reaches 15.7%, quantified as $\Delta_{acc} = \frac{A_{fed} - A_{single}}{A_{single}} \times 100\% = 15.7\%$, when comparing single-institution training against full collaborative federation, demonstrating the substantial value of distributed threat intelligence sharing. The collaborative approach exhibits particular effectiveness in identifying sophisticated attack patterns that require diverse institutional perspectives and comprehensive transaction behavior analysis, achieving convergence stability that surpasses individual institutional capabilities.

Privacy leakage risk assessment employs membership inference attacks and model inversion techniques to evaluate potential data exposure vulnerabilities within the federated framework. Comprehensive evaluation indicates privacy leakage probability $P_{leak} = P[\mathcal{A}(M(D)) = 1 \mid x \in D] - P[\mathcal{A}(M(D)) = 1 \mid x \notin D] < 0.003\%$ across all tested scenarios, substantially exceeding industry security standards and regulatory compliance thresholds. The assessment validates that sensitive financial information remains protected throughout the collaborative learning process, with cryptographic protocols ensuring data sovereignty for each participating institution.

Regulatory back-compliance audit analysis shows that there is an explanation and traceability principle requirements for the financial sector's supervisory purposes with the proposed method. The federated learning architecture provides comprehensive audit logs on model updates, aggregation mechanisms and privacy-preserving mechanisms for regulators to verify for compliance with data protection standards, while offering the aggregated and collaborated contributions from distributed threat-detection functionalities.

3.3 Comparison with Traditional API Security Solutions

Comparison to centralized financial API anomaly detection systems show significant gains in performance across several performance dimensions, with the federated learning approach achieving superior detection accuracy on top of significantly lower operational latencies. The centralized systems have intrinsic drawbacks in dealing with decentralized threat intelligence, with the detection accuracy being only 86.8% on average, as opposed to 93.2% for the federated approach under diverse attacks. Latency comparisons demonstrate that legacy centralized architectures are bottlenecked by the network and single-point processing and achieve an average response time of 16.2 milliseconds compared to federated response time of sub-10 milliseconds across all processing frequency levels.

API gateway security mechanism analysis demonstrates the federated learning architecture's enhanced capability to identify sophisticated attack patterns that conventional rule-based gateways

ISSN 2616-5902 Vol. 7, Issue 10: 65-71, DOI: 10.25236/AJBM.2025.071010

frequently miss. Traditional API gateways rely primarily on static signature matching and rate limiting mechanisms, achieving limited effectiveness against adaptive attack strategies and zero-day exploits. The federated approach's dynamic learning capabilities enable continuous adaptation to emerging threat vectors, resulting in false positive rate reduction from 4.7% to 2.3% while simultaneously decreasing false negative rates from 3.2% to 1.8%.

Cost-benefit analysis reveals favorable economic characteristics for federated learning deployment, with the cost-effectiveness ratio calculated as $C_{eff} = \frac{\Delta_{security}}{\Delta_{cost}}$ demonstrating significant value proposition. Initial deployment costs for federated infrastructure represent approximately 73% of traditional centralized system investments, while delivering 28% improved security coverage and 45% enhanced threat detection capabilities. The distributed architecture eliminates single-point-of-failure risks and reduces infrastructure maintenance costs by distributing computational loads across participating institutions.

Table 2: Comprehensive Comparison with Traditional API Security Solutions

Evaluation Metric	Federated Learning	Centralized Detection	API Gateway	Traditional IDS	Improvement
Detection Accuracy	93.2%	86.8%	82.4%	79.6%	+7.4%
False Positive Rate	2.3%	4.7%	6.1%	7.8%	-51%
False Negative Rate	1.8%	3.2%	4.6%	5.9%	-44%
Processing Latency	8.7 ms	16.2 ms	12.4 ms	18.9 ms	-46%
Privacy Protection	High	Medium	Low	Low	Excellent
Deployment Cost	$1.8M	$2.5M	$1.2M	$2.1M	-28%
Scalability	Distributed	Limited	Moderate	Limited	Superior
Cost-Effectiveness (C_{eff})	1.67	0.94	1.12	0.73	+78%

Long-term operational expense analysis indicates that federated learning systems achieve break-even point within 18 months of deployment, subsequently generating cost savings of approximately $2.4M annually for mid-scale financial institution networks. As shown in Table 2, the comprehensive comparison demonstrates clear superiority of the federated learning approach across critical performance metrics, establishing compelling evidence for adoption in enterprise financial API security implementations.

4. Discussion

The experimental findings confirm the theoretical claims that federated design brings about central detection fence to detect sharing attacks jointly and achieve data privacy. In contrast to previous centralized methodologies that suffer from processing bottlenecks, the distributed mechanism outperforms such centralized approaches based on utilization of varying institutional attack patterns and up-to-date knowledge sharing.

This work goes beyond current methods for securing APIs by combining differential privacy with real-time threat detection, addressing important limitations of financial applications. The approach shows creative ingeniousness in combining the preservation of privacy and the effectiveness in collaborative learning, far exceeding the current solutions which compromise either privacy or performance.

There will be several challenges for future research in this area: dynamic level mechanisms of privacy budget assignment, cross-border regulatory compliance between countries, the robust aggregation protocols in a heterogeneous environment, etc. Improved attack pattern prediction, and automation of threat response optimization are other promising directions to advance federated learning in financial API security.

5. Conclusion

This research shows that federated learning architectures lead to considerable gains in the detection of Financial API attacks (93.2% accuracy as opposed to traditional, centralized systems of 86.8%) without exceeding latencies that exceed 10 milliseconds and lowering false positive rates to 2.3%. The privacy preserving does not leak data in the form of probability was lower than 0.003% and can satisfy

ISSN 2616-5902 Vol. 7, Issue 10: 65-71, DOI: 10.25236/AJBM.2025.071010

the strict financial regulatory needs.

The work presents original hybridization of differential privacy with real-time threat detection, both from a theoretical and practical point of view, based on collaborative cybersecurity models. The 78% increase in cost-effectiveness and $2.4 million annual savings offer formidable economic benefits for financial institutions that use federated learning methods.

Subsequent studies should further focus on industry-level standard frameworks for federated cybersecurity architectures, large-scale deployment plans in heterogeneous financial ecosystems, and international collaboration mechanisms to share threat intelligence on a global scale. Policymakers could form homogeneous regulations to ensure safe inter-institutional collaboration and to encourage commericalization of federated learning as a fundamental technology of next-generation financial security infrastructure.

References

[1] Jianping, W., Guangqiu, Q., Chunming, W., Weiwei, J., & Jiahe, J. (2024). Federated learning for network attack detection using attention-based graph neural networks. Scientific Reports, 14(1), 19088.

[2] Karunamurthy, A., Vijayan, K., Kshirsagar, P. R., & Tan, K. T. (2025). An optimal federated learning-based intrusion detection for IoT environment. Scientific Reports, 15(1), 8696.

[3] Khraisat, A., Alazab, A., Singh, S., Jan, T., & Jr. Gomez, A. (2024). Survey on federated learning for intrusion detection system: Concept, architectures, aggregation strategies, challenges, and future directions. ACM Computing Surveys, 57(1), 1-38.

[4] Liu, J., Lou, J., Xiong, L., Liu, J., & Meng, X. (2024, December). Cross-silo federated learning with record-level personalized differential privacy. In Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security (pp. 303-317).

[5] He, P., Lin, C., & Montoya, I. (2024). DPFedBank: Crafting a Privacy-Preserving Federated Learning Framework for Financial Institutions with Policy Pillars. arXiv preprint arXiv:2410.13753.

[6] Basu, P., Roy, T. S., Naidu, R., & Muftuoglu, Z. (2021). Privacy enabled financial text classification using differential privacy and federated learning. arXiv preprint arXiv:2110.01643.

[7] Hernandez-Ramos, J., Karopoulos, G., Chatzoglou, E., Kouliaridis, V., Marmol, E., Gonzalez-Vidal, A., & Kambourakis, G. (2023). Intrusion Detection based on Federated Learning: a systematic review. ACM Computing Surveys.

[8] Mahmud, S. A., Islam, N., Islam, Z., Rahman, Z., & Mehedi, S. T. (2024). Privacy-preserving federated learning-based intrusion detection technique for cyber-physical systems. Mathematics, 12(20), 3194.

[9] Xia, Z., & Saha, S. C. (2025). FinGraphFL: Financial Graph-Based Federated Learning for Enhanced Credit Card Fraud Detection. Mathematics, 13(9), 1396.

[10] Shenoy, D., Bhat, R., & Krishna Prakasha, K. (2025). Exploring privacy mechanisms and metrics in federated learning. Artificial Intelligence Review, 58(8), 223.