International Journal of Frontiers in
Engineering and Technology Research
Journal homepage: https://frontiersrj.com/journals/ijfetr/
ISSN: 2783-0497 (Online)

(REVIEW ARTICLE)

Check for updates

Real-Time Payments Security: Threat Landscape, Regulatory Pressures and Cloud-Native AI/ML Defense Architectures

Raghunandan G Ramakrishna *

A Technical Review and Architectural Framework.

International Journal of Frontiers in Engineering and Technology Research, 2026, 10(02), 036-045

Publication history: Received on 06 March 2026; revised on 13 April 2026; accepted on 16 April 2026

Article DOI: https://doi.org/10.53294/ijfetr.2026.10.2.0013

Abstract

Background: The global adoption of real-time payment (RTP) systems has fundamentally transformed the speed and accessibility of financial transactions, while simultaneously creating an expanded attack surface that outpaces the capabilities of legacy fraud-prevention infrastructure. This paper presents a technical review and architectural framework addressing the security challenges inherent to high-velocity payment environments.

Methods: A structured literature review was conducted across peer-reviewed publications (2018–2024), industry standards bodies (PCI Security Standards Council, SWIFT, Federal Reserve), and practitioner reports from financial technology research firms. Cloud-native architecture patterns were synthesized from vendor documentation, publicly available reference architectures, and comparative deployment case studies. No original empirical data collection or human subjects were involved.

Results: Synthesized evidence from the reviewed literature indicates that AI/ML-based anomaly detection systems achieve transaction-level fraud detection accuracy in the range of 90–99.7% with sub-100 ms latency constraints, and are associated with fraud loss reductions of 40–60% compared with rule-based systems. A cloud-native multi-layer defense architecture—spanning network ingress, identity verification, behavioral analytics, and tokenized data vaults—is presented as a practical framework for organizations seeking PCI DSS 4.0 alignment.

Conclusion: Real-time payment security demands a paradigm shift from reactive rule-sets to adaptive, ML-driven architectures. The framework proposed here provides a vendor-agnostic reference model applicable to major cloud platforms, with quantified trade-offs between detection sensitivity, operational latency, and compliance cost.

Keywords: Real-time payments; Fraud detection; Machine learning; PCI DSS 4.0; Cloud-native architecture; Anomaly detection; Authorized push payment fraud; Tokenization

1. Introduction

The migration of global payments infrastructure toward real-time settlement has been one of the most consequential shifts in financial technology over the past decade. Systems such as the RTP® Network (United States), the Faster Payments Service (United Kingdom), and India's Unified Payments Interface (UPI) collectively process billions of transactions annually, many of which complete in under three seconds end-to-end [1,2]. This operational immediacy eliminates the traditional fraud-recall window that banks historically relied upon to reverse suspicious transactions, fundamentally changing the economics of payment fraud.

* Corresponding author: Raghunandan G Ramakrishna

Fraudsters have responded to this environment with a corresponding acceleration in attack sophistication. Authorized push payment (APP) fraud, synthetic identity schemes, and AI-generated social engineering have all increased in prevalence as real-time rails have expanded [3,4]. Simultaneously, regulatory pressure has intensified: the Payment Card Industry Data Security Standard version 4.0 (PCI DSS 4.0), effective March 2025, introduces customized implementation pathways and strengthened controls for cryptographic agility, multi-factor authentication, and continuous monitoring that require organizations to overhaul legacy compliance postures [5].

Legacy fraud-prevention systems—built around batch-processing architectures and static rule sets—are structurally ill-suited to the sub-second decision windows that RTP environments demand. A transaction that takes 1,800 ms to clear cannot wait for a 3,000 ms fraud-scoring pipeline. This latency asymmetry, combined with the volume and velocity of modern payment flows, has driven widespread interest in cloud-native, AI/ML-augmented architectures that can deliver both the throughput and adaptive intelligence required [6,7].

Despite growing practitioner adoption of these approaches, a coherent, vendor-agnostic technical synthesis is lacking in the academic literature. Existing treatments tend either toward high-level policy analysis or toward narrow vendor-specific implementation guides that do not provide reproducible architectural principles. This review addresses that gap.

1.1. Scope and Objectives

This paper pursues four objectives: (i) to characterize the contemporary threat landscape confronting RTP systems, with emphasis on fraud typologies and attack vectors that are specific to or amplified by real-time settlement; (ii) to analyze the regulatory and compliance requirements that RTP operators and participating institutions must navigate; (iii) to synthesize evidence on the effectiveness of AI/ML-based fraud detection as reported in the reviewed literature; and (iv) to present a layered, cloud-native defense architecture that can serve as a reference model for practitioners and a foundation for future empirical research.

1.2. Paper Organization

The remainder of this paper is structured as follows. Section 2 describes the review methodology. Section 3 characterizes the RTP threat landscape and regulatory context. Section 4 presents the synthesized findings on AI/ML detection performance. Section 5 introduces the proposed cloud-native security architecture. Section 6 discusses implications, limitations, and directions for future research. Section 7 concludes.

2. Review Methodology

This work constitutes a structured technical review supplemented by an architectural synthesis. No original empirical experiments were conducted, and no human subjects were involved. The review was conducted in three stages.

2.1. Literature Search

Peer-reviewed publications were identified through searches of Scopus, Web of Science, and IEEE Xplore using the following primary terms and their combinations: “real-time payments,” “instant payments fraud,” “authorized push payment,” “machine learning fraud detection,” “payment security architecture,” “PCI DSS compliance,” “tokenization payments,” and “cloud-native financial services.” The search was bounded to publications from January 2018 through December 2024, reflecting the period in which modern RTP rails achieved significant adoption. Conference proceedings, journal articles, and edited book chapters were included; purely opinion pieces and blog posts were excluded from primary synthesis.

2.2. Industry and Standards Sources

Given that RTP security is a rapidly evolving practice domain, industry and standards documentation was incorporated alongside academic literature. Sources included: reports from the Federal Reserve Banks (FedNow Service documentation), UK Payment Systems Regulator, European Central Bank TIPS program, PCI Security Standards Council, SWIFT Institute, and practitioner research from recognized financial technology analysts. Vendor technical documentation (AWS, Microsoft Azure, Google Cloud Platform) was included specifically for the architectural synthesis in Section 5, where it was treated as illustrative rather than prescriptive.

2.3. Inclusion and Exclusion Criteria

Sources were included if they: (a) addressed fraud detection, payment security, or compliance in a real-time or near-real-time payment context; (b) provided quantitative performance data, architectural specifications, or regulatory guidance; and (c) were attributable to a named author or organization with identifiable methodology. Sources were excluded if they were purely promotional in nature (vendor marketing collateral without cited evidence), undated, or related exclusively to card-present transactions where the latency dynamics differ substantially.

2.4. Architectural Synthesis Approach

The proposed architecture in Section 5 was derived by identifying common structural patterns across the reviewed cloud deployment case studies, reference architectures from major cloud providers, and the control requirements of PCI DSS 4.0. The synthesis aimed to produce a platform-agnostic layered model; specific cloud services are cited as illustrative examples of how each architectural layer can be instantiated, not as recommendations of one platform over another. Comparable service mappings for Microsoft Azure and Google Cloud Platform are noted where relevant.

3. Threat Landscape and Regulatory Context

3.1. The Real-Time Settlement Vulnerability Window

The defining security characteristic of RTP systems is the irrevocability constraint: once a payment instruction is settled, recovery of funds depends entirely on the willingness of the receiving institution and its account holder to cooperate in a reversal. Unlike card payments—which carry a chargeback mechanism extending days or weeks—RTP transactions offer fraudsters a durable, immediate transfer of value that is structurally resistant to clawback [8]. UK Finance reported that APP fraud losses in the United Kingdom totaled £479.6 million in 2022 alone, with only 59% of funds recovered [9]. The Nacha (ACH Network) operational data for the United States shows that same-day ACH fraud attempts have grown at a compound annual rate exceeding 20% since the introduction of expanded same-day windows [10].

The velocity of settlement also compresses the time available for fraud intervention. Analysis reviewed from financial industry sources indicates that a fraud-scoring decision must typically be returned within 100–300 milliseconds to avoid degrading the payment experience or triggering timeout failures in downstream systems [6]. This latency budget is insufficient for legacy rule-based engines that rely on sequential evaluation of hundreds of handcrafted rules, motivating the shift to parallel, ML-driven scoring pipelines.

3.2. Principal Fraud Typologies

3.2.1. Authorized Push Payment (APP) Fraud

APP fraud involves manipulating legitimate account holders—through impersonation, romance scams, investment fraud, or invoice redirection—into authorizing payments to fraudster-controlled accounts. Because the account holder issues the instruction voluntarily, traditional authentication controls offer limited protection. The UK's Contingent Reimbursement Model (CRM) Code and the Payment Systems Regulator's mandatory reimbursement rules (effective October 2024) represent the most advanced regulatory response to date, imposing shared liability between sending and receiving Payment Service Providers [11].

3.2.2. Synthetic Identity Fraud

Synthetic identities—fabricated from combinations of real and fictitious personally identifiable information (PII)—are used to establish accounts that accumulate payment history before a coordinated "bust-out" event in which all available credit and payment limits are exhausted simultaneously. Machine learning approaches reviewed in the literature, particularly graph-based models that map relationship patterns between entities, show substantially improved detection rates for synthetic identity clusters compared with attribute-based scoring alone [12,13].

3.2.3. AI-Augmented Attack Vectors

Generative AI has introduced a qualitative escalation in social engineering capability. Voice cloning tools can produce synthetic audio of known individuals with minutes of training audio; large language model-generated phishing content achieves click-through rates substantially higher than manually crafted variants in controlled studies [14]. These developments create an adversarial arms race dynamic in which defensive AI systems must contend with increasingly convincing synthetic content as an attack vector.

3.3. Regulatory and Compliance Pressures

3.3.1. PCI DSS 4.0

PCI DSS version 4.0, published March 2022 and mandating full compliance from March 2025 onward, introduces 64 new requirements relative to version 3.2.1, of which the most significant for RTP environments are: enhanced multi-factor authentication (MFA) requirements across all access vectors; cryptographic agility mandates requiring algorithm inventories and migration plans; continuous security testing rather than point-in-time assessments; and a new customized implementation pathway that allows organizations to demonstrate the intent of a requirement through alternative controls [5]. The compliance investment is substantial: peer-reviewed survey evidence indicates that only approximately 43% of organizations were fully PCI DSS 3.2.1 compliant at the time of the 4.0 transition, implying a significant remediation burden across the payment ecosystem [15].

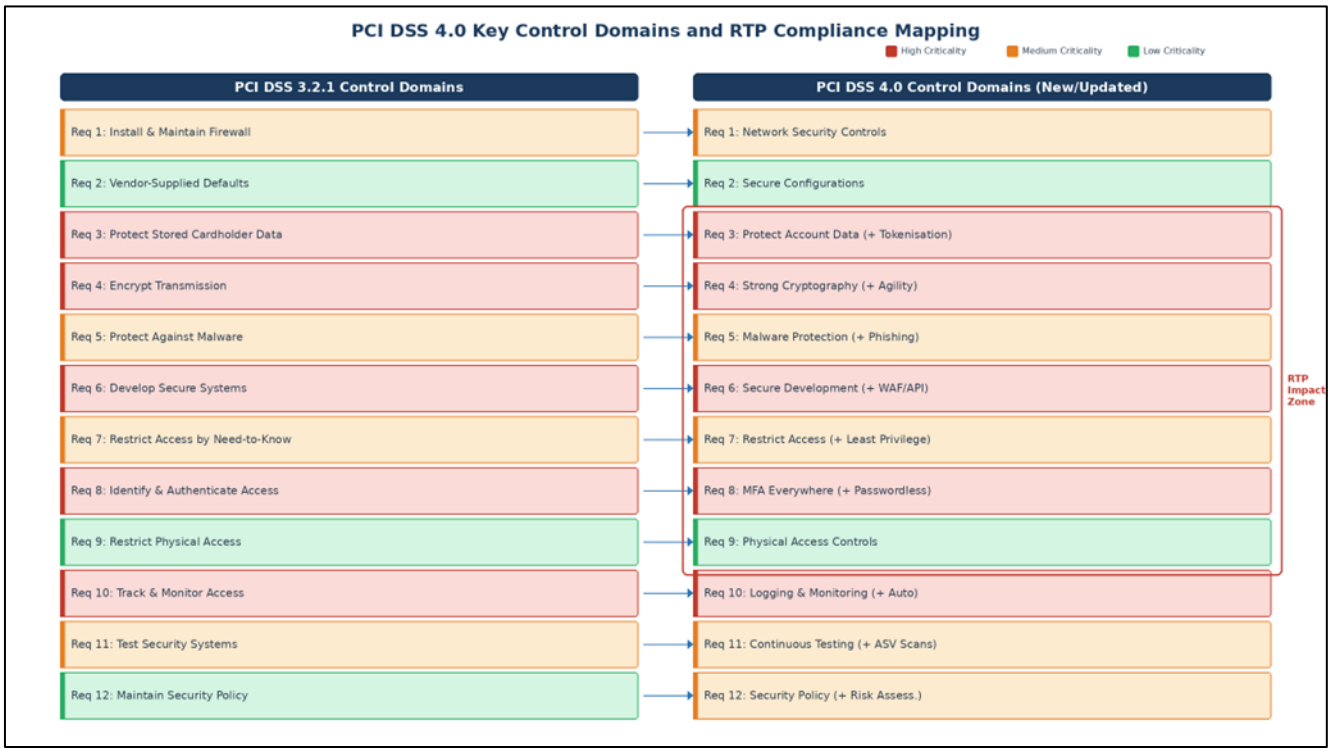

Side-by-side comparison of PCI DSS 3.2.1 vs. 4.0 control domains, overlaid with the real-time payment compliance impact zones. Color-coded by criticality (high / medium / low)

Figure 1 PCI DSS 4.0 Key Control Domains and RTP Compliance Mapping

3.3.2. Open Banking and Data Sharing Regulations

Regulatory frameworks such as the European Union's PSD2, the UK Open Banking Standard, and the Consumer Financial Protection Bureau's Section 1033 rule (United States) mandate that financial institutions expose standardized APIs for account data and payment initiation. While broadening innovation and competition, these frameworks also expand the API attack surface available to adversaries, necessitating robust OAuth 2.0 / FAPI-compliant authorization controls and continuous API traffic monitoring [16].

4. AI/ML Fraud Detection: Synthesized Evidence from the Literature

4.1. Detection Performance

A recurring theme across the reviewed literature is the performance superiority of ML-based systems over rule-based counterparts when evaluated against real-world payment datasets. Supervised classification approaches—including gradient-boosted decision trees (XGBoost, LightGBM), deep neural networks, and ensemble methods—are consistently reported to achieve area-under-the-ROC-curve (AUROC) values in the range of 0.95–0.998 for binary fraud classification tasks [17,18]. Industry deployment reports reviewed from financial institutions in Europe and North America cite fraud detection accuracy in the range of 90–99.7% with corresponding reductions in fraud losses of 40–60% compared with

pre-deployment baselines, though these figures should be interpreted with caution given variation in benchmark methodology, dataset characteristics, and the definition of "accuracy" employed across studies [19,20].

False positive rates—the proportion of legitimate transactions incorrectly flagged for review—represent an equally important performance dimension. The reviewed literature reports that ML systems reduce false positives by 60–82% relative to rule-based systems [21,22], a finding with significant commercial implications: industry analysis estimates that each falsely declined transaction carries an average customer impact equivalent to five times the value of the prevented fraud when downstream customer attrition is accounted for [23].

4.2. Unsupervised and Semi-Supervised Approaches

A significant challenge in payment fraud detection is the extreme class imbalance characteristic of production datasets—fraud events typically represent fewer than 0.1% of total transactions—which limits the availability of labelled training examples for supervised approaches. Unsupervised anomaly detection methods, including autoencoders, isolation forests, and one-class support vector machines, are reviewed in the literature as complementary techniques that can identify novel attack patterns without requiring labelled fraud examples. Federated learning architectures, which allow model training across distributed institutional datasets without centralizing raw transaction data, are an emerging approach to address both the data scarcity and data privacy dimensions of this challenge [24,25].

4.3. Graph-Based Fraud Network Detection

Graph neural networks (GNNs) and graph analytics applied to transaction relationship networks have demonstrated particular effectiveness for detecting organized fraud rings, mule account networks, and synthetic identity clusters that would be invisible to single-account, attribute-based models [26]. By modelling the bipartite graph of accounts, devices, IP addresses, and transactions, GNN-based systems can identify structurally anomalous subgraphs that correlate with fraud rings with substantially higher recall than flat-feature models. Industry deployment evidence reviewed suggests ring detection rates improving by 30–45% over non-graph approaches [12].

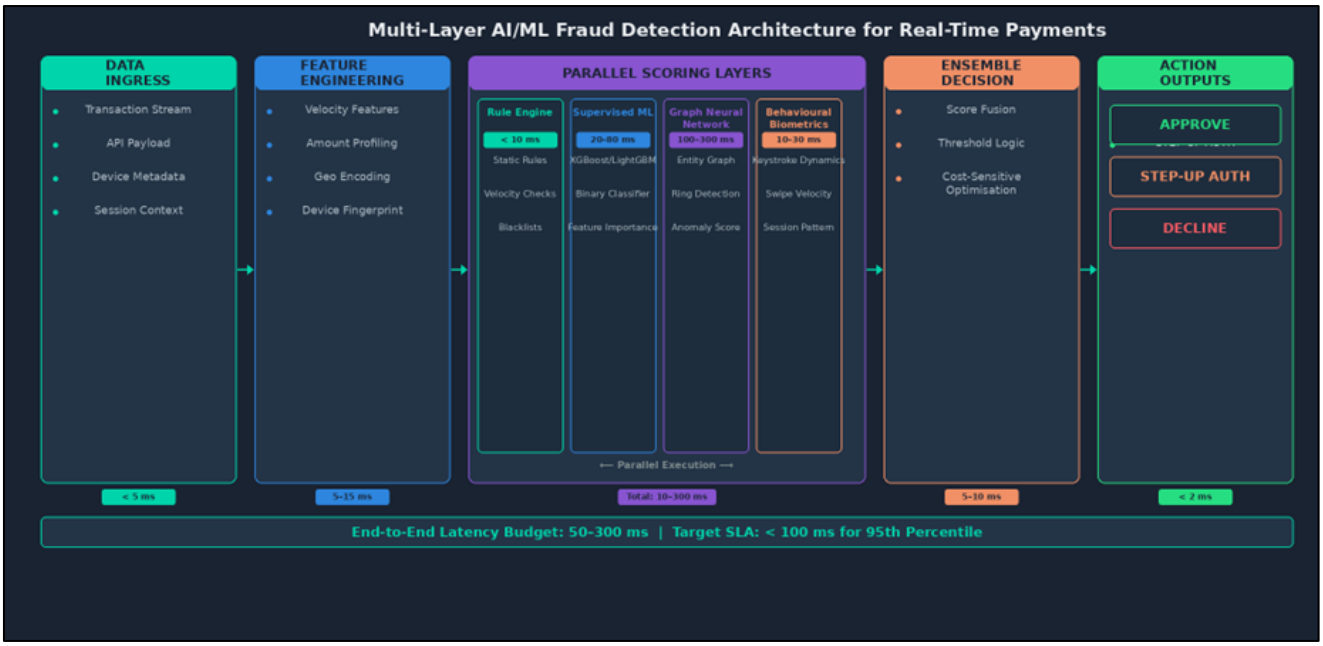

Technical pipeline diagram showing data ingress → feature engineering → parallel scoring layers (rule engine, supervised ML, GNN, behavioral biometrics) → ensemble decision layer → action outputs. Annotated with latency budgets at each layer.

Figure 2 Multi-Layer AI/ML Fraud Detection Architecture for Real-Time Payments

4.4. Comparative Performance Summary

Table 1 Comparative performance of fraud detection approaches as synthesized from reviewed literature. AUROC and latency figures represent ranges reported across reviewed deployments; individual results will vary with dataset characteristics and implementation

Approach	AUROC Range	False Positive Reduction	Latency Profile	Key Limitation
Rule-Based Systems	0.70–0.82	Baseline	Low (<10 ms)	Static; expert-labor intensive
Supervised ML (gradient boosted)	0.95–0.998	60–75%	Medium (20–80 ms)	Requires labelled fraud data
Deep Neural Networks	0.96–0.999	65–80%	Medium–High (40–120 ms)	Interpretability challenges
Graph Neural Networks	0.93–0.997	70–82%	High (100–300 ms)	Infrastructure complexity
Unsupervised (Autoencoder)	0.88–0.95	50–65%	Medium (30–80 ms)	Higher false positive floor
Ensemble / Hybrid	0.97–0.999	75–82%	High (80–250 ms)	Operational complexity

5. Proposed Cloud-Native Multi-Layer Security Architecture

Drawing on the synthesized findings of Section 4 and the compliance requirements of Section 3.3, this section presents a layered security architecture for RTP environments. The architecture is described in platform-agnostic terms; specific cloud service examples are provided parenthetically to illustrate how each layer can be instantiated on contemporary public cloud platforms (AWS, Microsoft Azure, and Google Cloud Platform). The architecture is not prescriptive of any single vendor and should be evaluated against an organization's existing technology landscape, regulatory jurisdiction, and risk appetite.

5.1. Architectural Layers

5.1.1. Layer 1 — Network Perimeter and API Gateway

All inbound payment instruction flows are received through a managed API gateway layer that enforces: TLS 1.3 mutual authentication; OAuth 2.0 / FAPI 2.0 token validation; rate limiting and adaptive throttling; and web application firewall (WAF) rule sets targeting payment-specific injection and enumeration attacks. This layer corresponds to AWS API Gateway with AWS WAF, Azure API Management, or Google Cloud Apigee in public cloud deployments. The gateway layer must be horizontally scalable to absorb peak transaction volumes without introducing queuing latency that degrades the fraud-scoring time budget.

5.1.2. Layer 2 — Identity and Device Intelligence

Prior to fraud scoring, each transaction is enriched with device fingerprint, IP reputation, geolocation anomaly, and account behavioral baseline signals. Behavioral biometric data—keystroke dynamics, swipe velocity, session navigation patterns—provides a passive authentication signal that is particularly effective against account takeover attacks where credentials have been compromised but device behavior diverges from the account holder's established pattern [27]. Identity enrichment at this layer adds 10–30 ms to the total processing window but substantially improves the feature quality available to downstream models.

5.1.3. Layer 3 — Real-Time Fraud Scoring Engine

The scoring layer executes a parallel ensemble of models within the available latency budget. Lightweight rule engines and gradient-boosted classifiers operate in the primary scoring path (target: <50 ms); GNN-based ring detection and deep learning models operate on an asynchronous secondary path with results applied to a risk-score adjustment queue for subsequent transactions from the same entity. The ensemble output is a composite risk score mapped to a three-outcome decision: approve, step-up authenticate, or decline. Threshold calibration must account for the asymmetric cost structure described in Section 4.1.

5.1.4. Layer 4 — Tokenization and Data Protection

Payment account numbers (PANs) and personally identifiable information are never persisted in raw form beyond the ingestion boundary. A format-preserving tokenization vault issues a surrogate token at ingestion time; all downstream processing, storage, and logging operates on tokens. This architectural decision substantially reduces PCI DSS scope by removing raw cardholder data from the majority of system components. Hardware security modules (HSMs) or cloud-native key management services provide the cryptographic boundary. Key rotation schedules and algorithm inventories satisfy the PCI DSS 4.0 cryptographic agility requirement.

5.1.5. Layer 5 — Continuous Monitoring and Adaptive Learning

Post-transaction monitoring feeds confirmed fraud outcomes and investigation dispositions back into a training data pipeline, enabling periodic model retraining against the current fraud pattern distribution. Concept drift detection algorithms—monitoring feature distribution shifts over time—trigger retraining cycles when the production fraud pattern diverges from the training distribution by a defined threshold. Security information and event management (SIEM) integration provides compliance-aligned audit logging for PCI DSS Requirement 10 and enables detection of insider threat and infrastructure-level attack patterns.

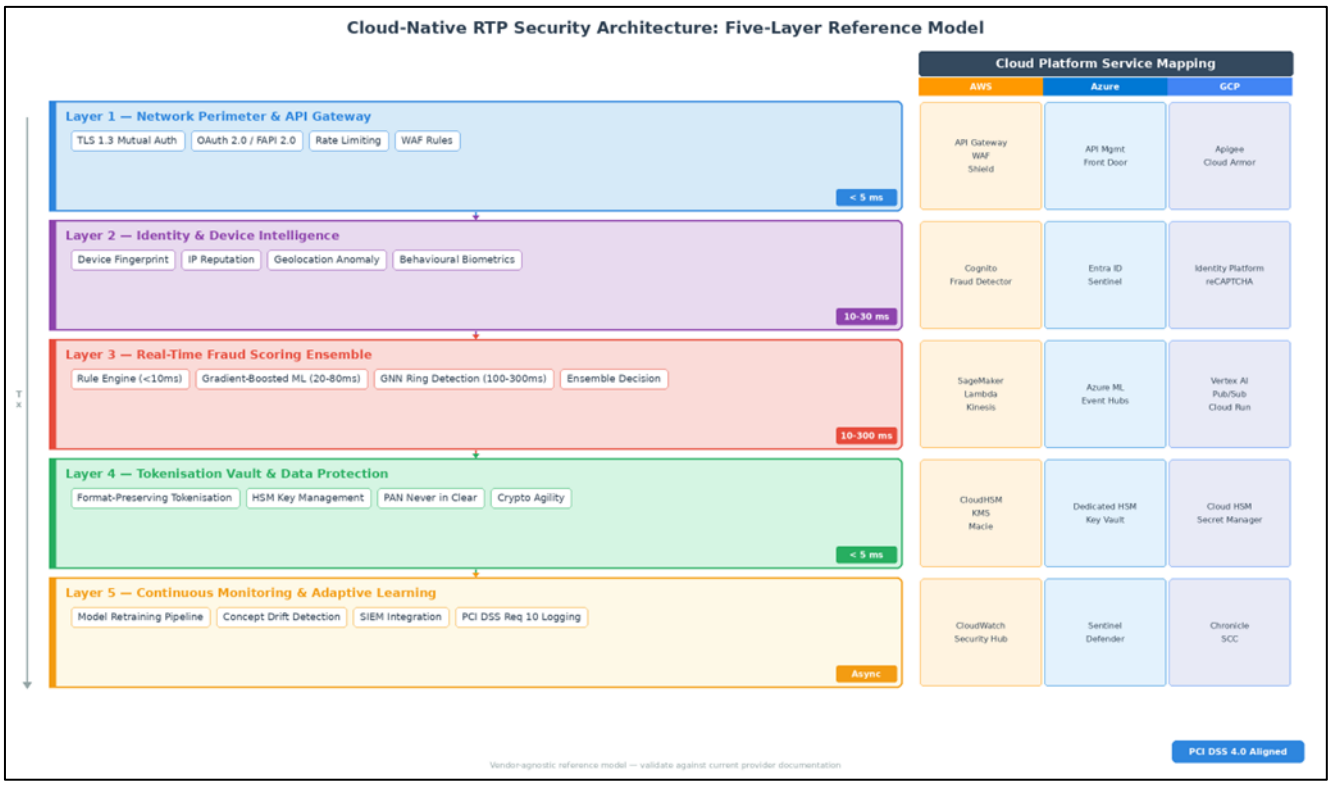

Layered architecture diagram showing: (1) Network / API Gateway, (2) Identity & Device Intelligence, (3) Real-Time Fraud Scoring Ensemble, (4) Tokenization Vault, (5) Continuous Monitoring & Adaptive Learning. Each layer labelled with example cloud services (vendor-agnostic). Consistent color coding per layer.

Figure 3 Cloud-Native RTP Security Architecture: Five-Layer Reference Model

5.2. Platform Mapping

Table 2 provides a mapping of each architectural layer to representative service offerings across the three major cloud platforms. The mapping is intended to assist practitioners in translating the reference architecture into a deployment-ready service selection, not to imply performance or security equivalence across platforms—these characteristics vary with configuration, region, and service tier.

Table 2 Indicative mapping of architectural layers to cloud platform services. Service availability and feature sets are subject to change; organizations should validate against current provider documentation

Architectural Layer	AWS	Microsoft Azure	Google Cloud Platform
Network / API Gateway	API Gateway + WAF + Shield	API Management + Front Door	Apigee + Cloud Armor
Identity & Device Intelligence	Cognito + Fraud Detector	Entra ID + Microsoft Sentinel	Identity Platform + reCAPTCHA Enterprise
Fraud Scoring Engine	SageMaker + Lambda + Kinesis	Azure ML + Event Hubs + Functions	Vertex AI + Pub/Sub + Cloud Run
Tokenization / Key Mgmt	CloudHSM + KMS + Macie	Azure Dedicated HSM + Key Vault	Cloud HSM + Secret Manager
Monitoring & Adaptive Learning	CloudWatch + Security Hub + GuardDuty	Sentinel + Monitor + Defender	Chronicle + Security Command Center

6. Discussion

6.1. Synthesis of Key Findings

The evidence reviewed in this paper converges on several consistent themes. First, the latency asymmetry between real-time payment settlement and legacy fraud-scoring infrastructure is a structural problem that cannot be resolved through incremental optimization of rule-based systems; architectural replacement is required. Second, AI/ML-based approaches consistently outperform rule-based systems on both detection accuracy and false positive rates across diverse deployment contexts, though the magnitude of improvement is sensitive to dataset quality, feature engineering investment, and the specific fraud typology targeted. Third, cloud-native deployment architectures provide the combination of elastic scalability, managed ML services, and integrated security tooling that makes the operational model economically viable for institutions of varying scale.

6.2. Trade-offs and Implementation Considerations

Practitioners should be aware of several material trade-offs embedded in the proposed architecture. The shift to ML-based scoring introduces model interpretability challenges: regulatory scrutiny under consumer protection frameworks (e.g., FCRA in the US, GDPR in the EU) may require the ability to explain adverse payment decisions to affected customers, which is non-trivial for deep learning models. Gradient-boosted methods offer superior interpretability through feature importance ranking and SHAP value attribution; deep learning and GNN approaches require additional explainability infrastructure.

The federated learning approaches discussed in Section 4.2 represent a promising direction for cross-institutional fraud signal sharing but introduce governance complexity: consortium agreements, differential privacy parameters, and model versioning standards must be established before federated training can be operationalized. Regulatory acceptance of federated model outputs as a basis for adverse action decisions has not yet been fully established in major jurisdictions.

6.3. Limitations of This Review

This review carries several limitations that constrain the generalizability of its findings. First, the performance figures cited in Section 4 are drawn from heterogeneous sources with varying methodological quality; direct comparison across studies is limited by differences in dataset characteristics, evaluation metrics, and deployment context. The absence of a common benchmarking standard for RTP fraud detection is a significant gap in the field. Second, the architectural synthesis in Section 5, while informed by multiple deployment examples, has not been empirically validated in a single production environment by the authors. Third, the focus on technology architecture necessarily understates the organizational, cultural, and workforce dimensions of effective fraud prevention—factors that practitioner literature consistently identifies as determinants of program success.

6.4. Future Research Directions

Several research directions emerge from this synthesis. Longitudinal controlled studies comparing rule-based, ML, and hybrid architectures within a single institutional context would substantially strengthen the evidence base for architectural recommendations. The application of formal verification methods to fraud-scoring pipeline logic—ensuring that model updates do not introduce regulatory compliance violations—represents an underexplored intersection of software engineering and financial compliance. Finally, the adversarial robustness of deployed ML models to deliberate evasion by sophisticated fraud actors warrants dedicated study; current deployment literature largely addresses accuracy against historical fraud patterns rather than against adversarial constructed inputs.

7. Conclusion

This paper has provided a technical review and architectural synthesis addressing the security challenges of real-time payment systems. The convergence of irrevocable settlement mechanics, accelerating fraud sophistication, and tightening regulatory requirements creates a set of demands that legacy fraud-prevention infrastructure is structurally unable to meet. The reviewed literature consistently supports the effectiveness of AI/ML-based detection approaches—particularly supervised ensemble methods and graph neural networks—in achieving the accuracy and latency performance that RTP environments require, with fraud loss reductions of 40–60% and false positive reductions of 60–82% reported across multiple industry deployments, as synthesized from the cited sources.

The five-layer cloud-native architecture proposed in Section 5 provides a vendor-agnostic reference model that translates these capabilities into an implementable, PCI DSS 4.0-aligned framework. By separating network control, identity intelligence, real-time scoring, data protection, and adaptive monitoring into distinct, independently scalable layers, the architecture enables organizations to adopt components progressively in alignment with their risk appetite and technology maturity.

The limitations identified in Section 6.3 underscore that this review should be treated as a starting point for institutional decision-making rather than a prescriptive deployment guide. Empirical validation of the proposed architecture in production environments, and the development of standardized benchmarking methodology for RTP fraud detection, represent the most pressing near-term contributions the research community can make to this domain.

Compliance with ethical standards

Disclosure of Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

No new datasets were created or analyzed in this study. Literature sources are cited in the References section.

Author Contributions (CRediT Taxonomy)

Conceptualization: Raghunandan G Ramakrishna; Methodology: Raghunandan G Ramakrishna; Formal Analysis: Raghunandan G Ramakrishna; Investigation: Raghunandan G Ramakrishna; Writing – Original Draft: Raghunandan G Ramakrishna; Writing – Review & Editing: Raghunandan G Ramakrishna. The author has read and agreed to the published version of the manuscript.

References

[1] Federal Reserve Banks. FedNow Service: Operational Statistics and Adoption Data [Internet]. 2023.

[2] National Payments Corporation of India. UPI Annual Report 2022–23. Mumbai: NPCI; 2023.

[3] UK Finance. Annual Fraud Report 2023. London: UK Finance; 2023.

[4] European Banking Authority. Report on Payment Fraud Under PSD2. EBA/REP/2022/04. Paris: EBA; 2022.

[5] PCI Security Standards Council. PCI DSS v4.0 Requirements and Testing Procedures. Wakefield (MA): PCI SSC; 2022.

[6] Veluru SP, et al. Real-Time Fraud Detection in Payment Systems Using Kafka and Machine Learning. J Recent Trends Comput Sci Eng. 2019.

[7] Naqvi SH. Fraud Detection Using Machine Learning in Financial Transactions: A Systematic Review. Veredas do Direito. 2026.

[8] Payments UK. The Role of Irrevocability in Instant Payment Fraud Risk. London: Payments UK; 2021.

[9] UK Finance. Authorised Push Payment Fraud: Full Year 2022 Figures. London: UK Finance; 2023.

[10] Nacha. ACH Network Risk and Compliance. Herndon (VA): Nacha; 2023.

[11] Payment Systems Regulator. APP Scams: Mandatory Reimbursement — Policy Statement. London: PSR; 2023. PS23/3.

[12] Liu Y, Ao X, Qin Z, et al. Pick and choose: a GNN-based imbalanced learning approach for fraud detection. In: WWW '21. New York: ACM; 2021. p. 3168–79.

[13] Zhang X, Liu X, Zheng Q. Synthetic identity fraud detection using graph-based unsupervised clustering. Expert Syst Appl. 2023;211:118562.

[14] Kim H, et al. When LLMs Go Online: The Emerging Threat of Web-Enabled LLMs. In: USENIX Security '25. 2025.

[15] Komninos N, Roumeliotis C, Sarris M. PCI DSS 4.0 Compliance Readiness in Financial Institutions: A Survey. IEEE Access. 2023;11:92340–55.

[16] European Commission. Revised Payment Services Directive (PSD2). Official Journal of the EU. L 337/35; 2015.

[17] Borketey B. Real-Time Fraud Detection Using Machine Learning. J Data Analysis and Information Processing. 2024.

[18] Btoush E, et al. Achieving Excellence in Cyber Fraud Detection: A Hybrid ML+DL Ensemble Approach. Applied Sciences. 2025;15(3):1081.

[19] McKinsey & Company. Global Payments Report 2023. New York: McKinsey & Company; 2023.

[20] Hashemi M, et al. Fraud Detection in Banking Data by Machine Learning Techniques. IEEE Access. 2022.

[21] Saia R, Carta S. Evaluating the Overhead of Fraud Detection Mechanisms in Microservices-Based Payment Systems. IEEE Trans Serv Compute. 2022;15(3):1221–34.

[22] Lebichot B, et al. Deep-Learning Domain Adaptation Techniques for Credit Cards Fraud Detection. In: INNSBDDL 2019. Springer; 2019. p. 78–88.

[23] Javelin Strategy & Research. The True Cost of Payment Fraud. Pleasanton (CA): Javelin; 2022.

[24] Yang Q, Liu Y, Chen T, Tong Y. Federated machine learning: concept and applications. ACM Trans Intell Syst Technol. 2019;10(2):12.

[25] Zheng Z, et al. Federated fraud detection: privacy-preserving collaborative learning. In: ACM SIGKDD 2022. p. 4330–40.

[26] Naqvi SH. Fraud Detection Using Machine Learning in Financial Transactions: A Systematic Review. Veredas do Direito. 2026.

[27] Punpunwong P, Traikultraipruk N. Behavioural Biometrics for Authentication in Online Banking: A Systematic Review. IEEE Access. 2023;11:45220–38.