iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

A HYBRID MACHINE LEARNING FRAMEWORK COMBINING BEHAVIORAL ANOMALY DETECTION AND RULE BASED THREAT INTELLIGENCE FOR REAL TIME INTRUSION DETECTION IN MULTI CLOUD AND FINTECH SYSTEMS

Adebayo Olufemi SOSANYA
Department of Computer Science, San Francisco Bay University, Fremont, California USA
asosanya166@student.sfbu.edu
https://orcid.org/0009-0002-0463-1726

Mary Oluwatoyin ALADE
Department of Business Administration, Lincoln University, Oakland, California USA
malade@lincolnucasf.edu

ABSTRACT
Financial technology (FinTech) enterprises and digital banking institutions increasingly operate across distributed multi-cloud infrastructures that expose microservices, payment application programming interfaces (APIs), and containerized workloads to high-velocity cyber intrusions. Conventional signature-based intrusion detection systems (IDS) excel at blocking known threat indicators with sub-millisecond latency but fail to detect novel zero-day exploits and low-and-slow behavioral lateral movement. Conversely, standalone unsupervised anomaly detectors generate high false-positive rates in dynamic cloud environments, creating severe alert fatigue for Security Operations Centers (SOCs). This study develops and evaluates a hybrid real-time intrusion detection architecture that fuses deterministic cyber threat intelligence (CTI) rule matching with unsupervised behavioral anomaly scoring and supervised gradient-boosted meta-classification. Grounded in design-science methodology, the proposed three-stage pipeline combines (1) high-throughput STIX/TAXII threat-intelligence rule filtering, (2) unsupervised Isolation Forest behavioral deviation scoring across east–west cloud API telemetry, and (3) a calibrated LightGBM meta-classifier that synthesizes rule-match severity, anomaly scores, and contextual FinTech asset criticality. The framework is benchmarked on canonical multi-attack intrusion datasets (UNSW-NB15 and CICIDS2017) augmented with multi-cloud API burst profiles. Results: The hybrid architecture achieves superior Precision-Recall Area Under the Curve (PR-AUC) and F1-score compared to standalone rule-based or purely unsupervised baselines, reducing false-positive SOC escalations by over 78% while preserving sub-10-millisecond inline inference latency. Integrating deterministic threat intelligence with behavioral machine learning provides FinTech institutions with an auditable, low-latency intrusion defense aligned with NIST Cybersecurity Framework (CSF) 2.0 continuous detection mandates.

Keywords:
intrusion detection system; hybrid machine learning; behavioral anomaly detection; cyber threat intelligence; multi-cloud security; FinTech infrastructure

1. INTRODUCTION

Modern financial technology (FinTech) platforms, payment gateways, and cloud-native banking institutions rely heavily on distributed multi-cloud architectures to achieve elastic scalability, high availability, and rapid service deployment across payment application programming interfaces (APIs) and core ledger microservices (U.S. Department of the Treasury, 2024). However, distributing sensitive financial workloads across multiple public and private cloud environments significantly expands the external and internal attack surface. Unlike traditional perimeter-bound enterprise networks, multi-cloud FinTech ecosystems process continuous east–west microservice traffic, ephemeral container sessions, and high-frequency external API calls, creating complex visibility blind spots that adversaries exploit for credential stuffing, API abuse, command-and-control (C2) beaconing, and lateral data exfiltration (Khraisat et al., 2019; U.S. Department of the Treasury, 2024).

Historically, financial-sector Security Operations Centers (SOCs) have depended on signature-based and rule-based Intrusion Detection Systems (IDS) enriched by structured Cyber Threat Intelligence (CTI) feeds such as malicious Internet Protocol (IP) blocklists, known exploit payload signatures, and Tactics, Techniques, and Procedures (TTPs) shared via Structured Threat Information Expression (STIX) and Trusted Automated Exchange

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

of Indicator Information (TAXII) protocols (Buczak & Guven, 2016; Tounsi & Rais, 2018). Rule-based CTI engines deliver deterministic accuracy, immediate regulatory auditability, and near-zero false positives for known attack signatures. Yet, their fundamental limitation is reactivity: polymorphic malware, zero-day API exploits, and compromised legitimate service accounts easily evade static indicators until threat feeds are updated post-incident (Khraisat et al., 2019; Sommer & Paxson, 2010).

To overcome signature blindness, researchers and practitioners have increasingly turned to machine learning (ML) and behavioral anomaly detection, modeling baseline network flow and API call distributions to flag statistical deviations indicative of novel intrusions (Chandola et al., 2009; Ahmed et al., 2016). Algorithms such as Isolation Forests and deep autoencoders can identify previously unseen reconnaissance or exfiltration patterns without prior attack labels (Liu et al., 2008; Ferrag et al., 2020). Nevertheless, as Sommer and Paxson (2010) famously demonstrated in their foundational critique of ML-based intrusion detection, purely unsupervised anomaly detectors deployed in high-variance production networks suffer from prohibitive false-positive rates, misclassifying benign cloud auto-scaling spikes, batch payment settlements, and software deployments as malicious intrusions.

Recognizing these complementary strengths and weaknesses, the National Institute of Standards and Technology (NIST) Cybersecurity Framework (CSF) 2.0 emphasizes continuous, risk-prioritized anomaly and threat detection (*DE.CM* and *DE.AE*) supported by supply-chain and threat-intelligence governance (National Institute of Standards and Technology [NIST], 2024). Yet, operationalizing a truly unified hybrid IDS in real-time FinTech environments poses three unresolved engineering challenges:

1. **Latency vs. Model Complexity Trade-Offs:** Deep neural architectures often incur inference latencies that bottleneck high-throughput payment API gateways, whereas ultra-fast tree ensembles require careful feature engineering to capture temporal behavioral drift (Ke et al., 2017; Ferrag et al., 2020).
2. **Naive Score Aggregation:** Simply taking a logical OR union of rule-based alerts and anomaly flags compounds false-positive alert fatigue, whereas taking a strict AND intersection sacrifices zero-day detection capability.
3. **Lack of Contextual FinTech Asset Weighting:** Generic network anomaly scores treat low-sensitivity public web endpoints identically to core payment-clearing or key-management microservices.

To address these gaps, this study designs and empirically evaluates a Hybrid Three-Stage Intrusion Detection Framework combining deterministic CTI rule matching, unsupervised Isolation Forest behavioral profiling, and supervised LightGBM meta-classification tailored to multi-cloud FinTech systems. Specifically, the paper investigates three research questions:

- **RQ1:** How can deterministic CTI rule-matching outputs and unsupervised behavioral anomaly scores be mathematically fused with FinTech microservice asset criticality to maximize detection sensitivity while suppressing false-positive SOC alert volume?
- **RQ2:** To what extent does the proposed hybrid meta-classification architecture outperform standalone rule-based IDS and standalone unsupervised anomaly detectors across canonical intrusion benchmarks (UNSW-NB15 and CICIDS2017) under multi-cloud traffic burst conditions?
- **RQ3:** Can the hybrid pipeline sustain sub-10-millisecond inline inference latency and generate structured, analyst-ready triage explanations aligned with NIST CSF 2.0 continuous monitoring requirements?

2. LITERATURE REVIEW AND MULTI-CLOUD FINTECH CONTEXT

2.1 Architectural Evolution of Multi-Cloud FinTech Environments

Financial technology (FinTech) platforms and digital banking institutions have increasingly migrated from monolithic on-premise data centers to distributed multi-cloud and hybrid-cloud architectures (U.S. Department of the Treasury, 2024). By deploying containerized payment gateways, real-time clearing engines, and identity microservices across multiple cloud service providers, FinTech operators gain geographic redundancy, elastic compute scaling during peak settlement windows, and reduced single-vendor lock-in (U.S. Department of the Treasury, 2024). However, this architectural decentralization fragments network visibility across heterogeneous virtual private clouds (VPCs), ephemeral container orchestrators, and external application programming interfaces (APIs; Khraisat et al., 2019).

In multi-cloud FinTech networks, adversarial campaigns rarely follow single-vector perimeter breaches. Instead, attackers exploit exposed API endpoints, compromised continuous-integration tokens, or misconfigured cloud storage buckets to establish an initial foothold, subsequently executing low-and-slow east–west lateral movement toward core ledger databases and cryptographic signing services (Khraisat et al., 2019; U.S. Department of the Treasury, 2024). Consequently, effective intrusion detection requires continuous, low-latency correlation of both

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

north–south ingress API traffic and east–west inter-service telemetry (National Institute of Standards and Technology [NIST], 2024).

2.2 Signature-Based IDS and Cyber Threat Intelligence (CTI) Feeds

Deterministic signature-based and rule-based Intrusion Detection Systems (IDS) remain the primary frontline defense across regulated financial networks (Buczak & Guven, 2016; Khraisat et al., 2019). Modern enterprise rule engines, augmented by automated Cyber Threat Intelligence (CTI) feeds formatted in Structured Threat Information Expression (STIX) and delivered via Trusted Automated Exchange of Indicator Information (TAXII) protocols, match incoming network flows against curated indicators of compromise (IoCs), including known command-and-control (C2) IP addresses, malicious TLS certificate hashes, and canonical exploit byte sequences (Tounsi & Rais, 2018).

Rule-based CTI architectures offer three operational advantages critical to FinTech compliance: sub-millisecond packet/flow evaluation, near-zero false-positive rates on verified indicators, and transparent audit trails mapping alerts directly to known adversary Tactics, Techniques, and Procedures (TTPs; Tounsi & Rais, 2018; NIST, 2024). Nevertheless, signature-based systems suffer from inherent structural blind spots. Because static rules depend on prior forensic observation and intelligence sharing, they cannot detect zero-day API parameter exploits, polymorphic malware variants, or credential-abuse attacks executed from clean cloud IP ranges (Buczak & Guven, 2016; Sommer & Paxson, 2010).

2.3 Unsupervised Behavioral Anomaly Detection and the False-Positive Bottleneck

To detect previously unseen intrusions without labeled attack signatures, extensive scholarship has investigated unsupervised and semi-supervised behavioral anomaly detection (Chandola et al., 2009; Ahmed et al., 2016). By modeling normal statistical baselines of packet inter-arrival times, byte-transfer ratios, API call frequencies, and session durations, unsupervised algorithms flag structural outliers as potential intrusions (Chandola et al., 2009). Among tree-based unsupervised methods, the Isolation Forest algorithm (Liu et al., 2008) has gained wide adoption in high-throughput network security due to its linear time complexity, low memory footprint, and ability to isolate high-dimensional anomalies without requiring expensive distance or density computations. Similarly, deep learning models such as autoencoders and recurrent neural networks capture complex non-linear dependencies across raw network flows (Ferrag et al., 2020).

Despite strong theoretical recall on zero-day anomalies, Sommer and Paxson (2010) identified a persistent operational barrier in production networks: the base-rate fallacy and false-positive alert fatigue. Because legitimate multi-cloud FinTech traffic exhibits extreme non-stationary variance driven by scheduled batch settlements, automated microservice scaling, and promotional retail payment surges standalone anomaly detectors routinely misclassify benign operational spikes as severe intrusions (Sommer & Paxson, 2010; Ahmed et al., 2016). In Security Operations Centers (SOCs), excessive false-positive alert volumes overwhelm human analysts and erode trust in automated blocking actions (Khraisat et al., 2019).

2.4 Benchmark Datasets and Hybrid Intrusion Detection Paradigms

Empirical evaluation of network intrusion frameworks relies heavily on standardized benchmark corpora that reflect realistic modern attack taxonomies (Ring et al., 2019). Early datasets such as KDD Cup 99 suffered from severe redundancy and outdated protocol distributions, prompting the development of UNSW-NB15 (Moustafa & Slay, 2015) which captures nine contemporary attack families (including fuzzers, backdoors, exploits, and reconnaissance) alongside realistic normal traffic and CICIDS2017 (Sharafaldin et al., 2018), which incorporates modern web-application, brute-force, infiltration, and distributed denial-of-service (DDoS) flows.

To reconcile the precision of signature engines with the zero-day sensitivity of unsupervised models, hybrid IDS architectures combine multiple detection paradigms (Buczak & Guven, 2016; Khraisat et al., 2019). Gradient-boosted decision trees, specifically LightGBM (Ke et al., 2017) have demonstrated exceptional performance as supervised meta-classifiers in tabular network security pipelines, offering rapid histogram-based training, native handling of skewed flow features, and low-latency inference compatible with inline payment gateways (Ke et al., 2017; Ferrag et al., 2020). However, existing hybrid studies rarely incorporate explicit FinTech microservice asset criticality or NIST Cybersecurity Framework (CSF) 2.0 governance controls (NIST, 2024; U.S. Department of the Treasury, 2024) into the real-time fusion layer; a gap addressed directly by the three-stage architecture developed in Section 3.

3. METHODOLOGY AND THREE-STAGE HYBRID INTRUSION DETECTION ARCHITECTURE

3.1 Design Science Formulation and Multi-Cloud Telemetry Pipeline

Following the design-science research methodology (Hevner et al., 2004), this study constructs a real-time, three-stage hybrid intrusion detection artifact engineered for distributed multi-cloud financial technology (FinTech) platforms. Rather than relying on a single monolithic classifier, the architecture decomposes real-time flow

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

inspection into three sequential, complementary stages: (1) a deterministic Cyber Threat Intelligence (CTI) Rule Engine for known indicators, (2) an unsupervised Behavioral Anomaly Detector for zero-day and east–west lateral deviations, and (3) a supervised Contextual Meta-Classifier that integrates rule matches, behavioral anomaly scores, and FinTech microservice asset criticality into an actionable Security Operations Center (SOC) disposition (Buczak & Guven, 2016; National Institute of Standards and Technology [NIST], 2024).

Telemetry ingestion normalizes north–south external payment gateway logs and east–west container-to-container network flow records into standardized feature vectors $x_i \in R^d$ within rolling 5-second micro-batches. Core extracted flow attributes include packet rate, byte asymmetry ratio, TCP flag entropy, TLS handshake anomalies, API endpoint invocation frequency, and inter-arrival time dispersion (Sharafaldin et al., 2018; Ring et al., 2019).

3.2 Stage 1: Deterministic STIX/TAXII Threat Intelligence Rule Filtering

In Stage 1, each incoming flow x_i is evaluated against an in-memory hash and signature index populated via Structured Threat Information Expression (STIX) and Trusted Automated Exchange of Indicator Information (TAXII) feeds (Tounsi & Rais, 2018). The rule engine evaluates four indicator classes: known command-and-control (C2) IP/autonomous-system blocklists, compromised TLS certificate fingerprints, canonical SQL/command-injection payload signatures, and unauthorized east–west port-scan sequences.

Each matched rule $r \in R$ carries an indicator confidence weight $c_r \in [0, 1]$ and an adversarial severity score $v_r \in [0, 1]$ calibrated to MITRE ATT&CK tactic criticality. The aggregate Stage 1 Rule-Match Threat Score $R(x_i) \in [0, 1]$ is formulated as:

$$R(x_i) = 1 - \prod_{r \in M(x_i)} (1 - c_r \cdot v_r)$$

where $M(x_i) \subseteq R$ denotes the subset of rules triggered by flow x_i, and $R(x_i) = 0$ when $M(x_i) = \emptyset$. This noisy-OR formulation ensures that multiple corroborating threat-intelligence hits monotonically approach 1.0 without exceeding unit bounds, while high-confidence deterministic matches ($R(x_i) \geq 0.90$) trigger immediate inline packet drop alongside parallel logging for downstream model retraining (Khraisat et al., 2019).

3.3 Stage 2: Unsupervised Behavioral Anomaly Detection (Isolation Forest)

Flows that do not trigger an immediate deterministic block ($R(x_i) < 0.90$) proceed concurrently to Stage 2, which detects novel zero-day exploits, credential stuffing, and low-and-slow lateral movement (Chandola et al., 2009; Ahmed et al., 2016). To preserve sub-10-millisecond inline latency across high-volume FinTech APIs, Stage 2 employs an ensemble of $T = 150$ Isolation Trees (Liu et al., 2008) trained exclusively on verified benign baseline traffic segmented by microservice tier.

For a flow vector x_i, let $h_t(x_i)$ denote the path length from the root node to the isolating external leaf in tree t. The raw unsupervised anomaly score $A_{raw}(x_i) \in (0, 1)$ is defined per Liu et al. (2008) as:

$$A_{raw}(x_i) = 2^{-E[h(x_i)] / c(\psi)}, \quad c(\psi) = 2H(\psi - 1) - 2(\psi - 1)/\psi$$

where $\psi = 256$ is the sub-sampling size, $H(k) \approx \ln(k) + 0.5772156649$ is the harmonic number, and $c(\psi)$ normalizes average path length. Because raw Isolation Forest scores exhibit compressed dynamic range in skewed cloud traffic, we apply min-max quantile calibration fitted on validation baseline flows to obtain a standardized Behavioral Anomaly Score $B(x_i) \in [0, 1]$, supplemented by local feature-deviation z-scores z_i capturing burst deviations in API call rate and outbound byte ratio (Sommer & Paxson, 2010).

3.4 Stage 3: Contextual LightGBM Meta-Classification and Asset Criticality Fusion

To eliminate the false-positive alert fatigue common to standalone anomaly detectors (Sommer & Paxson, 2010; Khraisat et al., 2019), Stage 3 fuses Stage 1 and Stage 2 outputs with FinTech Microservice Asset Criticality $\gamma_i \in [0.25, 1.00]$ (RQ1). Each cloud workload is assigned a structural criticality tier aligned with NIST Cybersecurity Framework (CSF) 2.0 asset management (*ID.AM*) and U.S. Department of the Treasury (2024) operational resilience guidelines:

- **Tier 1 ($\gamma_i = 1.00$):** Core payment clearing engines, HSM cryptographic signing services, and ledger databases.
- **Tier 2 ($\gamma_i = 0.75$):** Customer authentication gateways, KYC identity APIs, and card-tokenization vaults.
- **Tier 3 ($\gamma_i = 0.45$):** Internal developer telemetry, staging clusters, and public informational web frontends.

Stage 3 constructs a compact meta-feature vector $m_i = [R(x_i), B(x_i), R(x_i) \cdot B(x_i), \gamma_i, \gamma_i \cdot B(x_i), z_i^T]^T$ and passes m_i into a calibrated LightGBM gradient-boosted meta-classifier (Ke et al., 2017) trained via binary cross-entropy loss with class-imbalance weighting. The meta-classifier outputs a calibrated posterior intrusion probability $P_{meta}(y_i = 1 \mid m_i) \in [0, 1]$, mapped onto a Composite Cyber-Intrusion Risk Score $S_{IDS}(x_i) = 100 \times P_{meta}(y_i = 1 \mid m_i)$.

Operational routing follows three latency-bounded thresholds:

- **$S_{IDS}(x_i) \geq 80$ (Automated Inline Quarantine & SOC P1 Escalation):** Terminates anomalous API session tokens, isolates container network policies, and dispatches a structured STIX-enriched alert.

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

- **45 ≤ S_IDS(x_i) < 80 (Adaptive Rate-Limiting & SOC P2 Triage Queue):** Applies step-up mTLS/token verification and queues the flow for analyst review within 15 minutes.
- **S_IDS(x_i) < 45 (Permitted Inline Flow):** Passes traffic with zero user-facing latency penalty.

3.5 Experimental Benchmark Design and Ablation Protocol

To evaluate detection efficacy, false-positive suppression, and inference speed (RQ2–RQ3), the framework is benchmarked on two canonical, multi-attack intrusion datasets UNSW-NB15 (Moustafa & Slay, 2015) and CICIDS2017 (Sharafaldin et al., 2018) augmented with synthetic multi-cloud microservice tier labels (γ_i) and diurnal batch-settlement traffic bursts simulating FinTech production variance (Ring et al., 2019; Ferrag et al., 2020).

We conduct a systematic four-way ablation study comparing: (1) Baseline 1 — Standalone CTI Rule Engine (R(x_i) only), (2) Baseline 2 — Standalone Unsupervised Anomaly Detector (B(x_i) only), (3) Baseline 3 — Naive Logical-OR Hybrid (R OR B), and (4) Proposed Framework — Three-Stage Contextual LightGBM Hybrid (S_IDS(x_i)). All models are evaluated using Precision, Recall, F1-score, Precision-Recall Area Under the Curve (PR-AUC), False-Positive Rate (FPR), Zero-Day Recall (measured on held-out attack families excluded from Stage 1 rules), and median/99th-percentile inline inference latency in milliseconds (Ke et al., 2017; Ferrag et al., 2020).

4. EXPERIMENTAL RESULTS AND LATENCY EVALUATION

4.1 Comparative Detection and Ablation Performance (RQ1–RQ2)

Table 1 reports the classification performance across the multi-cloud FinTech intrusion benchmark (N = 200,000 network and API flows derived from UNSW-NB15 and CICIDS2017 attack profiles augmented with diurnal batch-settlement spikes; attack prevalence = 3.40%, including 28% held-out zero-day attack families excluded from Stage 1 Cyber Threat Intelligence [CTI] signatures). The proposed Three-Stage Contextual LightGBM Hybrid (S_IDS(x_i) ≥ 80) is benchmarked against the three ablation baselines defined in Section 3.5.

Table 1. Comparative Intrusion Detection and Zero-Day Recall Performance Across Multi-Cloud FinTech Flows (N = 200,000)

Baseline 1: Standalone CTI Rule Engine (R(x_i))	**0.942**	0.568	0.709	0.614	**0.12%**	0.041
Baseline 2: Standalone Isolation Forest (B(x_i))	0.418	0.784	0.545	0.521	3.84%	0.762
Baseline 3: Naive Logical-OR Hybrid (R OR B)	0.489	**0.916**	0.638	0.672	3.92%	0.771
Proposed: Three-Stage Contextual Hybrid (S_IDS ≥ 80)	0.896	0.894	**0.895**	**0.918**	0.36%	**0.815**

The ablation results confirm the structural trade-offs identified in prior intrusion detection scholarship (Sommer & Paxson, 2010; Khraisat et al., 2019). The Standalone CTI Rule Engine achieves high Precision (0.942) and minimal false positives (0.12%) on known indicators of compromise, yet fails almost entirely on held-out zero-day exploits (Zero-Day Recall = 0.041). Conversely, the Standalone Isolation Forest captures 76.2% of unseen

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

zero-day anomalies but generates a 3.84% false-positive rate during benign cloud auto-scaling and batch settlement bursts a classic manifestation of the base-rate fallacy (Sommer & Paxson, 2010; Ahmed et al., 2016). Taking a naive logical-OR union (R OR B) lifts overall Recall to 0.916 but compounds false-positive alerts to 3.92%.

By fusing Stage 1 rule confidence, Stage 2 behavioral deviation, and Stage 3 FinTech microservice asset criticality (γ_i) inside a calibrated LightGBM meta-classifier (Ke et al., 2017), the proposed framework achieves a 0.895 F1-score and 0.918 PR-AUC. Relative to the standalone anomaly detector and naive OR hybrid, the contextual meta-classifier suppresses false-positive Security Operations Center (SOC) escalations by 90.6% (reducing FPR from 3.84% to 0.36%) while lifting held-out zero-day recall to 81.5%.

4.2 Attack Family Disaggregation and Inline Latency Benchmarks (RQ3)

Table 2 disaggregates detection recall and inline inference latency across major multi-cloud FinTech attack vectors, spanning API exploitation, credential brute-forcing, east–west backdoor/lateral movement, and distributed denial-of-service (DDoS) floods (Moustafa & Slay, 2015; Sharafaldin et al., 2018).

Table 2. Detection Efficacy and Inline Inference Latency by Multi-Cloud Attack Vector

Known C2 & Payload Exploits (North–South API)	Stage 1 CTI + Stage 3	0.912	0.694	0.948	1.8 / 4.2 ms
Zero-Day API Parameter Fuzzing (Tier 1 Ledger)	Stage 2 IF + Stage 3 (γ_i)	0.038	0.781	0.842	2.6 / 5.9 ms
East–West Lateral Movement & Backdoor Beacons	Stage 2 IF + Stage 3	0.145	0.758	0.836	2.7 / 6.1 ms
Distributed Credential Stuffing & Volumetric DDoS	Stage 1 + Stage 2 Burst z_i	0.684	0.825	0.914	2.4 / 5.4 ms

Across all traffic classes, the end-to-end three-stage inference pipeline sustains a median latency of 2.4 ms and a 99th-percentile (P99) latency of 5.6 ms on standard multi-core cloud compute nodes, well within the sub-10-millisecond inline SLA required for real-time payment authorization gateways (Ke et al., 2017; Ferrag et al., 2020). Furthermore, each escalated P1 alert ($S_IDS \geq 80$) emits an auditable JSON triage payload detailing triggered STIX rule IDs, Isolation Forest percentile rank, asset criticality tier (γ_i), and top deviating flow features to accelerate SOC containment (Tounsi & Rais, 2018; National Institute of Standards and Technology [NIST], 2024).

5. DISCUSSION, SOC GOVERNANCE, AND CONCLUSION

5.1 Operationalizing Hybrid Detection in FinTech SOCs Under NIST CSF 2.0

The empirical findings provide FinTech security architects and cloud engineering teams with a practical blueprint for implementing continuous anomaly detection (*DE.CM*) and adverse event analysis (*DE.AE*) under NIST Cybersecurity Framework (CSF) 2.0 (NIST, 2024) and U.S. Department of the Treasury (2024) AI cybersecurity guidelines. Rather than deploying expensive deep neural packet inspectors that introduce latency bottlenecks and opaque black-box alerts (Ferrag et al., 2020), FinTech SOCs can pair deterministic STIX/TAXII feed ingestion with lightweight Isolation Forest baselines segmented by microservice criticality tier (γ_i). Crucially, weighting

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

behavioral deviations by asset criticality ensures that subtle east–west anomalies targeting Tier 1 payment clearing or cryptographic key-management services receive immediate automated quarantine (S_IDS ≥ 80), while minor fluctuations on Tier 3 staging clusters are routed to low-friction rate limiting.

5.2 Adversarial Resilience, Governance Controls, and Limitations

Because adversaries actively attempt to poison behavioral baselines or evade static rules in financial cloud networks (U.S. Department of the Treasury, 2024), production deployment requires strict lifecycle governance:

- **Clean-Baseline Training Isolation:** Stage 2 Isolation Forest models must be retrained exclusively on cryptographically verified, incident-free rolling windows so slow-drip reconnaissance cannot gradually shift anomaly thresholds.
- **STIX Feed Hygiene & Expiration:** Stage 1 threat-intelligence rules require automated time-to-live (TTL) decay to prevent stale cloud IP addresses from causing collateral blocking on legitimate payment partners (Tounsi & Rais, 2018).
- **Human-in-the-Loop Override Governance:** Inline session quarantines on Tier 1 payment APIs should incorporate rapid SOC analyst verification workflows and automated circuit breakers to prevent self-inflicted denial of service during legitimate peak settlement events (NIST, 2024).

Two limitations merit future investigation. First, while UNSW-NB15 and CICIDS2017 capture diverse network-layer exploits (Ring et al., 2019), encrypted TLS 1.3 application-layer business-logic abuse requires combining network flow features with authenticated API gateway telemetry (such as OAuth token scope and sequence-of-endpoint call graphs). Second, longitudinal evaluation across multi-region active-active cloud failovers is needed to quantify concept drift behavior during sudden cross-region traffic rerouting.

CONCLUSION

Protecting multi-cloud FinTech infrastructure against both known threat campaigns and novel zero-day intrusions requires reconciling the speed and precision of deterministic threat intelligence with the adaptability of behavioral machine learning. By uniting STIX/TAXII rule filtering, Isolation Forest behavioral anomaly scoring, and asset-criticality-aware LightGBM meta-classification into a single sub-10-millisecond inline architecture, this study reduces false-positive SOC escalations by over 90% while detecting more than 81% of unseen zero-day intrusions. The resulting framework delivers a scalable, auditable intrusion defense aligned with NIST CSF 2.0 continuous monitoring mandates for cloud-native financial institutions.

REFERENCES

1) Ahmed, M., Mahmood, A. N., & Hu, J. (2016). A survey of network anomaly detection techniques. *Journal of Network and Computer Applications*, 60, 19–31. https://doi.org/10.1016/j.jnca.2015.11.016
2) Buczak, A. L., & Guven, E. (2016). A survey of data mining and machine learning methods for cyber security intrusion detection. *IEEE Communications Surveys & Tutorials*, 18(2), 1153–1176. https://doi.org/10.1109/COMST.2015.2494502
3) Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), 1–58. https://doi.org/10.1145/1541880.1541882
4) Ferrag, M. A., Maglaras, L., Moschoyiannis, S., & Janicke, H. (2020). Deep learning for cyber security intrusion detection: Approaches, datasets, and comparative study. *Journal of Information Security and Applications*, 50, 102419. https://doi.org/10.1016/j.jisa.2019.102419
5) Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75–105. https://doi.org/10.2307/25148625
6) Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., & Liu, T.-Y. (2017). LightGBM: A highly efficient gradient boosting decision tree. In I. Guyon et al. (Eds.), *Advances in Neural Information Processing Systems* (Vol. 30, pp. 3146–3154). Curran Associates.
7) Khraisat, A., Gondal, I., Vamplew, P., & Kamruzzaman, J. (2019). Survey of intrusion detection systems: Techniques, datasets and challenges. *Cybersecurity*, 2(1), 20. https://doi.org/10.1186/s42400-019-0038-7
8) Liu, F. T., Ting, K. M., & Zhou, Z.-H. (2008). Isolation forest. In *2008 Eighth IEEE International Conference on Data Mining* (pp. 413–422). IEEE. https://doi.org/10.1109/ICDM.2008.17
9) Moustafa, N., & Slay, J. (2015). UNSW-NB15: A comprehensive data set for network intrusion detection systems (UNSW-NB15 network data set). In *2015 Military Communications and Information Systems Conference (MilCIS)* (pp. 1–6). IEEE. https://doi.org/10.1109/MilCIS.2015.7348942
10) National Institute of Standards and Technology. (2024). *The NIST cybersecurity framework (CSF) 2.0* (NIST CSWP 29). U.S. Department of Commerce. https://doi.org/10.6028/NIST.CSWP.29

 © Author(s) | CC BY 4.0 International License.

iJETRM
INTERNATIONAL JOURNAL OF ENGINEERING TECHNOLOGY RESEARCH & MANAGEMENT (IJETRM)
https://ijetrm.com/

11) Ring, M., Wunderlich, S., Scheuring, D., Landes, D., & Hotho, A. (2019). A survey of network-based intrusion detection data sets. *Computers & Security*, 86, 147–167. https://doi.org/10.1016/j.cose.2019.06.005
12) Sharafaldin, I., Lashkari, A. H., & Ghorbani, A. A. (2018). Toward generating a new intrusion detection dataset and intrusion traffic characterization. In *Proceedings of the 4th International Conference on Information Systems Security and Privacy (ICISSP)* (pp. 108–116). SciTePress. https://doi.org/10.5220/0006639801080116
13) Sommer, R., & Paxson, V. (2010). Outside the closed world: On using machine learning for network intrusion detection. In *2010 IEEE Symposium on Security and Privacy* (pp. 305–316). IEEE. https://doi.org/10.1109/SP.2010.25
14) Tounsi, W., & Rais, H. (2018). A survey on technical threat intelligence in the age of sophisticated cyber attacks. *Computers & Security*, 72, 212–233. https://doi.org/10.1016/j.cose.2017.09.001
15) U.S. Department of the Treasury. (2024). *Managing artificial intelligence-specific cybersecurity risks in the financial services sector*. U.S. Department of the Treasury. https://home.treasury.gov/system/files/136/Managing-Artificial-Intelligence-Specific-Cybersecurity-Risks-In-The-Financial-Services-Sector.pdf

 © Author(s) | CC BY 4.0 International License.