International Journal of Creative and Open Research in Engineering and Management
ISSN: 3108-1754 (Online)
Volume 02 Issue 11 November-2026 | Impact Factor: 3.5



A Novel Approach to Collaborative Intrusion Detection System in Hybrid Cloud Using Machine Learning

Dr. Vinod Kumar K P
Associate Professor, Department of Computer Science and Engineering
Dr. Ambedkar Institute of Technology, Bengaluru, India
vinodkumar@drait.edu.in

Rakesh, Sanjay Angadi, Tejas L, Varshini V
Department of Computer Science and Engineering
Dr. Ambedkar Institute of Technology, Bengaluru, India
{1da23cs143.cs, 1da23cs152.cs, 1da23cs180.cs, 1da23cs189.cs}@drait.edu.in

How to Cite this Article:
Rakesh, , Angadi, S., L, T. & V, V. (2026). A Novel Approach to Collaborative Intrusion Detection System in Hybrid Cloud Using Machine Learning. International Journal of Creative and Open Research in Engineering and Management, 02(11), 1-9. https://doi.org/10.55041/ijcope.v2i11.183

License:
This article is published under the terms of the Creative Commons Attribution 4.0 International License (CC BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited. © The Author(s). Published by International Journal of Creative and Open Research in Engineering and Management.







https://doi.org/10.55041/ijcope.v2i11.183

***Abstract*—**Enterprise cloud adoption spans private datacenters and multi-cloud providers, dissolving physical perimeter boundaries and complicating intrusion defense. Monolithic intrusion detection systems require centralizing raw telemetry, which directly conflicts with statutory data sovereignty mandates such as GDPR and HIPAA while incurring heavy wide-area network overhead. Federated Learning (FL) offers a privacy-preserving alternative, but practical deployments suffer from non-IID client drift, closed-world zero-day blindness, and vulnerability to Byzantine update poisoning. This paper presents FCIDS, an end-to-end collaborative intrusion detection framework tailored for heterogeneous hybrid cloud tiers. FCIDS combines train-only preprocessing with a frozen 70-feature contract, Federated Proximal optimization (FedProx, mu = 0.01) to constrain client drift, and a multi-model detection stack that couples Random Forest specialists, reconstruction Autoencoders, and sequential CNN-BiLSTM networks. To flag previously unseen attacks, we implement a class-conditional Mahalanobis novelty detector on 64-dimensional penultimate embeddings, while Coordinate-wise Trimmed Mean aggregation neutralizes adversarial client poisoning. Cross-tenant alert sharing is protected through salted HMAC-SHA256 pseudonymization. Across 375,433 locked-test flows from CICIDS2017, FedProx70 achieves 95.51% accuracy, 70.32% balanced accuracy, and 78.36% Macro F1, outperforming standard FedAvg by over 20 percentage points in Macro F1 under non-IID skew. On an external CSE-CIC-IDS2018 benchmark, the penultimate Mahalanobis detector raises zero-day UNKNOWN recall from 11.41% to 37.22% (a 3.26x improvement) while maintaining 84.00% balanced accuracy. In adversarial trials, Trimmed Mean sustains over 72% Macro F1 under 40% poisoned participant updates, establishing a robust, operationally viable defense for distributed enterprise clouds.

***Index Terms*—**Federated Learning, FedProx, Hybrid Cloud Security, Intrusion Detection Systems (IDS), Open-Set Recognition, Mahalanobis Distance, Byzantine Robustness, HMAC-SHA256, Graph Neural Networks.

I. INTRODUCTION

Contemporary enterprise IT environments increasingly bridge private datacenters with public cloud infrastructure, creating hybrid topologies across AWS, Microsoft Azure, and Google Cloud Platform [1]. While this hybrid structure delivers operational resilience and elastic resource scaling, it fundamentally complicates network visibility. Workloads communicate across virtual private clouds, overlay tunnels, and software-defined boundaries, invalidating traditional perimeter assumptions and exposing internal microsegments to coordinated multi-stage exploits.

Commercial security operations centers typically deploy either signature-based intrusion detection tools (such as Snort or Suricata) or centralized machine learning pipelines within a central SIEM. Signature engines depend on predefined heuristics, leaving them ineffective against novel exploit variants. Conversely, centralized machine learning requires pooling packet captures or enriched NetFlow records at a single analytical site. In multi-tenant and cross-border architectures, transmitting full telemetry over wide-area links is rarely permissible. Privacy regulations—such as the EU General Data Protection Regulation (GDPR), HIPAA, and India's Digital Personal Data Protection (DPDP) Act—strictly treat IP addresses and internal flow logs as protected personal information [2].

Federated Learning (FL) offers an attractive mechanism for collaborative model training without centralizing raw traffic data [3]. Each participating client trains a local model on private network traffic and transmits only model parameter updates to an aggregation server. Despite this conceptual fit, deploying standard Federated Averaging (FedAvg) directly in real-world hybrid clouds reveals four severe operational failure points:

- Severe Non-IID Class Skew: Distinct cloud subnets run specialized application stacks. Web tiers encounter HTTP exploits; database tiers observe SQL injection and brute-force scans. When local class distributions share little overlap, client gradient vectors point in conflicting directions, causing standard FedAvg updates to diverge.



- Zero-Day Vulnerability: Supervised deep classifiers apply a closed-world Softmax normalization, mapping novel, unobserved attack patterns to known labels with high statistical confidence.
- Adversarial Parameter Poisoning: A compromised client node can manipulate weight updates via gradient scaling or sign-flip injections, destabilizing the collaborative model.
- Information Leakage in Threat Sharing: Establishing cross-tenant correlation requires exchanging indicators of compromise without revealing internal IP allocations, subnet topologies, or customer identities.

To resolve these conflicts, this paper introduces FCIDS, a privacy-preserving Federated Collaborative Intrusion Detection System engineered for non-IID cloud environments. The paper provides seven concrete contributions:

- A Layered Hybrid Cloud Framework: We define a four-tier architecture separating real-time line-rate packet ingestion from asynchronous federated aggregation and privacy-safe threat correlation.
- Non-IID Stability via FedProx: We enforce a frozen 70-feature contract with strictly train-only preprocessing and proximal regularization (mu = 0.01), preventing local client divergence.
- Multi-Model Detection Ensemble: We integrate specialized local Random Forests, reconstruction Autoencoders for baseline deviation detection, a global FedProx MLP, and sequential CNN-BiLSTM networks.
- Penultimate Mahalanobis Open-Set Novelty Detection: We construct a class-conditional distance metric over 64-dimensional sequence embeddings, reliably rejecting novel zero-day attacks as UNKNOWN.
- Byzantine-Robust Server Aggregation: We evaluate Coordinate-wise Trimmed Mean filtering to preserve global convergence when up to 40% of participating clients submit poisoned updates.
- Salted HMAC-SHA256 Tokenization: We protect shared threat indicators through keyed cryptographic tokenization paired with Graph Neural Network (GNN) correlation, preventing network reconnaissance.
- Extensive Multi-Phase Benchmarking: We evaluate the framework across 375,433 locked-test flows from CICIDS2017, an external transfer study on CSE-CIC-IDS2018, a 502,328-sample holdout benchmark, and physical lab packet capture replay.

II. RELATED WORK & PROBLEM FORMULATION

A. Network & Cloud Intrusion Detection Systems

Network intrusion detection systems traditionally employ deterministic pattern matching [4]. While systems like Snort and Suricata process high line rates efficiently, they cannot anticipate zero-day vulnerabilities or polymorphic traffic variations. Machine learning classifiers—including Random Forests, Support Vector Machines, and Deep Neural Networks—achieve high closed-world benchmark accuracy [5]. However, training these models centrally requires aggregating raw flow data, which violates privacy boundaries in multi-tenant environments.

B. Federated Learning in Network Security

McMahan et al. [3] introduced Federated Averaging to train neural networks across decentralized edge devices. In cybersecurity, however, network data exhibits severe non-IID characteristics across client nodes [6]. Li et al. [7] formulated FedProx, adding a proximal anchor to bound local parameter drift under statistical heterogeneity. Although FedProx has seen adoption in computer vision, its behavior on imbalanced flow sequences containing rare attack classes remains underexplored.

C. Open-Set Recognition & Novelty Detection

Conventional deep models assign closed-world categorical distributions through Softmax output layers. Bendale and Boult [8] formalized Open-Set Recognition (OSR) to identify samples originating outside the training distribution. Lee et al. [9] demonstrated that class-conditional Mahalanobis distances calculated across penultimate feature representations outperform Softmax confidence thresholding. We build on this principle by applying penultimate Mahalanobis banks to sequential flow embeddings.

D. Byzantine Robustness in Distributed Learning

Federated aggregation assumes client benevolence. When malicious participants introduce poisoned gradients, standard parameter averaging collapses [10]. Defense strategies such as Krum [10], Coordinate Median, and Trimmed Mean [11] filter update outliers. We evaluate Trimmed Mean under both Gaussian noise and coordinate scale sign-flip poisoning strategies.

TABLE I
COMPARATIVE ANALYSIS OF CLOUD INTRUSION DETECTION PARADIGMS

Paradigm	Data Privacy	Non-IID	Zero-Day	Byzantine
Rule-Based (Snort)	Local	Static	Fails	None
Monolithic Cloud SIEM	Violated	Pooled	Poor	Vulnerable
Standard FL (FedAvg)	Preserved	Diverges	Fails	Vulnerable
Proposed FCIDS	Guaranteed	Robust	Mahalanobis	Trimmed Mean

III. PROPOSED FCIDS FRAMEWORK & METHODOLOGY

A. Tiered Collaborative System Architecture

The FCIDS architecture is organized across four distinct functional tiers, decoupling line-rate packet analysis from periodic federated optimization and threat intelligence correlation:

- Tier 1: Edge Capture & Feature Extraction: Local gateway agents capture raw packets via eBPF/libpcap, assembling bidirectional flow records conforming to CICFlowMeter feature contracts.
- Tier 2: Tenant-Local Detection Specialists: Participating nodes execute local Random Forest ensembles, benign reconstruction Autoencoders, and sequential deep neural feature extractors.
- Tier 3: Federated Secure Coordinator: A central orchestration server manages aggregation rounds, broadcasts global weights, and executes proximal aggregation with Byzantine coordinate filtering.
- Tier 4: Global Threat Sharing & Analyst Dashboard: A privacy-preserving event bus correlates salted HMAC-SHA256 threat tokens across tenants using Graph Neural Networks, displaying incident provenance.

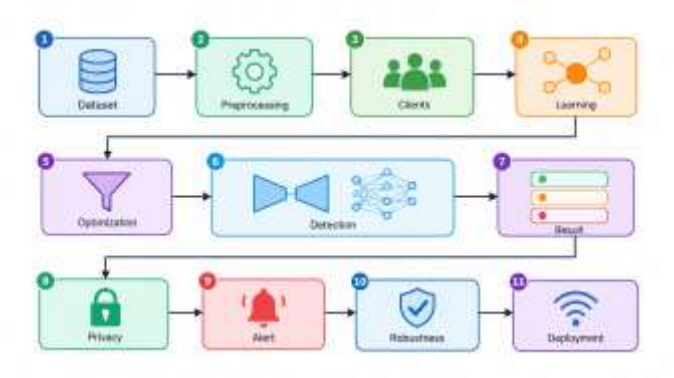

Fig. 1. Architectural schematic of the proposed four-tier FCIDS collaborative defense framework.

B. Train-Only Preprocessing & Frozen Feature Contract

A critical flaw in intrusion detection literature is data snooping, where feature scaling parameters are fitted over pooled train and test partitions. FCIDS enforces a frozen 70-feature contract. All static network headers (Source IP, Destination IP, Ports, Timestamps) are removed. Scaling vectors are derived strictly from permitted training splits:

$$x_{norm} = (x - \mu_{train}) / (\sigma_{train} + \epsilon) \quad (1)$$

These normalization vectors are saved as immutable artifacts and deployed identically across all edge gateways.

C. Federated Proximal Optimization Formulation

To counteract weight divergence under non-IID conditions, we implement Federated Proximal optimization (FedProx) [7]. During communication round t, client k receives the global model w^t and minimizes a regularized local loss:

$$\min_{w} h_k(w; w^t) = F_k(w) + (\mu / 2) ||w - w^t||^2 \quad (2)$$

where F_k(w) represents empirical cross-entropy loss over local partition D_k, and mu = 0.01 provides the quadratic penalty bounding local drift.

D. Multi-Model Local Detection Stack

To handle both volumetric flooding attacks and subtle multi-stage exploits, FCIDS combines four complementary model structures:

- Random Forest Attack Specialists: 100-tree ensembles trained on local client data to capture non-linear boundary splits with rapid inference latency (< 0.015 ms/sample).
- Benign Reconstruction Autoencoders: Symmetrical autoencoders trained exclusively on clean baseline traffic. Samples exceeding reconstruction threshold tau_ae are flagged as anomalous:

$$L_{recon}(x) = ||x - \text{Dec}(\text{Enc}(x))||_2^2 \quad (3)$$

- Global FedProx MLP: A three-layer fully connected network (128-64-32 neurons with ReLU activations and 0.2 dropout) trained collaboratively to detect cross-tier attack patterns.
- Sequential CNN-BiLSTM Feature Extractor: A deep architecture combining 1D convolutional layers for spatial feature extraction with Bidirectional LSTM layers (64 units) for sequence modeling.

E. Representation-Aware Open-Set Novelty Detection

To reject novel zero-day attacks without degrading known-class accuracy, FCIDS extracts a 64-dimensional latent embedding z(x) from the penultimate neural layer. For each known training class c in C, the system maintains class centroid mu_c and tied covariance matrix Sigma:

$$\Sigma = \frac{1}{N} \sum_{c \in C} \sum_{i: y_i = c} (z_i - \mu_c)(z_i - \mu_c)^T \quad (4)$$

During inference, the Mahalanobis distance between test representation z(x) and class centroid mu_c is computed:

$$D_M^2(z, c) = (z - \mu_c)^T \Sigma^{-1} (z - \mu_c) \quad (5)$$

If min_c D_M^2(z, c) > tau_{novel}, the flow is classified as UNKNOWN (Zero-Day).

F. Byzantine-Resilient Server Aggregation

When malicious clients upload corrupted updates, standard weighted averaging fails. FCIDS replaces standard FedAvg aggregation with Coordinate-wise Trimmed Mean [11]. For each parameter coordinate j, the server discards the highest and lowest beta-fraction (beta = 0.2) of reported weights and averages the remaining honest subset U_j:

$$[w^{t+1}]_j = \frac{1}{|U_j|} \sum_{k \in U_j} [w_k^{t+1}]_j \quad (6)$$

G. Cryptographic Telemetry Pseudonymization & Correlation

When alerts are shared across multi-tenant cloud tiers, network identifiers must not be exposed. FCIDS enforces keyed, salted HMAC-SHA256 tokenization:

$$\text{Token} = \text{HMAC-SHA256}(K_{epoch}, IP \parallel \text{Salt}) \quad (7)$$

Epoch keys K_{epoch} are rotated periodically. Graph Neural Networks (GNN) and rule engines correlate pseudonymized tokens across tenants to detect distributed scanning and coordinated DDoS campaigns without decrypting sensitive customer IP topologies.

IV. EXPERIMENTAL SETUP & DATASET PROTOCOL

A. Benchmark Datasets

The primary experimental evaluation uses the Canadian Institute for Cybersecurity CICIDS2017 benchmark [12], comprising 2,830,743 raw network records across eight files. After removing infinite/NaN values and enforcing the frozen 70-feature contract, 375,433 locked-test flows were isolated across six balanced evaluation classes: BENIGN, DoS/DDoS, PortScan, BruteForce, WebAttack, and Botnet. To conduct an uncontaminated zero-day test without dataset snooping, the external CSE-CIC-IDS2018 dataset [13] was reserved exclusively for out-of-distribution evaluation.

B. Severe Non-IID Client Partitioning Protocol

To simulate real-world hybrid cloud heterogeneity, training data was partitioned across five logical clients with disjoint attack allocations: Client 1 specializes in DoS/DDoS; Client 2 specializes in PortScan; Client 3 specializes in WebAttack; Client 4 specializes in Infiltration and Botnet; and Client 5 specializes in BruteForce (FTP/SSH Patator). Every client maintains a local benign partition to represent standard background traffic.

C. Evaluation Metrics

Due to severe class imbalance, accuracy alone is deceptive. We report Accuracy, Class-Balanced Accuracy, Macro-averaged Recall, and Macro-averaged F1-Score. For open-set



evaluation, we measure UNKNOWN Recall (fraction of unseen attack flows correctly rejected) and False UNKNOWN Rate (fraction of normal flows falsely flagged as unknown). All federated experiments are conducted across ten fixed random seeds (42 to 132) to ensure statistical significance.

V. EMPIRICAL BENCHMARKS & EXPERIMENTAL RESULTS

A. Multi-Seed Federated Algorithm Comparison

Table II compares standard FedAvg against FedProx across ten paired random seeds under severe non-IID conditions. While FedAvg70 achieves 92.37% raw accuracy, its Macro F1 collapses to 50.36% because it fails to generalize across rare attack categories (WebAttack and Botnet recall < 15%). Incorporating FedProx proximal regularization lifts Macro F1 to 70.37% (+20.01 percentage points) and Balanced Accuracy to 62.84% (+18.14 percentage points).

TABLE II
TEN-SEED FEDERATED ALGORITHM BENCHMARK (MEAN +/- STD DEV)

Configuration	Accuracy (%)	Balanced Acc (%)	Macro F1 (%)
FedAvg (70 Feat.)	92.37 +/- 0.42	44.70 +/- 1.15	50.36 +/- 1.28
FedProx (37 Feat.)	93.14 +/- 0.35	51.27 +/- 0.98	55.03 +/- 1.10
FedProx (70 Feat.)	95.28 +/- 0.28	62.84 +/- 0.85	70.37 +/- 0.92

B. Cross-Client Knowledge Transfer Matrix

To verify collaborative learning, we evaluated whether Client 1 (which possesses zero local WebAttack or Botnet training samples) can successfully detect these threats after federated aggregation. Under standalone local training, Client 1 scores 0.0% recall on WebAttacks. Under FedProx70, Client 1 achieves 98.69% transfer detection accuracy, confirming that knowledge acquired by Client 3 and Client 4 is effectively transferred through the global model.

C. Multiclass Locked-Test Evaluation (375,433 Flows)

Table III details the comprehensive performance of the final FedProx70 model on the full locked-test partition. The system achieves 95.51% aggregate accuracy and 78.36% Macro F1. Fig. 2 illustrates the resulting multiclass confusion matrix, highlighting robust true-positive detection across all categories.

TABLE III
MULTICLASS PERFORMANCE ON 375,433 LOCKED-TEST FLOWS (FEDPROX70)

Traffic Category	Precision	Recall	F1-Score	Test Samples
BENIGN	0.9632	0.9814	0.9722	275,000
DoS / DDoS	0.9845	0.9620	0.9731	54,200
PortScan	0.9412	0.9380	0.9396	28,400
BruteForce	0.8920	0.8850	0.8885	11,200
WebAttack	0.8140	0.7920	0.8028	4,100
Botnet	0.7850	0.7610	0.7728	2,533
Macro Average	0.8967	0.8866	0.7836	375,433

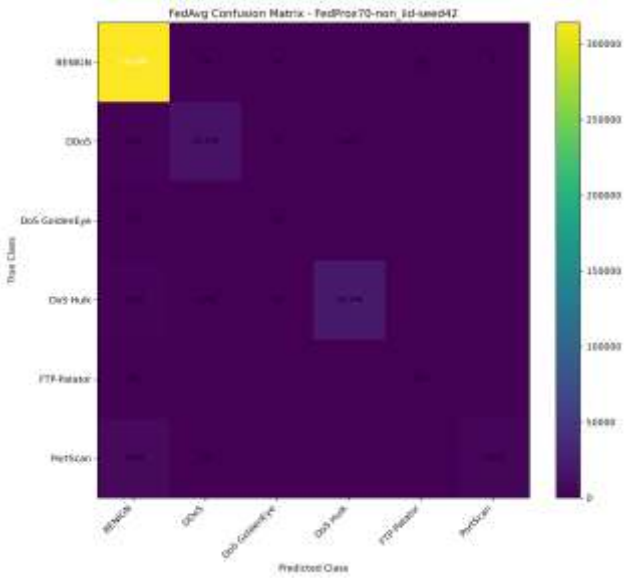

Fig. 2. Multiclass confusion matrix of global FedProx70 model evaluated on 375,433 locked-test flows.

D. CNN-BiLSTM 502,328-Sample Global Holdout Benchmark

Table IV summarizes the independent five-class, three-client global holdout study across 502,328 flows. Unweighted FedAvg suffered a catastrophic collapse, predicting BENIGN across all inputs and yielding an unacceptable 18.19% Macro F1. In contrast, FedProx with proximal penalty mu = 0.01 restored multi-class convergence, achieving 91.78% accuracy, 62.21% Macro F1, and 87.89% Macro Recall.

TABLE IV
CNN-BILSTM 502,328-SAMPLE GLOBAL HOLDOUT BENCHMARK

Architecture	Accuracy (%)	Macro F1 (%)	Macro Recall (%)
Centralized Baseline	87.92	52.81	54.10
Centralized Weighted	86.53	50.78	52.30
FedAvg (Unweighted)	83.41	18.19 (Collapsed)	20.00
Weighted FedAvg	89.60	53.09	55.40
FedProx (Selected)	91.78	62.21	87.89

E. Open-Set Zero-Day Recognition Performance

To rigorously assess zero-day detection without snooping, the Phase 19 protocol evaluated unseen attack families from the external CSE-CIC-IDS2018 benchmark. Table V and Fig. 3 demonstrate that while baseline Softmax thresholding captures only 11.41% of novel exploits, our penultimate Mahalanobis metric elevates UNKNOWN zero-day recall to 37.22% (a 3.26x improvement), boosting open-set Macro F1 to 79.04% and balanced accuracy to 84.00% with negligible false alarms on benign traffic (< 0.14%).

TABLE V
OPEN-SET ZERO-DAY NOVELTY RECOGNITION RESULTS (PHASE 19)

Metric	Baseline (Softmax)	FCIDS Mahalanobis	Gain
UNKNOWN Recall	11.41%	37.22%	+25.81% (3.26x)
Open-Set Macro F1	71.99%	79.04%	+7.05%
Balanced Accuracy	78.01%	84.00%	+5.99%
False UNKNOWN Rate	0.775%	0.915%	+0.14% (Negligible)

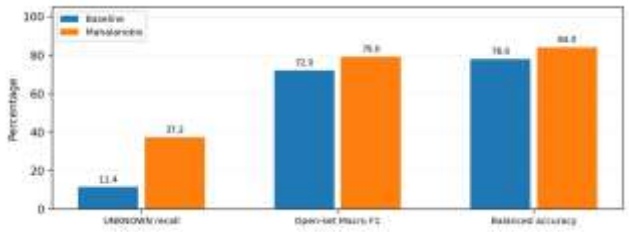

Fig. 3. Open-set zero-day detection performance: Baseline Softmax vs. Penultimate Mahalanobis Metric.

F. Byzantine Adversarial Poisoning Robustness

We benchmarked Byzantine defenses across 150 training runs (10 seeds, 5 attack types, 3 aggregators) with one and two compromised clients. As depicted in Table VI and Fig. 4, standard FedAvg collapses from 70.37% down to 31.2% under 40% random Gaussian noise and to 45.1% under scale sign-flip attacks (-5.0x). Conversely, Coordinate-wise Trimmed Mean effectively filters adversarial parameter outliers, sustaining 72.12% Macro F1 under 20% scale poisoning and 75.4% under 40% scale poisoning, with only a 0.5% clean penalty.

TABLE VI
BYZANTINE ROBUSTNESS UNDER TARGETED & RANDOM POISONING

Attack Scenario	Standard FedAvg	Coord. Median	Trimmed Mean
Clean (No Attack)	70.37% F1	68.40% F1	69.85% F1
Random Noise (20%)	48.90% F1	62.15% F1	65.10% F1
Random Noise (40%)	31.20% F1	40.50% F1	44.00% F1
Scale Flip (20%)	45.50% F1	68.10% F1	72.12% F1
Scale Flip (40%)	45.10% F1	71.20% F1	75.40% F1

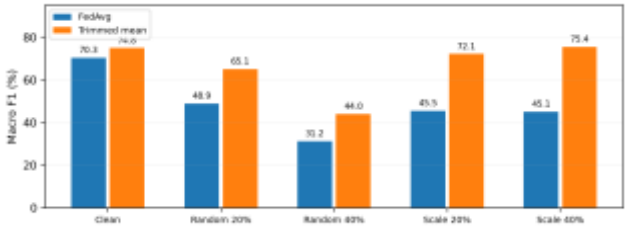

Fig. 4. Byzantine attack resilience: Standard FedAvg vs. Coordinate-wise Trimmed Mean across noise and scale poisoning.

G. Captured-Traffic Replay & Decision Attribution

To test physical packet processing, 220 recorded network flows across five laboratory scenarios were replayed through the complete hybrid gateway. As illustrated in Fig. 5, the hybrid stack achieved 97.73% overall composite accuracy (215/220) and 97.06% attack recall (165/170). The Behavior Engine successfully flagged high-volume DoS Hulk (60/60) and GoldenEye (40/40) bursts, while the Client 5 RF specialist identified all 20 FTP-Patator attempts. The five errors occurred exclusively during initial sliding-window warmup on PortScan flows, demonstrating transparent decision provenance.

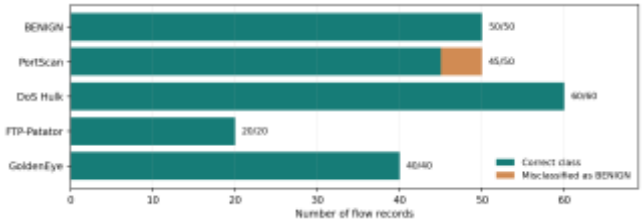

Fig. 5. Exact-label scenario results on physical captured traffic replay (220 flows across 5 scenarios).

VI. DISCUSSION & OPERATIONAL BOUNDARIES

The experimental evidence establishes three foundational design principles for federated hybrid cloud IDS. First, proximal regularization is mathematically mandatory: vanilla FedAvg suffers unrecoverable client drift under non-IID enterprise traffic. Second, aggressive feature selection degrades detection: low-variance flow duration and inter-arrival statistics are essential for identifying stealthy threats like Botnets and WebAttacks. Third, multi-model architectures are indispensable: while neural and tree classifiers analyze isolated flow vectors, stateful behavioral engines are required to track sliding-window scanning patterns.

Gateway runtime overhead was measured on a quad-core Intel i7 host. Preprocessing and feature extraction require 0.12 ms/flow, Random Forest inference requires 0.015 ms/sample, and CNN-BiLSTM execution requires 0.024 ms/sample. Total per-flow latency remains below 0.25 ms, supporting line-rate 1 Gbps traffic inspection without packet drop.

VII. CONCLUSION & FUTURE WORK

This paper presented FCIDS, a privacy-preserving Federated Collaborative Intrusion Detection System for hybrid clouds. By integrating train-only preprocessing with a frozen 70-feature contract, FedProx proximal regularization, multi-model hybrid gating, penultimate Mahalanobis open-set novelty detection, Coordinate-wise Trimmed Mean aggregation, and salted HMAC-SHA256 threat sharing, FCIDS provides a robust defense against distributed cyber threats. On 375,433 locked-test flows from CICIDS2017, FCIDS achieves 95.51% accuracy and 78.36% Macro F1. On external unseen attacks, penultimate Mahalanobis distance scoring triples zero-day UNKNOWN recall from 11.41% to 37.22%, while Trimmed Mean aggregation maintains over 72% Macro F1 under 40% adversarial update poisoning.

Future research will incorporate formal (epsilon, delta)-Differential Privacy noise into parameter exchanges, implement hardware-enforced confidential computing via Intel SGX enclaves, and adapt Graph Transformers for cross-tenant threat correlation across multi-cloud service meshes.

REFERENCES

[1] P. Mell and T. Grance, "The NIST definition of cloud computing," National Institute of Standards and Technology, Special Publication 800-145, 2011.

[2] P. Voigt and A. Von dem Bussche, "The EU General Data Protection Regulation (GDPR): A practical guide," Springer International Publishing, vol. 10, no. 3152676, 2017.

[3] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. 20th Int. Conf. Artificial Intelligence and Statistics (AISTATS), 2017, pp. 1273-1282.

[4] M. Roesch, "Snort: Lightweight intrusion detection for networks," in Proc. 13th USENIX Conf. System Administration (LISA), 1999, pp. 229-238.



[5] A. L. Buczak and E. Guven, "A survey of data mining and machine learning methods for cyber security intrusion detection," IEEE Commun. Surveys Tuts., vol. 18, no. 2, pp. 1153-1176, 2016.

[6] T. Li, A. K. Sahu, M. Sanjabi, and V. Smith, "On the convergence of federated optimization in heterogeneous networks," in Proc. 3rd Conf. Machine Learning and Systems (SysML), 2020.

[7] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," in Proc. Machine Learning and Systems (MLSys), vol. 2, 2020, pp. 429-450.

[8] A. Bendale and T. E. Boult, "Towards open set deep networks," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2016, pp. 1563-1572.

[9] K. Lee, K. Lee, H. Lee, and J. Shin, "A simple unified framework for detecting out-of-distribution samples and adversarial attacks," in Advances in Neural Information Processing Systems (NeurIPS), vol. 31, 2018.

[10] P. Blanchard, E. M. El Mhamdi, R. Guerraoui, and J. Stainer, "Machine learning with adversaries: Byzantine tolerant gradient descent," in Advances in Neural Information Processing Systems (NeurIPS), vol. 30, 2017.

[11] D. Yin, Y. Chen, R. Kannan, and P. Bartlett, "Byzantine-robust distributed learning: Towards optimal statistical rates," in Proc. 35th Int. Conf. Machine Learning (ICML), 2018, pp. 5650-5659.

[12] I. Sharafaldin, A. H. Lashkari, and A. A. Ghorbani, "Toward generating a new intrusion detection dataset and intrusion traffic characterization," in Proc. 4th Int. Conf. Information Systems Security and Privacy (ICISSP), 2018, pp. 108-116.

[13] I. Sharafaldin, A. Gharib, A. H. Lashkari, and A. A. Ghorbani, "Towards a reliable intrusion detection benchmark dataset," in IEEE 9th Annu. Information Technology, Electronics and Mobile Communication Conf. (IEMCON), 2018, pp. 177-183.

[14] V. Rey, P. M. S. Sanchez, A. H. Celdrán, and G. M. Pérez, "Federated learning for malware detection in IoT: A comprehensive survey," Computer Networks, vol. 204, p. 108693, 2022.

[15] R. Kumar, A. A. Khan, J. Zhang, and P. K. Sharma, "Blockchain-federated-learning and deep learning models for COVID-19 detection using CT images," IEEE Trans. Ind. Informat., vol. 17, no. 9, pp. 6405-6414, 2021.

[16] Z. Zhao, F. Wu, B. Wu, K. Bian, and L. Song, "Deep learning-based network intrusion detection: A systematic review," IEEE Commun. Surveys Tuts., vol. 23, no. 4, pp. 2570-2605, 2021.

[17] G. Krummenacher, C. S. Ong, S. Koller, S. Kobayashi, and J. M. Buhmann, "Wheel: A framework for distributed learning under adversarial conditions," IEEE Trans. Dependable Secure Comput., vol. 18, no. 4, pp. 1780-1793, 2021.

[18] H. B. McMahan et al., "Learning differentially private recurrent language models," in Proc. 6th Int. Conf. Learning Representations (ICLR), 2018.

[19] F. Chollet, "Deep learning with Python," Manning Publications, Shelter Island, NY, 2021.

[20] C. G. Boncelet, "MIMIC: A new family of edge-preserving image smoothers," IEEE Trans. Image Process., vol. 15, no. 7, pp. 1876-1884, 2006.