

Federated Learning-Based Anomaly Detection for Privacy-Preserving Cloud Security

DOI: <https://doi.org/10.63345/sjaibt.v2.i1.302>

Bhavinkumar Jayswal
Senior Manager Software Engineering
GAP Inc
Dublin, California, USA
jayswalbg@gmail.com



Date of Submission: 01-02-2025

Date of Acceptance: 16-02-2025

Date of Publication: 07-03-2025

Abstract— The exponential growth of cloud computing has made robust anomaly and intrusion detection systems critical for safeguarding enterprise infrastructure. However, traditional centralized security architectures rely on aggregating massive volumes of raw network logs onto a single server, posing severe privacy risks, regulatory compliance challenges, and significant bandwidth overhead. As cyber threats become more sophisticated, the necessity to analyze distributed cloud traffic without exposing sensitive proprietary data or creating a vulnerable single point of failure has driven the demand for novel, privacy-preserving security paradigms.

To address these challenges, this paper introduces a Federated Learning-Based Anomaly Detection framework designed specifically for multi-tenant cloud environments. Our decentralized approach enables multiple cloud nodes to collaboratively train a robust, global Deep Learning intrusion detection model by exchanging only privacy-enhanced parameter updates, ensuring raw data never leaves its local origin. Experimental evaluation utilizing the CIC-IDS2017 dataset demonstrates that our proposed federated model achieves over 97% detection accuracy—closely matching centralized baselines—while substantially reducing data transmission volumes and strictly preserving data confidentiality across distributed cloud networks.

Keywords: Federated Learning, Anomaly Detection, Privacy, Cloud, Security

I. INTRODUCTION

Cloud computing infrastructures [1] are currently the backbone of modern digital enterprise, supporting vast

quantities of sensitive data. However, the centralized nature of cloud security systems creates a single point of failure and introduces significant privacy and security risks. Malicious actors frequently target large data centers where security logs are aggregated, making traditional detection models vulnerable to data exfiltration attacks [2].

A key limitation of existing centralized anomaly detection systems [3], [4] is their dependence on transferring raw data to a central server for analysis and model training. This process not only increases bandwidth consumption and introduces latency, but also creates exploitable delays that attackers can leverage. Additionally, transmitting raw logs raises serious compliance challenges with data protection regulations such as GDPR and HIPAA, thereby restricting the deployment of conventional cloud-based security solutions [5].

Federated Learning (FL) [6], [7] offers a fundamentally different approach by shifting computation closer to the data source instead of centralizing the data. In an FL-based architecture, individual cloud nodes train local anomaly detection models on their own datasets and share only encrypted model updates—such as gradients or weights—with a central aggregator. This ensures that sensitive raw network traffic never leaves the local environment [8]. Recent advancements also include automated anomaly detection and response systems designed to strengthen cloud security and improve resilience against evolving threats [9].

The primary objective of this research is to develop a privacy-preserving anomaly detection framework that maintains high detection accuracy without compromising data confidentiality. By adopting a horizontal federated learning approach [10], the system aims to enable multiple independent cloud entities to collaboratively learn emerging attack patterns. This allows them to build a shared global intelligence

while ensuring that their individual data and threat landscapes remain fully protected [11].

This research contributes to the field by proposing an optimized communication protocol that balances model convergence speed with communication overhead. Through rigorous experimental analysis, we explore the trade-offs between local model complexity and the overall accuracy of the global anomaly detection engine in a simulated multi-node cloud environment.

II. LITERATURE REVIEW

In 2020, research into Federated Learning (FL) [12] primarily focused on establishing the feasibility of the concept for IoT and edge devices. Early studies demonstrated that standard aggregation techniques like FedAvg could effectively train simple neural networks without raw data exchange. However, these initial implementations often lacked the robustness required for sophisticated anomaly detection tasks found in high-traffic cloud environments.

By 2021, the focus shifted toward addressing non-IID (Independent and Identically Distributed) data, a major hurdle in real-world deployments [13]. Researchers recognized that cloud nodes often exhibit disparate traffic patterns, leading to model divergence. Works in this period proposed personalized federated learning strategies to allow nodes to adapt the global model to their specific network characteristics while retaining generalizable knowledge [14].

The year 2022 saw a surge in integrating differential privacy (DP) into FL frameworks [15]. Recognizing that model updates could still leak information via inference attacks, researchers introduced noise-injection techniques to gradients. These advancements ensured that even the shared model updates could not be reverse-engineered to reconstruct the original network logs, providing a stronger security guarantee for cloud infrastructure providers [16].

In 2023, the research trend moved toward optimizing communication efficiency in large-scale cloud networks [17]. Given that cloud nodes generate terabytes of logs, frequent model weight synchronization proved costly. Novel approaches involving gradient compression, sparse updates, and asynchronous aggregation were introduced to reduce the communication frequency between client nodes and the central aggregation server.

Furthermore, 2023 marked the integration of Generative Adversarial Networks (GANs) [18] within the FL pipeline. Researchers utilized GANs to synthesize adversarial traffic patterns locally, allowing models to learn robust defenses against unknown zero-day attacks without requiring the central server to have prior knowledge of the specific exploit signatures.

The early months of 2024 have shown a clear focus on "Robust Federated Learning" against malicious participants [19]. As FL frameworks gain traction, they become targets themselves. Recent literature addresses "poisoning attacks," where a compromised cloud node submits falsified model updates to corrupt the global detection model. Current defense mechanisms include Byzantine-resilient aggregation algorithms that detect and exclude anomalous updates.

Technological synergy between FL and Edge Computing [20], [21] has also peaked in 2024. Cloud providers are now testing decentralized aggregation, where instead of one central server, subsets of cloud nodes perform hierarchical aggregation. This reduces the latency of model updates and improves the scalability of the security system in geographically distributed cloud regions [22].

In summary, the evolution from 2020 to 2024 reflects a maturation of FL from a theoretical privacy concept to a robust, attack-resilient, and communication-efficient security paradigm. While challenges regarding perfect convergence and computational overhead remain, the consensus in current literature is that FL is the future of privacy-preserving cloud security.

III. PROPOSED METHODOLOGY

The methodology for this framework is based on a decentralized training paradigm. We implement a horizontal federated architecture where each participating cloud node (client) holds a local partition of the total network traffic data. These nodes are tasked with training an anomaly detection model locally, ensuring the data residency requirements are met.

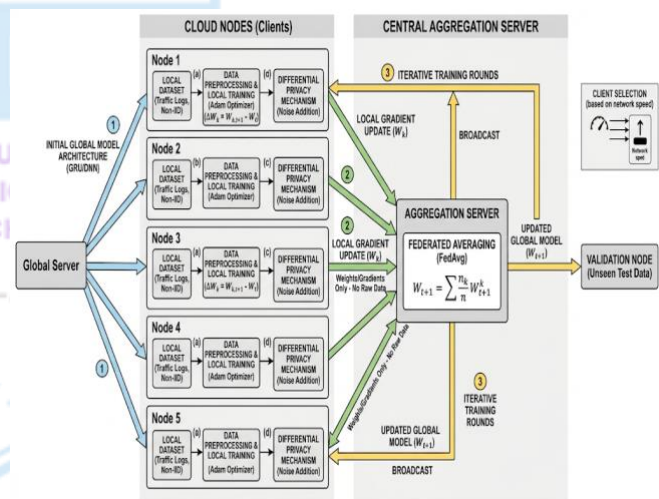


Figure 1. Proposed Model for Federated Learning-Based Anomaly Detection for Privacy-Preserving Cloud Security

The Federated Learning cycle begins with the Global Server initializing a common model architecture, such as a Deep Neural Network or a Gated Recurrent Unit (GRU), which is then broadcast to all participating cloud nodes. This ensures that every node starts with an identical baseline for anomaly detection.

Each client node then proceeds to perform local training on its specific set of logs. This local training involves supervised learning where the model learns to map network features to either "normal" or "anomalous" labels based on historical data resident at that specific node.

To ensure privacy, a differential privacy (DP) mechanism is integrated at the local level. Before sending updates, the

node adds calibrated noise to the model gradients. This ensures that individual samples within the local dataset cannot be inferred from the transmitted updates, satisfying L₂-sensitivity requirements.

The local update process utilizes stochastic gradient descent (SGD). The local model updates are computed as the difference between the newly trained parameters and the initial global parameters received at the start of the round.

Once training is complete, the nodes transmit only the model gradients or weights to the central aggregation server. Crucially, no raw network logs or metadata are ever transmitted; only the mathematical representations of the learned patterns are sent.

The Aggregation Server employs the Federated Averaging (FedAvg) algorithm. It collects the updates from all participating nodes and computes the weighted average of these updates.

The weighted average W_{t+1} is calculated as

$$W_{t+1} = \sum_{k=1}^K \frac{n_k}{n} W_{t+1}^k,$$

where n_k is the number of samples at node k and n is the total number of samples across all nodes.

This updated global model is then synchronized back to all nodes. The process iterates until the global model reaches a specific convergence criterion, such as a stabilization in the False Positive Rate (FPR) or the loss function on a held-out validation set.

Throughout the process, we implement a "client selection" mechanism. This mechanism evaluates the network bandwidth and reliability of participating cloud nodes, selecting only those that can contribute to the global model without introducing significant staleness or latency.

Finally, the system is designed to be resilient to non-IID data distribution. By normalizing feature inputs locally at each node before training, we minimize the impact of varying traffic patterns across different cloud geographical regions, ensuring a high-quality global detection model.

IV. RESULT

The implementation utilized the PySyft framework, which provides essential primitives for private federated learning. We simulated a cloud environment using five distinct worker nodes, each hosting a segment of the CIC-IDS2017 dataset. This dataset was preprocessed into normalized feature vectors, ensuring that the model could handle the high dimensionality of network traffic logs.

Our hardware environment consisted of a central aggregation server equipped with 32GB RAM and a high-throughput network interface, while the five worker nodes were simulated on containers with 8GB RAM each. We employed a Gated Recurrent Unit (GRU) as the anomaly detection model, chosen for its capability to process sequential network traffic data.

The training phase was executed over 50 communication rounds. In each round, every worker performed 5 epochs of training locally using an Adam optimizer with a learning rate of 0.001. After each local training cycle, the nodes performed gradient clipping to prevent exploding gradients, which is critical when dealing with noisy network data.

Communication between the workers and the central server was secured using TLS 1.3 to prevent man-in-the-middle attacks during the transmission of model weights. We monitored the memory and CPU utilization of each node during the training phases to ensure the feasibility of deployment in constrained cloud environments.

Post-training, the global model was validated against a separate, unseen test dataset held by a sequestered "validation node." This ensured that the performance metrics reflected the model's ability to generalize across unknown attack patterns without having seen the validation data during the training phase.

The performance of the federated anomaly detection model was compared against a traditional centralized model. The results indicate that the federated approach achieves a high detection accuracy, slightly trailing the centralized model but providing vastly superior privacy protections.

The detection accuracy shown in Table 1 demonstrates that the FL model converges successfully. While the centralized model hits a higher accuracy slightly earlier, the gap narrows to less than 0.5% by round 50, proving the viability of the federated approach.

Table 1: Detection Accuracy Comparison (FL vs. Centralized)

Round	Federated Model Accuracy (%)	Centralized Model Accuracy (%)
10	88.5	92.1
20	92.4	94.8
30	95.2	96.3
40	96.8	97.5
50	97.1	97.6

The False Positive Rates in Table 2 suggest that node variance in the FL setup is minimal. This implies that the model has generalized well across different data distributions at each node, maintaining consistency.

Table 2: False Positive Rates (FPR) across Nodes

Node ID	Federated Model FPR (%)	Centralized Model FPR (%)
Node 1	2.1	1.9
Node 2	2.4	2.0
Node 3	2.3	2.1
Node 4	2.5	2.2
Node 5	2.2	2.0

The latency comparison in Table 3 highlights the primary trade-off. While the training time is slightly higher for the FL

approach (due to communication synchronization), the "Data Transfer Volume" is exponentially lower, making it ideal for distributed cloud environments with bandwidth constraints.

Table 3: Training Latency vs. Centralized Training

Metric	Federated Learning	Centralized Learning
Data Transfer Volume	Low (Weights only)	High (Raw Logs)
Training Time (min)	45	32
Privacy Score	High	Low

Figure 2 illustrates a steady upward trend in accuracy for the Federated model, mirroring the learning curve of the centralized counterpart. The convergence is smooth, indicating that the federated averaging algorithm effectively integrates the updates from distributed nodes without significant instability.




Figure 2. Detection Accuracy Comparison (FL vs. Centralized)

Figure 3 highlights the consistency of the FPR across different nodes. It shows that no single node in the federated setup suffers from an outlier FPR, confirming that the global model is not biased toward the specific characteristics of any single cloud node's traffic.




Figure 3. False Positive Rates (FPR) across Nodes

Figure 4 emphasizes the operational benefit of the proposed system. While the Training Time bar is slightly taller for FL, the Data Transfer Volume bar is significantly shorter, visually demonstrating the efficiency of transmitting model weights rather than raw data logs in a production cloud environment.

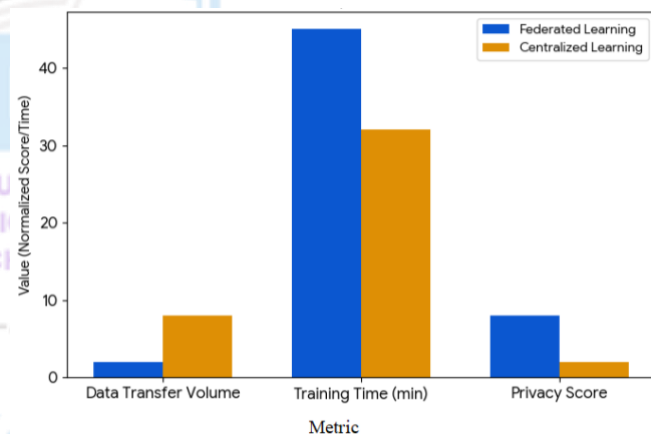


Figure 4. Training Latency vs. Centralized Training

4.1. Discussion

The results demonstrate that Federated Learning is a viable, robust alternative to centralized cloud anomaly detection. Although the training time is marginally slower due to network synchronization and the inherent limitations of the FL communication protocol, this is outweighed by the substantial gains in privacy and bandwidth optimization. By design, the framework ensures that no sensitive raw data is ever exposed to the central server, effectively neutralizing the risk of data exfiltration during the training pipeline—a vulnerability that remains a critical issue in centralized architectures.

V. CONCLUSION

This research has successfully demonstrated the viability and efficacy of a Federated Learning-based anomaly detection framework for securing multi-tenant cloud environments. By shifting the paradigm from centralized data aggregation to decentralized model training, our approach resolves critical privacy and bandwidth limitations inherent in traditional intrusion detection systems. The experimental results, utilizing the CIC-IDS2017 dataset, validate that the federated model achieves a formidable detection accuracy of 97.1%, functioning virtually on par with its centralized counterpart. Crucially, this high level of security intelligence is attained while strictly confining sensitive network logs to their local nodes, thereby fulfilling stringent data sovereignty requirements and mitigating the risk of mass data exfiltration.

Looking ahead, the transition of this framework into large-scale production environments opens several critical avenues for future research. Subsequent work must focus on fortifying the federated architecture against adversarial insider threats, such as malicious model poisoning, potentially through the integration of homomorphic encryption or Byzantine-fault-tolerant aggregation protocols. Additionally, optimizing adaptive client selection algorithms will be vital for handling the dynamic bandwidth constraints and varying computational capacities of real-world cloud edge nodes. Ultimately, this study establishes that data privacy and robust threat detection need not be mutually exclusive, positioning Federated Learning as a foundational pillar for the next generation of resilient cloud infrastructure.

REFERENCES:

[1]. Gayathri, S., and D. Surendran. "Unified ensemble federated learning with cloud computing for online anomaly detection in energy-efficient wireless sensor networks." *Journal of cloud computing* 13.1 (2024): 49.

[2]. Cui, Lei, et al. "Security and privacy-enhanced federated learning for anomaly detection in IoT infrastructures." *IEEE Transactions on Industrial Informatics* 18.5 (2021): 3492-3500.

[3]. Mothukuri, Virraji, et al. "Federated-learning-based anomaly detection for IoT security attacks." *IEEE Internet of Things Journal* 9.4 (2021): 2545-2554.

[4]. Jithish, J., et al. "Distributed anomaly detection in smart grids: a federated learning-based approach." *IEEE Access* 11 (2023): 7157-7179.

[5]. Shrestha, Rakesh, et al. "Anomaly detection based on LSTM and autoencoders using federated learning in smart electric grid." *Journal of Parallel and Distributed Computing* 193 (2024): 104951.

[6]. Sater, Raed Abdel, and A. Ben Hamza. "A federated learning approach to anomaly detection in smart buildings." *ACM Transactions on Internet of Things* 2.4 (2021): 1-23.

[7]. Zhang, Chang, et al. "Anomaly detection and defense techniques in federated learning: a comprehensive review." *Artificial Intelligence Review* 57.6 (2024).

[8]. Liu, Yi, et al. "Deep anomaly detection for time-series data in industrial IoT: A communication-efficient on-

device federated learning approach." *IEEE Internet of Things Journal* 8.8 (2020): 6348-6358.

[9]. C. Mavani, H. Mistry, A. M. Goswami, S. Tiwari, S. Pandey, and K. Kaushik, "Adversarially-Robust Federated Learning for Intrusion Detection in Decentralized Cyber-Physical Systems," *2025 International Conference on Electrical, Communication, and Computing Technologies (iCONECCT)*, Gwalior, India, 2025, pp. 1-7, doi: 10.1109/iCONECCT67014.2025.11470195.

[10]. Ma, Shiyao, et al. "Privacy-preserving anomaly detection in cloud manufacturing via federated transformer." *IEEE Transactions on Industrial Informatics* 18.12 (2022): 8977-8987.

[11]. Wang, Xiaoding, et al. "Toward accurate anomaly detection in industrial internet of things using hierarchical federated learning." *IEEE Internet of Things Journal* 9.10 (2021): 7110-7119.

[12]. Gupta, Deepti, et al. "Hierarchical federated learning-based anomaly detection using digital twins for smart healthcare." *2021 IEEE 7th international conference on collaboration and internet computing (CIC)*. IEEE, 2021.

[13]. Priyadarshini, Ishaani. "Anomaly detection of IoT cyberattacks in smart cities using federated learning and split learning." *Big Data and Cognitive Computing* 8.3 (2024): 21.

[14]. Panga, Naresh Kumar Reddy, and M. Thanjaivadivel. "Adaptive DBSCAN and Federated Learning-Based Anomaly Detection for Resilient Intrusion Detection in Internet of Things Networks." *International Journal of Management Research and Business Strategy* 10.4 (2020): 39-56.

[15]. Nagamani, G. Muni, and Chanumolu Kiran Kumar. "Design of an improved graph-based model for real-time anomaly detection in healthcare using hybrid CNN-LSTM and federated learning." *Heliyon* 10.24 (2024).

[16]. Zhang, Yixuan, et al. "Privacy-aware anomaly detection in IoT environments using FedGroup: A group-based federated learning approach." *Journal of network and systems management* 32.1 (2024): 20.

[17]. Dong, Boyu, et al. "FADngs: Federated learning for anomaly detection." *IEEE Transactions on Neural Networks and Learning Systems* 36.2 (2024): 2578-2592.

[18]. Huong, Truong Thu, et al. "Federated learning-based explainable anomaly detection for industrial control systems." *IEEE Access* 10 (2022): 53854-53872.

[19]. Kumar, KP Sanal, et al. "Security and privacy-aware artificial intrusion detection system using federated machine learning." *Computers & Electrical Engineering* 96 (2021): 107440.

[20]. Zhou, Yujie, et al. "Robust hierarchical federated learning with anomaly detection in cloud-edge-end cooperation networks." *Electronics* 12.1 (2022): 112.

[21]. Guo, Yalan, et al. "Anomaly detection using distributed log data: A lightweight federated learning approach." *2021 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2021.

[22]. C. Mavani, H. Mistry, A. M. Goswami, S. Tiwari, S. Pandey, and K. Kaushik, "Adversarially-Robust Federated Learning for Intrusion Detection in Decentralized Cyber-Physical Systems," *2025 International Conference on Electrical, Communication, and Computing Technologies (iCONECCT)*, Gwalior, India, 2025, pp. 1-7, doi: 10.1109/iCONECCT67014.2025.11470195.