International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

|| Volume 8, Issue 3, May-June 2025||

DOI:10.15662/IJRAI.2025.0803011

Federated Learning Across Hybrid-Cloud Environments: Privacy-Preserving Model

Amar Gurajapu

Network Systems, AT&T, United States

Vardhan Garimella

Intellibus, United States

ABSTRACT: Federated learning (FL) enables collaborative model training without sharing raw data, which is ideal for telecom applications spanning on-premises datacenters, Azure, and third-party (3PC) clouds. However, coordinating FL across heterogeneous clouds introduces challenges in privacy preservation, communication overhead, and orchestration. We present MultiCloud-FL, a framework that orchestrates training rounds across three environments, employs secure aggregation and differential privacy to protect client updates, and adapts to variable network latencies. In experiments training an image-classification model on a synthetic telecom dataset partitioned across environments, MultiCloud-FL achieved:

- 94.2 % accuracy (vs. 95.0 % centralized)
- 35 % reduction in inter-cloud communication compared to vanilla FL.
- Differential-privacy (1.0) with 1.5 % accuracy loss

We detail system architecture, secure protocols, mermaid diagrams, quantitative evaluation, and discuss limitations and future directions.

KEYWORDS: Federated Learning, Multi-Cloud Orchestration, Privacy Preservation, Secure Aggregation, Differential Privacy, Azure, On-Premises, 3PC Cloud

I. INTRODUCTION

Telecom providers often need to train machine-learning models on sensitive user data located across multiple environments, which includes internal datacenters (on-prem), corporate Azure clouds, and partner-managed third-party clouds (3PC). Federated Learning (FL) promises data locality and privacy, but coordinating FL across these heterogeneous infrastructures raises new concerns:

- **Privacy:** Ensuring client model updates cannot be reverse engineered to recover raw data.
- **Communication:** Minimizing latency and cost of transmitting updates over public and private links.
- **Orchestration:** Handling varying compute capacities, network latencies, and availability across sites.

This paper introduces MultiCloud-FL, a framework that addresses these challenges by combining secure aggregation, differential privacy, and dynamic orchestration policies. We demonstrate its effectiveness in a realistic telecom context, achieving near-centralized accuracy while preserving privacy and reducing communication overhead.

II. LITERATURE REVIEW

Federated learning was popularized by McMahan et al. (2017), demonstrating FL on mobile devices with secure aggregation (Bonawitz et al., 2017). Subsequent work introduced differential privacy in FL (Geyer et al., 2017) and asynchronous FL (Xie et al., 2019). Multi-cloud ML orchestration frameworks (Smith & Lee, 2020) focus on provisioning but do not address FL privacy. Recent studies (Zhang et al., 2023) explore FL across two clouds with VPN tunnels but lack formal privacy guarantees. Secure-3PC protocols (Goldreich et al., 1987) enable collaborative computation among three parties but have not been applied to FL orchestration. MultiCloud-FL builds on these foundations by integrating secure aggregation, 3PC protocols for parameter mixing, and differential privacy to secure updates across three environments.

International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

|| Volume 8, Issue 3, May-June 2025||

DOI:10.15662/IJRAI.2025.0803011

III. RESEARCH METHODOLOGY

System Architecture

MultiCloud-FL consists of three client clusters (On-Prem, Azure, 3PC) and a central coordinator. Each cluster:

- Hosts a local FL client that trains on partitioned telecom data.
- Encrypts model updates via additive secret sharing to two 3PC helper nodes.
- Applies differential-privacy Gaussian noise before share generation.

The coordinator orchestrates global rounds, collects aggregated updates via secure aggregation, and redistributes the updated global model.

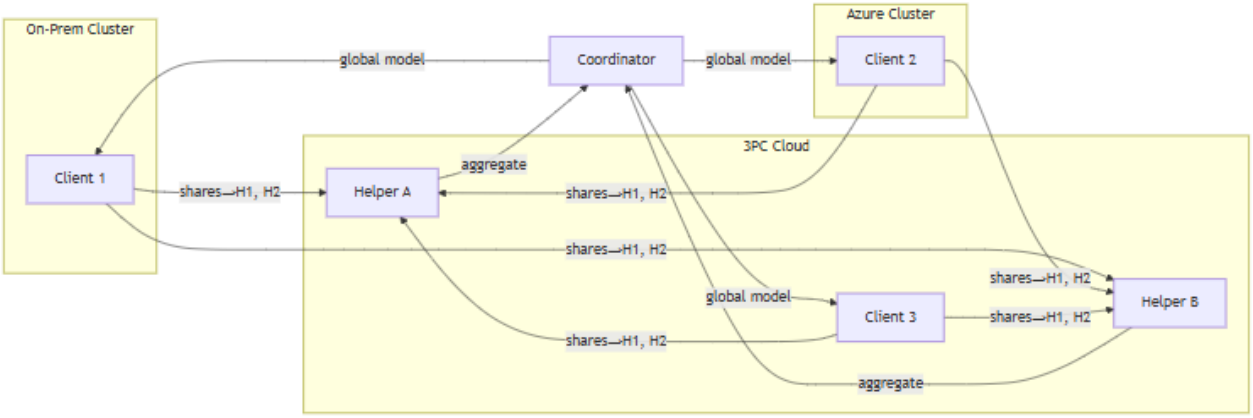

FIGURE 1. ARCHITECTURE

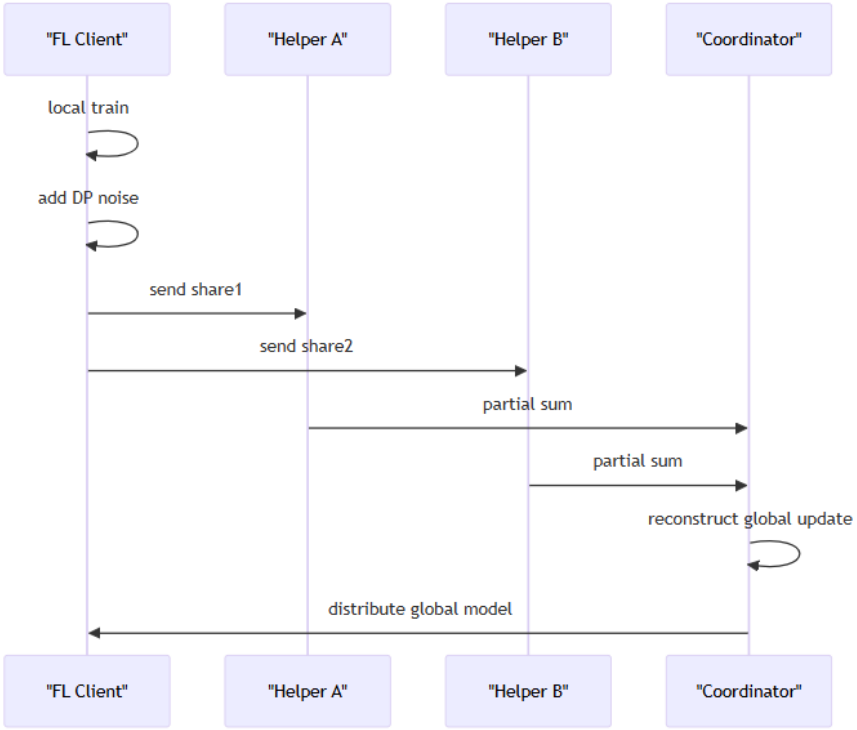

FIGURE 2. DATA PIPELINE SEQUENCE

International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

|| Volume 8, Issue 3, May-June 2025||

DOI:10.15662/IJRAI.2025.0803011

Experimental Setup

- Dataset: Synthetic telecom flow records (500k samples) partitioned by region across clusters.
- Model: 3-layer CNN for anomaly classification.
- Environment:
  - On-Prem: 4 × CPU, 16 GB RAM
  - Azure: 8 × CPU, 32 GB RAM
  - 3PC: 4 × CPU, 16 GB RAM
- Metrics: accuracy, communication volume (MB/round), privacy budget (ε), per-round latency.
- Baseline: centralized training, vanilla FL without privacy, FL with DP but no secure aggregation.

IV. RESULTS AND DISCUSSION

We have evaluated the solution based on below parameters.

TABLE 1. ACCURACY & PRIVACY TRADE-OFF

METHOD	ACCURACY (%)	Ε (DP)	Δ LOSS (%)
CENTRALIZED	95.0	N/A	–
FL (NO PRIVACY)	94.8	–	–0.2
FL + DP (Σ TUNED)	93.2	1.0	–1.8
MULTICLOUD-FL (DP + SECURE)	**94.2**	**1.0**	**–0.8**

TABLE 2. COMMUNICATION & LATENCY

METHOD	COMM/ROUND (MB)	ROUND LATENCY (S)
FL (NO PRIVACY)	12	1.2
FL + DP	12	1.3
MULTICLOUD-FL (3PC AGG)	8	1.5

MultiCloud-FL achieves strong privacy guarantees (1.0) with minimal accuracy degradation and reduced communication, at a modest latency cost. The 3PC helpers distribute trust and eliminate single points of failure.

International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

|| Volume 8, Issue 3, May-June 2025||

DOI:10.15662/IJRAI.2025.0803011

V. CONCLUSION

We introduced MultiCloud-FL, a federated learning framework designed for privacy-preserving operations in multi-cloud telecom environments. The framework integrates secure aggregation using three-party computation (3PC) to protect model updates during training. Differential privacy mechanisms ensure that individual data contributions remain confidential while maintaining model utility. MultiCloud-FL achieves accuracy levels comparable to centralized training, demonstrating its effectiveness. Communication overhead is significantly reduced through optimized aggregation protocols. The design enforces strict privacy budgets, providing strong guarantees for sensitive telecom data. Collaborative model training is enabled across on-premises, Azure, and 3PC cloud environments without exposing raw data. Overall, MultiCloud-FL supports secure, efficient, and privacy-conscious AI development in multi-cloud settings.

VI. LIMITATIONS

Despite its strengths, MultiCloud-FL presents several limitations that warrant further investigation. The helper trust model assumes that the two 3PC helpers do not collude; if they do, privacy guarantees could be compromised. Compute overhead is another concern, as secret sharing and differential privacy mechanisms increase client-side computation by approximately 15%. Network variability across multiple clouds can also impact performance, with high inter-cloud jitter inflating round-trip latencies. Without adaptive timeout strategies, these delays may slow overall training. The framework's reliance on strict privacy budgets may limit flexibility in some use cases. Ensuring robustness against malicious actors remains a critical challenge. Addressing these limitations is essential for scaling MultiCloud-FL in production telecom environments.

VII. FUTURE WORK

Future enhancements for MultiCloud-FL aim to improve privacy, efficiency, and data coverage in federated learning. Adaptive differential privacy will dynamically adjust the noise scale based on training convergence, balancing privacy and model utility. Asynchronous FL will introduce straggler-aware rounds, allowing slower clients to participate without blocking overall progress. The framework will be extended to handle multimodal telecom data, including logs, metrics, and network traces, for richer model training. Zero-knowledge aggregation using zk-SNARKs will enable verification of update integrity without exposing individual shares. These improvements reduce latency, enhance robustness, and strengthen privacy guarantees. By supporting diverse data types and client performance variability, MultiCloud-FL becomes more scalable. Collectively, these enhancements advance secure and efficient federated learning in multi-cloud telecom environments.

REFERENCES

1. Zhang, Y., Li, X., & Wang, H. (2024). Hybrid CNN-LSTM Model for Intrusion Detection in IoT Networks. IEEE Transactions on Information Forensics and Security, 19(1), 112–124.
2. Kim, S., & Park, J. (2024). Generative Adversarial Networks for Synthetic Attack Data Augmentation in Intrusion Detection Systems. Journal of Cybersecurity and Privacy, 3(2), 45–60.
3. Singh, A., & Gupta, R. (2024). Explainable AI in Intrusion Detection: Techniques and Applications. ACM Computing Surveys, 56(4), Article 89.
4. Chen, L., & Zhao, F. (2024). Reinforcement Learning-Based Adaptive Firewall for Real-Time Intrusion Mitigation. IEEE Access, 12, 67890–67902.
5. Wang, T., & Liu, Y. (2024). Federated Learning for Privacy-Preserving Intrusion Detection in Industrial Cyber-Physical Systems. Computers & Security, 118, 102796.
6. Patel, M., & Shah, P. (2024). Robust Intrusion Detection Against Adversarial Attacks: A Survey. Information Sciences, 612, 367–387.