



PDF Download
3744969.3748413.pdf
17 March 2026
Total Citations: 0
Total Downloads: 418

 Latest updates: <https://dl.acm.org/doi/10.1145/3744969.3748413>

Published: 10 September 2025

[Citation in BibTeX format](#)

SIGCOMM '25: ACM SIGCOMM 2025
Conference

September 8 - 11, 2025
Coimbra, Portugal

Conference Sponsors:
SIGCOMM

SHORT-PAPER

NEBULA - Decentralized Federated Learning for Heterogeneous Networks

[ENRIQUE TOMÁS MARTÍNEZ BELTRÁN](#), University of Murcia, Murcia, Murcia, Spain

[GÉRÔME BOVET](#), Armasuisse, Switzerland, Bern, BE, Switzerland

[GREGORIO MARTÍNEZ PÉREZ](#), University of Murcia, Murcia, Murcia, Spain

[ALBERTO HUERTAS CELDRÁN](#), University of Murcia, Murcia, Murcia, Spain

Open Access Support provided by:

[Armasuisse, Switzerland](#)

[University of Murcia](#)

DEMO: NEBULA – Decentralized Federated Learning for Heterogeneous Networks

Enrique Tomás Martínez Beltrán¹, G r me Bovet², Gregorio Mart nez P rez¹, Alberto Huertas Celdr n¹

¹Department of Information and Communications Engineering, University of Murcia, Spain

[enriquetomas,gregorio,alberto.huertas]@um.es

²Cyber-Defence Campus, armasuisse Science & Technology, Switzerland
gerome.bovet@armasuisse.ch

Abstract

Federated learning (FL) enables collaborative model training without sharing raw data, which is pivotal for maintaining privacy. However, existing FL frameworks often rely on a central coordinator, posing risks in heterogeneous networks. This work presents NEBULA, a decentralized FL platform that unifies centralized and peer-to-peer FL paradigms, integrating network awareness and autonomous adaptation for improved resilience and efficiency. Key contributions include: (1) a unified architecture supporting both server-coordinated and fully decentralized operation; (2) network-aware orchestration for dynamic communication and aggregation optimization; and (3) built-in mechanisms for robust operation. The demonstration will showcase real-time performance, defense against adversarial attacks, and adaptive client participation in challenging network scenarios.

CCS Concepts

- **Networks** → **Network simulations**; **Network mobility**;
- **Security and privacy** → **Malware and its mitigation**.

Keywords

Decentralized Systems, Resilience, Visualization

ACM Reference Format:

Enrique Tom s Mart nez Beltr n¹, G r me Bovet², Gregorio Mart nez P rez¹, Alberto Huertas Celdr n¹. 2025. DEMO: NEBULA – Decentralized Federated Learning for Heterogeneous Networks. In *ACM SIGCOMM 2025 Posters and Demos (SIGCOMM Posters and Demos '25)*, September 8–11, 2025, Coimbra, Portugal. ACM, New York, NY, USA, 3 pages.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

SIGCOMM Posters and Demos '25, Coimbra, Portugal

  2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2026-0/2025/09.

<https://doi.org/10.1145/3744969.3748413>

1 Introduction

Federated learning (FL) facilitates collaborative machine learning model training across distributed data sources without centralizing sensitive user data [9]. This paradigm is crucial for privacy-sensitive applications in edge and cloud computing. However, many prominent FL solutions, such as TensorFlow Federated (TFF) [11], FATE [6], Flower [3], PySyft [12], and FedML [5], often depend on a central server. This centralized architecture can introduce a single point of failure and performance bottlenecks, particularly in dynamic or resource-constrained heterogeneous network environments. Decentralized Federated Learning (DFL) [1] mitigates these issues via peer-to-peer communication, removing the central coordinator. However, existing DFL frameworks, such as BrainTorrent [10] or Fedstellar [8], often overlook real-world network dynamics or lack integrated adaptation. Furthermore, while secure aggregation protocols are essential for cryptographic privacy, they do not inherently address system resilience or efficient operation over unreliable networks, aspects often underemphasized in FL demonstrations.

This paper introduces NEBULA, a novel FL platform that unifies centralized and decentralized paradigms while being *network-aware* and *autonomous* (publicly available in [2]). NEBULA is engineered to maintain training progress and adapt performance despite network disruptions, node failures, or adversarial behavior. The key contributions are:

- **Unified FL Architecture:** A flexible platform supporting operation in server-coordinated and DFL modes, adaptable to diverse network topologies and application needs.
- **Network-Aware Orchestration:** Integration of real-time network context into training decisions, dynamically optimizing communication patterns and node selection.
- **Robustness and Adaptation Mechanisms:** Strategies for resilience against common adversarial attacks and adaptive responses to dynamic client states, such as mobility and fluctuating network connectivity.

2 NEBULA Platform

Architecture Overview: NEBULA employs a modular architecture comprising three main components, as illustrated

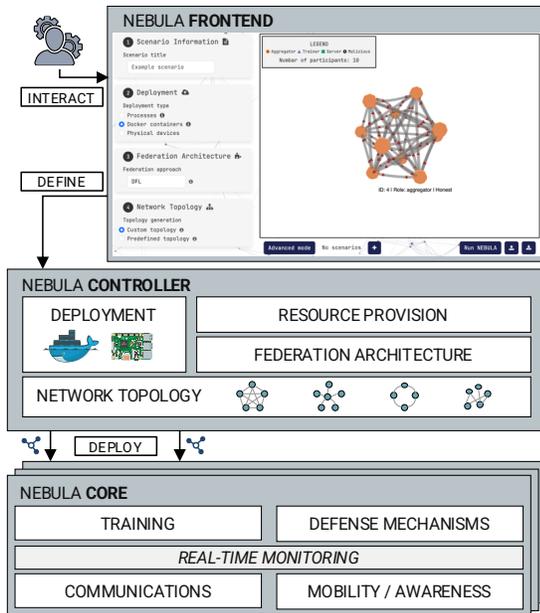


Figure 1: NEBULA Platform Architecture

in Figure 1. The *Frontend* provides a web-based dashboard and REST/WebSocket APIs for experiment configuration, real-time monitoring, and results visualization. It interacts with the *Controller*, orchestrating the FL experiment lifecycle. The *Controller* manages initial setup, including session instantiation, client registration, and initial model distribution. The *Core* executes on each physical and virtual client. It is responsible for local model training, network condition sensing, node discovery, and a configurable peer-to-peer protocol for robust model dissemination and aggregation among nodes. The system is designed for asynchronous operations to enhance tolerance to network delays and straggler nodes. **Features and Implementation:** NEBULA supports a versatile range of FL settings. The platform implements multiple aggregation algorithms, including FedAvg [9], and robust aggregation algorithms like Krum [4] to defend against malicious updates. Additional defense mechanisms, such as reputation based on update consistency and gradient norm bounding, are integrated. Network awareness is achieved through a combination of active probing (e.g., round-trip time measurements to potential nodes) and passive monitoring of communication statistics. This network intelligence informs dynamic node selection in DFL mode, aggregation frequency adjustments, and potentially adaptive model compression strategies. NEBULA is implemented in Python 3.11 with *asyncio* for efficient and scalable asynchronous message passing. The platform is containerized using Docker and is compatible with Kubernetes orchestration.

3 Demonstration

The demonstration showcases NEBULA capabilities across two distinct testbed configurations: (i) 10 containerized clients and (ii) 10 resource-constrained Raspberry Pi 4 devices. For decentralized operations within each configuration, clients establish a peer-to-peer network based on a random Erdős–Rényi topology. Network variability is emulated using a custom wrapper of the Linux traffic control (*tc*) utility. Nodes collaboratively train models for image classification (CIFAR-10 with a MobileNetV3-Small architecture) and sentiment classification (IMDb [7] with a distilled BERT model). The demonstration is structured around three scenarios:

- *Unified Architecture Showcase:* This scenario demonstrates NEBULA orchestrating an FL task on the heterogeneous testbed, showing centralized and decentralized operating modes. The dashboard visualizes node participation, communication links, and model convergence.
- *Adversarial Attacks and Defense:* Building upon an ongoing federated task, one or more clients are configured to act maliciously via label-flipping or model-poisoning attacks. The impact is visualized on the live dashboard. Attendees will observe the system detecting and mitigating malicious contributions by filtering or down-weighting harmful updates, followed by stabilization of model performance.
- *Mobile Client Adaptation:* A client simulates mobility by dynamically changing network conditions, altering its set of reachable nodes. Clients use their integrated network sensing module to detect these real-time changes. The dashboard will highlight adaptive behaviors, such as dynamically altering node selection or adjusting update frequency based on the detected network context.

During the scenarios, a web dashboard displays telemetry for attendees to access and engage with the demonstrators.

4 Conclusion

NEBULA offers a unified and robust platform for FL in heterogeneous networks. By integrating centralized and decentralized approaches with network awareness and adaptive mechanisms, it addresses key limitations of existing frameworks. The demonstration highlights capabilities in handling diverse hardware, dynamic network conditions, and adversarial threats. The platform is a valuable tool for research in networking, distributed systems, and FL. Future work includes larger-scale evaluations and enhanced security.

Acknowledgments

This work has been partially supported by (a) the Swiss Federal Office for Defense Procurement (armasuisse) with the DECIMAL project, (b) 21629/FPI/21, Fundación Séneca, and (c) the European Commission through the Horizon Europe/JU SNS project ROBUST-6G (Grant no. 101139068).

References

- [1] Enrique Tomás Martínez Beltrán et al. 2023. Decentralized Federated Learning: Fundamentals, State of the Art, Frameworks, Trends, and Challenges. *IEEE Communications Surveys & Tutorials* 25, 4 (2023), 2983–3013. doi:10.1109/COMST.2023.3315746
- [2] Enrique Tomás Martínez Beltrán, G r me Bovet, Gregorio Mart nez P rez, and Alberto Huertas Celdr n. 2025. NEBULA Platform: Web Application. Available at <https://nebula-dfl.com>. Last accessed on 20/06/25.
- [3] Daniel J Beutel et al. 2020. Flower: A Friendly Federated Learning Research Framework. *arXiv preprint arXiv:2007.14390* (2020).
- [4] Peva Blanchard, El Mahdi El Mhamdi, Rachid Guerraoui, and Julien Stainer. 2017. Machine Learning with Adversaries: Byzantine Tolerant Gradient Descent. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2017/file/f4b9ec30ad9f68f89b29639786cb62ef-Paper.pdf
- [5] Chaoyang He et al. 2020. FedML: A Research Library and Benchmark for Federated Machine Learning.
- [6] Yang Liu, Tao Fan, Tianjian Chen, Qian Xu, and Qiang Yang. 2021. FATE: An Industrial Grade Platform for Collaborative Learning With Data Protection. *Journal of Machine Learning Research* 22, 226 (2021), 1–6. <http://jmlr.org/papers/v22/20-815.html>
- [7] Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. Learning Word Vectors for Sentiment Analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, Portland, Oregon, USA, 142–150. <http://www.aclweb.org/anthology/P11-1015>
- [8] Enrique Tomás Martínez Beltrán,  ngel Luis Perales G mez, Chao Feng, Pedro M guel S nchez S nchez, Sergio L pez Bernal, G r me Bovet, Manuel Gil P rez, Gregorio Mart nez P rez, and Alberto Huertas Celdr n. 2024. Fedstellar: A Platform for Decentralized Federated Learning. *Expert Systems with Applications* 242 (2024), 122861. doi:10.1016/j.eswa.2023.122861
- [9] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 54)*. PMLR, 1273–1282.
- [10] Abhijit Guha Roy, Shayan Siddiqui, Sebastian P lsterl, Nassir Navab, and Christian Wachinger. 2019. BrainTorrent: A Peer-to-Peer Environment for Decentralized Federated Learning. arXiv:1905.06731 [cs.LG] <https://arxiv.org/abs/1905.06731>
- [11] TensorFlow Federated Authors. 2019. TensorFlow Federated: Machine Learning on Decentralized Data. <https://www.tensorflow.org/federated>. Software framework, accessed 2 May 2025.
- [12] Alexander Ziller et al. 2021. *PySyft: A Library for Easy Federated Learning*. Springer International Publishing, Cham, 111–139. doi:10.1007/978-3-030-70604-3_5