SUCCESS-6G: VERIFY

WP4 Deliverable E9

Decentralized methods for predictive diagnosis of vehicles condition

Project Title:	SUCCESS-6G: VERIFY
Title of Deliverable:	Decentralized methods for predictive diagnosis of vehicles condition
Status-Version:	v1.0
Delivery Date:	17/03/2025
Contributors:	Guillermo Candela Belmonte, Antonio Fernandez (Optare), Pavol Mulinka, Charalampos Kalalas, Roshan Sedar (CTTC), Michail Dalgitsis, Eftychia Datsika, Angelos Antonopoulos (NBC)
Lead editor:	Charalampos Kalalas (CTTC)
Reviewers:	Miquel Payaro (CTTC)
Keywords:	Federated learning, decentralized condition diagnostics, service orchestration

Document revision history

Version	Date	Description of change
v0.1	15/12/2024	Table of contents created
v0.2	25/02/2025	Main content added
v0.3	10/03/2025	Final edits
v1.0	17/03/2025	Final version uploaded to the website

Acknowledgment

The research conducted by SUCCESS-6G - TSI-063000-2021-39/40/41 receives funding from the Ministerio de Asuntos Económicos y Transformación Digital and the European Union-NextGenerationEU under the framework of the “Plan de Recuperación, Transformación y Resiliencia” and the “Mecanismo de Recuperación y Resiliencia”.

[1] http://creativecommons.org/licenses/by-nc-nd/3.0/deed.en_US

Executive Summary

Decentralized learning methods, particularly federated learning (FL), are increasingly being explored for vehicle condition monitoring within an edge-based monitoring infrastructure. This Deliverable investigates the application of FL in a distributed system where multiple edge nodes, serving as clients, collaboratively train an autoencoder model without sharing raw data. The deployment of this system leverages Kubernetes, offering scalability, fault tolerance, and efficient resource management. The decentralized components ensuring the orchestration process of the vehicle condition monitoring service are also explained.

Table of Contents

List of Figures

List of Tables

1 Introduction

Decentralized learning methods, e.g., federated learning (FL) [1], [2], are increasingly being explored for vehicle condition monitoring within an edge-based monitoring infrastructure, which is particularly valuable in high-mobility vehicular environments. In such dynamic settings, vehicles continuously move between different locations, interact with various road conditions, and generate vast amounts of sensor data, e.g., related to engine health, braking efficiency, fuel consumption, and tire wear.

Traditional centralized machine learning approaches face challenges such as latency, high communication overhead, and data privacy risks due to the need for constant data transmission to a remote cloud [4]. In decentralized setups, edge nodes—such as roadside units, onboard vehicle sensors, and edge servers—serve as clients (agents) that collaboratively train a machine learning model without sharing raw data. Instead of transmitting vast amounts of potentially sensitive vehicle data to a centralized server, each edge node processes and trains local models based on the data it collects from vehicles in its vicinity. Periodically, these edge nodes share only model updates (such as gradients or weights) with a central server or coordinator, which aggregates these updates to improve a global model. This decentralized approach enhances data privacy, reduces communication costs, and allows for real-time learning across different geographical locations and environmental conditions.

In the context of vehicle condition monitoring, FL enables vehicles to share insights about equipment status without compromising data security. Moreover, as edge nodes form an interconnected network of learning agents, they can adapt to localized driving conditions, traffic variations, and climate effects, making the predictive maintenance models more robust and region-specific. In high-mobility scenarios, FL provides seamless adaptability, allowing vehicles to contribute to and benefit from an evolving global model while on the move. This is crucial for ensuring real-time fault detection, predictive maintenance, and system optimization in connected ecosystems [3], [6], [7]. The use of FL in such an architecture also ensures scalability, as new edge nodes (such as new vehicles, sensor stations, or monitoring hubs) can seamlessly join the network without overburdening a central server [8]. Furthermore, advanced techniques such as secure aggregation, differential privacy, and blockchain-based consensus mechanisms can be integrated to improve the security and reliability of FL in such a dynamic environment [5]. Ultimately, this decentralized learning paradigm enables a more intelligent, autonomous, and responsive vehicle condition monitoring system, crucial for fleet management, smart cities, and connected vehicle ecosystems.

This Deliverable investigates the application of FL in a distributed system where multiple edge nodes, serving as clients, collaboratively train an autoencoder (AE) model without sharing raw data. The deployment of this system leverages Kubernetes, ensuring scalability, fault tolerance, and efficient resource management. The decentralized features for service orchestration are also explained.

The objectives are the following:

1. Evaluate FL in a distributed environment with isolated clients:
    - Implement an FL framework that enables multiple clients to train local models on their respective datasets while contributing to a shared global model.
    - Examine the effectiveness of FL in scenarios where data remains decentralized, ensuring a practical, real-world approach to distributed machine learning.
2. Ensure data privacy, avoiding centralized storage:
    - Unlike traditional machine learning paradigms where data is aggregated into a central repository, FL allows data to remain localized on each client's device or server.
    - Secure model updates (e.g., gradients or weights) are shared instead of raw data,

preserving privacy and reducing the risk of data breaches.

3. Analyze the convergence of a global model compared to models trained locally:
    - Compare the performance of a globally aggregated model (trained via FL) with independent local models trained on individual datasets.
    - Measure key metrics to determine the effectiveness of FL in achieving a well-generalized model.

2 Implementation

2.1 General Description

The workflow consists of an FL system where multiple clients (workers) collaboratively train an AE model in a distributed manner. The considered architecture is illustrated in Figure 1.

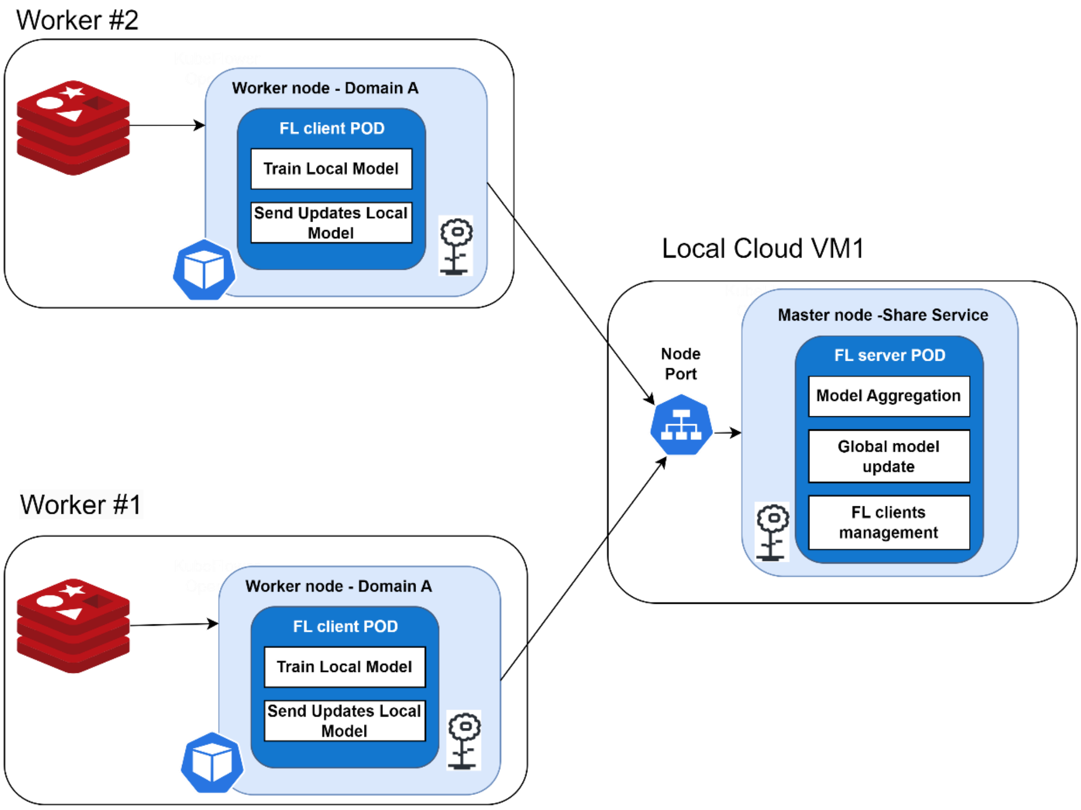

Figure 1 Considered decentralized architecture for monitoring measurement streams

The core components of this implementation include:

- Federated Clients (Workers):
  - Each client has access to its own dataset, ensuring that data remains private and isolated.
  - Clients train an AE model locally using their respective datasets.
  - They participate in the FL process by periodically sending model updates (e.g., weights or gradients) to a central Federated Server instead of sharing raw data.
- Federated Server:
  - Acts as the central orchestrator of the FL training process.
  - Aggregates local model updates from all clients using Federated Averaging (FedAvg) [9] to create a global model.
  - Coordinates training rounds, ensuring that model convergence is achieved efficiently across distributed clients.
- Deployment Environment:
  - All components are deployed in a Kubernetes cluster, providing a containerized and scalable infrastructure.
  - Kubernetes ensures data isolation by restricting access between clients, preventing any unintended data sharing.

- It enables automated scaling, resource allocation, and fault tolerance for efficient FL execution.

2.2 Decentralized Monitoring Setup

The system incorporates a decentralized monitoring approach, ensuring that data remains isolated while allowing federated model training to proceed effectively. Each component plays a crucial role that can be summarized as follows:

- **FL Server**:
  - Orchestrates FL rounds, determining when clients should train and send updates.
  - Aggregates local models from multiple clients using the Federated Averaging (FedAvg) algorithm.
  - Consolidates the final global model, incorporating learned representations from all clients.
- **FL Clients**:
  - Maintain local autonomy, each working with its own dataset stored in a private Redis Stack instance.
  - Train an initial AE model on their dataset, creating a local baseline model before federated training begins.
  - Participate in FL training rounds, periodically training and sharing only model parameters (weights/gradients) with the FL server—never sharing raw data.
- **Redis Stack**:
  - Ensures data isolation by assigning a separate Redis instance per client.
  - Acts as a local storage solution for each client, housing time-series sensor data used for AE training.
  - Facilitates fast and efficient data retrieval for training without exposing raw data outside the client environment.

2.3 Service orchestration with NearbyOne

The decentralized implementation of vehicular condition monitoring service relies on the advanced capabilities of a service orchestrator. Emerging applications and services, such as Tactile Internet and autonomous driving, demand stringent requirements for throughput, latency, and reliability; challenges that traditional monolithic architectures cannot address. Within the scope of fifth-generation (5G) and forthcoming sixth-generation (6G) networks, Multi-Access Edge Computing (MEC) is envisioned as a fundamental component in network and service architecture. By bringing computational and storage resources closer to the network edge, MEC minimizes latency and enables real-time processing [10]. Furthermore, the Control and User Plane Separation (CUPS) architecture enhances network flexibility and scalability by decoupling control plane and user plane functions [11]. This separation also aids in latency reduction, as user plane functions can be distributed to the network edge. Supporting these paradigms are key technologies such as Software-Defined Networking (SDN) and Network Functions Virtualization (NFV). SDN enables network administrators to programmatically configure and manage network services by decoupling the control plane from the data plane, thereby enhancing flexibility and ease of management [12]. Meanwhile, NFV facilitates the virtualization of network functions, allowing them to operate on standard server hardware instead of dedicated equipment, reducing costs and improving scalability.

The transition toward software-based and virtualized networks has pioneered the shift to the cloud-native era, characterized by a microservices architecture, where applications are decomposed into smaller, autonomous services [13]. This era also promotes standardized communication through Application Programming Interfaces (APIs) over HyperText Transfer Protocol (HTTP) and embraces containerization, where lightweight containers replace traditional virtual machines, offering superior efficiency and scalability [14]. As networks continue to evolve, future systems, including 5G, are expected to embrace softwarization, virtualization, and cloudification across their core, Radio Access Network (RAN), and transport domains. The core and RAN have increasingly adopted the cloud-native paradigm, while the transport network primarily relies on the SDN framework, with a growing incorporation of cloud-native principles. In the 5G core network, the service-based architecture (SBA) is employed for control plane functions, utilizing Service-Based Interfaces (SBI) to facilitate communication between these functions [15]. Ongoing efforts focus on the integration of the User Plane Function (UPF) (Figure 2). The UPF, which can operate with the core functions or be distributed at the edge, connects to the Session Management Function (SMF) via the N4 interface. This distributed architecture enhances scalability and reduces latency. Although the UPF is not directly integrated into the SBI architecture, it still follows cloud-native techniques and approaches, ensuring flexibility and efficiency in its operation. With the shift towards cloud-native network functions, Communication Service Providers (CSPs) now require edge-cloud infrastructure to host and run these functions. This need has given rise to the concept of Edge-Cloud Telco Continuum}, encompassing physical or virtual nodes capable of supporting cloud-native platforms to host containers and microservices.

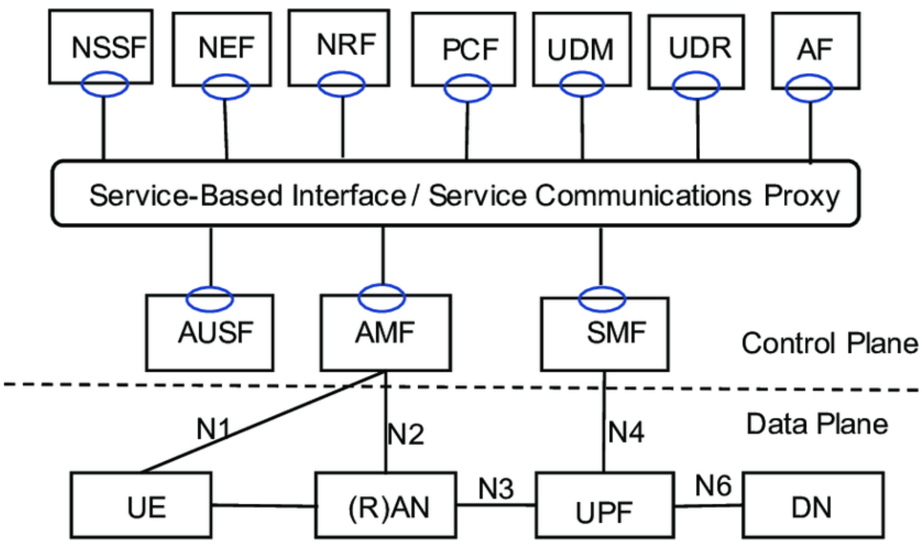

Figure 2: The Service Based Architecture (SBA) for the 5G core network

As latency-sensitive applications increase, network and compute functions need to shift towards the network edge, closer to the end user. This shift optimizes traffic distribution but also necessitates the integration of edge computing and edge networking. However, unlike the centralized cloud model, the edge lacks the necessary real estate, power budget, and operational overhead to host these services. In addition, the transition to a disaggregated, cloud-native model highlights the need for efficient scaling solutions. Effective orchestration and automation solutions are essential to address these challenges, ensuring seamless operation, resource optimization, and efficient service delivery. For example, consider edge applications and distributed UPF (d-UPF) collocated in the same edge node. In scenarios such as the d-UPF use case, the resources required for one service can conflict with or impact the resources and performance of the other. Additionally, resources at the edge are inherently more scarce than those in the cloud, necessitating careful management to avoid bottlenecks and ensure optimal performance.

Service orchestration emerges as a critical solution to manage and automate the complexities of the 5G and upcoming 6G environments. It plays a vital role in coordinating the diverse network functions, slices, applications, and API-centric services across the Edge-Cloud Telco Continuum. Its role is paramount in optimizing operations, enhancing the Quality of Service (QoS) and Quality of Experience (QoE) for end users, while also optimizing the cost-efficiency for CSPs. In essence, service orchestration is the key that unlocks the full potential of next-generation networks, turning the promise of high-speed, low-latency, cost-effective services into a reality. A high-level overview of the service orchestration within the Edge-Cloud Telco Continuum is depicted in Figure 3.

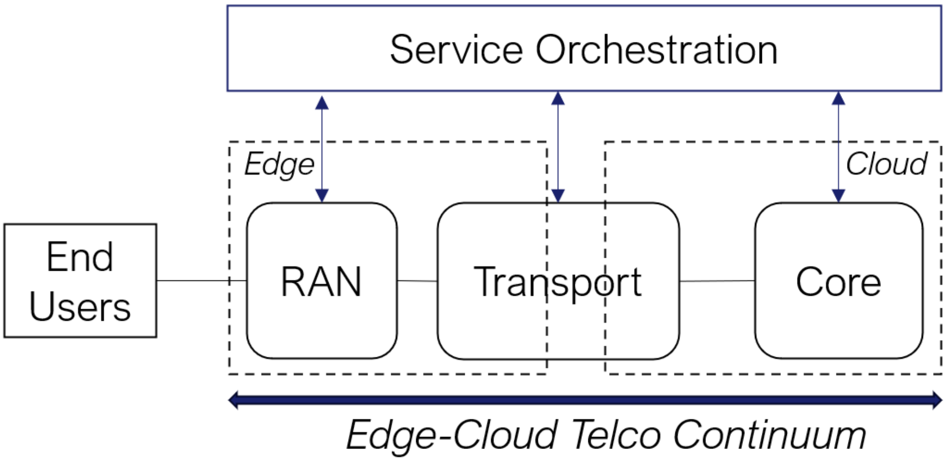

Figure 3: Service Orchestration in the Edge-Cloud Telco Continuum

A critical component of service orchestration in edge computing is the service migration among edge clusters. This is essential for maintaining seamless connectivity and low-latency performance as vehicles move across different network regions. As a vehicle transitions from one geographic area to another, its associated edge services must dynamically relocate to the nearest available edge cluster. This migration process can be triggered by mobility events detected by the 5G core network, ensuring that computational tasks and data processing remain close to the vehicle. By efficiently transferring services, service disruptions can be minimized and resource utilization in highly dynamic vehicular environments can be optimized [16].

In the context of this project, NearbyOne functions as a service orchestrator, responsible for managing the lifecycle of the vehicular condition monitoring service at the edge. It can observe vehicle mobility, initiating and performing service migrations accordingly. More specifically, during handovers, the NearbyOne orchestrator facilitates seamless service migration between available edge nodes, guided by event triggers from the 5G core network. The NearbyOne solution primarily consists of: a) the Nearby edge orchestration platform, i.e., the core component responsible for managing all tasks related to the orchestration of applications and infrastructure, and b) the Nearby Blocks, which are distributed components that encapsulate logic and code to support various application-specific functionalities. The NearbyOne multi-site orchestrator manages application services as containerized applications, packaged using Helm, the widely adopted package manager for Kubernetes. These cloud-native applications are then encapsulated within Nearby Blocks to cater to ecosystems that demand seamless inter-application communication, precise placement decisions, and fine-tuned execution platform configurations, similar to the ecosystem shaped by the use cases of this project. It manages the onboarding and lifecycle of cloud-native applications and infrastructure on a global scale and it can operate across edge and cloud sites within the SUCCESS-6G architecture.

3 Training Flow

3.1 Detailed Workflow Explanation

Figure 4 describes the step-by-step process of the FL workflow using AEs, designed to ensure data privacy and collaborative model improvement across distributed clients.

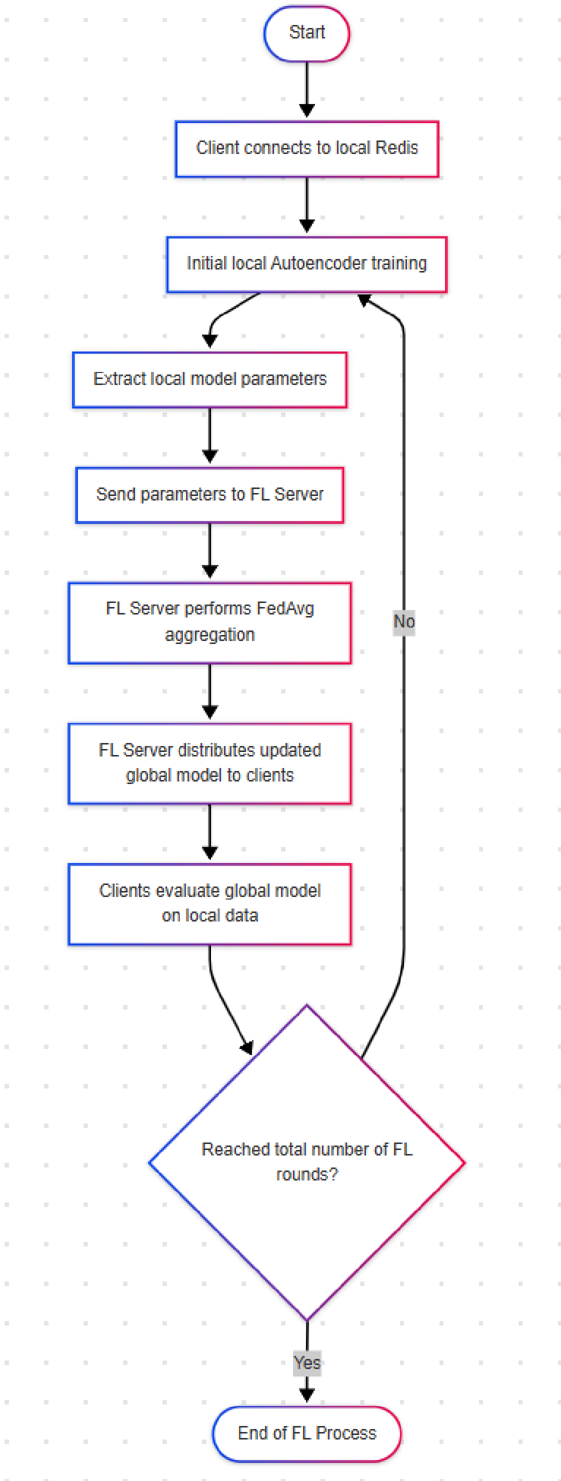

Figure 4 Training flow

3.2 Step-by-Step Process

1. Start of the Process

- The FL system is initialized and deployed in a distributed Kubernetes environment.

2. Client Connection to Local Redis

- Each FL client (worker) connects to its own private Redis Stack instance, where its own local dataset is stored.
- Data typically includes sensor time series or customer-specific data, ensuring data isolation.

3. Initial Local Training of AE

- Before participating in FL rounds, each client performs local training of an AE using its own data.
- This local base model is used as:
    - A reference model for comparison.
    - A starting point to assess how FL improves over purely local training.

4. Extraction of Local Model Parameters

- After local training, each client extracts model parameters (weights and biases).
- Only model parameters are shared — raw data remains private and local.

5. Sending Parameters to FL Server

- Each client sends its model parameters to the FL server for aggregation.

6. FedAvg Aggregation on the Server

- The FL server receives all model parameters from clients.
- FedAvg is performed:
    - The parameters are average, weighted by each client's data contribution.
- The result is a global aggregated model that combines knowledge from all clients.

7. Distribution of the Global Model to Clients

- The FL server sends the updated global model back to all clients.
- Each client updates its local model with this global version to continue training.

8. Evaluation of Global Model on Local Data

- Each client evaluates the global model on its own local dataset:
    - For example, calculating the reconstruction error (MSE) as a measure of anomaly detection.
- This assessment helps determine if the global model performs better than local-only models.

9. Check Number of Rounds

- The system checks whether the configured number of FL rounds has been completed:
    - If not completed, clients resume training using the latest global model (go back to local training).
    - If completed, the process terminates.

10. End of Process

- The FL process concludes.

- Results, including final global model performance, can now be analyzed and compared against local-only models.

4 Autoencoder Modeling

4.1 Purpose of the AE model

The AE model is specifically designed for anomaly detection in time-series sensor data. It functions by learning a compact representation of normal operating conditions and reconstructing input data with minimal error. The AE is trained on normal (non-anomalous) data, learning to encode and decode the input effectively. During inference, if the model encounters anomalous data that deviates significantly from the learned patterns, it struggles to reconstruct it accurately. The reconstruction error (difference between input and reconstructed output) serves as an anomaly detection metric—higher errors indicate potential anomalies.

4.2 AE Architecture

The AE consists of two primary components:

1. Encoder – Compresses the input data into a lower-dimensional latent representation.
2. Decoder – Reconstructs the original input from the encoded representation.

The architecture is structured as follows:

- **Input Layer**: The size of the input layer is determined by the number of input features (e.g., sensor readings).
- **Encoder**:
  - First Linear Layer:
    - Input → 32 neurons
    - Activation: ReLU (Rectified Linear Unit) to introduce non-linearity
  - Second Linear Layer:
    - 32 → 16 neurons
    - Activation: ReLU
  - Third Linear Layer (Latent Space):
    - 16 → 8 neurons (compressed representation, lowest-dimensional encoding)
- **Decoder**:
  - First Linear Layer:
    - 8 → 16 neurons
    - Activation: ReLU
  - Second Linear Layer:
    - 16 → 32 neurons
    - Activation: ReLU
  - Third Linear Layer (Output Reconstruction):
    - 32 → Original Input Size (same as input layer)

5 Training Results

5.1 Initial Local Training (Before FL)

Before engaging in FL, each client (Worker 1 and Worker 2) trained an AE model locally on their respective datasets. This served as a baseline performance metric to compare against the FL-trained global model.

Client	Initial Local Model MSE
Worker 1	0.9812
Worker 2	0.9971

Table 1 Initial local model training

As shown in Table 1, both clients exhibit relatively high Mean Squared Error (MSE), indicating that their initial local models struggled to generalize well. This high MSE highlights the potential benefits of collaborative learning via FL to improve model performance.

5.2 Evolution of the Global Model during FL Rounds

The FL process was conducted over 10 rounds, where clients trained their local models and sent updates to the FL Server, which aggregated them into a global model using FedAvg.

Round	Global MSE
1	0.9891
2	0.9757
3	0.9543
4	0.9226
5	0.8779
6	0.8220
7	0.7604
8	0.6946
9	0.6217
10	0.5465

Table 2 Evolution of the global model during FL rounds

From Table 2, it can be observed that the global MSE consistently decreased across training rounds, indicating effective convergence of the FL model. By round 10, the global MSE was reduced by nearly 50% compared to the initial local models. FedAvg successfully aggregated knowledge from both clients, improving overall model performance without sharing raw data.

5.3 Final MSE Comparison (Global Model Evaluated Locally)

After completing 10 rounds of FL, each client evaluated the final global model on their own local datasets:

Client	Initial Local Model MSE	Global Model MSE on Local Data
Worker 1	0.9812	0.5916
Worker 2	0.9971	0.5292

Table 3 Final MSE comparisons Local vs FL Model

From Table 3, it can be observed that the FL-trained global model significantly outperformed the local models for both clients:

- Worker 1: MSE reduced from 0.9812 → 0.5916 (~39% improvement).
- Worker 2: MSE reduced from 0.9971 → 0.5292 (~47% improvement).

This demonstrates that FL effectively leverages distributed learning, allowing clients to benefit from collective knowledge even without sharing raw data.

5.4 Detailed Client Metrics During FL Rounds

The performance of each client was tracked throughout the FL rounds, showing how local models contributed to global model aggregation. During the 10 FL rounds, each client's performance was tracked, showing individual MSEs contributing to the global model aggregation. Below are some key snapshots:

Round	Client 1 MSE	Client 2 MSE	Global MSE
1	0.9921	0.9921	0.9921
5	0.8795	0.8795	0.8795
10	0.5916	0.5916	0.5916

Table 4 Client's performance

From Table 4, it can be observed that the client MSEs steadily decreased in sync with the global MSE, showing successful FL training convergence. Both clients contributed equally to model improvement, demonstrating a balanced FL setup. The final global model performed significantly better than the initial local models, confirming the effectiveness of FL for anomaly detection.

6 Conclusions

FL demonstrated a substantial reduction in error compared to models trained individually by each client. This improvement is evident from the drop in MSE observed across FL rounds. Each client trained their AE model independently, but the high MSE values indicated that the models struggled to generalize well using only their own datasets. After 10 rounds of FL, the global model's MSE decreased significantly, showing that collaborative learning helped refine the model more effectively than isolated training. FL enabled seamless knowledge sharing among distributed clients while ensuring data privacy. Both clients participated equally in FL rounds, contributing their locally trained model updates to the central server. Even though each client trained on its own dataset, the global model incorporated patterns from both clients, allowing each client to benefit from the other's data insights without ever accessing raw data.

Additionally, the FL-trained global model demonstrated superior generalization compared to locally trained models. While local models struggled with dataset-specific biases, limiting their ability to detect anomalies effectively across different conditions, the FL model learned a broader representation by integrating diverse data distributions from multiple clients, resulting in better performance across all local datasets. Finally, the FL framework provided enhanced model performance without compromising data security. Raw data never left the clients' environments, ensuring data confidentiality and compliance with privacy regulations. Clients only exchanged model parameters (weights, gradients) instead of full datasets, reducing communication overhead while achieving significant performance gains. FL can be further expanded to more clients without increasing privacy risks, making it a viable long-term solution for distributed learning.

References

[1] X. Li, L. Cheng, C. Sun, K.-Y. Lam, X. Wang, and F. Li, “Federated-Learning-Empowered Collaborative Data Sharing for Vehicular Edge Networks,” IEEE Network, vol. 35, no. 3, pp. 116–124, 2021.

[2] S. Savazzi, M. Nicoli, M. Bennis, S. Kianoush, and L. Barbieri, “Opportunities of Federated Learning in Connected, Cooperative, and Automated Industrial Systems,” IEEE Communications Magazine, vol. 59, no. 2, pp. 16–21, 2021.

[3] D. Gutierrez-Rojas et al., "Three-layer Approach to Detect Anomalies in Industrial Environments based on Machine Learning," 2020 IEEE Conference on Industrial Cyberphysical Systems (ICPS), Tampere, Finland, 2020, pp. 250-256, doi: 10.1109/ICPS48405.2020.9274780.

[4] W. Y. B. Lim, N. C. Luong, D. T. Hoang, Y. Jiao, Y. C. Liang, Q. Yang, D. Niyato, and C. Miao, “Federated Learning in Mobile Edge Networks: A Comprehensive Survey,” IEEE Communications Surveys & Tutorials, vol. 22, no. 3, pp. 2031–2063, 2020.

[5] R. Sedar, C. Kalalas, F. Vázquez-Gallego, L. Alonso and J. Alonso-Zarate, "A Comprehensive Survey of V2X Cybersecurity Mechanisms and Future Research Paths," in *IEEE Open Journal of the Communications Society*, vol. 4, pp. 325-391, 2023, doi: 10.1109/OJCOMS.2023.3239115.

[6] T. D. Nguyen, S. Marchal, M. Miettinen, H. Fereidooni, N. Asokan, and A. Sadeghi, “D¨IoT: A Federated Self-learning Anomaly Detection System for IoT,” in 2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS), 2019, pp. 756–767.

[7] D. Gutierrez-Rojas et al., "Detection and Classification of Anomalies in WSN-Enabled Cyber-Physical Systems," in IEEE Sensors Journal, vol. 25, no. 4, pp. 7193-7204, 15 Feb.15, 2025, doi: 10.1109/JSEN.2024.3520507.

[8] Z. Yu, J. Hu, G. Min, Z. Zhao, W. Miao, and M. S. Hossain, “Mobility-Aware Proactive Edge Caching for Connected Vehicles Using Federated Learning,” IEEE Transactions on Intelligent Transportation Systems, vol. 22, no. 8, pp. 5341–5351, 2021

[9] Jakub Konečný, H. Brendan McMahan, Daniel Ramage, Peter Richtárik, “Federated Optimization: Distributed Machine Learning for On-Device Intelligence”, arXiv preprint arXiv:1610.02527, 2016

[10] Taleb, Tarik, et al. "On multi-access edge computing: A survey of the emerging 5G network edge cloud architecture and orchestration." IEEE Communications Surveys & Tutorials 19.3 (2017): 1657-1681.

[11] Valente, Pedro, et al. "Edge-based IPFS in a Disaggregated Mobile Core." 2023 IEEE 24th International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM). IEEE, 2023.

[12] Shah, Syed Danial Ali, et al. "SDN-based service mobility management in MEC-enabled 5G and beyond vehicular networks." IEEE Internet of Things Journal 9 (15), (2022): 13425-13442.

[13] Li, Nan, et al. "Transforming the 5G RAN with innovation: The confluence of cloud native and intelligence." IEEE Access 11 (2023): 4443-4454.

[14] Queiroz, Rui, et al. "Container-based virtualization for real-time industrial systems—a systematic review." ACM Computing Surveys 56.3 (2023): 1-38.

[15] Zhao, Qianli, et al. "Customizable cloud-native infrastructure for private 5G." *2023 26th Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN)*. IEEE, 2023.

[16] M. Zakarya, et al. "ApMove: A Service Migration Technique for Connected and Autonomous Vehicles," in IEEE Internet of Things Journal, vol. 11, no. 17, pp. 28721-28733, 1 Sept.1, 2024

Appendix A Federated Server Logs

This appendix contains the full log output of the FL server during the 10 rounds of training, showing model aggregation and evaluation steps.

Excerpt:

kubectl logs -l app=fl-server -f

INFO : Received initial parameters from one random client

2025-03-11 12:18:22,564 | INFO | Received initial parameters from one random client

INFO : Starting evaluation of initial global parameters

2025-03-11 12:18:22,564 | INFO | Starting evaluation of initial global parameters

INFO : Evaluation returned no results (`None`)

2025-03-11 12:18:22,564 | INFO | Evaluation returned no results (`None`)

INFO :

2025-03-11 12:18:22,564 | INFO |

INFO : [ROUND 1]

2025-03-11 12:18:22,564 | INFO | [ROUND 1]

INFO : configure_fit: strategy sampled 2 clients (out of 2)

2025-03-11 12:19:08,814 | INFO | configure_fit: strategy sampled 2 clients (out of 2)

INFO : aggregate_fit: received 2 results and 0 failures

2025-03-11 12:19:09,046 | INFO | aggregate_fit: received 2 results and 0 failures

WARNING : No fit_metrics_aggregation_fn provided

2025-03-11 12:19:09,049 | WARNING | No fit_metrics_aggregation_fn provided

INFO : configure_evaluate: strategy sampled 2 clients (out of 2)

2025-03-11 12:19:09,049 | INFO | configure_evaluate: strategy sampled 2 clients (out of 2)

INFO : aggregate_evaluate: received 2 results and 0 failures

2025-03-11 12:19:09,065 | INFO | aggregate_evaluate: received 2 results and 0 failures

2025-03-11 12:19:09,065 | INFO | [Server] Métricas individuales de los clientes en esta ronda:

2025-03-11 12:19:09,065 | INFO | Cliente 1: Ejemplos=429, MSE=0.9921

2025-03-11 12:19:09,065 | INFO | Cliente 2: Ejemplos=1113, MSE=0.9880

2025-03-11 12:19:09,065 | INFO | [Server] MSE Loss Global Agregado: 0.9891

INFO :

2025-03-11 12:19:09,065 | INFO |

INFO : [ROUND 2]

2025-03-11 12:19:09,065 | INFO | [ROUND 2]

INFO : configure_fit: strategy sampled 2 clients (out of 2)

2025-03-11 12:19:09,066 | INFO | configure_fit: strategy sampled 2 clients (out of 2)

INFO : aggregate_fit: received 2 results and 0 failures

2025-03-11 12:19:09,201 | INFO | aggregate_fit: received 2 results and 0 failures

```
INFO :      configure_evaluate: strategy sampled 2 clients (out of 2)
2025-03-11 12:19:09,205 | INFO | configure_evaluate: strategy sampled 2 clients (out of 2)
INFO :      aggregate_evaluate: received 2 results and 0 failures
2025-03-11 12:19:09,241 | INFO | aggregate_evaluate: received 2 results and 0 failures
2025-03-11 12:19:09,243 | INFO | [Server] Métricas individuales de los clientes en esta ronda:
2025-03-11 12:19:09,244 | INFO |     Cliente 1: Ejemplos=429, MSE=0.9802
2025-03-11 12:19:09,244 | INFO |     Cliente 2: Ejemplos=1113, MSE=0.9740
2025-03-11 12:19:09,245 | INFO | [Server] MSE Loss Global Agregado: 0.9757
INFO :
2025-03-11 12:19:09,247 | INFO |
INFO :      [ROUND 3]
2025-03-11 12:19:09,247 | INFO | [ROUND 3]
INFO :      configure_fit: strategy sampled 2 clients (out of 2)
2025-03-11 12:19:09,247 | INFO | configure_fit: strategy sampled 2 clients (out of 2)
INFO :      aggregate_fit: received 2 results and 0 failures
2025-03-11 12:19:09,361 | INFO | aggregate_fit: received 2 results and 0 failures
INFO :      configure_evaluate: strategy sampled 2 clients (out of 2)
2025-03-11 12:19:09,363 | INFO | configure_evaluate: strategy sampled 2 clients (out of 2)
INFO :      aggregate_evaluate: received 2 results and 0 failures
2025-03-11 12:19:09,376 | INFO | aggregate_evaluate: received 2 results and 0 failures
2025-03-11 12:19:09,376 | INFO | [Server] Métricas individuales de los clientes en esta ronda:
2025-03-11 12:19:09,377 | INFO |     Cliente 1: Ejemplos=429, MSE=0.9600
2025-03-11 12:19:09,377 | INFO |     Cliente 2: Ejemplos=1113, MSE=0.9522
2025-03-11 12:19:09,377 | INFO | [Server] MSE Loss Global Agregado: 0.9543
INFO :
2025-03-11 12:19:09,377 | INFO |
INFO :      [ROUND 4]
2025-03-11 12:19:09,377 | INFO | [ROUND 4]
INFO :      configure_fit: strategy sampled 2 clients (out of 2)
2025-03-11 12:19:09,377 | INFO | configure_fit: strategy sampled 2 clients (out of 2)
INFO :      aggregate_fit: received 2 results and 0 failures
2025-03-11 12:19:09,508 | INFO | aggregate_fit: received 2 results and 0 failures
INFO :      configure_evaluate: strategy sampled 2 clients (out of 2)
2025-03-11 12:19:09,511 | INFO | configure_evaluate: strategy sampled 2 clients (out of 2)
INFO :      aggregate_evaluate: received 2 results and 0 failures
2025-03-11 12:19:09,523 | INFO | aggregate_evaluate: received 2 results and 0 failures
2025-03-11 12:19:09,524 | INFO | [Server] Métricas individuales de los clientes en esta ronda:
```

2025-03-11 12:19:09,524 | INFO | Cliente 1: Ejemplos=429, MSE=0.9277

2025-03-11 12:19:09,524 | INFO | Cliente 2: Ejemplos=1113, MSE=0.9206

2025-03-11 12:19:09,524 | INFO | [Server] MSE Loss Global Agregado: 0.9226

INFO :

2025-03-11 12:19:09,524 | INFO |

INFO :   [ROUND 5]

2025-03-11 12:19:09,524 | INFO | [ROUND 5]

INFO :   configure_fit: strategy sampled 2 clients (out of 2)

2025-03-11 12:19:09,524 | INFO | configure_fit: strategy sampled 2 clients (out of 2)

INFO :   aggregate_fit: received 2 results and 0 failures

2025-03-11 12:19:09,626 | INFO | aggregate_fit: received 2 results and 0 failures

INFO :   configure_evaluate: strategy sampled 2 clients (out of 2)

2025-03-11 12:19:09,631 | INFO | configure_evaluate: strategy sampled 2 clients (out of 2)

INFO :   aggregate_evaluate: received 2 results and 0 failures

2025-03-11 12:19:09,645 | INFO | aggregate_evaluate: received 2 results and 0 failures

2025-03-11 12:19:09,647 | INFO | [Server] Métricas individuales de los clientes en esta ronda:

2025-03-11 12:19:09,652 | INFO | Cliente 1: Ejemplos=429, MSE=0.8795

2025-03-11 12:19:09,653 | INFO | Cliente 2: Ejemplos=1113, MSE=0.8772

2025-03-11 12:19:09,653 | INFO | [Server] MSE Loss Global Agregado: 0.8779

INFO :

2025-03-11 12:19:09,653 | INFO |

INFO :   [ROUND 6]

2025-03-11 12:19:09,653 | INFO | [ROUND 6]

INFO :   configure_fit: strategy sampled 2 clients (out of 2)

2025-03-11 12:19:09,653 | INFO | configure_fit: strategy sampled 2 clients (out of 2)

INFO :   aggregate_fit: received 2 results and 0 failures

2025-03-11 12:19:09,908 | INFO | aggregate_fit: received 2 results and 0 failures

INFO :   configure_evaluate: strategy sampled 2 clients (out of 2)

2025-03-11 12:19:09,911 | INFO | configure_evaluate: strategy sampled 2 clients (out of 2)

INFO :   aggregate_evaluate: received 2 results and 0 failures

2025-03-11 12:19:09,930 | INFO | aggregate_evaluate: received 2 results and 0 failures

2025-03-11 12:19:09,931 | INFO | [Server] Métricas individuales de los clientes en esta ronda:

2025-03-11 12:19:09,931 | INFO | Cliente 1: Ejemplos=1113, MSE=0.8224

2025-03-11 12:19:09,931 | INFO | Cliente 2: Ejemplos=429, MSE=0.8212

2025-03-11 12:19:09,931 | INFO | [Server] MSE Loss Global Agregado: 0.8220

INFO :

2025-03-11 12:19:09,931 | INFO |

INFO : [ROUND 7]

2025-03-11 12:19:09,931 | INFO | [ROUND 7]

INFO : configure_fit: strategy sampled 2 clients (out of 2)

2025-03-11 12:19:09,931 | INFO | configure_fit: strategy sampled 2 clients (out of 2)

INFO : aggregate_fit: received 2 results and 0 failures

2025-03-11 12:19:10,118 | INFO | aggregate_fit: received 2 results and 0 failures

INFO : configure_evaluate: strategy sampled 2 clients (out of 2)

2025-03-11 12:19:10,123 | INFO | configure_evaluate: strategy sampled 2 clients (out of 2)

INFO : aggregate_evaluate: received 2 results and 0 failures

2025-03-11 12:19:10,148 | INFO | aggregate_evaluate: received 2 results and 0 failures

2025-03-11 12:19:10,148 | INFO | [Server] Métricas individuales de los clientes en esta ronda:

2025-03-11 12:19:10,148 | INFO | Cliente 1: Ejemplos=1113, MSE=0.7600

2025-03-11 12:19:10,148 | INFO | Cliente 2: Ejemplos=429, MSE=0.7615

2025-03-11 12:19:10,148 | INFO | [Server] MSE Loss Global Agregado: 0.7604

INFO :

2025-03-11 12:19:10,148 | INFO |

INFO : [ROUND 8]

2025-03-11 12:19:10,148 | INFO | [ROUND 8]

INFO : configure_fit: strategy sampled 2 clients (out of 2)

2025-03-11 12:19:10,148 | INFO | configure_fit: strategy sampled 2 clients (out of 2)

INFO : aggregate_fit: received 2 results and 0 failures

2025-03-11 12:19:10,284 | INFO | aggregate_fit: received 2 results and 0 failures

INFO : configure_evaluate: strategy sampled 2 clients (out of 2)

2025-03-11 12:19:10,288 | INFO | configure_evaluate: strategy sampled 2 clients (out of 2)

INFO : aggregate_evaluate: received 2 results and 0 failures

2025-03-11 12:19:10,310 | INFO | aggregate_evaluate: received 2 results and 0 failures

2025-03-11 12:19:10,321 | INFO | [Server] Métricas individuales de los clientes en esta ronda:

2025-03-11 12:19:10,321 | INFO | Cliente 1: Ejemplos=429, MSE=0.7052

2025-03-11 12:19:10,321 | INFO | Cliente 2: Ejemplos=1113, MSE=0.6904

2025-03-11 12:19:10,322 | INFO | [Server] MSE Loss Global Agregado: 0.6946

INFO :

2025-03-11 12:19:10,322 | INFO |

INFO : [ROUND 9]

2025-03-11 12:19:10,322 | INFO | [ROUND 9]

INFO : configure_fit: strategy sampled 2 clients (out of 2)

2025-03-11 12:19:10,323 | INFO | configure_fit: strategy sampled 2 clients (out of 2)

INFO : aggregate_fit: received 2 results and 0 failures

```
2025-03-11 12:19:10,482 | INFO | aggregate_fit: received 2 results and 0 failures
INFO :      configure_evaluate: strategy sampled 2 clients (out of 2)
2025-03-11 12:19:10,487 | INFO | configure_evaluate: strategy sampled 2 clients (out of 2)
INFO :      aggregate_evaluate: received 2 results and 0 failures
2025-03-11 12:19:10,517 | INFO | aggregate_evaluate: received 2 results and 0 failures
2025-03-11 12:19:10,520 | INFO | [Server] Métricas individuales de los clientes en esta ronda:
2025-03-11 12:19:10,520 | INFO |     Cliente 1: Ejemplos=429, MSE=0.6479
2025-03-11 12:19:10,520 | INFO |     Cliente 2: Ejemplos=1113, MSE=0.6116
2025-03-11 12:19:10,520 | INFO | [Server] MSE Loss Global Agregado: 0.6217
INFO :
2025-03-11 12:19:10,520 | INFO |
INFO :      [ROUND 10]
2025-03-11 12:19:10,520 | INFO | [ROUND 10]
INFO :      configure_fit: strategy sampled 2 clients (out of 2)
2025-03-11 12:19:10,520 | INFO | configure_fit: strategy sampled 2 clients (out of 2)
INFO :      aggregate_fit: received 2 results and 0 failures
2025-03-11 12:19:10,662 | INFO | aggregate_fit: received 2 results and 0 failures
INFO :      configure_evaluate: strategy sampled 2 clients (out of 2)
2025-03-11 12:19:10,674 | INFO | configure_evaluate: strategy sampled 2 clients (out of 2)
INFO :      aggregate_evaluate: received 2 results and 0 failures
2025-03-11 12:19:10,708 | INFO | aggregate_evaluate: received 2 results and 0 failures
2025-03-11 12:19:10,708 | INFO | [Server] Métricas individuales de los clientes en esta ronda:
2025-03-11 12:19:10,708 | INFO |     Cliente 1: Ejemplos=1113, MSE=0.5292
2025-03-11 12:19:10,708 | INFO |     Cliente 2: Ejemplos=429, MSE=0.5916
2025-03-11 12:19:10,709 | INFO | [Server] MSE Loss Global Agregado: 0.5465
INFO :
2025-03-11 12:19:10,709 | INFO |
INFO :      [SUMMARY]
2025-03-11 12:19:10,709 | INFO | [SUMMARY]
INFO :      Run finished 10 round(s) in 48.14s
2025-03-11 12:19:10,709 | INFO | Run finished 10 round(s) in 48.14s
INFO :          History (loss, distributed):
2025-03-11 12:19:10,709 | INFO |        History (loss, distributed):
INFO :                  round 1: 0.9891064766548023
2025-03-11 12:19:10,710 | INFO |                round 1: 0.9891064766548023
INFO :                  round 2: 0.9757183024855439
2025-03-11 12:19:10,710 | INFO |                round 2: 0.9757183024855439
```

INFO : round 3: 0.9543428694691639

2025-03-11 12:19:10,710 | INFO | round 3: 0.9543428694691639

INFO : round 4: 0.9225645219537535

2025-03-11 12:19:10,710 | INFO | round 4: 0.9225645219537535

INFO : round 5: 0.8778633264252184

2025-03-11 12:19:10,710 | INFO | round 5: 0.8778633264252184

INFO : round 6: 0.8220354211933418

2025-03-11 12:19:10,710 | INFO | round 6: 0.8220354211933418

INFO : round 7: 0.7604289689184627

2025-03-11 12:19:10,710 | INFO | round 7: 0.7604289689184627

INFO : round 8: 0.6945559690898494

2025-03-11 12:19:10,710 | INFO | round 8: 0.6945559690898494

INFO : round 9: 0.6217126347675398

2025-03-11 12:19:10,710 | INFO | round 9: 0.6217126347675398

INFO : round 10: 0.5465218512928439

2025-03-11 12:19:10,710 | INFO | round 10: 0.5465218512928439

INFO : History (metrics, distributed, evaluate):

2025-03-11 12:19:10,710 | INFO | History (metrics, distributed, evaluate):

INFO : {'mse_loss': [(1, 0.9891064766548023),

2025-03-11 12:19:10,711 | INFO | {'mse_loss': [(1, 0.9891064766548023),

INFO : (2, 0.9757183024855439),

2025-03-11 12:19:10,711 | INFO | (2, 0.9757183024855439),

INFO : (3, 0.9543428694691639),

2025-03-11 12:19:10,711 | INFO | (3, 0.9543428694691639),

INFO : (4, 0.9225645219537535),

2025-03-11 12:19:10,711 | INFO | (4, 0.9225645219537535),

INFO : (5, 0.8778633264252184),

2025-03-11 12:19:10,711 | INFO | (5, 0.8778633264252184),

INFO : (6, 0.8220354211933418),

2025-03-11 12:19:10,711 | INFO | (6, 0.8220354211933418),

INFO : (7, 0.7604289689184627),

2025-03-11 12:19:10,711 | INFO | (7, 0.7604289689184627),

INFO : (8, 0.6945559690898494),

2025-03-11 12:19:10,711 | INFO | (8, 0.6945559690898494),

INFO : (9, 0.6217126347675398),

2025-03-11 12:19:10,711 | INFO | (9, 0.6217126347675398),

INFO : (10, 0.5465218512928439)]}

2025-03-11 12:19:10,713 | INFO | (10, 0.5465218512928439)]}

INFO :

2025-03-11 12:19:10,713 | INFO |

2025-03-11 12:19:10,722 | INFO | ********** RESUMEN FINAL DEL ENTRENAMIENTO **********

2025-03-11 12:19:10,722 | INFO | Ronda 1: MSE Global = 0.9891

2025-03-11 12:19:10,722 | INFO | Ronda 2: MSE Global = 0.9757

2025-03-11 12:19:10,722 | INFO | Ronda 3: MSE Global = 0.9543

2025-03-11 12:19:10,722 | INFO | Ronda 4: MSE Global = 0.9226

2025-03-11 12:19:10,722 | INFO | Ronda 5: MSE Global = 0.8779

2025-03-11 12:19:10,722 | INFO | Ronda 6: MSE Global = 0.8220

2025-03-11 12:19:10,722 | INFO | Ronda 7: MSE Global = 0.7604

2025-03-11 12:19:10,722 | INFO | Ronda 8: MSE Global = 0.6946

2025-03-11 12:19:10,722 | INFO | Ronda 9: MSE Global = 0.6217

2025-03-11 12:19:10,722 | INFO | Ronda 10: MSE Global = 0.5465

2025-03-11 12:19:10,722 | INFO | **

2025-03-11 2:19:10,722 | INFO | Servidor FL ejecutándose. A la espera de clientes...

Appendix B Appendix B: FL Client Logs (Worker 1)

Detailed logs of FL Client Worker 1, including local AE training before FL rounds and participation in each federated round.

Excerpt:

kubectl logs -l app=fl-client-worker1 -f

2025-03-11 12:19:07,281 | INFO | [Init] Cliente ID: da72caf7-fdd8-4294-95c1-e42d29c42b62

2025-03-11 12:19:07,284 | INFO | Conexión a Redis exitosa.

2025-03-11 12:19:08,687 | INFO | [Local] Entrenamiento previo del modelo local...

2025-03-11 12:19:08,701 | INFO | [Local Training] Epoch 1/20: Loss: 1.0193

2025-03-11 12:19:08,704 | INFO | [Local Training] Epoch 2/20: Loss: 1.0166

2025-03-11 12:19:08,706 | INFO | [Local Training] Epoch 3/20: Loss: 1.0140

2025-03-11 12:19:08,707 | INFO | [Local Training] Epoch 4/20: Loss: 1.0116

2025-03-11 12:19:08,709 | INFO | [Local Training] Epoch 5/20: Loss: 1.0094

2025-03-11 12:19:08,712 | INFO | [Local Training] Epoch 6/20: Loss: 1.0072

2025-03-11 12:19:08,714 | INFO | [Local Training] Epoch 7/20: Loss: 1.0052

2025-03-11 12:19:08,718 | INFO | [Local Training] Epoch 8/20: Loss: 1.0033

2025-03-11 12:19:08,721 | INFO | [Local Training] Epoch 9/20: Loss: 1.0015

2025-03-11 12:19:08,726 | INFO | [Local Training] Epoch 10/20: Loss: 0.9998

2025-03-11 12:19:08,728 | INFO | [Local Training] Epoch 11/20: Loss: 0.9980

2025-03-11 12:19:08,729 | INFO | [Local Training] Epoch 12/20: Loss: 0.9964

2025-03-11 12:19:08,732 | INFO | [Local Training] Epoch 13/20: Loss: 0.9947

2025-03-11 12:19:08,739 | INFO | [Local Training] Epoch 14/20: Loss: 0.9931

2025-03-11 12:19:08,741 | INFO | [Local Training] Epoch 15/20: Loss: 0.9915

2025-03-11 12:19:08,745 | INFO | [Local Training] Epoch 16/20: Loss: 0.9898

2025-03-11 12:19:08,748 | INFO | [Local Training] Epoch 17/20: Loss: 0.9882

2025-03-11 12:19:08,751 | INFO | [Local Training] Epoch 18/20: Loss: 0.9865

2025-03-11 12:19:08,754 | INFO | [Local Training] Epoch 19/20: Loss: 0.9848

2025-03-11 12:19:08,757 | INFO | [Local Training] Epoch 20/20: Loss: 0.9830

2025-03-11 12:19:08,763 | INFO | [Local] MSE Local Baseline: 0.9812

2025-03-11 12:19:08,763 | INFO | [FL] Iniciando entrenamiento federado...

2025-03-11 12:19:08,769 | INFO | Opened insecure gRPC connection (no certificates were passed)

INFO flower 2025-03-11 12:19:08,769 | connection.py:102 | Opened insecure gRPC connection (no certificates were passed)

DEBUG flower 2025-03-11 12:19:08,773 | connection.py:39 | ChannelConnectivity.IDLE

2025-03-11 12:19:08,773 | DEBUG | ChannelConnectivity.IDLE

DEBUG flower 2025-03-11 12:19:08,781 | connection.py:39 | ChannelConnectivity.CONNECTING

2025-03-11 12:19:08,781 | DEBUG | ChannelConnectivity.CONNECTING

```
DEBUG flower 2025-03-11 12:19:08,812 | connection.py:39 | ChannelConnectivity.READY
2025-03-11 12:19:08,812 | DEBUG | ChannelConnectivity.READY
2025-03-11 12:19:08,821 | INFO | [FL Training] Epoch 1/10: Loss: 0.9994
2025-03-11 12:19:08,898 | INFO | [FL Training] Epoch 2/10: Loss: 0.9986
2025-03-11 12:19:08,920 | INFO | [FL Training] Epoch 3/10: Loss: 0.9978
2025-03-11 12:19:08,943 | INFO | [FL Training] Epoch 4/10: Loss: 0.9970
2025-03-11 12:19:08,975 | INFO | [FL Training] Epoch 5/10: Loss: 0.9961
2025-03-11 12:19:08,978 | INFO | [FL Training] Epoch 6/10: Loss: 0.9952
2025-03-11 12:19:08,993 | INFO | [FL Training] Epoch 7/10: Loss: 0.9943
2025-03-11 12:19:08,995 | INFO | [FL Training] Epoch 8/10: Loss: 0.9933
2025-03-11 12:19:09,009 | INFO | [FL Training] Epoch 9/10: Loss: 0.9923
2025-03-11 12:19:09,011 | INFO | [FL Training] Epoch 10/10: Loss: 0.9912
2025-03-11 12:19:09,012 | INFO | [FL Training] Modelo entrenado.
2025-03-11 12:19:09,060 | INFO | [FL Evaluation] MSE: 0.9921
2025-03-11 12:19:09,109 | INFO | [FL Training] Epoch 1/10: Loss: 0.9921
2025-03-11 12:19:09,131 | INFO | [FL Training] Epoch 2/10: Loss: 0.9908
2025-03-11 12:19:09,154 | INFO | [FL Training] Epoch 3/10: Loss: 0.9895
2025-03-11 12:19:09,162 | INFO | [FL Training] Epoch 4/10: Loss: 0.9882
2025-03-11 12:19:09,171 | INFO | [FL Training] Epoch 5/10: Loss: 0.9868
2025-03-11 12:19:09,180 | INFO | [FL Training] Epoch 6/10: Loss: 0.9854
2025-03-11 12:19:09,183 | INFO | [FL Training] Epoch 7/10: Loss: 0.9838
2025-03-11 12:19:09,188 | INFO | [FL Training] Epoch 8/10: Loss: 0.9822
2025-03-11 12:19:09,194 | INFO | [FL Training] Epoch 9/10: Loss: 0.9805
2025-03-11 12:19:09,198 | INFO | [FL Training] Epoch 10/10: Loss: 0.9787
2025-03-11 12:19:09,198 | INFO | [FL Training] Modelo entrenado.
2025-03-11 12:19:09,238 | INFO | [FL Evaluation] MSE: 0.9802
2025-03-11 12:19:09,267 | INFO | [FL Training] Epoch 1/10: Loss: 0.9802
2025-03-11 12:19:09,291 | INFO | [FL Training] Epoch 2/10: Loss: 0.9782
2025-03-11 12:19:09,296 | INFO | [FL Training] Epoch 3/10: Loss: 0.9761
2025-03-11 12:19:09,299 | INFO | [FL Training] Epoch 4/10: Loss: 0.9739
2025-03-11 12:19:09,301 | INFO | [FL Training] Epoch 5/10: Loss: 0.9716
2025-03-11 12:19:09,303 | INFO | [FL Training] Epoch 6/10: Loss: 0.9691
2025-03-11 12:19:09,320 | INFO | [FL Training] Epoch 7/10: Loss: 0.9666
2025-03-11 12:19:09,350 | INFO | [FL Training] Epoch 8/10: Loss: 0.9639
2025-03-11 12:19:09,354 | INFO | [FL Training] Epoch 9/10: Loss: 0.9610
2025-03-11 12:19:09,357 | INFO | [FL Training] Epoch 10/10: Loss: 0.9580
2025-03-11 12:19:09,357 | INFO | [FL Training] Modelo entrenado.
```

2025-03-11 12:19:09,368 | INFO | [FL Evaluation] MSE: 0.9600

2025-03-11 12:19:09,393 | INFO | [FL Training] Epoch 1/10: Loss: 0.9600

2025-03-11 12:19:09,426 | INFO | [FL Training] Epoch 2/10: Loss: 0.9568

2025-03-11 12:19:09,438 | INFO | [FL Training] Epoch 3/10: Loss: 0.9535

2025-03-11 12:19:09,444 | INFO | [FL Training] Epoch 4/10: Loss: 0.9499

2025-03-11 12:19:09,455 | INFO | [FL Training] Epoch 5/10: Loss: 0.9462

2025-03-11 12:19:09,469 | INFO | [FL Training] Epoch 6/10: Loss: 0.9423

2025-03-11 12:19:09,476 | INFO | [FL Training] Epoch 7/10: Loss: 0.9382

2025-03-11 12:19:09,479 | INFO | [FL Training] Epoch 8/10: Loss: 0.9339

2025-03-11 12:19:09,495 | INFO | [FL Training] Epoch 9/10: Loss: 0.9294

2025-03-11 12:19:09,499 | INFO | [FL Training] Epoch 10/10: Loss: 0.9247

2025-03-11 12:19:09,499 | INFO | [FL Training] Modelo entrenado.

2025-03-11 12:19:09,518 | INFO | [FL Evaluation] MSE: 0.9277

2025-03-11 12:19:09,560 | INFO | [FL Training] Epoch 1/10: Loss: 0.9277

2025-03-11 12:19:09,565 | INFO | [FL Training] Epoch 2/10: Loss: 0.9229

2025-03-11 12:19:09,566 | INFO | [FL Training] Epoch 3/10: Loss: 0.9179

2025-03-11 12:19:09,581 | INFO | [FL Training] Epoch 4/10: Loss: 0.9127

2025-03-11 12:19:09,604 | INFO | [FL Training] Epoch 5/10: Loss: 0.9075

2025-03-11 12:19:09,608 | INFO | [FL Training] Epoch 6/10: Loss: 0.9020

2025-03-11 12:19:09,615 | INFO | [FL Training] Epoch 7/10: Loss: 0.8964

2025-03-11 12:19:09,617 | INFO | [FL Training] Epoch 8/10: Loss: 0.8907

2025-03-11 12:19:09,620 | INFO | [FL Training] Epoch 9/10: Loss: 0.8848

2025-03-11 12:19:09,623 | INFO | [FL Training] Epoch 10/10: Loss: 0.8787

2025-03-11 12:19:09,623 | INFO | [FL Training] Modelo entrenado.

2025-03-11 12:19:09,636 | INFO | [FL Evaluation] MSE: 0.8795

2025-03-11 12:19:09,661 | INFO | [FL Training] Epoch 1/10: Loss: 0.8795

2025-03-11 12:19:09,670 | INFO | [FL Training] Epoch 2/10: Loss: 0.8733

2025-03-11 12:19:09,729 | INFO | [FL Training] Epoch 3/10: Loss: 0.8669

2025-03-11 12:19:09,746 | INFO | [FL Training] Epoch 4/10: Loss: 0.8605

2025-03-11 12:19:09,797 | INFO | [FL Training] Epoch 5/10: Loss: 0.8540

2025-03-11 12:19:09,853 | INFO | [FL Training] Epoch 6/10: Loss: 0.8474

2025-03-11 12:19:09,895 | INFO | [FL Training] Epoch 7/10: Loss: 0.8406

2025-03-11 12:19:09,898 | INFO | [FL Training] Epoch 8/10: Loss: 0.8338

2025-03-11 12:19:09,901 | INFO | [FL Training] Epoch 9/10: Loss: 0.8270

2025-03-11 12:19:09,903 | INFO | [FL Training] Epoch 10/10: Loss: 0.8200

2025-03-11 12:19:09,903 | INFO | [FL Training] Modelo entrenado.

2025-03-11 12:19:09,916 | INFO | [FL Evaluation] MSE: 0.8212

2025-03-11 12:19:09,989 | INFO | [FL Training] Epoch 1/10: Loss: 0.8212

2025-03-11 12:19:10,005 | INFO | [FL Training] Epoch 2/10: Loss: 0.8140

2025-03-11 12:19:10,007 | INFO | [FL Training] Epoch 3/10: Loss: 0.8069

2025-03-11 12:19:10,023 | INFO | [FL Training] Epoch 4/10: Loss: 0.7998

2025-03-11 12:19:10,089 | INFO | [FL Training] Epoch 5/10: Loss: 0.7927

2025-03-11 12:19:10,094 | INFO | [FL Training] Epoch 6/10: Loss: 0.7857

2025-03-11 12:19:10,106 | INFO | [FL Training] Epoch 7/10: Loss: 0.7787

2025-03-11 12:19:10,111 | INFO | [FL Training] Epoch 8/10: Loss: 0.7718

2025-03-11 12:19:10,113 | INFO | [FL Training] Epoch 9/10: Loss: 0.7650

2025-03-11 12:19:10,116 | INFO | [FL Training] Epoch 10/10: Loss: 0.7582

2025-03-11 12:19:10,116 | INFO | [FL Training] Modelo entrenado.

2025-03-11 12:19:10,129 | INFO | [FL Evaluation] MSE: 0.7615

2025-03-11 12:19:10,154 | INFO | [FL Training] Epoch 1/10: Loss: 0.7615

2025-03-11 12:19:10,155 | INFO | [FL Training] Epoch 2/10: Loss: 0.7542

2025-03-11 12:19:10,157 | INFO | [FL Training] Epoch 3/10: Loss: 0.7470

2025-03-11 12:19:10,161 | INFO | [FL Training] Epoch 4/10: Loss: 0.7400

2025-03-11 12:19:10,168 | INFO | [FL Training] Epoch 5/10: Loss: 0.7330

2025-03-11 12:19:10,172 | INFO | [FL Training] Epoch 6/10: Loss: 0.7260

2025-03-11 12:19:10,186 | INFO | [FL Training] Epoch 7/10: Loss: 0.7190

2025-03-11 12:19:10,188 | INFO | [FL Training] Epoch 8/10: Loss: 0.7120

2025-03-11 12:19:10,193 | INFO | [FL Training] Epoch 9/10: Loss: 0.7050

2025-03-11 12:19:10,254 | INFO | [FL Training] Epoch 10/10: Loss: 0.6979

2025-03-11 12:19:10,254 | INFO | [FL Training] Modelo entrenado.

2025-03-11 12:19:10,301 | INFO | [FL Evaluation] MSE: 0.7052

2025-03-11 12:19:10,368 | INFO | [FL Training] Epoch 1/10: Loss: 0.7052

2025-03-11 12:19:10,376 | INFO | [FL Training] Epoch 2/10: Loss: 0.6976

2025-03-11 12:19:10,384 | INFO | [FL Training] Epoch 3/10: Loss: 0.6900

2025-03-11 12:19:10,387 | INFO | [FL Training] Epoch 4/10: Loss: 0.6825

2025-03-11 12:19:10,388 | INFO | [FL Training] Epoch 5/10: Loss: 0.6750

2025-03-11 12:19:10,399 | INFO | [FL Training] Epoch 6/10: Loss: 0.6675

2025-03-11 12:19:10,412 | INFO | [FL Training] Epoch 7/10: Loss: 0.6600

2025-03-11 12:19:10,441 | INFO | [FL Training] Epoch 8/10: Loss: 0.6524

2025-03-11 12:19:10,447 | INFO | [FL Training] Epoch 9/10: Loss: 0.6447

2025-03-11 12:19:10,452 | INFO | [FL Training] Epoch 10/10: Loss: 0.6369

2025-03-11 12:19:10,453 | INFO | [FL Training] Modelo entrenado.

2025-03-11 12:19:10,498 | INFO | [FL Evaluation] MSE: 0.6479

2025-03-11 12:19:10,533 | INFO | [FL Training] Epoch 1/10: Loss: 0.6479

2025-03-11 12:19:10,559 | INFO | [FL Training] Epoch 2/10: Loss: 0.6396

2025-03-11 12:19:10,570 | INFO | [FL Training] Epoch 3/10: Loss: 0.6314

2025-03-11 12:19:10,587 | INFO | [FL Training] Epoch 4/10: Loss: 0.6232

2025-03-11 12:19:10,591 | INFO | [FL Training] Epoch 5/10: Loss: 0.6151

2025-03-11 12:19:10,600 | INFO | [FL Training] Epoch 6/10: Loss: 0.6070

2025-03-11 12:19:10,610 | INFO | [FL Training] Epoch 7/10: Loss: 0.5990

2025-03-11 12:19:10,612 | INFO | [FL Training] Epoch 8/10: Loss: 0.5910

2025-03-11 12:19:10,628 | INFO | [FL Training] Epoch 9/10: Loss: 0.5830

2025-03-11 12:19:10,645 | INFO | [FL Training] Epoch 10/10: Loss: 0.5750

2025-03-11 12:19:10,647 | INFO | [FL Training] Modelo entrenado.

2025-03-11 12:19:10,682 | INFO | [FL Evaluation] MSE: 0.5916

DEBUG flower 2025-03-11 12:19:10,739 | connection.py:121 | gRPC channel closed

INFO flower 2025-03-11 12:19:10,739 | app.py:149 | Disconnect and shut down

2025-03-11 12:19:10,739 | DEBUG | gRPC channel closed

2025-03-11 12:19:10,739 | INFO | Disconnect and shut down

2025-03-11 12:19:10,758 | INFO | [FL Final] MSE del modelo global en datos locales: 0.5916

2025-03-11 12:19:10,759 | INFO | [Metrics Summary] Cliente ID: da72caf7-fdd8-4294-95c1-e42d29c42b62 | MSE Local: 0.9812 | MSE Global en Local: 0.5916

Appendix C Appendix C: FL Client Logs (Worker 2)

Detailed logs of FL Client Worker 2, including local AE training before FL and its participation in federated training rounds.

Excerpt:

kubectl logs -l app=fl-client-worker2 -f

2025-03-11 12:18:22,503 | INFO | [Local] MSE Local Baseline: 0.9971

2025-03-11 12:18:22,503 | INFO | [FL] Iniciando entrenamiento federado...

INFO flower 2025-03-11 12:18:22,509 | connection.py:102 | Opened insecure gRPC connection (no certificates were passed)

2025-03-11 12:18:22,509 | INFO | Opened insecure gRPC connection (no certificates were passed)

DEBUG flower 2025-03-11 12:18:22,512 | connection.py:39 | ChannelConnectivity.IDLE

2025-03-11 12:18:22,512 | DEBUG | ChannelConnectivity.IDLE

DEBUG flower 2025-03-11 12:18:22,514 | connection.py:39 | ChannelConnectivity.CONNECTING

2025-03-11 12:18:22,514 | DEBUG | ChannelConnectivity.CONNECTING

2025-03-11 12:18:22,559 | DEBUG | ChannelConnectivity.READY

2025-03-11 12:19:08,927 | INFO | [FL Training] Epoch 1/10: Loss: 0.9971tivity.READY

2025-03-11 12:19:08,982 | INFO | [FL Training] Epoch 2/10: Loss: 0.9962

2025-03-11 12:19:09,026 | INFO | [FL Training] Epoch 3/10: Loss: 0.9953

2025-03-11 12:19:09,030 | INFO | [FL Training] Epoch 4/10: Loss: 0.9944

2025-03-11 12:19:09,032 | INFO | [FL Training] Epoch 5/10: Loss: 0.9934

2025-03-11 12:19:09,034 | INFO | [FL Training] Epoch 6/10: Loss: 0.9925

2025-03-11 12:19:09,036 | INFO | [FL Training] Epoch 7/10: Loss: 0.9914

2025-03-11 12:19:09,038 | INFO | [FL Training] Epoch 8/10: Loss: 0.9903

2025-03-11 12:19:09,040 | INFO | [FL Training] Epoch 9/10: Loss: 0.9892

2025-03-11 12:19:09,043 | INFO | [FL Training] Epoch 10/10: Loss: 0.9880

2025-03-11 12:19:09,043 | INFO | [FL Training] Modelo entrenado.

2025-03-11 12:19:09,062 | INFO | [FL Evaluation] MSE: 0.9880

2025-03-11 12:19:09,081 | INFO | [FL Training] Epoch 1/10: Loss: 0.9880

2025-03-11 12:19:09,097 | INFO | [FL Training] Epoch 2/10: Loss: 0.9867

2025-03-11 12:19:09,114 | INFO | [FL Training] Epoch 3/10: Loss: 0.9854

2025-03-11 12:19:09,133 | INFO | [FL Training] Epoch 4/10: Loss: 0.9840

2025-03-11 12:19:09,136 | INFO | [FL Training] Epoch 5/10: Loss: 0.9826

2025-03-11 12:19:09,138 | INFO | [FL Training] Epoch 6/10: Loss: 0.9811

2025-03-11 12:19:09,139 | INFO | [FL Training] Epoch 7/10: Loss: 0.9796

2025-03-11 12:19:09,153 | INFO | [FL Training] Epoch 8/10: Loss: 0.9779

2025-03-11 12:19:09,155 | INFO | [FL Training] Epoch 9/10: Loss: 0.9762

2025-03-11 12:19:09,163 | INFO | [FL Training] Epoch 10/10: Loss: 0.9744

2025-03-11 12:19:09,164 | INFO | [FL Training] Modelo entrenado.

2025-03-11 12:19:09,213 | INFO | [FL Evaluation] MSE: 0.9740

2025-03-11 12:19:09,270 | INFO | [FL Training] Epoch 1/10: Loss: 0.9740

2025-03-11 12:19:09,284 | INFO | [FL Training] Epoch 2/10: Loss: 0.9721

2025-03-11 12:19:09,287 | INFO | [FL Training] Epoch 3/10: Loss: 0.9700

2025-03-11 12:19:09,304 | INFO | [FL Training] Epoch 4/10: Loss: 0.9679

2025-03-11 12:19:09,307 | INFO | [FL Training] Epoch 5/10: Loss: 0.9657

2025-03-11 12:19:09,322 | INFO | [FL Training] Epoch 6/10: Loss: 0.9634

2025-03-11 12:19:09,326 | INFO | [FL Training] Epoch 7/10: Loss: 0.9610

2025-03-11 12:19:09,329 | INFO | [FL Training] Epoch 8/10: Loss: 0.9584

2025-03-11 12:19:09,343 | INFO | [FL Training] Epoch 9/10: Loss: 0.9557

2025-03-11 12:19:09,346 | INFO | [FL Training] Epoch 10/10: Loss: 0.9530

2025-03-11 12:19:09,346 | INFO | [FL Training] Modelo entrenado.

2025-03-11 12:19:09,370 | INFO | [FL Evaluation] MSE: 0.9522

2025-03-11 12:19:09,391 | INFO | [FL Training] Epoch 1/10: Loss: 0.9522

2025-03-11 12:19:09,427 | INFO | [FL Training] Epoch 2/10: Loss: 0.9493

2025-03-11 12:19:09,433 | INFO | [FL Training] Epoch 3/10: Loss: 0.9463

2025-03-11 12:19:09,436 | INFO | [FL Training] Epoch 4/10: Loss: 0.9432

2025-03-11 12:19:09,442 | INFO | [FL Training] Epoch 5/10: Loss: 0.9400

2025-03-11 12:19:09,449 | INFO | [FL Training] Epoch 6/10: Loss: 0.9367

2025-03-11 12:19:09,452 | INFO | [FL Training] Epoch 7/10: Loss: 0.9332

2025-03-11 12:19:09,455 | INFO | [FL Training] Epoch 8/10: Loss: 0.9297

2025-03-11 12:19:09,465 | INFO | [FL Training] Epoch 9/10: Loss: 0.9260

2025-03-11 12:19:09,505 | INFO | [FL Training] Epoch 10/10: Loss: 0.9221

2025-03-11 12:19:09,505 | INFO | [FL Training] Modelo entrenado.

2025-03-11 12:19:09,520 | INFO | [FL Evaluation] MSE: 0.9206

2025-03-11 12:19:09,532 | INFO | [FL Training] Epoch 1/10: Loss: 0.9206

2025-03-11 12:19:09,534 | INFO | [FL Training] Epoch 2/10: Loss: 0.9165

2025-03-11 12:19:09,536 | INFO | [FL Training] Epoch 3/10: Loss: 0.9124

2025-03-11 12:19:09,537 | INFO | [FL Training] Epoch 4/10: Loss: 0.9081

2025-03-11 12:19:09,541 | INFO | [FL Training] Epoch 5/10: Loss: 0.9037

2025-03-11 12:19:09,562 | INFO | [FL Training] Epoch 6/10: Loss: 0.8991

2025-03-11 12:19:09,564 | INFO | [FL Training] Epoch 7/10: Loss: 0.8944

2025-03-11 12:19:09,569 | INFO | [FL Training] Epoch 8/10: Loss: 0.8896

2025-03-11 12:19:09,582 | INFO | [FL Training] Epoch 9/10: Loss: 0.8846

2025-03-11 12:19:09,598 | INFO | [FL Training] Epoch 10/10: Loss: 0.8795

2025-03-11 12:19:09,599 | INFO | [FL Training] Modelo entrenado.

2025-03-11 12:19:09,643 | INFO | [FL Evaluation] MSE: 0.8772

2025-03-11 12:19:09,774 | INFO | [FL Training] Epoch 1/10: Loss: 0.8772

2025-03-11 12:19:09,785 | INFO | [FL Training] Epoch 2/10: Loss: 0.8719

2025-03-11 12:19:09,791 | INFO | [FL Training] Epoch 3/10: Loss: 0.8664

2025-03-11 12:19:09,822 | INFO | [FL Training] Epoch 4/10: Loss: 0.8608

2025-03-11 12:19:09,828 | INFO | [FL Training] Epoch 5/10: Loss: 0.8551

2025-03-11 12:19:09,831 | INFO | [FL Training] Epoch 6/10: Loss: 0.8492

2025-03-11 12:19:09,844 | INFO | [FL Training] Epoch 7/10: Loss: 0.8432

2025-03-11 12:19:09,848 | INFO | [FL Training] Epoch 8/10: Loss: 0.8370

2025-03-11 12:19:09,869 | INFO | [FL Training] Epoch 9/10: Loss: 0.8307

2025-03-11 12:19:09,885 | INFO | [FL Training] Epoch 10/10: Loss: 0.8242

2025-03-11 12:19:09,885 | INFO | [FL Training] Modelo entrenado.

2025-03-11 12:19:09,917 | INFO | [FL Evaluation] MSE: 0.8224

2025-03-11 12:19:09,967 | INFO | [FL Training] Epoch 1/10: Loss: 0.8224

2025-03-11 12:19:09,970 | INFO | [FL Training] Epoch 2/10: Loss: 0.8158

2025-03-11 12:19:09,984 | INFO | [FL Training] Epoch 3/10: Loss: 0.8092

2025-03-11 12:19:10,013 | INFO | [FL Training] Epoch 4/10: Loss: 0.8025

2025-03-11 12:19:10,016 | INFO | [FL Training] Epoch 5/10: Loss: 0.7958

2025-03-11 12:19:10,030 | INFO | [FL Training] Epoch 6/10: Loss: 0.7890

2025-03-11 12:19:10,051 | INFO | [FL Training] Epoch 7/10: Loss: 0.7820

2025-03-11 12:19:10,062 | INFO | [FL Training] Epoch 8/10: Loss: 0.7750

2025-03-11 12:19:10,065 | INFO | [FL Training] Epoch 9/10: Loss: 0.7679

2025-03-11 12:19:10,068 | INFO | [FL Training] Epoch 10/10: Loss: 0.7607

2025-03-11 12:19:10,069 | INFO | [FL Training] Modelo entrenado.

2025-03-11 12:19:10,143 | INFO | [FL Evaluation] MSE: 0.7600

2025-03-11 12:19:10,166 | INFO | [FL Training] Epoch 1/10: Loss: 0.7600

2025-03-11 12:19:10,182 | INFO | [FL Training] Epoch 2/10: Loss: 0.7527

2025-03-11 12:19:10,198 | INFO | [FL Training] Epoch 3/10: Loss: 0.7453

2025-03-11 12:19:10,216 | INFO | [FL Training] Epoch 4/10: Loss: 0.7377

2025-03-11 12:19:10,234 | INFO | [FL Training] Epoch 5/10: Loss: 0.7301

2025-03-11 12:19:10,236 | INFO | [FL Training] Epoch 6/10: Loss: 0.7224

2025-03-11 12:19:10,254 | INFO | [FL Training] Epoch 7/10: Loss: 0.7145

2025-03-11 12:19:10,270 | INFO | [FL Training] Epoch 8/10: Loss: 0.7066

2025-03-11 12:19:10,275 | INFO | [FL Training] Epoch 9/10: Loss: 0.6986

2025-03-11 12:19:10,280 | INFO | [FL Training] Epoch 10/10: Loss: 0.6904

2025-03-11 12:19:10,281 | INFO | [FL Training] Modelo entrenado.

2025-03-11 12:19:10,306 | INFO | [FL Evaluation] MSE: 0.6904

2025-03-11 12:19:10,332 | INFO | [FL Training] Epoch 1/10: Loss: 0.6904

2025-03-11 12:19:10,334 | INFO | [FL Training] Epoch 2/10: Loss: 0.6821

2025-03-11 12:19:10,350 | INFO | [FL Training] Epoch 3/10: Loss: 0.6737

2025-03-11 12:19:10,403 | INFO | [FL Training] Epoch 4/10: Loss: 0.6652

2025-03-11 12:19:10,439 | INFO | [FL Training] Epoch 5/10: Loss: 0.6565

2025-03-11 12:19:10,457 | INFO | [FL Training] Epoch 6/10: Loss: 0.6477

2025-03-11 12:19:10,461 | INFO | [FL Training] Epoch 7/10: Loss: 0.6387

2025-03-11 12:19:10,468 | INFO | [FL Training] Epoch 8/10: Loss: 0.6296

2025-03-11 12:19:10,472 | INFO | [FL Training] Epoch 9/10: Loss: 0.6204

2025-03-11 12:19:10,478 | INFO | [FL Training] Epoch 10/10: Loss: 0.6110

2025-03-11 12:19:10,479 | INFO | [FL Training] Modelo entrenado.

2025-03-11 12:19:10,508 | INFO | [FL Evaluation] MSE: 0.6116

2025-03-11 12:19:10,539 | INFO | [FL Training] Epoch 1/10: Loss: 0.6116

2025-03-11 12:19:10,551 | INFO | [FL Training] Epoch 2/10: Loss: 0.6022

2025-03-11 12:19:10,565 | INFO | [FL Training] Epoch 3/10: Loss: 0.5927

2025-03-11 12:19:10,568 | INFO | [FL Training] Epoch 4/10: Loss: 0.5833

2025-03-11 12:19:10,588 | INFO | [FL Training] Epoch 5/10: Loss: 0.5739

2025-03-11 12:19:10,605 | INFO | [FL Training] Epoch 6/10: Loss: 0.5645

2025-03-11 12:19:10,622 | INFO | [FL Training] Epoch 7/10: Loss: 0.5551

2025-03-11 12:19:10,634 | INFO | [FL Training] Epoch 8/10: Loss: 0.5457

2025-03-11 12:19:10,642 | INFO | [FL Training] Epoch 9/10: Loss: 0.5363

2025-03-11 12:19:10,656 | INFO | [FL Training] Epoch 10/10: Loss: 0.5270

2025-03-11 12:19:10,658 | INFO | [FL Training] Modelo entrenado.

2025-03-11 12:19:10,702 | INFO | [FL Evaluation] MSE: 0.5292

2025-03-11 12:19:10,724 | DEBUG | gRPC channel closed

DEBUG flower 2025-03-11 12:19:10,724 | connection.py:121 | gRPC channel closed

2025-03-11 12:19:10,724 | INFO | Disconnect and shut down

INFO flower 2025-03-11 12:19:10,724 | app.py:149 | Disconnect and shut down

2025-03-11 12:19:10,729 | INFO | [FL Final] MSE del modelo global en datos locales: 0.5292

2025-03-11 12:19:10,730 | INFO | [Metrics Summary] Cliente ID: 938a2525-d7a2-4718-9915-d9920b001025 | MSE Local: 0.9971 | MSE Global en Local: 0.5292