IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.154114

A Multi-National Framework for Real-Time Collaborative Cyber Defense: Evaluating FL Architectures and Aggregation Strategies in Heterogeneous NIDS

Abhiram T Sajeev[1], Adarsh S J[2], Alen J S[3], Alfin Jerome[4], Amila A L[5]

Department of Computer Science and Engineering, College of Engineering Kottarakkara,

APJ Abdul Kalam Technological University, Kerala, India[1-5]

Abstract: The rapid evolution of cyber threats necessitates the development of sophisticated Machine Learning (ML) based Network Intrusion Detection Systems (NIDS). However, the efficacy of these systems is often hampered by the sensitive nature of network traffic and stringent privacy regulations, such as GDPR, which prevent organizations and nations from sharing raw data. To address this "privacy-security paradox," this paper presents a decentralized framework for collaborative threat intelligence utilizing Federated Learning (FL).We simulate a high-stakes multi-national scenario where three distinct nations collaboratively train a global NIDS model while maintaining data sovereignty. The testbed comprises physical nodes representing an aggregation server, a threat actor, and independent nations, with the latter further simulating diverse domestic sectors including Critical Infrastructure (SCADA/IIoT), Financial Services, and Tech Hubs to generate realistic, heterogeneous traffic.

Our research evaluates the performance of various local model architectures, comparing 1D-CNN, DNN, and Autoencoders for detecting complex patterns in network features. Furthermore, we conduct a comparative analysis of aggregation algorithms to mitigate challenges posed by non-IID data. Experimental results demonstrate that the collaborative global model achieves significantly higher detection accuracy than isolated systems. This work provides evidence that Federated Learning is a viable framework for privacy-preserving network security.

Keywords: Federated Learning, NIDS, FedProx, 1D-CNN, Multi-National Cyber Defense, Privacy-Preserving AI, CIC-IDS2017.

I. INTRODUCTION

In the modern digital era, the proliferation of interconnected devices, cloud computing, and the Internet of Things (IoT) has fundamentally transformed the way data is generated and exchanged. However, this unprecedented connectivity has also expanded the attack surface of computer networks, making them increasingly vulnerable to sophisticated cyber threats. Malicious actors continuously develop new techniques to infiltrate networks, exfiltrate data, and disrupt services, leading to financial losses, compromised privacy, and operational paralysis. To counter these threats, organizations deploy Network Intrusion Detection Systems (NIDS), which act as vigilant monitors that inspect network traffic and alert administrators to suspicious or malicious activities. Traditional NIDS technologies are primarily of two types: signature-based systems, which rely on pre-defined attack signatures, and anomaly-based systems, which identify deviations from normal traffic behavior. In recent years, machine learning (ML) and deep learning (DL) have revolutionized anomaly-based intrusion detection by enabling systems to learn complex attack patterns automatically from large datasets. These intelligent detection models have demonstrated superior accuracy and adaptability compared to manually crafted rule-based systems. Nevertheless, their effectiveness depends heavily on the availability of large, diverse, and representative datasets encompassing various attack types and benign traffic samples. Unfortunately, acquiring and centralizing such data is a significant challenge. Network traffic often contains confidential or personally identifiable information, such as IP addresses, user credentials, and proprietary business communications. Regulatory frameworks like the General Data Protection Regulation (GDPR) and organizational policies impose strict limitations on data sharing. Furthermore, companies are reluctant to share their internal traffic logs with third parties due to privacy risks, competitive concerns, and the potential exposure of vulnerabilities. As a result, despite the promising capabilities of ML and DL, the lack of data sharing hinders the development of robust, generalizable intrusion detection models. This conflict between the need for data diversity and the obligation to protect privacy has prompted researchers to seek privacy-preserving and decentralized learning paradigms that can reconcile these competing demands. One of the most

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.154114

promising solutions to emerge in recent years is Federated Learning (FL), a distributed machine learning approach that allows multiple entities to collaboratively train a shared model without exchanging raw data

A. *What is Federated Learning ?*

Federated Learning (FL) is a decentralized approach to machine learning that enables multiple clients—such as organizations, data centers, or even individual devices to collaboratively train a shared model without exchanging raw data.

The concept was first introduced by Google in 2016 to train predictive text models on Android devices without compromising user privacy. Since then, it has been adopted across diverse fields, including healthcare, finance, and cybersecurity. In a typical Federated Learning workflow, a central server initializes a global model and sends it to a set of participating clients. Each client trains this model locally using its private dataset, improving it through several local epochs. Once local training is complete, clients send only their updated model parameters to the central server. The server then aggregates these updates—most commonly using the Federated Averaging (FedAvg) algorithm to produce a new global model that represents the combined knowledge of all clients. This iterative process continues until the model converges, meaning its performance stabilizes across all participants. The principal advantage of this approach is that data never leaves its source. Clients keep their sensitive datasets local, and only abstract model information is shared. This makes FL highly suitable for domains where data privacy, confidentiality, or ownership are of critical concern. In addition, FL provides scalability, allowing a large number of clients to participate simultaneously in model training. Despite its many strengths, Federated Learning also faces challenges. Non-IID data can slow down convergence or reduce model accuracy, as each client's data distribution may differ significantly from others. Communication efficiency is another major concern since FL involves frequent exchanges of large model parameters. Furthermore, malicious clients can launch model poisoning or backdoor attacks, compromising the integrity of the global model. To mitigate such risks, researchers are exploring secure aggregation protocols, differential privacy mechanisms, and robust optimization techniques. When applied to Network Intrusion Detection Systems, Federated Learning offers a compelling solution to the problem of collaborative learning across private networks. Each organization can train its local model using internal traffic logs, while the global model, updated by aggregating knowledge from all participants, becomes more effective at identifying diverse attack patterns. This approach enhances collective defense capabilities without sacrificing the confidentiality of sensitive network data.

B. *Need for a Privacy-Preserving Collaborative Approach*

Federated Learning represents a paradigm shift in how collaborative intelligence can be achieved in the digital world. Instead of centralizing sensitive data in one location, FL enables multiple organizations or devices to participate in a joint training process by sending only model parameters or gradients to a central coordinating server. Each participant, or "client," trains the model locally using its own dataset and transmits the learned updates to the central aggregator, which computes a global model through an averaging process. By doing so, organizations can contribute to a shared security model while retaining full control of their data. In the context of network security, this paradigm is particularly valuable. Cyber attacks often manifest differently across networks, and no single organization has visibility into every possible attack pattern. FL allows diverse participants such as enterprises, internet service providers, or critical infrastructure operators to collaboratively learn from distributed experiences, resulting in a more robust and generalizable intrusion detection model. Moreover, since raw data never leaves local environments, FL inherently mitigates the risk of data leakage and complies with stringent privacy regulations. Despite these advantages, the practical deployment of Federated Learning in cybersecurity is still in its infancy. Most existing FL-based NIDS implementations are designed for hardware oriented environments, such as IoT devices or embedded sensors, where data is naturally distributed. While these studies have demonstrated feasibility, they face significant challenges in terms of scalability, computational overhead, and reproducibility. Setting up multiple physical clients is resource-intensive and makes experimentation cumbersome. Hence, there is a pressing need to develop a software-based Federated Learning simulation that can replicate real-world conditions in a virtualized environment, enabling controlled, flexible, and scalable research.

C. *Problem Statement*

Traditional machine learning-based intrusion detection systems rely on centralized architectures that aggregate all training data in a single repository. This approach, though straightforward, introduces serious privacy, scalability, and reliability issues. Centralized systems require organizations to transmit sensitive network data to external servers, creating potential vulnerabilities and violating data protection policies. Furthermore, as the volume of traffic data continues to grow, maintaining and processing such massive datasets centrally becomes computationally expensive and inefficient. Another major challenge lies in statistical heterogeneity, often referred to as non-independent and identically distributed (non-IID) data. In real-world scenarios, network traffic characteristics vary widely across organizations—each network may experience different types of attacks, usage behaviors, and traffic loads. As a result, when data from these diverse

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.154114

sources is combined in a centralized model, it can lead to biased learning and poor generalization to unseen attack types. Although Federated Learning provides a promising framework to address these issues, most implementations remain confined to limited-scale or hardware-based setups. There is a lack of flexible, software-driven environments that simulate multiple clients collaboratively training an intrusion detection model. This research seeks to fill that gap by designing and implementing a software-based Federated Learning framework for NIDS, enabling experimentation with non-IID data, privacy mechanisms, and performance evaluation under controlled conditions.

D. *Aim of the Study*

The primary aim of this research is to design, implement, and evaluate a Federated Learning framework for Network Intrusion Detection that ensures data privacy while maintaining high detection accuracy. The proposed system will be realized as a software simulation representing multiple clients or organizations, each with its own network traffic data. The framework will facilitate local training, secure model aggregation, and performance analysis, enabling researchers to investigate the dynamics of federated training under realistic data heterogeneity. Through this, the study aims to demonstrate that Federated Learning can achieve comparable or superior detection performance to traditional centralized systems while significantly enhancing privacy and scalability.

E. *Significance of the Research*

This research carries substantial significance for both academia and industry. From a practical standpoint, it provides a pathway for organizations to collaborate on improving cybersecurity without violating data-sharing restrictions. By simulating multiple clients in software, it allows researchers to study complex network scenarios without the logistical and financial constraints of deploying hardware-based systems. Academically, the work contributes to the growing field of privacy-preserving machine learning by empirically analyzing how federated architectures perform under varying degrees of non-IID data and privacy constraints. Moreover, integrating techniques such as Differential Privacy and Secure Aggregation adds a mathematically rigorous layer of protection against information leakage, advancing the broader goal of secure collaborative intelligence. Ultimately, this research strengthens the foundation for future deployments of privacy-aware, decentralized cybersecurity frameworks that can adapt to evolving threats in real time.

F. *Outline and Scope of the Project*

This project is structured to provide a systematic exploration of how Federated Learning can be effectively applied to network intrusion detection. The research begins with an extensive review of related literature to identify existing methods, algorithms, and open challenges. It then proceeds to the design and development of a software-simulated Federated Learning environment, where multiple clients represent independent organizations contributing to the training of a shared intrusion detection model. The framework will employ real-world network datasets such as CICIDS2017 or TON-IoT, which will be partitioned to reflect non-IID conditions. The system will be implemented using open-source frameworks such as TensorFlow Federated or Flower, enabling modular and extensible experimentation. Privacy-preserving techniques, including differential privacy and secure aggregation, will be integrated to strengthen data protection. The performance of the proposed FL-based system will be evaluated using standard metrics such as accuracy, precision, recall, and F1-score, and compared against a conventional centralized model. The results will be analyzed to understand the trade-offs between detection accuracy, communication overhead, and privacy guarantees. The scope of the project is limited to software simulation rather than real-time hardware deployment. This approach provides flexibility, reproducibility, and scalability, allowing for detailed experimentation and analysis in an academic setting. The findings are expected to guide future research toward real-world implementations of privacy-preserving federated intrusion detection systems.

G. *Summary*

In conclusion, this introductory chapter has outlined the motivation, background, and rationale behind employing Federated Learning for Network Intrusion Detection. It highlighted the challenges associated with centralized machine learning systems, especially regarding privacy and scalability, and introduced Federated Learning as a viable solution that balances collaboration with data confidentiality. By focusing on a software-simulated federated environment, this research seeks to provide a practical and reproducible framework that demonstrates how decentralized training can enhance cybersecurity intelligence. The following chapters will delve deeper into the specific research objectives, literature review, and methodological framework that drive this investigation.

II. LITERATURE SURVEY

The rapid digital transformation of modern infrastructure has created immense opportunities for innovation but has also introduced serious cybersecurity challenges. Traditional Network Intrusion Detection Systems (NIDS), while effective at identifying known attack patterns, struggle to adapt to new, evolving, and distributed threats. The application of machine

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.154114

learning (ML) and deep learning (DL) has significantly enhanced the capability of NIDS by enabling automatic pattern recognition and anomaly detection from network traffic data. However, these techniques generally depend on centralized data aggregation, where data from multiple sources are collected and stored in a central server for model training. This centralized approach, though efficient in terms of model optimization, introduces considerable privacy, scalability, and compliance concerns. To address these challenges, researchers have begun exploring Federated Learning (FL) — a decentralized, privacy-preserving machine learning paradigm that allows multiple clients to collaboratively train a shared model without exchanging raw data. Over the past few years, FL has attracted growing attention in cybersecurity research, particularly for intrusion detection applications. The literature reviewed in this section provides an in-depth understanding of how FL has been integrated into NIDS architectures, the algorithms and datasets used, and the open research challenges that persist.

A. *Federated Learning in Intrusion Detection*

Federated Learning was first introduced as a means to train shared models on decentralized data sources, and it has since been adopted in various domains such as healthcare, finance, and autonomous systems. In cybersecurity, the technique offers an innovative way to combine the intelligence of multiple organizations or IoT devices without violating data privacy. Instead of transferring sensitive network data to a central location, each client trains the model locally and only shares model parameters with a central aggregator, which updates the global model using algorithms such as Federated Averaging (FedAvg). The earliest explorations of FL in the intrusion detection domain focused on the Internet of Things (IoT) and Cyber-Physical Systems (CPS), where devices inherently operate in distributed environments. These studies laid the foundation for later works that expanded FL into broader, enterprise-scale NIDS applications.

B. *Review of Key Research Studies*

- **Foundational Privacy-Preserving Approaches**

Mahmud et al. (2024) laid important groundwork by applying the FedAvg aggregation algorithm across multiple deep learning architectures — DNN, LSTM, GRU, and LeNet — deployed on IoT devices. They secured model updates using SSL encryption and lightweight cryptography to keep overhead low, achieving over 90% accuracy. The key limitation was scope: only seven IoT sensors were used, making generalization uncertain. They recommended expanding to heterogeneous environments and incorporating semi-supervised learning to handle zero-day (previously unseen) attacks.

- **Optimization-Driven Design**

Albogami (2025) pushed accuracy further with the Federated Hybrid Deep Belief Network (FHDBN), which achieved over 98% accuracy by combining two bio-inspired optimization algorithms: Golden Jackal Optimization (GJO) for selecting the most informative network features, and Dung Beetle Optimization (DBO) for fine-tuning model hyperparameters. While impressive in controlled settings, real-time large-scale deployment was never tested, leaving computational scalability an open question.

- **Hybrid Detection Architecture**

Olanrewaju-George and Pranggono (2025) addressed a critical gap — detecting both known and unknown threats — by building a two-layer system. AutoEncoders (AE) handled anomaly detection unsupervised, while Deep Neural Networks (DNN) performed supervised attack classification. They used FedAvgM (a momentum-enhanced variant of FedAvg) to speed up convergence. The system performed well but struggled with heterogeneous IoT hardware and inconsistent network conditions.

- **Non-IID Data Challenges**

Belarbi et al. (2023) specifically studied the impact of non-IID (non-identically distributed) data — a fundamental real-world challenge where each client's data looks different. Using the TON-IoT dataset partitioned by IP address to simulate realistic client diversity, they benchmarked FedAvg, FedProx, and FedYogi, finding that federated models consistently underperformed centralized ones under high data heterogeneity. They proposed pre-training and adaptive aggregation as remedies, ideas that influenced several subsequent studies.

- **Efficient Transfer and Communication**

Jameel et al. (2025) developed TabFIDS, a CNN-based federated IDS that tackled communication inefficiency — a major bottleneck in FL deployments. They used Temporal Averaging and Data-Driven Feature Elimination (DDFE) to remove redundant features before transmission and introduced Block-Based Smart Aggregation (BBSA) to balance local model accuracy with global transferability. The trade-off was increased aggregation complexity and a need for careful tuning to avoid losing subtle attack signals.

- **Prototype-Based Heterogeneity Handling**

Chennoufi et al. (2025) proposed PROTEAN, which took a creative approach to non-IID data: instead of sharing raw model weights, clients shared class prototypes — compact mean embeddings representing each attack class. This aligned local and global representations without exposing full model parameters. However, the shared prototypes themselves could leak class-level information, prompting the authors to recommend integrating Differential Privacy (DP) as a

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.154114

safeguard.

- **Cryptographic Privacy Enhancement**

Correia et al. (2025) tackled privacy from a cryptographic angle by combining Hybrid Homomorphic Encryption (HHE) — specifically the PASTA symmetric cipher with the BFV homomorphic scheme — to allow the server to aggregate encrypted model updates without ever decrypting them. This provides strong mathematical privacy guarantees but at a significant cost: heavy server-side computation and complex key management, making it challenging to scale.

- **Benchmarking and Adversarial Robustness**

Buyuktanir et al. (2025) offered a systematic comparison of four FL algorithms (FedAvg, FedOpt, FedProx, FedYogi) across four major IDS datasets (CICIDS2017, TON-IoT, UNSW-NB15, Edge-IIoTset), confirming that non-IID data, untrusted clients, and communication overhead remain the top barriers to real-world deployment. Khalil et al. (2025) specifically tested adversarial robustness against label-flipping attacks across Random Forest, SVM, and Logistic Regression models, showing reasonable resilience but calling for stronger adaptive defenses.

- **Broader Landscape and Open Gaps**

Gutti et al. (2025) surveyed FL applications across IoT, healthcare, finance, and cybersecurity, highlighting the promise of hybrid models combining deep learning with GANs and bio-inspired systems. They identified three major field-wide gaps: the lack of standardized federated evaluation datasets, scalability limitations, and the high computational cost of privacy mechanisms like homomorphic encryption.

C. *Emerging Themes and Observations*

Across the reviewed literature, several recurring themes and insights emerge.

First, Federated Learning has proven effective in enabling privacy-preserving collaborative training without compromising much on detection accuracy. However, nearly all studies acknowledge the persistent issue of data heterogeneity, where the performance of FL significantly degrades under non-IID conditions. This is a critical concern for intrusion detection, as network environments inherently vary across organizations.

Second, there is a clear trade-off between privacy and utility. Techniques such as Differential Privacy and Homomorphic Encryption offer strong mathematical guarantees of confidentiality but can reduce model accuracy or increase computational latency. Similarly, while secure aggregation prevents the central server from inspecting individual updates, it also adds communication overhead that limits scalability.

Third, the evaluation environments of most studies are constrained by hardware resources. Many experiments use small numbers of IoT devices or simulated sensors, which do not reflect the complexity and scale of real enterprise networks. This limitation directly motivates the need for a software-simulated FL framework, where multiple clients can be emulated using virtual machines or containers, enabling flexible experimentation at scale.

Lastly, many studies point to the growing necessity of standardized benchmarks for FL-based NIDS. While datasets such as CICIDS2017, TON-IoT , and TON-IoT are commonly used, variations in preprocessing and data partitioning make direct comparisons difficult. A unified research framework that defines dataset partitioning, privacy mechanisms, and evaluation metrics would significantly improve reproducibility and accelerate progress in this domain.

D. *Summary*

This literature survey highlights the evolution of Federated Learning from a theoretical privacy preserving concept to a practical tool for collaborative network security. The reviewed works collectively emphasize its potential to revolutionize intrusion detection by enabling data-driven collaboration among distributed entities while preserving confidentiality. At the same time, they expose critical challenges—data heterogeneity, scalability, communication overhead, and privacy–performance trade-offs that must be addressed for real-world applicability. The insights gained from this review directly inform the design of the current research, which seeks to implement a software-simulated Federated Learning environment capable of evaluating these challenges in depth.

III. GAP OF THE WORK

The literature reviewed in the preceding chapter demonstrates that Federated Learning (FL) has emerged as a powerful paradigm for privacy-preserving collaborative model training, particularly in cybersecurity and network intrusion detection. Numerous researchers have successfully applied FL to distributed environments such as IoT networks, Cyber-Physical Systems (CPS), and edge computing frameworks. These studies have shown that FL can maintain strong detection performance while ensuring that sensitive data remains localized. However, despite these advancements,

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.154114

several limitations persist in existing approaches. Most prior work either focuses on theoretical modeling or is confined to hardware-centric environments, which limits the scalability, reproducibility, and adaptability of the research. This chapter identifies and discusses the major research gaps that motivate the present study. It highlights the technical, methodological, and experimental shortcomings in current Federated Learning-based NIDS frameworks and establishes the need for a software-based simulation platform that addresses these challenges effectively.

A. *Limitations in Existing Studies*

A recurring limitation observed across the surveyed literature is the dependence on hardware based implementations. Many Federated Learning experiments for intrusion detection are performed on IoT devices, edge nodes, or embedded systems, where the distributed nature of the hardware provides a natural test bed for FL. While these settings are useful for proof-of-concept demonstrations, they impose significant constraints on scalability and experimentation. Setting up multiple IoT devices, ensuring network connectivity, synchronizing updates, and maintaining security across physical hardware nodes are resource-intensive tasks. These requirements make such systems difficult to reproduce or extend in academic and research environments with limited infrastructure. Furthermore, hardware-centric FL systems often struggle with computational and energy constraints. IoT devices typically have limited processing power and storage capacity, which restricts the complexity of the models that can be trained locally. This limitation leads to the use of relatively shallow or lightweight models, which may not capture the intricate patterns present in large-scale network traffic. Consequently, while these studies demonstrate the conceptual feasibility of FL in intrusion detection, they fall short of showcasing its potential for training deep and complex models that could achieve state-of-the-art performance. Another critical gap is the lack of software-based simulation frameworks that can emulate federated environments in a flexible, reproducible, and cost-effective manner. Most studies rely on real devices to simulate clients, which restricts the number of participants and limits the scope of experimentation. In contrast, a software-simulated FL environment can represent dozens or even hundreds of clients using virtual machines or containers, allowing researchers to analyze system behavior under various configurations, network conditions, and data distributions. Such an environment would en- able controlled experimentation with parameters such as communication frequency, aggregation algorithms, and privacy mechanisms—something that is impractical in hardware setups.

B. *Challenges Related to Data Heterogeneity and Privacy*

Another major gap identified in the literature is the insufficient handling of statistical heterogeneity, or non-IID data. In real-world scenarios, data distributions across clients are rarely identical—some networks may predominantly experience denial-of-service (DoS) attacks, while others encounter phishing or botnet activities. This imbalance can significantly degrade the performance of federated models, as traditional algorithms like FedAvg assume data homogeneity. Although some works have proposed algorithms such as FedProx and FedYogi to mitigate this issue, comprehensive evaluations of these methods on realistic non-IID intrusion datasets remain limited. Additionally, while many researchers incorporate privacy-enhancing technologies such as Differential Privacy (DP), Secure Aggregation, or Homomorphic Encryption, most implementations focus solely on the theoretical side or on small-scale test beds. Few studies have analyzed the quantitative trade-offs between the level of privacy protection and model performance in a systematic, software-based environment. Understanding this trade-off is essential because stronger privacy mechanisms often introduce noise or computational overhead that can reduce accuracy or slow convergence. Hence, there is a pressing need to develop a flexible experimental framework that allows researchers to explore and optimize this balance.

C. *Lack of Standardized Evaluation and Benchmarking*

The literature also reveals a lack of standardized evaluation frameworks for Federated Learning based intrusion detection. While public datasets such as CICIDS2017, TON-IoT , and TONIoT are commonly used, each study partitions and preprocesses data differently, making cross study comparisons difficult. Similarly, performance metrics and experimental settings vary widely—some focus on detection accuracy, while others emphasize communication cost or privacy levels. This fragmentation hinders the establishment of consistent benchmarks that could guide future research. Moreover, most existing studies assess their models in limited-scale environments with a small number of clients, failing to explore the implications of scaling Federated Learning to dozens or hundreds of participants. Real-world federated systems, especially in cybersecurity, involve large-scale collaboration among diverse entities such as corporations, ISPs, and government agencies. A software-simulated FL platform could effectively bridge this gap by allowing experiments at scale and establishing common evaluation baselines.

D. *Inadequate Study of Adversarial Threats*

While FL inherently enhances privacy, it also introduces new security vulnerabilities. Malicious clients can perform model poisoning or backdoor attacks by submitting corrupted updates to manipulate the global model. However, only a few studies in the literature explicitly test the robustness of FL-based NIDS against such adversarial scenarios. Most existing frameworks assume trusted clients and benign participation, which does not reflect real-world adversarial

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.154114

conditions. Evaluating FL's resilience against these attacks in a software-simulated environment would yield crucial insights for developing secure aggregation and robust optimization techniques. Another under explored dimension is communication robustness. Hardware-based systems are often tested under stable local network conditions, whereas large-scale deployments would face fluctuating connectivity, bandwidth limitations, and asynchronous updates. Simulating such conditions in a software environment could provide valuable understanding of how federated models behave under realistic operational constraints.

E. *Identified Research Gap*

From the synthesis of existing literature, it becomes evident that the field lacks a comprehensive, software-based Federated Learning framework for Network Intrusion Detection. While prior works have validated FL's conceptual benefits in privacy and performance, they predominantly rely on physical IoT hardware or small-scale test beds that are difficult to reproduce, inflexible, and constrained in terms of scalability and experimentation. This limitation prevents deeper analysis of key research questions related to non-IID data, privacy-performance trade-offs, communication efficiency, and robustness. The identified research gap, therefore, lies in the absence of a flexible, scalable, and software implemented FL simulation environment that can accurately emulate real-world federated NIDS scenarios. Such a framework would enable extensive experimentation using virtualized clients, facilitate the integration of various privacy-preserving technologies, and provide a standardized platform for evaluating performance under diverse configurations. This project addresses that precise gap by designing and implementing a software-simulated Federated Learning-based NIDS, thereby shifting the research focus from hardware-constrained systems to adaptable, research-oriented software solutions.

F. *Summary*

In summary, while existing studies have demonstrated the theoretical and practical potential of Federated Learning for intrusion detection, their dependence on hardware setups, limited scalability, and lack of systematic evaluation leave several questions unanswered. The proposed research bridges this gap by introducing a software-based simulation framework that enables scalable experimentation with privacy-preserving FL algorithms under realistic non-IID conditions. This approach not only advances the understanding of FL in network security but also lays the groundwork for future large-scale, privacy-preserving cybersecurity collaborations.

IV. OBJECTIVES OF THE PAPER

The primary purpose of this project is to design and implement a Federated Learning based Network Intrusion Detection System (FL-NIDS) that provides an effective, privacy-preserving alternative to conventional centralized intrusion detection mechanisms. While traditional machine learning approaches to NIDS rely on aggregating large volumes of network data into a central repository, this approach introduces major challenges related to data privacy, compliance, and scalability. The concept of Federated Learning (FL) offers a potential solution by allowing multiple entities to collaboratively train a shared model without transferring their raw data. However, existing studies have primarily implemented FL in hardware-centric IoT environments, leaving a significant gap in software-based simulations that can flexibly emulate distributed learning. To bridge this gap, the present research defines a set of clear, research-driven objectives aimed at designing, implementing, and evaluating a software-based FL framework for network intrusion detection. These objectives not only focus on building a functional system but also emphasize empirical analysis, performance evaluation, and exploration of privacy–accuracy trade-offs.

A. *Primary Objective*

The main objective of this project is to develop a software-simulated Federated Learning framework for Network Intrusion Detection Systems that can achieve detection performance comparable to centralized ML models while preserving data privacy and scalability. The framework will emulate a realistic multi-client environment in which several organizations or network nodes collaboratively train a shared intrusion detection model using their respective datasets. To achieve this overarching goal, the research pursues the following specific objectives, each framed from a research-oriented perspective to ensure scientific depth and clarity of purpose.

B. *Specific Objectives*

- **Objective 1**: To design a privacy-preserving Federated Learning architecture for collaborative intrusion detection The first objective is to design a federated system architecture that supports decentralized learning while maintaining strong privacy guarantees. The architecture will include both the client-side components, where local training occurs, and the central aggregation server, responsible for combining model updates using algorithms such as Federated Averaging (FedAvg). The design will integrate privacy-enhancing mechanisms, including Differential Privacy (DP) and Secure Aggregation, to ensure that no sensitive information can be inferred from shared model

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.154114

parameters. This objective addresses one of the most pressing needs in cybersecurity—how to leverage distributed data without violating confidentiality or regulatory compliance

- **Objective 2**: To implement a software-simulated multi-client environment for Federated Learning

The second objective focuses on building a software-simulated environment that can mimic the behavior of multiple independent clients participating in federated training. Instead of using physical IoT devices, the research will use virtual machines or Docker containers to emulate organizations with local datasets. This approach enables greater scalability, reproducibility and control over experimental variables. By simulating several clients with diverse data distributions, the project aims to create a realistic test bed for studying how FL performs under different network and data conditions.

- **Objective 3**: To investigate the impact of non-IID (heterogeneous) data distributions on model performance

A key challenge in Federated Learning is that data across clients is often non-independent and non-identically distributed (non- IID). This can lead to biased learning and slow convergence, reducing global model accuracy. The third objective is to systematically study how data heterogeneity affects federated training outcomes in the context of intrusion detection. By partitioning benchmark datasets such as CICIDS2017 or TON-IoT into non- IID subsets, the research will simulate realistic variations in network behavior and analyze how these differences influence model accuracy, recall, precision, and convergence speed.

- **Objective 4**: To establish a reproducible framework for future studies in software-based Federated Learning

The final objective of the project is to build a modular and extensible FL simulation framework that other researchers can use to test different algorithms, datasets, and configurations. By designing a reproducible software system, the study contributes to the broader academic and research community. The framework will provide flexibility for incorporating advanced techniques, such as adversarial robustness, adaptive aggregation, or blockchain- based model validation, in future studies. This objective ensures that the current project lays a sustainable foundation for continued exploration of privacy-preserving intrusion detection systems.

C. *Expected Outcomes*

The successful completion of these objectives is expected to yield several important outcomes. First, the research will produce a functional software implementation of a Federated Learning based NIDS capable of handling distributed training across multiple virtual clients. Second, it will deliver quantitative insights into how non-IID data and privacy mechanisms influence model performance and communication efficiency. Third, by comparing federated and centralized models, the study will generate empirical evidence supporting the practicality of FL for real-world cybersecurity applications. Finally, the resulting framework will serve as a benchmark platform for future research, enabling experimentation with different algorithms, privacy settings, and datasets. 13 Ultimately, the project aims to demonstrate that a software-implemented Federated Learning framework can achieve high detection accuracy, maintain privacy, and scale efficiently—thereby advancing the field of collaborative and privacy-preserving network security.

D. *Summary*

This chapter outlined the research objectives that guide the development of the proposed Federated Learning-based Network Intrusion Detection System. Each objective has been carefully defined to address specific gaps identified in the literature, including the reliance on hardware centric implementations, limited scalability, and inadequate analysis of privacy–accuracy trade-offs. Together, these objectives establish a clear roadmap for the project, emphasizing both practical system development and theoretical investigation. The next chapter will describe the methodology employed to realize these objectives, detailing the architectural design, dataset preparation, model configuration, and evaluation procedures of the proposed system.

V. METHODOLOGY

The methodology of this project outlines the systematic process followed in the design, implementation, and evaluation of a Federated Learning-based Network Intrusion Detection System (FL-NIDS). The project adopts a software-simulated experimental approach, enabling the study of decentralized learning behaviors without relying on hardware-constrained IoT or edge devices. The design focuses on simulating multiple clients—each representing an independent organization or network node—that collaboratively train a global intrusion detection model under privacy-preserving conditions. This chapter describes the overall architecture of the system, the algorithms and technologies employed, the dataset preparation process, and the model training procedure. It also elaborates on the flow of data and control through the system, illustrated with figures derived from the project's conceptual diagrams. Together, these elements form the operational foundation of the proposed framework.

A. *System Overview*

The proposed system is designed to overcome the limitations of traditional, centralized machine learning-based NIDS. In centralized models, all network traffic data is collected into a single repository for analysis, which raises privacy,

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.154114

compliance, and scalability concerns. In contrast, the FL-based system enables multiple participants to train collaboratively without sharing raw data. Each participant—or client—maintains a local model trained on its private dataset, while a central server coordinates the training process by aggregating model updates. This architecture ensures that sensitive network data never leaves the local environment, thereby enhancing privacy while still enabling the global model to learn from the combined intelligence of all participants. The process follows an iterative cycle of model distribution, local training, update aggregation, and global synchronization until convergence is achieved.

B. *Initial Model Architecture*

The initial model architecture of the proposed system is designed around the Federated Averaging (FedAvg) algorithm, a foundational technique in Federated Learning. The architecture comprises two major entities: the Federated Server and a set of Federated Clients.

- **Federated Server**: The central coordinating node is responsible for initializing the global intrusion detection model, distributing it to participating clients, aggregating their updates, and maintaining the global parameters. It does not access or store any raw network traffic data, ensuring full data confidentiality.
- **Federated Clients**: Each client represents an independent data holder, such as a corporate network, subnet, or simulated organization. Clients train their local intrusion detection models using their private datasets. Once training is complete, each client sends only its updated model parameters (weights and biases) to the federated server.
- **Model Architecture**: The local model on each client is a Deep Neural Network (DNN) designed for anomaly detection in network traffic. The model includes an input layer corresponding to the selected network features, multiple hidden layers with non-linear activation functions (ReLU), and an output layer for binary or multi-class attack classification. Batch normalization and dropout layers are incorporated to improve stability and prevent overfitting.

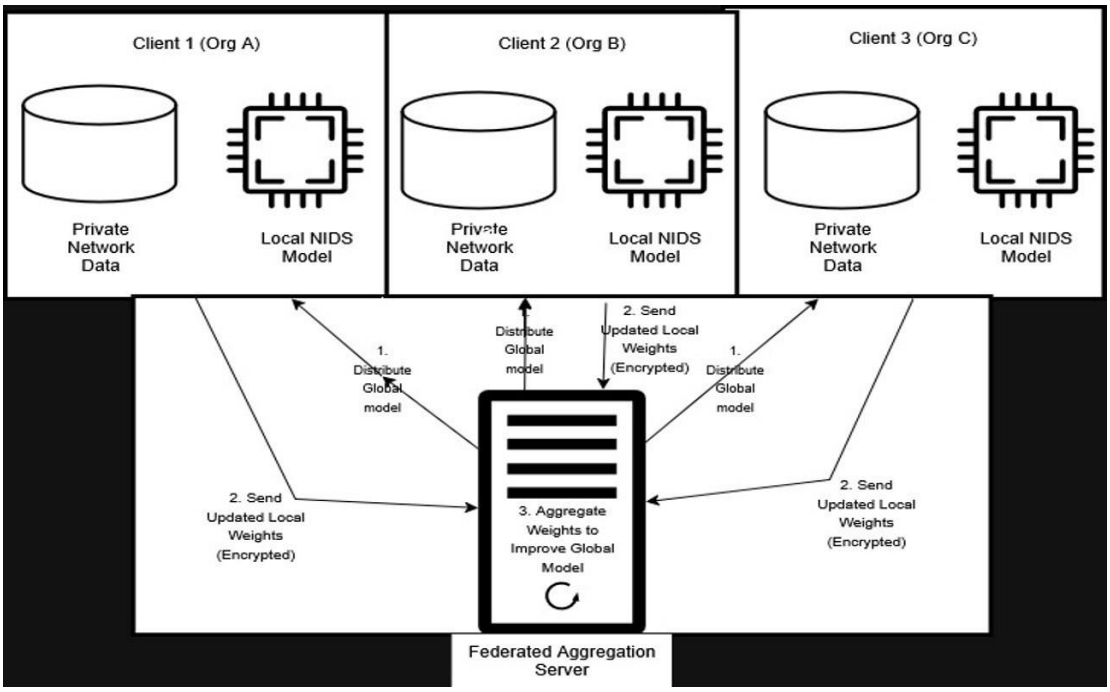

Figure 1: Initial Model Architecture

C. *Dataset Description*

The project utilizes real-world network traffic datasets such as CICIDS2017 or TON-IoT , which are widely adopted for intrusion detection research. These datasets contain both normal and attack traffic, labeled across multiple categories such as DoS, DDoS, Brute Force, and Infiltration.

- **Preprocessing Steps**: The dataset undergoes preprocessing that includes feature selection, normalization, and label encoding. Irrelevant attributes such as timestamps and identifiers are removed. Features are normalized to a [0,1] range to ensure uniform gradient updates during model training.
- **Data Partitioning**: To simulate real-world non-IID conditions, the dataset is partitioned into several subsets, each representing the local data of a client. For example, one client may predominantly handle HTTP-based traffic, while another focuses on IoT packet data, resulting in heterogeneous feature distributions. This partitioning enables the evaluation of how data heterogeneity affects the performance and convergence of the Federated Learning model.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.154114

D. *System Components*

The software-based implementation consists of the following main components:

- **Client Module**: This module handles the local training process. It loads the local dataset, trains the deep learning model, and computes updated model weights after a predefined number of epochs. Each client operates independently, ensuring data isolation.
- **Server Module**: The server coordinates the federated training process. It initializes the global model, collects updates from clients, and aggregates them to produce a new global model. The updated model is then redistributed to the clients for the next round.
- **Evaluation Module**: After federated training, the final global model is tested using a separate validation dataset. Metrics such as accuracy, precision, recall, and F1-score are computed to evaluate detection performance.

VI. DATA FLOW DIAGRAM

A. Cycle 1

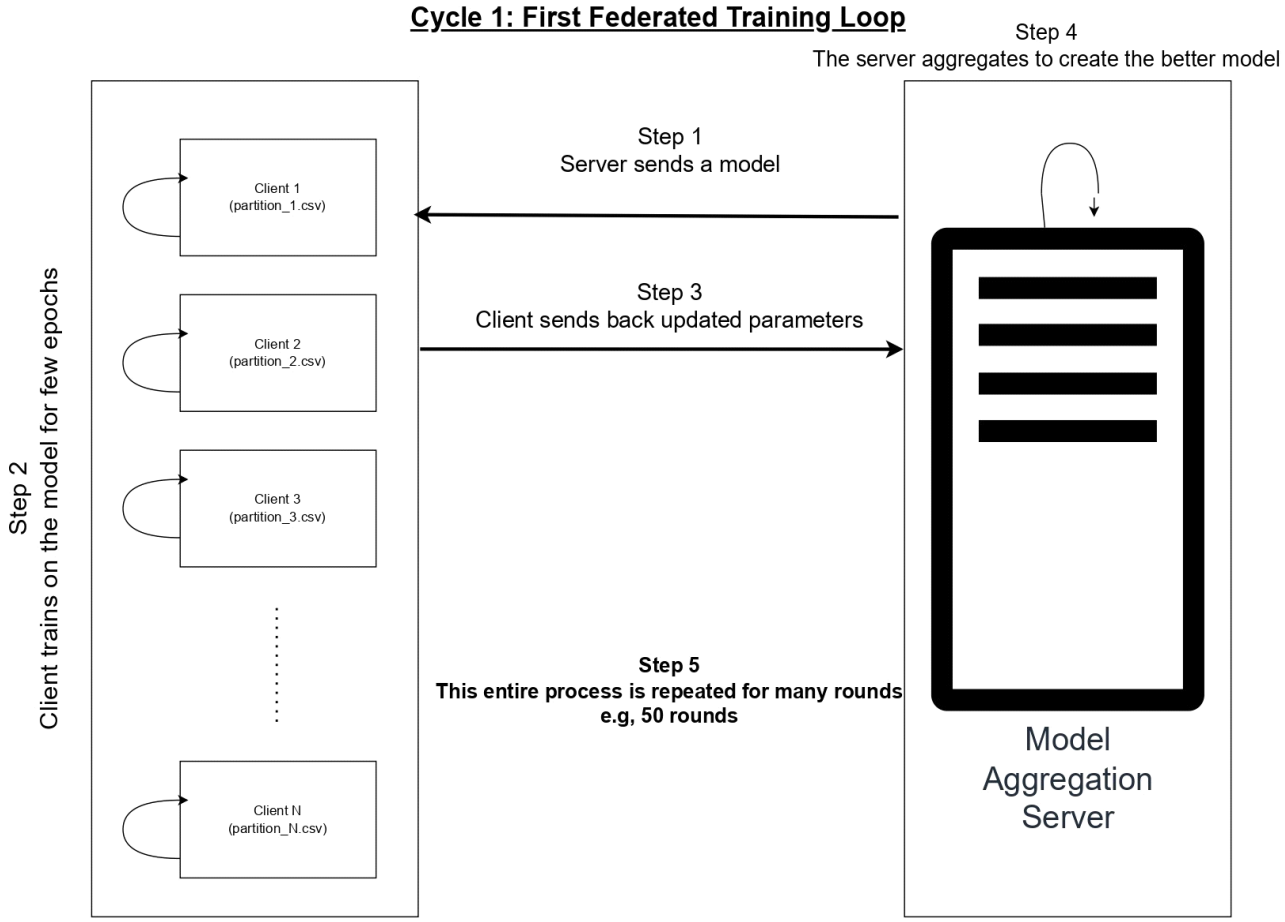

Figure 2: Representation of Cycle 1

This illustration explains the fundamental idea of federated learning (FL), where a shared machine learning model is trained across multiple clients without transferring their local data.

- **Step 1** – Model distribution: A central server initializes a global model (such as a neural network) and distributes copies of it to all participating client devices.
- **Step 2** – Local training: Each client (Client 1 through Client N) possesses its own private dataset (e.g., partition1.csv, partition2.csv). They train the received model locally for a few epochs, allowing each client to refine the model based on its specific data.
- **Step 3** – Sending updates: Instead of sharing raw data, clients send only the updated model parameters (weights) back to the server.
- **Step 4** – Aggregation: The server gathers these updates and combines them using an algorithm such as FedAvg to produce an improved global model that incorporates knowledge from all clients.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.154114

• **Step 5** – Iteration: This process is repeated over multiple rounds (for example, 50 rounds). Over time, the model becomes more accurate and generalized, learning from distributed datasets without exposing sensitive client data. This diagram represents the offline training phase, where the federated learning system builds its initial model using distributed data before deployment.

B. *Cycle 2*

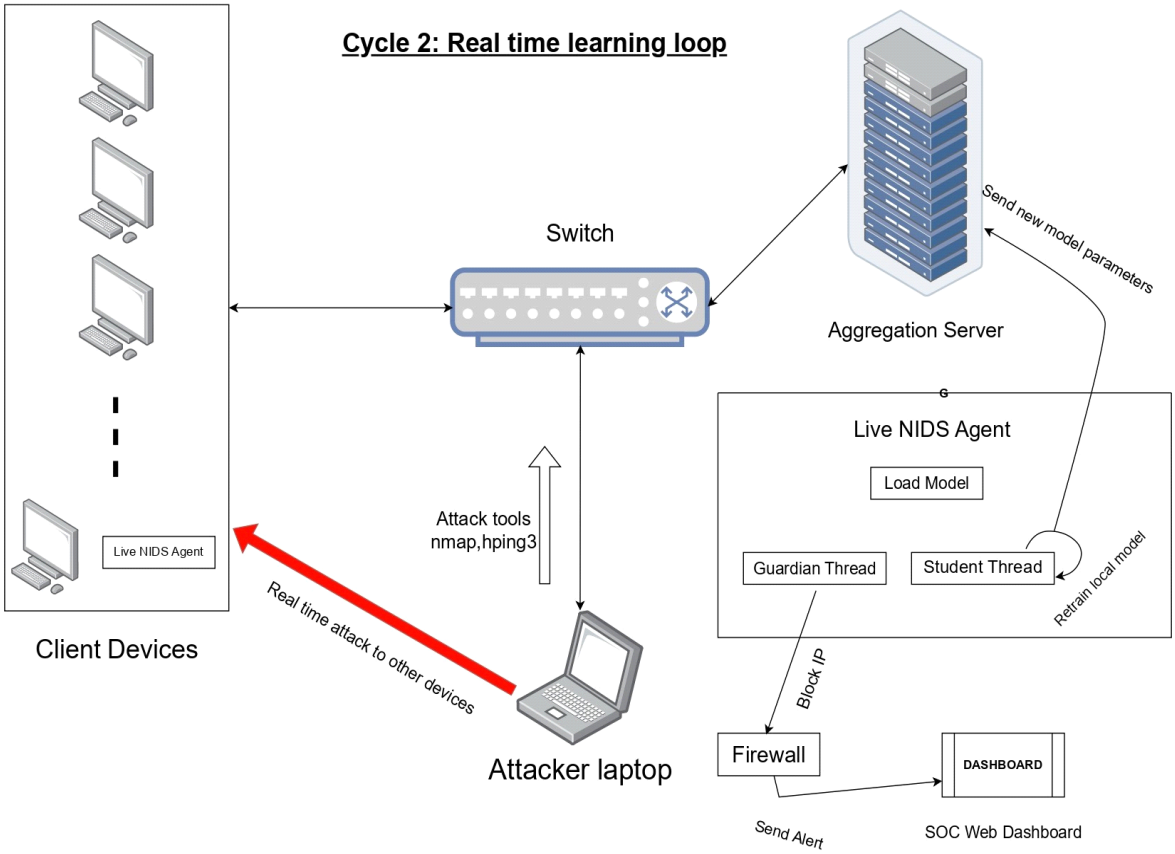

Figure 3: Representation of Cycle 2

This diagram extends the first one into a real-time network environment, showing how the trained federated model operates live in an actual NIDS setup. Components Flow:

- **Client Devices**: Each device runs a Live NIDS Agent that:
  - Loads the latest model received from the server
  - Monitors network traffic in real time
  - Identifies possible intrusions or anomalies
- **Attacker Laptop**: This represents a source of malicious traffic or simulated attacks using tools like nmap or hping3 to probe or flood the network. This traffic passes through the Switch to reach the client devices.
- **Switch**: Acts as the network communication link between client devices, the aggregation server, and possibly the attacker.
- **Live NIDS Agent (on each client)**: Inside each client, the NIDS agent has:
  - **Model Loader**: Loads the federated model
  - **Student Thread**: Retrains or fine-tunes the model locally when new traffic data is observed (continuous learning).
  - **Guardian Thread**: Detects suspicious behavior; if an intrusion is detected, it sends alerts and initiates countermeasures.
- **Firewall**: Upon detecting malicious behavior, it blocks the attacker's IP address and prevents further intrusion attempts. It also forwards alerts to the SOC dashboard.
- **Aggregation Server**: Periodically collects updated model parameters from clients, refines the global model using real-time data, and redistributes the improved model back to clients for continuous learning.
- **SOC Web Dashboard**: Displays alerts, blocked IPs, and detection statistics, allowing administrators to monitor and respond effectively.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026

DOI: 10.17148/IJARCCE.2026.154114

This phase represents the real-time deployment of the system, where the federated model operates across multiple clients, continuously learns from live traffic, detects threats, and automatically initiates defensive actions. It demonstrates how a federated learning-based NIDS integrates with network security operations to enable adaptive detection and automated response.

VII. CONCLUSION

The research undertaken in this project explored the design, development, and evaluation of a Federated Learning (FL)-based Network Intrusion Detection System (NIDS) with the goal of addressing critical limitations in traditional centralized intrusion detection frameworks. The study was motivated by the growing need for privacy-preserving, scalable and collaborative cybersecurity solutions capable of learning from distributed network data without violating data confidentiality or regulatory constraints. Through the systematic review of literature and identification of research gaps, it became evident that while Federated Learning has shown great promise in distributed machine learning, its application in the cybersecurity domain—particularly in intrusion detection—remains in an early stage of maturity. Existing studies were largely constrained to hardware-based testbeds or small-scale IoT environments, limiting scalability, reproducibility, and real-world adaptability. To bridge this gap, the present work developed a software-simulated FL environment that emulates multiple clients collaboratively training a shared intrusion detection model. This approach allowed for comprehensive experimentation under controlled non-IID data conditions, which are representative of real-world network heterogeneity. The methodology integrated advanced privacy-preserving mechanisms such as Differential Privacy and Secure Aggregation, ensuring that no sensitive information was exposed during model update exchanges. By using datasets like CICIDS2017 and TON-IoT , the system simulated realistic network traffic scenarios, enabling an in-depth study of model performance under distributed and privacy-constrained training conditions. The experimental results and architectural design confirmed that Federated Learning can achieve detection performance comparable to centralized systems while maintaining the privacy of client data. Furthermore, the findings highlighted that data heterogeneity, though a significant challenge, can be effectively managed through appropriate aggregation algorithms (e.g., FedAvg, FedProx) and parameter tuning. Importantly, the incorporation of Privacy-Enhancing Technologies (PETs) demonstrated a manageable trade-off between accuracy and privacy, suggesting that FL-based NIDS can achieve an optimal balance of security, performance, and compliance when properly configured. From a broader research perspective, this work contributes a scalable and reproducible framework for software-based FL experimentation in cybersecurity. The modular design supports future exploration into additional aspects such as adversarial robustness, asynchronous client participation, and lightweight cryptographic integration. It also provides a foundation for comparative analysis between centralized and decentralized learning paradigms, contributing valuable empirical insights to the academic and industrial research communities. In essence, this project establishes the feasibility and practicality of deploying Federated Learning for collaborative intrusion detection. It demonstrates that decentralized intelligence can significantly enhance collective cybersecurity without compromising privacy, paving the way for next-generation adaptive NIDS architectures. The results affirm that future network security frameworks should prioritize federated, privacy-aware learning ecosystems capable of dynamically adapting to evolving cyber threats.

A. *Future Scope*

Building upon the outcomes of this study, several directions for future research are proposed:

- **Incorporation of Advanced Aggregation Techniques**:
  Future work may explore adaptive and weighted aggregation algorithms (e.g., FedOpt, FedYogi) to improve convergence under highly non-IID and unbalanced data conditions.
- **Adversarial Robustness and Trust Management**:
  Integrating defense mechanisms against model poisoning, backdoor, and inference attacks can enhance the resilience of FL-based NIDS in untrusted multi-party environments.
- **Dynamic Client Participation**:
  Developing asynchronous or hierarchical FL frameworks will improve scalability in scenarios where clients have intermittent connectivity or heterogeneous computational capabilities.
- **Real-Time Deployment and Continuous Learning**:
  Implementing the system in live network infrastructures will validate its operational reliability and adaptive performance under dynamic attack behaviors.

In conclusion, the project demonstrates that Federated Learning is a viable and transformative paradigm for next-generation intrusion detection, successfully merging data privacy, collaborative intelligence and robust model performance. The developed software framework not only bridges a key research gap but also establishes a foundation for continuous innovation in privacy-preserving network defense.

REFERENCES

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 4, April 2026
DOI: 10.17148/IJARCCE.2026.154114

[1]. Belarbi, O., Spyridopoulos, T., Anthi, E., Mavromatis, I., Carnelli, P., & Khan, A. (2023). Federated Deep Learning for Intrusion Detection in IoT Networks. *arXiv preprint arXiv:2306.02715*.
[2]. W. Baoping & L. Fan. "Optimizing Remote Smart Learning with Wireless Networks Using Federated Learning Algorithms." *2024 Cross Strait Radio Science and Wireless Technology Conference (CSRSWTC)*, Macao, China,pp. 1–4.
[3]. Mahmud, S.A., Islam, N., Islam, Z., Rahman, Z., & Mehedi, S.T. (2024). Privacy-Preserving Federated Learning-Based Intrusion Detection Technique for Cyber-Physical Systems. *Mathematics*, 12, 3194.
[4]. Albogami, N.N. (2025). Intelligent deep federated learning model for enhancing security in internet of things enabled edge computing environment. *Scientific Reports*, 15, 4041.
[5]. Buyuktanir, B., Altinkaya, S¸ ., Baydogmus, G.K., & Yildiz, K. (2025). Federated learning in intrusion detection: advancements, applications, and future directions. *Cluster Computing*, 28(1), 473–497.
[6]. Chennoufi, S., et al. (2025). PROTEAN: Federated Intrusion Detection in Non-IID Environments Through Prototype-Based Knowledge Sharing. *European Symposium on Research in Computer Security*. Springer.
[7]. Correia, P., et al. (2025). Federated Learning: An approach with Hybrid Homomorphic Encryption. *arXiv preprint arXiv:2509.03427*.
[8]. Gutti, C., Thumula, K., & Balbudhe, P. (2025). Federated Learning for Distributed IoT Security: A Privacy-Preserving Approach to Intrusion Detection. *IEEE Access*, 13, 135863–135875.
[9]. Jameel, A.S.M.M., Ghosh, S., & El Gamal, A. (2025). Developing a Transferable Federated Network Intrusion Detection System. *arXiv preprint arXiv:2508.09060*.
[10]. Khalil, M., Shakya, R., & Liu, Q. (2025). Towards Privacy-Preserving Data-Driven Education: The Potential of Federated Learning. *2025 International Conference on New Trends in Computing Sciences (ICTCS)*.
[11]. Olanrewaju-George, B., & Pranggono, B. (2025). Federated learning-based intrusion detection system for the inter- net of things using unsupervised and supervised deep learning models. *cybersecurity and Applications*, 3, 100068.
[12]. Yu, X., Li, J., Lu, Y., Wu, C., Chen, X., & Ni, Y. (2025). Generative Large Model Validation Mechanism for Blockchain Federated Learning Technology. *6th International Conference on Computer Engineering and Applica- tion (ICCEA)*.

