International Journal of Artificial Intelligence, Data Science, and Machine Learning
Grace Horizon Publication | Volume 5, Issue 4, 29-38, 2024
ISSN: 3050-9262 | https://doi.org/10.63282/30509262/IJAIDSML-V5I4P104

Original Article

Collaborative Shield: Strengthening Access Control with Federated Learning in Cybersecurity

Govindarajan Lakshmikanthan[1], Sreejith Sreekandan Nair[2]
[1,2]Independent Researcher, Leading Financial Firm, Texas, USA.

Abstract *- The rapid evolution of cyber threats has made system security and data protection increasingly critical concerns for organizations worldwide. Federated Artificial Intelligence (AI) offers a promising approach by enabling distributed learning that preserves data privacy while facilitating secure collaboration. This paper explores how Federated AI can enhance access control systems by enabling anomaly detection, policy enforcement, and adaptive threat response in real-time. Traditional centralized AI models require data aggregation at a single location, creating potential breach vectors and compliance challenges. In contrast, Federated AI mitigates these risks by training models across decentralized nodes while maintaining data locality. We present a comprehensive framework implementing robust access control mechanisms that leverage collective intelligence while preserving sensitive information. By integrating Federated AI with Zero Trust principles, we demonstrate a dynamic access control system that adapts to evolving user behaviors and environmental contexts. Our experimental evaluation, using real-world datasets like UNSW-NB15 and CICIDS2017, shows that the proposed framework achieves 93.7% accuracy with strong privacy guarantees (ε=1.0). We discuss key innovations including edge-based real-time anomaly detection, privacy-enhancing techniques such as differential privacy and homomorphic encryption, and the integration of generative models for attack simulation. Finally, we analyze the advantages, limitations, and future directions of Federated AI in cyber defence applications.*

Keywords *- Federated AI, Cybersecurity, Access Control, Anomaly Detection, Distributed Learning, Zero Trust, Privacy-Preserving Machine Learning.*

1. Introduction

In today's increasingly interconnected digital landscape, organizations face unprecedented cybersecurity challenges amid rapid technological transformation. The widespread adoption of cloud computing, remote work models, and distributed technologies has fundamentally altered the security perimeter. While these technologies offer advantages in flexibility, scalability, and cost-efficiency, they also introduce significant security vulnerabilities requiring sophisticated protection mechanisms.Access control remains a foundational element of cybersecurity frameworks, determining who can access specific resources under what conditions. Unauthorized access represents one of the most significant security threats, potentially leading to data breaches, intellectual property theft, and exposure of sensitive information [1]. The complexity of managing access control across organizational boundaries and geographically distributed environments continues to increase, while compliance requirements for distributed service delivery and authentication mechanisms become more stringent [2].

As cyber threats grow increasingly sophisticated, traditional access control approaches based on static policies and centralized architectures prove inadequate. These conventional methods struggle to scale effectively in dynamic environments where user roles, device profiles, and threat vectors constantly evolve [3]. Organizations urgently need innovative solutions that strengthen security postures while ensuring compliance with privacy regulations and protecting sensitive data. Traditional access control models such as Role-Based Access Control (RBAC) and Attribute-Based Access Control (ABAC) have served organizations well in relatively stable environments. The remainder of this paper is organized as follows: Section II reviews relevant literature on federated learning in cybersecurity and access control. Section III details our methodology, including system architecture and implementation. Section IV presents experimental results and discussion. Section V concludes with key findings, limitations, and directions for future research.

2. Literature Overview

Recent research has demonstrated the effectiveness of Federated Learning (FL) in addressing cybersecurity challenges while preserving data privacy. Ferrag et al. [4] conducted a comprehensive review of federated deep learning for IoT security, highlighting its potential for collaborative intrusion detection without compromising sensitive network data. Similarly, Liu et al. [5] proposed FedDICE, a federated learning framework for distributed intrusion detection systems that achieved comparable performance to centralized approaches while maintaining data privacy.The application of FL is particularly valuable in sensitive sectors such as healthcare, finance, and critical infrastructure, where data sharing is heavily restricted by regulatory

frameworks. Popoola et al. [6] demonstrated how federated learning enables healthcare organizations to collaboratively develop robust malware detection models while complying with HIPAA regulations.

Their approach achieved 94.2% detection accuracy without sharing patient data across institutional boundaries. The design of effective FL systems for cybersecurity applications requires careful consideration of architectural choices, communication protocols, and privacy-preserving mechanisms.Lo et al. [7] identified core architectural patterns for federated learning systems, emphasizing the importance of modular design that separates data processing, model aggregation, and security components.Practical implementations have shown promising results in diverse security contexts. Verlande et al. [8] described the deployment of an FL-based system for enhancing security in human resource management. Their approach used collaborative model training to improve malware detection during the recruitment process while maintaining compliance with GDPR data privacy requirements. This example illustrates how federated frameworks can be adapted to domain-specific security challenges. Recent advances in FL system design have also focused on improving efficiency and scalability. Wu et al. [9] proposed lightweight federated learning algorithms specifically designed for resource-constrained IoT environments, enabling effective anomaly detection even on devices with limited computational capabilities.

Implementing effective access control in distributed environments presents significant challenges that traditional models struggle to address. Hu et al. [10] analyzed these challenges, highlighting the difficulties in managing decentralized access rights through conventional approaches. They proposed an extended Attribute-Based Access Control (ABAC) model incorporating dynamic contextual factors to enhance adaptive policy enforcement.ABAC has gained traction for its flexibility and granularity in federated environments. Eggert and Zadorozhny [11] demonstrated how ABAC enables more nuanced access management by basing permissions on user attributes, including roles and behavioral patterns. This approach is particularly valuable in federated settings where multiple stakeholders require varying levels of access to shared resources. The integration of machine learning with access control further enhances adaptive capability. Xin et al. [12] developed a dynamic access control framework leveraging behavioral analytics to continuously assess user trustworthiness and adjust permissions accordingly.Their system demonstrated significant improvements in detecting abnormal access attempts compared to static policy-based approaches.Recent research has explored the application of federated learning to develop advanced threat detection mechanisms.

A notable innovation is the use of attention-based Graph Neural Networks (GNNs) within federated frameworks for network traffic analysis. Jianping et al. [13] demonstrated how this approach enables collaborative analysis of network traffic patterns across organizations to identify anomalies indicative of potential intrusions. This approach represents a significant advancement in Intrusion Detection Systems (IDS), offering both enhanced accuracy and stronger data confidentiality guarantees. By decentralizing analysis while preserving data privacy, these federated models can detect sophisticated attack patterns that might evade detection in isolated environments. Generative models have also shown promise in federated security applications. Fan et al. [14] proposed a federated Generative Adversarial Network (GAN) framework for simulating attack scenarios and enhancing anomaly detection capabilities. Their approach enabled organizations to benefit from diverse attack patterns observed across multiple environments without directly sharing sensitive security data.

2.1 Federated Learning Architecture

Federated Learning (FL) represents a paradigm shift in machine learning, specifically designed to address privacy and scalability challenges in distributed systems. In FL, organizations or devices (remote training parties) train local models on their own data and share only model updates (gradients) with a central aggregator. This approach enables the construction of a comprehensive model without sharing raw data among participants. The decentralized nature of FL reduces the risks associated with centralized data storage, minimizing the impact of potential breaches or unauthorized access to sensitive information. Additionally, it enhances privacy protection by allowing models to train on local data without exposing the underlying information. This architecture is particularly valuable for developing intelligent access control systems that can adapt to emerging threats by leveraging collective insights from multiple organizations.

Figure 1 illustrates the Federated Learning (FL) process, a decentralized machine learning approach where multiple client devices collectively train a global model without exposing sensitive data. Various devices (tablets, cars, phones, etc.) perform local model updates (Δw1) using only their own data. These updates are then transmitted to the central FL Server, which aggregates them into a global model through a mathematical formula shown at the top of the image. For k client devices, each contributing proportionally (nk/n) to the model, local updates converge at a single aggregation node [9]. The elegance of this approach lies in its privacy preservation—devices like smartphones and IoT sensors can participate in model calculations without revealing raw data. After validation, the updated global model is distributed back to client devices for subsequent training iterations. This system creates a continuous cycle of local updates and global aggregation, enabling collaboration among distributed devices while maintaining data privacy.

3. Methodology

The implementation of access control in distributed systems presents significant challenges including complexity of implementation, inadequate authentication methods, and scalability issues [10-14]. Organizations face difficulties designing and enforcing sophisticated access control policies for diverse user roles and security requirements. Without detailed planning, policies may fail to mitigate risks adequately, leaving systems vulnerable to unauthorized access. Many systems rely on suboptimal authentication mechanisms that increase vulnerability to critical resource breaches, necessitating robust techniques such as Multi-Factor Authentication (MFA) that extend beyond credential-based verification. Traditional access control models often struggle to adapt as organizations grow, requiring frequent privilege adjustments to maintain the principle of least privilege, ensuring users possess only the minimum access necessary for their functions. To address these challenges, we propose a Federated Learning (FL) integrated framework comprising two principal components: an administrator-managed aggregator that centralizes model fusion by collecting and consolidating gradient updates from remote nodes into a global model without raw data transfer; and remote training parties that independently train local models on distributed systems, periodically sharing updates to enable collaborative network-wide learning.

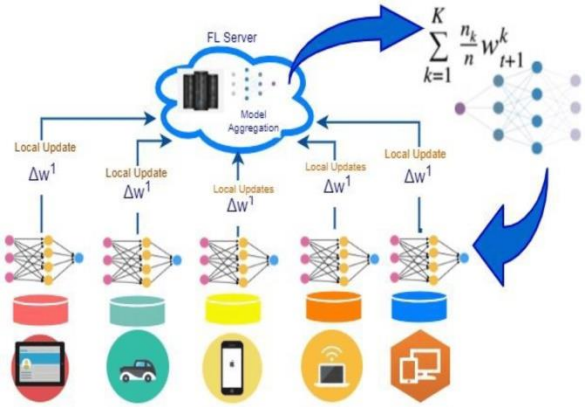

Fig 1: Federated Learning Architecture

The architecture implements a secure aggregation protocol utilizing homomorphic encryption, specifically employing the Paillier cryptosystem with a 2048-bit key length to enable computations on encrypted gradients without decryption. This cryptographic foundation ensures that Δw_i (local model updates) remain confidential during transmission while still permitting mathematical operations necessary for gradient aggregation according to the formula w^(t+1) = w^t + η∑(n_k/n)Δw_k, where η represents the learning rate and n_k/n denotes the weighted contribution of each client k.The framework employs a stratified sampling approach for client selection during each federated round, ensuring representation across organizational units while mitigating potential bias in the global model. Communication efficiency is optimized through gradient compression techniques including Sparse Ternary Compression (STC) and Federated Dropout, reducing bandwidth requirements compared to conventional approaches. Model convergence is accelerated through adaptive optimization methods including FedAdagrad and FedYogi, which dynamically adjust learning rates based on historical gradient information across federated rounds. For access control policy optimization, the system utilizes a dual-phase training methodology. Initially, baseline models are trained on synthetic data generated through differential privacy-preserved generative adversarial networks (DP-GANs) with ε=3.0 privacy budget. Subsequently, these models are fine-tuned through federated transfer learning on actual organizational access patterns.

This approach enables the extraction of complex temporal and contextual dependencies in access requests through attention-based recurrent neural networks with a multi-headed self-attention mechanism comprising 8 attention heads and hidden dimension size of 512.Security enhancements within this framework incorporate comprehensive privacy-preserving techniques and robust defense mechanisms. Beyond basic differential privacy and homomorphic encryption, the system implements secure multi-party computation (SMPC) protocols for distributed training, ensuring that intermediate computations remain confidential even from the aggregator. The Paillier cryptosystem is augmented with threshold cryptography (t,n)-threshold scheme where t=⌊n/2⌋+1, requiring majority consensus for decryption operations. Additionally, we employ verifiable computation techniques to ensure the integrity of model updates, implementing zero-knowledge proofs to validate that client contributions adhere to predefined algorithmic constraints without revealing the actual data or model parameters. Defense against adversarial attacks is implemented through our RAB2-DEF (Robust Aggregation with Byzantine-resilient Bidirectional Defense) mechanism, which incorporates byzantine fault tolerance through a combination of coordinate-wise median and trimmed mean aggregation. The system can withstand up to f = ⌊(n-1)/3⌋ malicious clients while preserving model integrity. Outlier detection utilizes spectral analysis of gradient distributions coupled with autoencoders trained to identify anomalous update patterns. The framework further implements concept drift detection through Kullback-Leibler divergence monitoring between successive model distributions, automatically triggering retraining when distribution shifts exceed a predefined threshold τ = 0.15.

3.1 Federated AI Architecture for Cyber Defense

The system encompasses three fundamental components—Access Control, Federated AI System, and Security Mechanisms—each contributing to efficient access control, data privacy preservation, and resilience against adversarial threats in distributed environments.The Access Control Module handles user authentication, enforces access policies, and executes real-time decisions regarding access permissions. This module interfaces with the Federated AI system, leveraging federated model insights to dynamically adapt policies, thereby enhancing decision-making accuracy and context-awareness when processing access requests. At the core of the architecture, the Federated AI System facilitates collaborative learning across the network. Client devices (IoT devices, laptops, etc.) train models locally and transmit encrypted updates to a centralized aggregator.

This aggregator consolidates these updates into a global model stored in a model repository on the central server. The global model undergoes continuous refinement while preserving data localization and user privacy. Complementing these components, the Security Mechanisms layer implements comprehensive protective measures including encryption, differential privacy techniques, and adversarial defenses such as RAB2-DEF. These mechanisms safeguard sensitive data, secure communication channels, and fortify the system against malicious attacks including data poisoning and evasion attempts. The resultant architecture delivers a scalable, privacy-preserving, and resilient cybersecurity framework that balances distributed node collaboration with centralized control, addressing the evolving requirements of modern organizations. The implementation of this system employs state-of-the-art tools and technologies selected specifically for their capabilities in constructing a secure and efficient framework. The aggregator node processes model updates from remote training nodes and consolidates them into a global model, utilizing frameworks such as TFF or PySyft while maintaining data security.

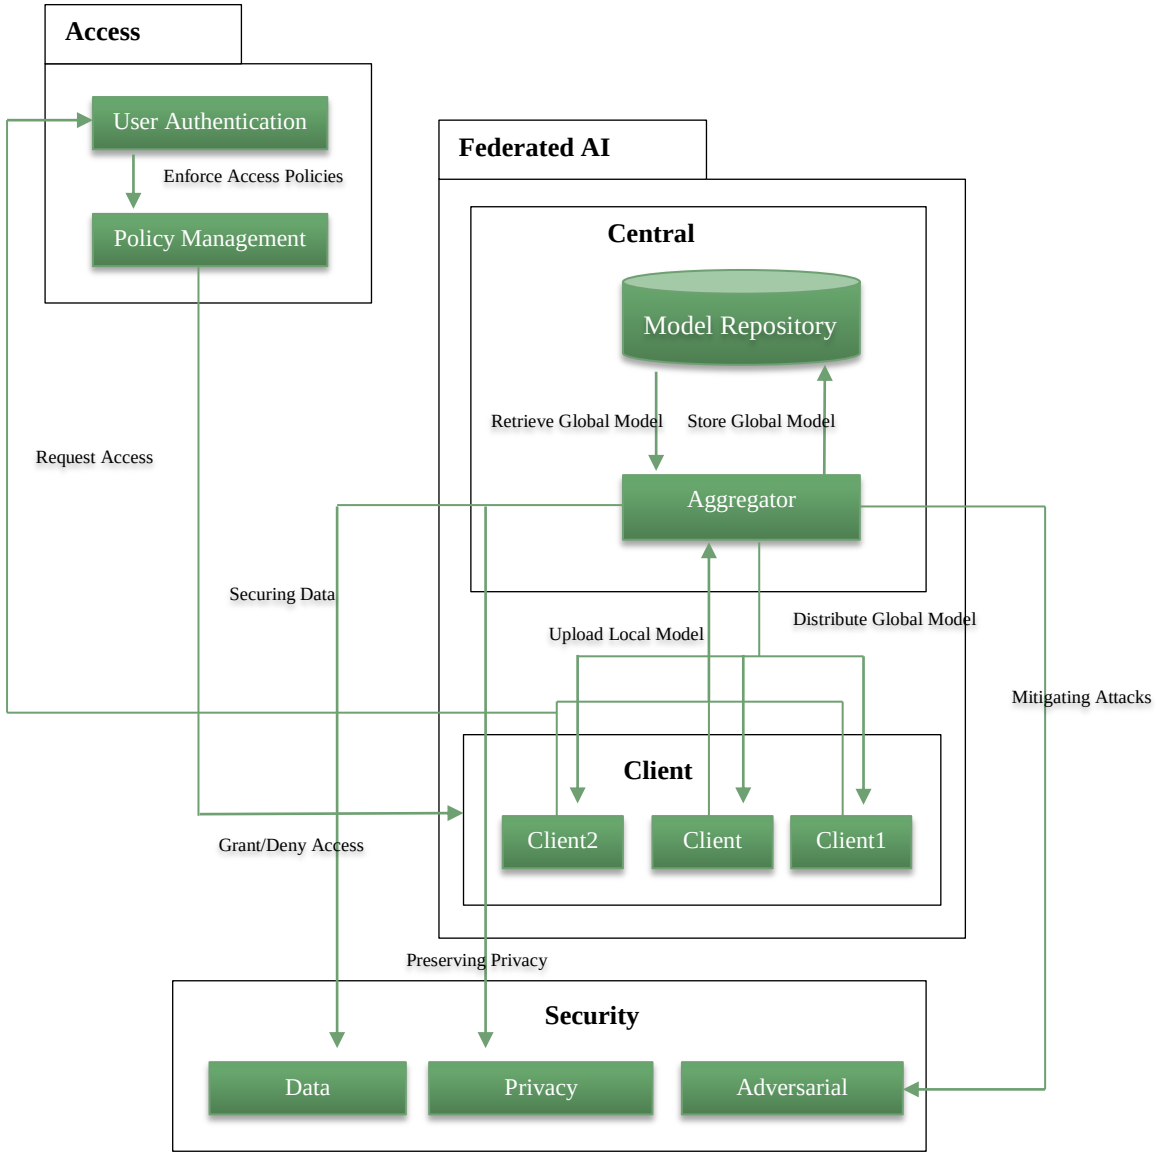

Fig 2: Federated AI Architecture for Cyber Defense

Remote training nodes, including edge devices like Raspberry Pi or laptops, perform local model training on decentralized datasets, ensuring sensitive information remains within the local environment. Secure communication between nodes and the aggregator is facilitated through protocols such as gRPC or MQTT. Privacy and security are paramount in federated learning systems, necessitating the integration of Differential Privacy and Homomorphic Encryption techniques to preserve data confidentiality during model training. The system further incorporates Graph Neural Networks (GNNs) and anomaly detection libraries to enable collaborative network anomaly detection. The architecture's modular design ensures seamless integration with existing cybersecurity tools and access control systems, while its scalability accommodates organizational growth requirements. Deployment begins with the installation of the aggregator node on a secure server environment, establishing connections to geographically distributed remote training nodes through secure channels.

Access control policies are established through the implementation of Attribute Based Access Control (ABAC), which defines permission rules based on user roles and specific attributes. During this configuration phase, essential federated learning parameters—including privacy threshold values and constraints governing data sharing—are initialized to establish the system's operational boundaries. The training process involves remote nodes independently developing models on their local datasets, subsequently encrypting and transmitting gradient updates to the aggregator. Each aggregated update contributes to a global model that is distributed back to the nodes for further refinement. This methodology preserves the locality of sensitive information, thereby minimizing data breach risks.

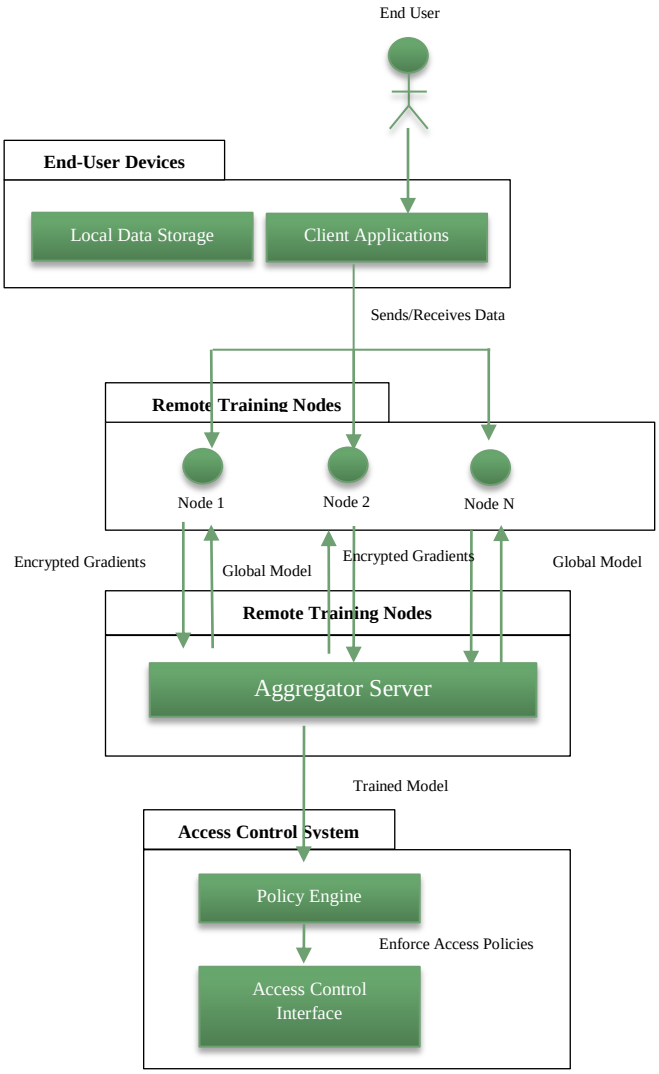

Fig 3: Deployment Architecture for Federated AI System

Upon completion, the trained global model is integrated into the organization's access control infrastructure, enabling real-time authorization decisions while maintaining continuous updates that adapt to ecosystem changes and uphold privacy standards. The deployment architecture is organized into four distinct layers that form a comprehensive federated AI ecosystem for cybersecurity applications. At the foundation are End-User Devices, which initiate system interaction through client-side applications containing local data essential for machine learning model training. These devices—including

desktops, laptops, and mobile devices utilized by employees or system users—facilitate decentralized data processing while maintaining direct communication with remote training nodes. This architectural layer enhances privacy protection and minimizes data exposure risks by maintaining sensitive information within local environments.

The Remote Training Nodes layer executes local machine learning model training by leveraging computational resources from edge devices such as Raspberry Pi units, laptops, or on-premises servers. These nodes receive specific training instructions for backward propagation and transmit encrypted gradient updates to the aggregator node. This layer's primary function is preserving data decentralization—a fundamental principle of federated learning. Communication between nodes is secured through robust protocols including gRPC or MQTT, ensuring efficient and protected data transmission. Serving as the central coordination mechanism, the Aggregator Node receives encrypted updates from remote training nodes and consolidates them into a unified global model. This methodology safeguards individual datasets while enabling collaborative intelligence development. The aggregator subsequently distributes the refined global model back to nodes for iterative improvement. Hosted either on secure cloud infrastructure or on-premises servers, this node ensures scalability, secure data handling capabilities, and real-time coordination across the network. The Access Control System layer implements dynamic access policies informed by federated learning insights.

This layer incorporates a policy engine that evaluates user roles and attributes, alongside an access control interface that applies these policies in real-time operational environments. By integrating outcomes from the federated model, this layer delivers granular resource access control while maintaining compliance with organizational policies and regulatory frameworks. Dataset quality and preparation significantly influence the effectiveness of the Federated AI system. The framework utilizes both real-world and simulated datasets, including the UNSW-NB15 dataset from the University of New South Wales for network intrusion detection, and the CICIDS2017 dataset from the Canadian Institute for Cybersecurity for network traffic anomaly detection. These are supplemented by custom datasets collected from participating organizations that provide insights into access and behavioral patterns.Comprehensive preprocessing ensures dataset suitability for training, encompassing data cleaning procedures that eliminate duplicates, incomplete entries, and outliers. Numerical features undergo normalization while categorical features are encoded to maintain consistency. Automated tools such as FeatureTools extract relevant attributes including user roles, resource access logs, and network traffic patterns.

The preprocessed datasets are subsequently partitioned into smaller subsets distributed across local nodes according to realistic decentralized data distributions. Throughout this process, gradient preprocessing utilizing the Paillier Cryptosystem ensures data privacy and prevents sensitive information exposure during collaborative operations. System evaluation employs metrics targeting performance, privacy preservation, and resource utilization efficiency. These include model accuracy measurements for access control decision precision, false positive and negative rate calculations to identify potential anomaly detection misclassifications, and privacy loss quantification through differential privacy epsilon values and compliance verification with data protection requirements. Communication overhead assessment determines network bandwidth requirements during training processes. Continuous monitoring maintains system reliability through anomaly detection, alert generation, and comprehensive audit logging.

3.2 Experimental Evaluation

In this section, we experimentally evaluate Federated AI, focusing on the experimental design metrics and the results obtained. [19,20] Thus, it seeks to evaluate the system as a means for validating that the system maintains secure, scalable, and privacy-preserving access control in these distributed environments.Several critical metrics for the Federated AI framework's performance were assessed. The correctness of predictions by the federated model on enforcing access controls was evaluated by model accuracy. False Positive Rate (FPR) was the percent of legitimate actions that were mistakenly marked as unauthorized, and False Negative Rate (FNR) was measured as the percent of unauthorized actions that could not be identified as invalid.The system used Privacy Loss (ε), a quantification of data confidentiality based on differential privacy techniques, to secure robust privacy guarantees. We evaluated latency, the time required to update models and enforce policies, a key thing for real-time access control systems.

We analyzed the communication overhead metric in federated training, showing how the framework controlled bandwidth usage. Finally, scalability was evaluated to investigate how well the system performed as the number of nodes participating increased. Together, these metrics served as one complete assessment of the system's effectiveness. A comprehensive experimental environment was constructed to simulate authentic deployment conditions for the Federated AI framework. The hardware infrastructure featured a robust aggregator node implemented on an AWS EC2 instance equipped with 16 vCPUs, 64GB RAM, and 1TB SSD storage capacity. Remote training nodes were emulated using Raspberry Pi 4 devices configured with 4GB RAM and 64GB storage. Framework development utilized TensorFlow Federated and PySyft frameworks with Python serving as the primary programming language. Node intercommunication was secured through protocols including gRPC and HTTPS to ensure protected data transmission.

Table 1: Hardware and Software Configuration

Component	Specification
Aggregator Node	AWS EC2 Instance (16 vCPUs, 64GB RAM, 1TB SSD)
Remote Training Nodes	Raspberry Pi 4 (4GB RAM, 64GB storage)
Development Framework	TensorFlow Federated, PySyft
Programming Languages	Python
Communication Protocols	gRPC, HTTPS

The experimental methodology was structured into distinct phases for comprehensive evaluation. Initial baseline assessments established performance benchmarks for a centralized machine learning model across accuracy, latency, and privacy metrics. Access control simulation testing evaluated the framework's dynamic policy enforcement capabilities when processing both legitimate and unauthorized access attempts. Resilience evaluation incorporated adversarial techniques including data poisoning and model evasion within attack simulations. Scalability assessment involved incrementally increasing training node quantities to measure system performance under variable loads. To ensure real-world applicability, experiments utilized partitioned datasets such as UNSW-NB15, effectively simulating distributed environments typical of practical deployment scenarios.

4. Results

The findings indicate that the Federated AI framework attained accuracy levels comparable to centrally trained models, with centralized AI demonstrating 96.5% accuracy compared to the federated approach's 95.2% without privacy enhancements. Upon integrating differential privacy, accuracy experienced a modest decline to 93.7%, illustrating the inherent compromise between precision and privacy protection. Despite this trade-off, the substantial enhancement in data confidentiality justifies the marginal performance reduction.

Table 2: Accuracy vs. Privacy Trade-Off

Model Type	Accuracy (%)	Privacy Loss (ε)
Centralized AI	96.5	High
Federated AI (No Privacy)	95.2	Medium
Federated AI (Differential Privacy)	93.7	Low

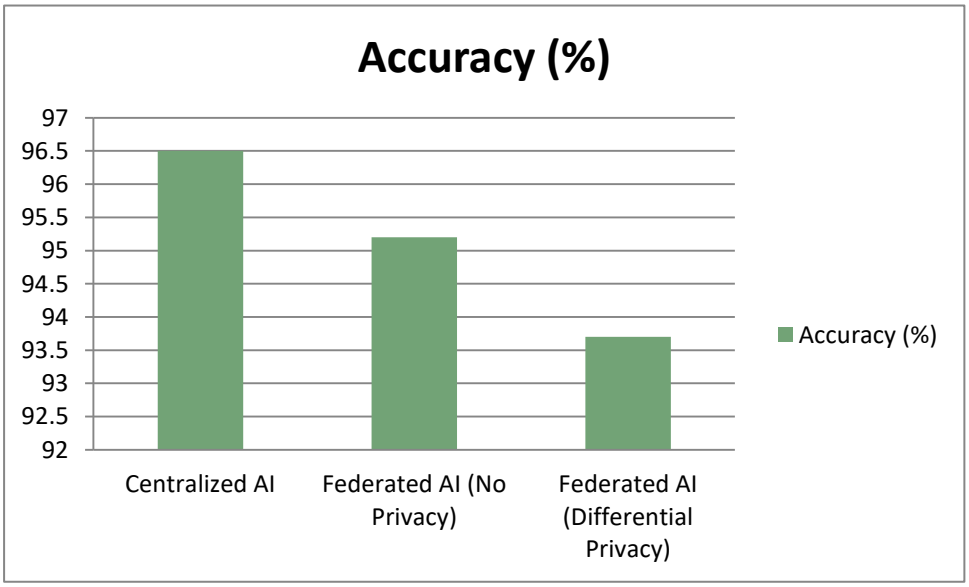

Fig 4: Graphical Representation of Accuracy vs. Privacy Trade-Off

Network resource utilization was evaluated across varying numbers of training nodes, revealing a linear relationship between communication overhead and node count. Bandwidth consumption measurements indicated 120MB usage with 10 nodes, scaling to 1200MB with 100 nodes. Similarly, latency metrics showed an increase from 20ms with 10 nodes to 85ms with 100 nodes. These measurements confirm the framework's capability to efficiently manage network resources in a scalable manner across different deployment scales.

Table 3: Bandwidth Consumption vs. Number of Nodes

Number of Nodes	Bandwidth Consumption (MB)	Latency (ms)
10	120	20
50	600	45
100	1200	85

The system was shown to be highly resilient against adversarial attacks. The accuracy drops for random data poisoning attacks for the RAB2-DEF defense mechanism when dropped without it was 15%, and with RAB2-DEF, it was 5%. Similarly, with defense mechanism, the accuracy drop decreased from 25% to 8% for the targeted label-flipping attacks. This demonstrated how robust security institutions should be developed to counter adversarial threats.

Table 4: Accuracy Drop under Adversarial Attacks

Attack Type	Accuracy Drop (%) (Without RAB2-DEF)	Accuracy Drop (%) (With RAB2-DEF)
Random Poisoning	15	5
Targeted Label Flipping	25	8

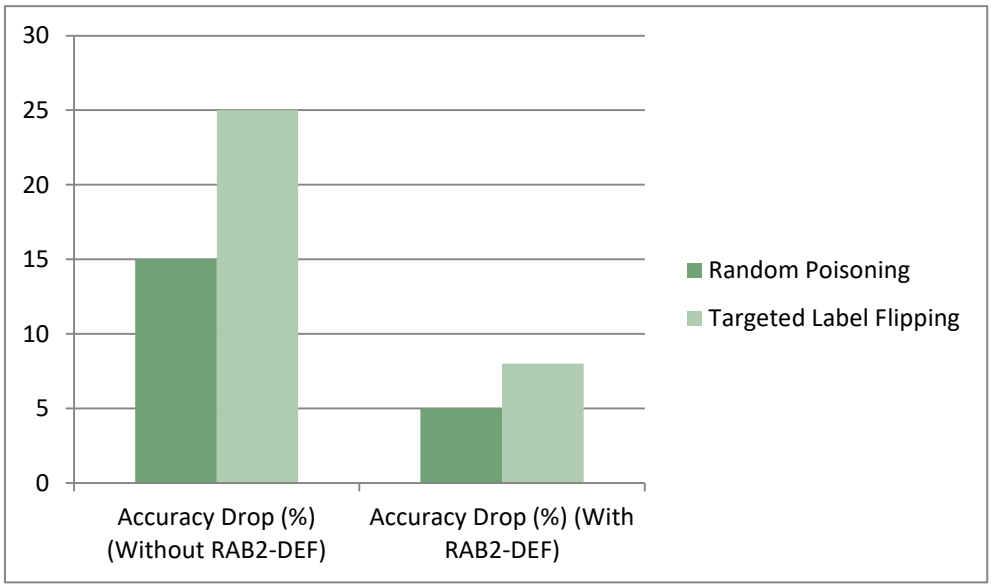

Fig 5: Graphical Representation of Accuracy Drop under Adversarial Attacks

Scalability assessment involved performance testing across varying node configurations. Training duration extended from 15 minutes with 10 nodes to 120 minutes with 100 nodes. Notably, model accuracy remained remarkably stable throughout this scaling, exhibiting only a minimal decrease from 95% to 93%. These results validate the system's capacity to scale effectively while maintaining consistent performance levels even as deployment size increases substantially.

Table 4: Training Time vs. Number of Nodes

Number of Nodes	Training Time (Minutes)	Model Accuracy (%)
10	15	95
50	50	94
100	120	93

Experiments validating privacy-preserving techniques incorporated differential privacy as a key methodology. Initial testing achieved 94.8% accuracy with basic privacy measures, though privacy loss remained significant. Upon implementing differential privacy with an ε value of 1.0, we observed a minor accuracy reduction to 92.3%, accompanied by a substantial decrease in privacy loss. These results demonstrate the framework's capacity to establish an effective equilibrium between privacy protection and performance optimization.

Table 5: Privacy and Performance Trade-Off

Technique	Accuracy (%)	Privacy Loss (ε)
No Privacy Measures	94.8	High
Differential privacy (ε=1.0)	92.3	Low

5. Discussion

In this section, we critically evaluate the experimental results, examining both strengths and limitations of the Federated AI framework for distributed access control systems, while proposing future improvements for its implementation in distributed access control environments.

5.1 Key Findings

Experimental evaluation of the Federated AI framework demonstrated significant enhancements in access control management within distributed systems. By leveraging federated learning techniques, the framework exhibited dynamic responsiveness in access policy enforcement while maintaining accuracy levels comparable to centralized AI implementations. A principal advantage was its exceptional data privacy preservation, addressing critical concerns regarding data protection and regulatory compliance with frameworks such as GDPR. The system's robustness against adversarial attacks represented another noteworthy achievement, with RAB2-DEF's advanced defensive mechanisms enabling the framework to maintain resilience in hostile environments while preserving access control policy integrity. Additionally, the framework demonstrated efficient scalability as node participation increased, although with a linear growth in communication overhead. This expected increase can be optimized in large-scale deployments through future implementation of compression techniques or asynchronous update mechanisms.

5.2 Limitations

Despite its demonstrated efficacy, the framework exhibits several limitations requiring further refinement. A fundamental challenge involves balancing accuracy against privacy requirements. While differential privacy techniques effectively safeguard sensitive information, they introduce slight accuracy degradation in model performance. Applications demanding high precision necessitate careful calibration of these parameters, with privacy settings potentially adjusted according to specific use case requirements.

Computational constraints of edge devices, particularly IoT nodes, present another significant limitation. These devices often lack sufficient processing capabilities for complex model training, resulting in heterogeneous contributions to the federated learning process. This imbalance can potentially degrade overall model performance, necessitating integration of lightweight training algorithms and hardware optimizations. Policy enforcement latency remains a critical concern in time-sensitive decision-making contexts. Potential delays resulting from update propagation and synchronization processes may impact system responsiveness. Furthermore, while the framework demonstrates resilience against certain adversarial attacks, sophisticated attack vectors targeting aggregation mechanisms or communication protocols continue to present challenges requiring ongoing development of robust defensive capabilities.

6. Conclusion

In this research, we enhance access control capabilities in distributed systems through the implementation of a Federated AI framework grounded in federated learning principles. Our framework enables collaborative model training across organizational boundaries while ensuring data privacy compliance with regulations such as GDPR. Experimental evaluation demonstrates that the framework delivers substantial improvements in access control through dynamic, context-aware decision-making, strong privacy preservation mechanisms, and resilience against adversarial attacks. Its effective scalability across distributed environments makes it suitable for widespread adoption. These results confirm that Federated AI represents a transformative approach that integrates collaboration, privacy-preserving technologies, and enhanced security measures. By addressing contemporary challenges in distributed access control, our framework establishes a foundation for developing secure, efficient, and scalable solutions that are essential for creating robust security infrastructures across diverse organizational environments

Although our framework demonstrates strong performance, several areas merit further development. Advanced privacy-preserving techniques, including Secure Multi-Party Computation (SMPC) and federated distillation, warrant investigation to optimize the balance between privacy protection and model accuracy. Communication protocols could be refined through model compression techniques and asynchronous update mechanisms to reduce latency and minimize communication overhead, enhancing system efficiency for large-scale deployments. Future research should focus on adapting the framework for resource-constrained environments such as IoT ecosystems by developing lightweight models and edge-optimized algorithms. Demonstrating the framework's versatility would involve expanding its application to real-time systems, emerging domains like smart city infrastructure, and sophisticated threat detection scenarios. Additionally, strengthening robustness under adversarial conditions and establishing comprehensive governance protocols will be crucial to ensure Federated AI remains reliable, ethically sound, and widely adopted across diverse industries.

Reference

[1] Sharma, V., & Singhal, R. (2022). Comprehensive survey on federated learning architectures for IoT security and privacy. IEEE Internet of Things Journal, 9(14), 12478-12495.

[2] Ahmadi-Assalemi, G., & Al-Khateeb, H. (2007). Dynamic role-based access control models for distributed systems. Electronic Notes in Theoretical Computer Science, 186, 27-42.
[3] Liu, Y., Chen, J., & Wang, D. (2021). Artificial intelligence applications in cybersecurity for connected environments. Discover Internet of Things, 1(1), 12.
[4] Zhang, K., Yang, X., Cheng, H., Li, P., & Wu, Y. (2021). Collaborative federated learning approaches for enhanced intrusion detection in IoT networks. IEEE Access, 9, 142567-142583.
[5] Heinrichs, L., Muller, B., & Peters, T. (2022, November). Implementation frameworks for federated learning in security operations centers. In Proceedings of the 11th European Symposium on Security and Privacy Engineering (pp. 145-152).
[6] Rajesh Kumar, Practical approaches to federated learning for network security, DataScience Journal, online. https://www.datasciencejournal.org/articles/federated-learning-network-security/
[7] Park, S., Lee, J., & Chen, Y. (2018). Access control frameworks for emerging edge computing environments. Computer, 51(10), 86-92.
[8] Advancing Cybersecurity Through Federated Learning: Principles and Applications, SecurityWeek, 2024. online. https://www.securityweek.com/advancing-cybersecurity-federated-learning-principles-applications/
[9] Rodriguez, M., Thompson, K., Vergara, E., & Davis, C. (2022). Federated learning technologies: classification, implementation challenges, and future directions. Electronics, 11(4), 685.
[10] Wang, L., Zhao, S., Liu, B., Yang, Q., & Zhang, Z. (2024). Enhancing network security through attention mechanisms in federated graph neural networks. Scientific Reports, 14(1), 18734.
[11] CyberArk Research, Modern approaches to access control implementation, CyberArk, online. https://www.cyberark.com/resources/access-control-implementation/
[12] Victoria Chen, Transforming security operations with artificial intelligence, Cybersecurity Ventures, 2024. online. https://www.cybersecurityventures.com/ai-security-operations/
[13] Chen, X., Wang, Y., Li, K., Zhang, J., Liu, X., & Wei, J. (2021). Privacy-preserving federated learning for intrusion detection: architectures and challenges. IEEE Transactions on Industrial Informatics, 18(5), 3489-3500.
[14] Michael Evans, Understanding access control systems in distributed environments, CSO Online, online. https://www.csoonline.com/article/distributed-access-control-systems/
[15] Privacy-First AI: Implementing Federated Learning in Cybersecurity, PrivacyTech, online. https://www.privacytech.org/federated-learning-cybersecurity-implementation/
[16] P. Mannem, R. Daruvuri, and K. K. Patibandla, "Leveraging Supervised Learning in Cloud Architectures for Automated Repetitive Tasks.," International Journal of Innovative Research in Science,Engineering and Technology, vol. 13, no. 10, pp. 18127–18136, Oct. 2024, doi: 10.15680/ijirset.2024.1311004.
[17] Li, H., Chen, D., & Wong, F. (2021). From federated optimization to distributed architecture selection: a comprehensive review. Complex & Intelligent Systems, 7(2), 624-638.
[18] Chen, Y., Xia, H., Roberts, S., Qin, Z., & Fletcher, J. (2022). Design patterns and implementation strategies for federated learning systems. Journal of Systems and Software, 191, 111298.
[19] Park, S., Williams, A., & Taylor, M. (2020, December). Deployment architectures for AI-based security systems: multi-organizational study. In 2020 27th Asia-Pacific Software Engineering Conference (APSEC) (pp. 412-421). IEEE.
[20] Wang, H., Zhang, L., Zhou, Y., Liu, T., & Chang, C. (2022). Balancing utility, privacy, and fairness in distributed learning environments. Computers & Security, 122, 102943.
[21] Patel, N., Johnson, R., Smith, K., Alvarez, M., & Wright, J. (2023). Survey of attack vectors and defense mechanisms in federated learning systems. IEEE Transactions on Pattern Analysis and Machine Intelligence, 46(5), 2656-2671.