

Federated Learning in Cybersecurity: Enhancing Decentralized Threat Detection

Sai Bhuvana Kurada

Arizona State University

Quality Assurance Analyst Technology Data Solution Tempe, Arizona, United States

Abstract: Federated Learning (FL) is an emerging machine learning paradigm with transformative potential in the field of cybersecurity. By enabling decentralized model training across distributed devices without transferring raw data, FL addresses critical privacy concerns while supporting real-time threat detection. This review synthesizes recent advances in the application of FL to cybersecurity, focusing on its deployment in domains such as the Internet of Things (IoT), intrusion detection systems, and threat intelligence networks. The review highlights key advantages, including enhanced data privacy, collaborative learning across diverse sources, and improved anomaly detection. It also identifies significant implementation challenges, such as communication overhead, computational inefficiencies, and susceptibility to adversarial attacks. Furthermore, the paper outlines major research gaps, including the lack of standardized benchmarks, limited real-world deployment studies, and the need for personalized federated models that address non-IID data and heterogeneous device capabilities. To bridge these gaps, the review proposes future research directions such as developing secure and efficient aggregation methods, prototyping FL systems in live environments, and advancing lightweight, adaptive FL frameworks. Overall, this review underscores the potential of FL to become a foundational technology in next-generation cybersecurity systems, enabling scalable and privacy-preserving threat mitigation across distributed infrastructures.

Keywords: Federated Learning, Cybersecurity, Intrusion Detection, Privacy-Preserving Machine Learning, Internet of Things (IoT) Security, Threat Intelligence

I. Introduction:

Federated learning, as an emerging method in the field of cybersecurity, presents itself as an opportunity to improve threat detection while minimizing risks associated with leaking sensitive information. As devices learn from each other's data point patterns without sharing relevant information, decentralized federated learning methods create a functional architecture for detecting and responding to potential cyber threats (Ghimire & Rawat, 2022). Through this architecture, data can remain private while federated learning enhances threat detection services and their functionality across varying environments to improve system resilience to attacks. Although privacy and efficiency improvements are the most prominent advantages of federated learning for architectural designs, its impact on cybersecurity and system privacy requires evaluation of the method's applications, challenges, and successful integration into existing systems.

Recent studies have highlighted the dual role of federated learning in both enhancing cybersecurity and addressing the security challenges posed by its own deployment in Internet of Things (IoT) environments (Ghimire & Rawat, 2022). By categorizing federated learning models into specific applications, researchers have been able to pinpoint its effectiveness in real-time threat detection and response strategies. This dual focus not only improves the security posture of IoT devices but also provides a framework for protecting federated learning systems themselves from potential vulnerabilities. For instance, the layered architecture of federated learning ensures that even if one part of the system is compromised, the overall network remains resilient and operational. As these advances continue, the integration of federated learning into cybersecurity protocols promises to enhance the robustness of threat detection mechanisms, while ensuring that user data privacy is maintained across decentralized networks.

Background and Motivation:

Federated Learning (FL) enables decentralized model training across multiple edge devices while ensuring that raw data never leaves local systems. This distributed paradigm is particularly significant in cybersecurity, where data sensitivity and privacy regulations demand minimal data exposure. By allowing devices to collaboratively train models on-site, FL reduces the risk of data breaches inherent in centralized systems and supports compliance with frameworks like GDPR and HIPAA.

In the context of cybersecurity, this architecture introduces transformative potential. It facilitates real-time anomaly detection, intrusion prevention, and adaptive threat responses at the network edge. As such, FL has been explored in various high-risk domains, including Internet of Things (IoT) security (Ferrag et al., 2021), spectrum monitoring (Sánchez et al., 2022), and cyber-physical systems (Consul et al., 2022). For example, Preuveneers et al. (2018) proposed chained anomaly detection models under a federated architecture for intrusion detection, demonstrating improved scalability and robustness while preserving data privacy.

However, the very properties that make FL attractive also present new challenges. Communication across distributed nodes increases network overhead, making FL vulnerable to latency, bandwidth constraints, and power limitations—especially in resource-constrained environments like UAV-based military systems (Consul et al., 2022). Moreover, the collaborative nature of FL opens up avenues for adversarial behavior. Attacks such as model poisoning, inference attacks, and compromised update aggregation pose serious risks to system integrity (Ghimire & Rawat, 2022; Koike et al., 2024).

Despite these limitations, the core architecture of FL aligns well with modern cybersecurity goals: privacy preservation, distributed resilience, and adaptive learning. As threats become more sophisticated and decentralized themselves—such as the rise of ransomware variants using polymorphic behavior—FL provides a scalable foundation for continuous, privacy-aware defense systems (Koike et al., 2024). Its ability to harness localized intelligence without compromising confidentiality makes it an ideal candidate for next-generation cyber defense strategies.

Application of Federated Learning in Cybersecurity:

The proposed application of federated learning in the field of cybersecurity can be described as a complex fusion of machine learning and decentralized computing that works towards enhancing existing security systems. Instead of leveraging centralized model that aggregates sensitive information into a common repository, federated learning models work in a distributed compute environment where it can access information from a wider array of data sources while ensuring the information remains private (Alazab et al., 2021). Such operational framework supports the need for real time threat detection as it can access information and data trends at a local level, which can prove to be useful in determining anomalous patterns that may compromise the integrity of different systems. Moreover, the use of this learning model within existing cybersecurity infrastructure makes it a suitable candidate for systems that would continue to evolve as new threats emerge without sacrificing information privacy or security of user data. Of course, the addition of federated learning to existing models would not be without challenges, particularly in the areas of network communication overhead and model convergence, which are still being studied actively by researchers (Ferrag et al., 2021).

Case studies and Examples:

Federated learning application in cybersecurity systems has started to reveal fruitful results in actual scenarios. For instance, the use of federated learning models in IoT security systems showed better performance for threat detections in distributed IoT networks (Ferrag et al., 2021). Since the IoT devices intervene with each other to detect threats and anomalies, federated learning enabled accurate and fast anomaly detections while keeping data privacy safe due to not centralized data. Also, applications in authentication methods have shown that information preciseness and security methods are enhanced with the use of federated learning due to less vulnerability against cyber threats since it uses common machine learning methods (Alazab et al., 2021). These studies suggest that there is potential for further applications of federated learning knowledge in cybersecurity systems and overall insights that can be used as fundamentals in future implementations.

II. Literature Review and Comparative Analysis

Federated Learning (FL) has gained significant traction as a viable approach to secure and privacy-aware cybersecurity solutions. Between 2021 and 2025, numerous studies have investigated FL's applicability in domains such as IoT security, intrusion detection, mobile networks, and data privacy enhancement. This section synthesizes five representative studies that collectively span technical diversity, application scope, and system architecture depth, offering a critical comparative view of FL in the context of cybersecurity.

Ferrag et al. (2021) present one of the earliest and most cited experimental studies demonstrating the effectiveness of federated deep learning for cybersecurity in IoT environments. Using custom IoT datasets, the authors implemented a decentralized threat detection architecture and reported improved detection accuracy, scalability, and privacy preservation compared to centralized approaches. However, the study also acknowledges communication overhead as a limiting factor in practical IoT deployments.

Ghimire and Rawat (2022) provide a dual-perspective review by analyzing both how FL contributes to cybersecurity and how FL itself needs protection against adversarial threats. The authors propose layered security mechanisms and discuss resilience to model poisoning attacks. Their analysis emphasizes the need for

robust encryption and secure aggregation techniques within FL protocols, making it a foundational work in securing federated architectures.

Anastasakis et al. (2022) move the research further by implementing a federated intrusion detection system (FL-IDS) integrated with differential privacy in IoT networks. Their experimental setup using synthetic data emphasizes privacy enhancement without compromising detection performance. Notably, they address the challenge of training on non-IID data and suggest that differential privacy adds noise that mitigates reverse-engineering risks, albeit at a small cost to accuracy.

Blika et al. (2024) provide a comprehensive survey exploring FL's role in securing 5G and 6G networks. The study categorizes FL frameworks across communication layers and emphasizes the importance of trustworthiness and scalability. It introduces taxonomy-based comparisons of FL architectures, identifies open challenges in cross-silo vs cross-device FL, and advocates for FL-powered trust management systems in heterogeneous, high-mobility environments.

Chitimoju (2025) explores FL in the context of digital transformation and innovation. The study proposes FL-driven privacy-preserving threat detection models and advocates for adaptive FL strategies that dynamically respond to changing threat landscapes. While primarily conceptual, the paper highlights the importance of real-time FL orchestration and proposes a high-level architecture for smart grid applications.

These studies collectively indicate that FL is effective in addressing core cybersecurity challenges such as privacy, decentralization, and scalability. However, real-world implementation barriers remain, particularly regarding high communication cost, attack vulnerability, and system personalization under heterogeneous data conditions. The next section elaborates on these open challenges and formulates future research directions aimed at making FL a deployable and secure solution for cybersecurity in critical systems.

Key Findings from Literature:

The synthesis of recent research reveals that Federated Learning (FL) presents a transformative opportunity for enhancing cybersecurity infrastructure while simultaneously preserving user data privacy. Several key insights have emerged from the literature:

Privacy Preservation:

FL enables local training of machine learning models on user or edge devices, thereby eliminating the need to transfer raw data to a central server. This architectural advantage aligns with modern data protection regulations and significantly mitigates the risk of privacy breaches. Studies by Ferrag et al. (2021) and Ghimire and Rawat (2022) confirm that this decentralized paradigm is especially beneficial in sensitive environments like IoT systems, where device-level data exposure could have critical consequences.

Real-Time Anomaly Detection:


By operating at the edge, FL facilitates real-time detection of threats and anomalous behaviors. Ferrag et al. (2021) demonstrated that FL can achieve low-latency response times in IoT security applications, effectively identifying distributed attacks without requiring centralized data collection. This capability is further validated by Anastasakis et al. (2022), who incorporated differential privacy to enhance security while maintaining performance.

Collaborative Knowledge Sharing:

FL fosters a collective intelligence framework in which multiple nodes contribute to a shared model. This approach enables robust generalization across diverse environments without sharing sensitive information. Blika et al. (2024) highlight this as a key advantage in securing heterogeneous networks like 5G/6G, where cross-device learning is essential for scalability and resilience. Ghimire and Rawat (2022) also emphasize the value of FL in enabling trust-based model sharing across different stakeholders in cybersecurity ecosystems.

Implementation Challenges:

Despite its promise, FL is not without limitations. The iterative exchange of model updates leads to significant communication overhead, which can be particularly problematic in bandwidth-constrained or mobile environments (Ferrag et al., 2021). Additionally, the challenge of training across non-identically distributed (non-IID) data complicates model convergence, leading to potential performance degradation (Anastasakis et al., 2022). Security risks such as model poisoning and adversarial updates remain active concerns, as discussed by Ghimire and Rawat (2022) and Blika et al. (2024), highlighting the necessity of robust defense mechanisms like secure aggregation and anomaly-aware client filtering.



Research Gaps:

Despite promising advancements in applying Federated Learning (FL) to cybersecurity, several unresolved challenges continue to hinder its widespread adoption in operational environments. A closer examination of recent literature from 2021 to 2025 reveals critical areas that require further investigation. These research gaps, if addressed, could enable FL to evolve into a scalable and secure solution for next-generation cybersecurity systems.

1. Standardized Benchmarks and Evaluation Protocols:

A persistent challenge across existing studies is the lack of standardized datasets and evaluation metrics to assess FL-based cybersecurity systems. As seen in works by Ferrag et al. (2021) and Anastasakis et al. (2022), many experiments are based on custom or synthetic datasets, which complicates performance comparisons and reduces reproducibility. Darwish and Roy (2025) further emphasize this issue in their comparative analysis of FL, deep learning, and traditional methods for malware detection, highlighting inconsistent dataset use as a major obstacle to model benchmarking. This inconsistency undermines confidence in the scalability and generalizability of current solutions.

2. Security of FL Architectures:

While FL enhances data privacy, it also introduces new vulnerabilities. Attack vectors such as model poisoning, backdoor injection, and inference attacks are increasingly relevant in FL-based deployments. Ghimire and Rawat (2022) call attention to the lack of mature mitigation mechanisms against these threats. Sánchez et al. (2022) provide further evidence of this concern, showing that many anti-adversarial FL implementations are not robust against adaptive cyberattacks in real-time IoT environments. The field lacks a standardized taxonomy and testing methodology for adversarial resilience, making security validation fragmented and incomplete.

3. Communication Efficiency and Resource Constraints:

Communication overhead continues to be a critical bottleneck in federated environments, especially when deployed across resource-constrained IoT devices or edge networks. The iterative model synchronization inherent in FL often results in excessive latency and energy consumption. Ferrag et al. (2021) and Blika et al. (2024) identify the need for model compression and efficient update schemes. Bodagala and Priyanka (2022) similarly observe that FL's benefits for IoT security can only be realized if communication protocols are optimized to fit real-world bandwidth limitations.

4. Lack of Real-World Deployments:

Many FL implementations remain confined to laboratory simulations or academic testbeds. There is a notable absence of real-world pilot deployments in sectors such as critical infrastructure, healthcare, and smart grid networks. Chitimoju (2025) stresses the importance of transitioning from theoretical models to live environments to ensure operational robustness. Zemel et al. (2023) further emphasize that cyber threats to critical infrastructure systems are growing in both complexity and impact, yet FL-based cybersecurity solutions remain largely untested in such high-stakes, real-world settings.

5. Personalization and Non-IID Data Handling:

One of the core challenges of FL is its struggle to generalize across clients with non-identically distributed (non-IID) data—a condition commonly found in cybersecurity applications. Standard FL algorithms assume homogeneous data across clients, leading to poor convergence and biased predictions when applied in practice. Anastasakis et al. (2022) and Blika et al. (2024) advocate for adaptive or personalized FL approaches that can accommodate local heterogeneity without degrading overall model performance. However, few implementations have successfully validated such strategies under adversarial pressure.

6. Robustness Against Sophisticated Malware and Evolving Threats:

As threat actors increasingly use polymorphic and metamorphic malware to evade traditional detection, FL models must adapt to continuously evolving attack patterns. Darwish and Roy (2025) argue that FL's current learning frameworks do not provide sufficient agility for rapidly shifting threat landscapes. More work is needed on continual learning and model retraining strategies that enable FL to respond to zero-day threats and unknown malware families.

Proposed Conceptual Framework

To address the implementation gaps identified in the previous section, we propose a conceptual framework that outlines the integration of Federated Learning (FL) into decentralized cybersecurity systems. This framework emphasizes privacy-preserving data training, collaborative threat detection, and continuous model optimization while preserving the core tenets of data confidentiality and system adaptability.




Fig 2 Conceptual Framework: Federated Learning for Decentralized Threat Detection

Figure 2 illustrates a multi-stage process that begins at the edge level with local devices (e.g., IoT sensors, mobile endpoints) performing on-device training without transmitting raw data. Instead, only encrypted or aggregated model updates are shared with a central federated aggregator (e.g., cloud server or blockchain-enabled node), which generates a refined global model.

This global model is then redistributed to all participating devices, enhancing their local detection capabilities. A dedicated threat detection module integrates the aggregated intelligence to identify intrusions and anomalies in real time. The lower part of the framework emphasizes security enhancements, including secure aggregation, defense against adversarial attacks, and anomaly-aware model filtering. Finally, the system enters a collaborative intelligence loop that supports continuous learning and adaptation as new threats evolve.

This architecture effectively balances data privacy, scalability, and real-time responsiveness, making it suitable for cybersecurity deployment in diverse, heterogeneous environments such as smart cities, critical infrastructure, and industrial IoT ecosystems.

III. Conclusion

Federated Learning (FL) offers a paradigm shift in the design of secure, privacy-aware cybersecurity systems. By decentralizing model training and preserving sensitive data at the edge, FL aligns with contemporary privacy regulations and enables scalable real-time threat detection. As this review has demonstrated, FL's applications span across IoT, Industrial IoT (IIoT), authentication systems, and blockchain-enabled environments, providing a flexible defense architecture capable of adapting to evolving cyber threats. Recent advances validate FL's potential across various threat detection models and frameworks. For instance, Markovic et al. (2022) successfully implemented a random forest-based FL model for intrusion detection, showing promising performance in distributed setups. Meanwhile, Yazdinejad et al. (2022) introduced "Block Hunter," a federated framework for cyber threat hunting in blockchain-integrated IIoT systems. Similarly, the integration of blockchain with FL, as surveyed by Ali et al. (2024), highlights a growing consensus around combining distributed trust mechanisms with decentralized learning to enhance cybersecurity resilience. Kumar et al. (2024) further underscore the importance of efficient secure aggregation and trust management in FL to build privacy-preserving infrastructures for critical IoT environments.

However, challenges such as computational overhead, communication inefficiencies, and the need for standardization continue to limit real-world adoption. The variability in datasets, lack of unified evaluation protocols, and vulnerability to adversarial attacks emphasize the need for more robust and scalable FL models. Additionally, personalization of models in non-IID environments and deployment in live, multi-platform infrastructures remain underexplored.

In summary, FL represents a critical innovation for future-proofing cybersecurity in decentralized digital ecosystems. To fully unlock its potential, ongoing research must address implementation bottlenecks while fostering collaboration between academia, industry, and regulatory bodies. A well-optimized federated framework—integrating trust, scalability, and security—can serve as the backbone for next-generation cybersecurity systems equipped to defend against the increasingly sophisticated and distributed nature of modern cyber threats.

References:

- [1]. Alazab, M., Rm, S. P., Maddikunta, P. K. R., Gadekallu, T. R., & Pham, Q. V. (2021a). Federated learning for cybersecurity: Concepts, challenges, and future Transactions Informatics, directions. on IEEE Industrial 18(5), 3501–3509. <https://ieeexplore.ieee.org/abstract/document/9566732/>
- [2]. Alazab, M., Rm, S. P., Maddikunta, P. K. R., Gadekallu, T. R., & Pham, Q. V. (2021b). Federated learning for cybersecurity: Concepts, challenges, and future Transactions Informatics, directions. on IEEE Industrial 18(5), 3501–3509. <https://ieeexplore.ieee.org/abstract/document/9566732/>
- [3]. Ferrag, M. A., Friha, O., Maglaras, L., Janicke, H., & Shu, L. (2021). Federated deep learning for cyber security in the internet of things: Concepts, applications, and experimental analysis. IEE Access, 9, 138509–138542. <https://ieeexplore.ieee.org/abstract/document/9562531/>
- [4]. Ghimire, B., & Rawat, D. B. (2022). Recent advances on federated learning for cybersecurity and cybersecurity for federated learning for internet of things. IEEE Internet of Things Journal, 9(11), 8229–8249. <https://ieeexplore.ieee.org/abstract/document/9709603/>
- [5]. Katari, P., Bonam, V. S. M., Bojja, S. G. R., Ravi, Venkataramanan, Decentralized C. S. S., & (2021). Cybersecurity: Implementing Federated Learning in Threat Intelligence Networks. Journal of Informatics Education and Research, 1, 29–40. https://www.researchgate.net/profile/Srinivasan-Venkataramanan-2/publication/389660179_Decentralized_Cybersecurity_Implementing_Federated_Learning_in_Threat_Intelligence_Networks/links/67cbc585cc055043ce6f45bc/Decentralized-Cybersecurity-Implementing-Federated-Learning-in-Threat-Intelligence-Networks.pdf
- [6]. Kumar, M. A., Mohammed, A., Sumanth, S., & Sivanantham, V. (2025). Enhancing Cybersecurity Through Federated Learning: A Critical Evaluation of Strategies and Implications. In Model Optimization Methods for Efficient and Edge AI: Federated Learning Architectures, Frameworks and Applications (pp. 281–297). Wiley Online Library. <https://onlinelibrary.wiley.com/doi/abs/10.1002/9781394219230.ch14>
- [7]. Sarhan, M., Layeghy, S., Moustafa, N., & Portmann, M. (2023). Cyber threat intelligence sharing scheme based on federated learning for network intrusion detection. Journal of Network Management, and 31(1), Systems 3. <https://link.springer.com/article/10.1007/s10922-022-09691-3>
- [8]. Ghimire, B., & Rawat, D. B. (2022). Recent advances on federated learning for cybersecurity and cybersecurity for federated learning for internet of things. IEEE Internet of Things Journal, 9(11), 8229–8249.
- [9]. Anastasakis, Z., Psychogios, K., Velivassaki, T., Bourou, S., Voulkidis, A., Skias, D., ... & Zahariadis, T. (2022, September). Enhancing cyber security in IoT systems using FL-based IDS with differential privacy. In 2022 Global Information Infrastructure and Networking Symposium (GIIS) (pp. 30–34). IEEE.
- [10]. Blika, A., Palmos, S., Doukas, G., Lamprou, V., Pelekis, S., Kontoulis, M., ... & Askounis, D. (2024). Federated Learning For Enhanced Cybersecurity And Trustworthiness In 5G and 6G Networks: A Comprehensive Survey. IEEE Open Journal of the Communications Society.

Biographies

Sai Bhuvana Kurada received her Bachelor of Technology (B.Tech) in Computer Science in 2018 and completed her Master's in Information Technology Project Management in 2023. Between her academic pursuits, she gained industry experience working with multiple organizations as a Business Analyst. She is currently serving as a Quality Assurance Analyst. Her professional interests include software quality assurance, project coordination, and IT process optimization, cyber security, AI and ML including data analysis.

