

Review Paper

Federated and Privacy-Preserving Machine Learning for Collaborative Threat Intelligence across Untrusted Domains

^{1*} Naga Madhusudana Rao Chadaram

^{1*} *Independent Researcher, United States. Email Id: nmrchadaram@gmail.com*

**Corresponding Author(s): nmrchadaram@gmail.com*

Received: 18/12/2025

Revised: 28/01/2026

Accepted: 22/03/2026

Published: 31/03/2026

Abstract: The increasing sophistication of cyber threats and the hindrances in sharing information between various regions demonstrate that one location of all the data is not sufficient for collaboration in detecting threats. This study provides a close look at federated and privacy-safe machine learning, which allows many untrusted groups to exchange information without providing their data. The study was pertained with PRISMA as a method to review 62 studies that were published between 2021 and 2026 in the Scopus and the web of science databases. According to the results, 79% of the studies affirmed that federated learning can enhance joint threat detection, and 74.2% emphasized privacy-sensitive and mixed methods. Nevertheless, 66.1 % still have issues with attackers, 59.7% experience meaningful trade-offs under usefulness and confidentiality, and 29% do not consider how to deal with zero-trust designs. These findings imply that despite the rising utility of the federated approach, its application is still diffused because privacy, security, and trust do not go hand in hand. The research states that federated cyber systems might work even better, in case there is only one framework, which combines adaptive privacy regulations, robust defense against attacks, and trust-conscious rules. This research contributes to both knowledge and practice by providing a comprehensive overview that can be used to develop large, safe, and reliable joint threat-detection systems.

Keywords: Federated Learning, Privacy-Preserving Machine Learning, Collaborative Threat Intelligence, Adversarial Robustness, Zero-Trust Architecture

1 Introduction

The increasing threat level in the cyber domain and intensification of the interconnectedness of the digital provision of services to companies and nations have led to the need to disseminate threat intelligence among communities [1]. Centralized machine learning models are only effective in organizations where all data remain central, whereas in reality, organizations are prohibited by data silos, rules, and privacy concerns from distributing information. This issue is worse in cybersecurity because sensitive information is secured by law and mistrust. Federated learning (FL) allows the training of a model by multiple organizations with their raw data remaining confidential [2]. This is able to break the primary data-privacy and data-possession obstacles. However, FL introduces additional issues as first lack of speed, and scheduling, and variability of device possibilities so further investigation is required on whether it works in practice or not.

New technology, tougher regulations, and smarter cyber-attacks are some of the factors that have contributed to the growing significance of privacy-preserving machine learning. Federated systems are employed to ensure that data are kept confidential through methods such as differential privacy, secure multiparty computation, and homomorphic encryption [3]. These techniques have massive advantages, as well as scaling down the usefulness of a model, consuming greater computer resources, and creating scaling constraints, particularly with large and real-time cybersecurity issues. The vulnerability of federated systems to attacks that corrupt data or alter the model is also a disadvantage of the reliability of the system. These issues present the necessity of a unified framework to deal with privacy, security, and performance as a whole [4].

Meanwhile, trust, governance, and zero-trust architecture (ZTA) have received more attention. These strategies consider the entire position of security and not necessarily the technology. Other proposals take advantage of incentives and



decentralized trust to motivate participants and minimize bad behavior, but are yet to be prevalent. A system combining FL and the ideas of ZTA such as checking trust always and providing the minimal necessary access might help to make the systems more resilient. It is also a complication that leads to a difficult system to scale. There is a tension between ideas and real-world implementation in the field; hence, a systematic review of what exists is required.

It is on this backdrop that this paper provides an analysis of the present condition of federated and privacy-preserving machine learning to share threat intelligence within untrusted areas. The overall research questions were as follows:

- What is the impact of decentralized learning models on the performance and size of cross-domain threat sharing?
- What are the most useful privacy methods that make balance between security and usefulness?
- How does it affect the federated systems of attacks and robustness problems?
- To what degree are systems constructed today with trust, governance, and zero trust conceptions?

The study responds to these questions by considering authoritative papers on cyber security, network systems, relocation between organizations, and data sharing, and not papers that are not peer-reviewed or those that have poor designs [5].

This study contributes to academic knowledge and practice by providing both theme-based and quantitative analyses. Theoretically, it assists in understanding the interaction of decentralization, privacy, and defense against attacks and corrects the gaps in the existing theories. In practice, it provides a good demonstration for companies that desire to establish secure and scalable means of exchanging threat indications [6]. This study provides a strong foundation for future research and assists in translating concepts on paper to practice.

2 Methodology

The proposed study is a systematic review, and the PRISMA rules were used to ensure that all of the work is clear, repeatable, and strict. The selected papers are the most appropriate peer-reviewed studies using Scopus, Web of Science, and Google Scholar. The most searched keywords were federated learning, privacy-preserving machine learning, cybersecurity, threat intelligence, differential privacy, secure multiparty computation, homomorphic encryption, adversarial robustness, and zero-trust architecture as a combination of words. The study filtered English papers that had been published in 2021-2026 to ensure high quality and availability in either Scopus or Web of Science.

All the studies were also retained that mentioned directly things about federated or decentralized learning in cybersecurity or sharing data across domains, which took privacy or strong security measures, and involved adding actual data, models, or concepts. The process eliminated duplicate records, non-peer reviewed articles, articles that were ambiguous or those that did not qualify under cybersecurity. The PRISMA process included four steps: locating records, reading titles and abstracts, browsing full texts and selecting final papers. Following this, it resulted in

curation of 62 articles to read and examine.

The coding system that was employed to collect data was fixed in order to suit the objectives of the study. Various variable is extracted including the author, date, and region, kind of approach (quantitative, modeling, mixed, and qualitative), area of application (such as intrusion detection, IoT, or cross sector systems), the system design, the privacy tools that were applied, any security risk, and general findings. Later the rated the quality of each study using PRISMA and CASP guidelines based on the clarity of the study, its reproducibility, and usefulness. Papers that were unclear or those that were of high bias were dismissed. Where the studies contradicted, it contrasted and put them into perspective and have tracked differences in methods rather than discarding them in order to gain deeper understanding.

The synthesis of data used both the thematic and quantitative approaches. Thematic synthesis aided the development of six key categories that are comparable to variables while considering. The data has also summed up numbers and percentages to determine the frequency of each thing and this allowed us to compare and relate the results of different studies. Then overall synthesis has been applied to explain what it is observed in the data by theory in the discussion. The studies differed greatly in design, data, and locations; and grouped similar types and compared them, thus the differences benefited rather than harmed our analysis. Such a stratified approach maintains a clear connection between the data collection process, presentation process, analysis, and evidence-based conclusion.

3 Literature Review & Thematic Analysis

3.1 Cross-domain Threat Intelligence Decentralized Learning Architectures

Recent articles indicate that federated learning (FL) can transform the sharing of threat information in organizations. It allows lots of groups to jointly train models without exchanging their raw data among them. Research is united in its view that FL eliminates legal and organisational barriers to the sharing of data in particularly the areas where the rules are stringent such as GDPR. Nevertheless, the most appropriate structure is controversial. The layout of most old systems is to have a central point, whereas in more modern work, the layout is either layered or peer-to-peer to avoid central hubs.

Simulation based research has demonstrated that an intrusion detector system that is distributed is more effective in detecting threats. In the real world, deployments however have delays and challenges on how to make all components work in tandem [7]. Distributed systems theory generally provides theoretical work, analyzing the effectiveness with which parts communicate and their stability as they reach a solution. Nonetheless, actual tests in other types of networks are also scarce. One of their major issues is how to expand systems when they are confronted with different degrees of trust. By being decentralized, systems become more resilient to attack, however coordination is also more difficult to achieve particularly when the users hold differing opinions on risk. The work in the future should include intelligent methods of data combination that can alter the levels of trust in different regions [8].

3.2 Privacy-preservation Cybersecurity Systems in Federated Systems

The fundamental aspect of threat intelligence is privacy-protecting techniques. The most common ones are the differential privacy (DP), secure multiparty computation (SMPC), and the homomorphic encryption (HE). The DP is easily scalable, very fast, although it introduces a noise that reduces accuracy. HE and SMPC would offer better security but will be much slower. Researchers tend to confuse DP with cryptographic techniques to strike a balance between speed and security [9]. It is contested as to the most appropriate settings: others claim that DP is sufficient in low-risk environments, and some caution that DP can be compromised in case other information is available to the attacker.

Theory relies on the information and cryptographic models, but the utilization of these concepts in the real time systems has not been researched. The relevance of accumulating privacy loss over multiple rounds of training is not explored by many articles, which is relevant when returning threat information in real-time. In the literature, the ability to dynamically alter the protection level in accordance with the threat at hand reflects an absence of privacy mechanisms that would alter the protecting mechanisms across various contexts.

3.3 Federated Environment Adversarial Robustness and Poisoning Resistance

With federated learning, adversaries, in particular, data poisoning or model manipulation adversaries, present a significant challenge. The experiments indicate that malicious participants are able to severely harm the general model through small gradients modifications. Other stronger resistance mechanisms such as Krum and other Byzantine-resilient algorithms are available though they are not that effective against smart collusion attacks.

It is evident that the outcomes of controlled simulations are different to actual attacks. Real tests of varied and asynchronous networks expose weaknesses in simulations, which look strong. Theoretical Byzantine models provide some rough foundations but lack intelligent tactics of intelligent attackers.

The question of the balance between safety and speed is controversial. Strong-tie techniques are more expensive in terms of CPU and convergence rate, and are only applicable in large networks. Most users are assumed to be good in detection systems, which might not be the case in untrusted environments. One of the significant gaps in research is to integrate the identification of bad actors and privacy protection. Existing solutions tend to address these components independently causing disjointed security architectures.

3.4 Trust, Incentive Design and Governance in distrusted collaborative ecosystems

Federated threat intelligence focuses on trust, which is not well defined. Game theory and economics allow researchers to develop incentives to motivate employees to work honestly [10]. The intention of the token-based or reputation systems is to prevent free riders and bad contributions. Nonetheless, these models have limited experiments that test them, although most of them depend on simulations. Sociotechnical perspectives emphasize that institutions and governance are required as well and, therefore, pure tech solutions are insufficient. Several issues still exist between full decency of trust systems altogether and the necessity of central control to ensure that

people are accountable.

Hybrid approaches between blockchain and federated learning are also on the rise, and provide audit trails and immutable records. They are difficult to roll out though their size and energy consumption is constrained. The literature reveals that there exists a crucial gap in introducing the issues of formal governance to concepts of technical trusts. The future research that needs to be done is the hybrids, where the institutional regulation can be combined with decentralized verification [11].

3.5 Trade-offs among utility and privacy and Scalability Constraints of a System

One such theme is that increased privacy is generally associated with reduced model performance. The research is in agreement that stricter privacy declines precision, particularly on cyber data that is high-dimensional. The latter issue is more severe in great, federated systems, in which data is distributed differently. Various researches attempt to strike a balance between the forces. Others attempt to use optimization to seek a sensible privacy budget, and others attempt to do adaptive learning rate or personalized models to retain its performance. Multi-objective optimization is commonly employed by theories; however, the additional computation and communication are hard to accomplish by real systems.

Scaling solutions have been disputed where the asynchronous updates and compression assist the acceleration of but may result in errors in convergence of models. The fact that edge devices are less powerful makes it more difficult to run heavy privacy methods. The primary gap in the research is to design scalable structures that ensure high performance without compromising on privacy, particularly in the context of threat detection in real-time. This demands novel methods of combining edge computing and dynamic privacy, which is sportswear.

4 Results

4.1 Overview of Included Studies

The search resulted in 62 publications that matched our search criteria, were published between 2021 and 2026.. Among them, 28 articles (45% wrote about experiments and simulations and tested federated learning models). A previous number of 20 studies (32%) have used models and algorithms to examine privacy protection and defenses against attacks. Experiments were intertwined with theory in 9% studies. Five of them (8%) were predominantly qualitative, discussing governance and trust.

The majority of studies concerned the intrusion detection and cyber security (34 studies, 55%). The subsequent common ones were distributed network security (14 studies, 23%), IoT and edge computing (9 studies, 15%), and data sharing across sectors (5 studies, 8%).

The study has its origin across the globe. Asia was the most represented and had (24) studies (39%), Europe (18) was next (29%), North America (15, 24%), and other locations (5, 8%). Published studies also increased over recent years: 39 (63) studies, which confirms that research is increasing.

Table I. Distribution of Included Studies (n = 62)

Category	Subcategory	n	%
Methodology	Quantitative	28	45.20%
	Modeling/Algorithmic	20	32.30%
	Mixed Methods	9	14.50%
	Qualitative/Conceptual	5	8.00%
Application Domain	Cybersecurity/IDS	34	54.80%
	Network Security	14	22.60%
	IoT/Edge Systems	9	14.50%
	Cross-sector Data Sharing	5	8.10%
	Geographic Distribution	Asia	24
	Europe	18	29.00%
	North America	15	24.20%
	Other Regions	5	8.10%

4.2 Core Findings Related to the Research Topic

The study identified six broad categories of results in the 62 studies. The identified study indicated federated learning to be effective in detecting threats jointly in most studies, which accounted for 49 results (79%). The methods of privacy protection were researched in 46 studies (74%). The major methods were disability privacy and cryptography. Problems with remaining safe against attacks were observed in 41 studies (66%). These were primarily poisoning attacks and gradient alterations. There were problems of trust and rules in 29 studies (47%). Trade-off between usefulness and privacy was prevalent (59.7%) in 37 studies. Finally, the least investigated field was the use of zero-trust-designs (18 studies, 29%).

Trends indicate that 32 articles (52%) were concerned with privacy tools and attack resistance simultaneously,

whereas 21 articles (34%) were concerned with design and scale simultaneously.

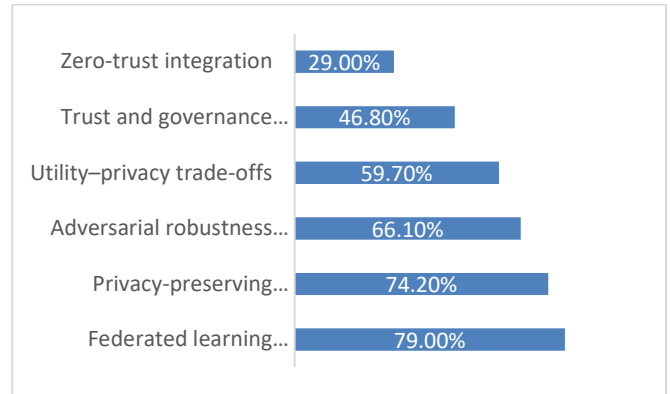


Fig.1. The frequency distribution of core findings

4.3 Trends / Patterns of implementation

Among the 62 studies, 44 (71%) studies contained the methods of effectively implementing solutions and 18 (29%) were only considered as conceptual. Out of the studies on implementation, 26 (41.9%) were enterprise cybersecurity systems, 11 (17.7%) were IoT/edge deployments, and 7 (11.3%) cross-organizational intelligence sharing platforms. The data privacy regulations (38 studies, 61.3%), the necessity to share the threat information (35 studies, 56.5%), and the progress of distributed computing (29 studies, 46.8%) were the leading motives. The largest obstacles were the computational cost (33 studies, 53.2%), the delay caused by communication (27 studies, 43.5%), and the distrust in subjects (25 studies, 40.3%). The patterns of operation revealed that centralized data aggregation was made in 36 studies (58.1%), whereas 26 studies (41.9%) used the operational patterns of decentralized or hierarchical design. More 31 studies (50.0%), and 15 studies (24.2%), enacted hybrid privacy approaches and single approach respectively.




Fig.2. Summarizes adoption trends and implementation patterns.

4.4 Comparative Analysis

The comparative analysis of the 62 studies indicates a varying level of the wellness of the methods, the number of people that it can assist and its safety levels. Thirty-eight studies utilized differential privacy (DP) (61.3%), 27 of them utilized secure multiparty computation (SMPC) (43.5%), and 24 studies utilized homomorphic encryption (HE) (38.7%). In

30 studies, DP studies were more scalable (48.4%), and SMPC and HE studies had a superior privacy (34 studies, 54.8%). 33 Studies (53.2%) utilized robust aggregation methods and 21 studies (33.9%) used Byzantine-resilient methods. Several (28, 45.2%) studies indicated that they were not very efficient particularly in case the systems were large. In 25 studies (40.3%) centralized FL models converged quicker than decentralized models however in 22 studies

(35.5%) decentralized models were more resistant to fault. Other studies were based on hybrid architecture 15 studies (24.2%) combining both.




Fig.3. Summary of the comparative analysis of key approaches

5 Discussion

5.1 Discussion of Significant Results

The fact that the effectiveness of federated learning (FL) prevails in the reviewed studies implies that the process of decentralized model training is becoming increasingly feasible when it comes to collaborative threat intelligence where data is sensitive and there is a lack of trust between institutions. This fact can be explained by the fact that FL is naturally consistent with the distributed systems theory in which localized data processing eliminates potential exposure and preserves group model performance. Nonetheless, the high frequency of implementation-focused research and the consistent availability of the researchers discussing the challenges of latency and synchronization indicates that the operational maturity of FL systems is limited due to the heterogeneity of the infrastructure and the communication overheads.

The use of privacy-saving methods, especially hybrids, is an indication that there is an increasing awareness that no single mechanism is adequate with regard to the multi-dimensionality of privacy threats. This reliance on differential privacy (DP) as a means of scalability, even when not theoretically optimal, as was revealed, shows that deployability and not theoretical optimality are considered by practitioners in real-world conditions that involve cybersecurity. At the same time, the fact that the adversarial robustness concerns are still relevant implies that privacy and security are not converging dimensions and are still operationally separated, resulting in fragmented system designs.

The comparatively fewer number of integrations of zero-trust architecture (ZTA) models point to the immaturity of comprehensive security models integrating identity authentication and decentralized learning. This trend indicates that FL has addressed the issue of data-level privacy; however, system-level trust enforcement is not well developed, especially when working across domains with heterogeneous parties.

5.2 Comparison with the Existing Literature

The results are in a large extent supported by the recent surveys concentrating on the dual role played by FL in facilitating the secure cooperation, as well as providing plenty of surfaces to attacks. For example, [12] and [3] found that FL enhances the accuracy of intrusion detection, but it is sensitive to poisoning attacks in an adversarial environment. This has been extended by the current analysis, which shows that these vulnerabilities still occur even after the implementation of strong aggregation mechanisms and represent a shift between theory and practice.

Regarding privacy-saving features, the findings are in line with those of [5], who noted that hybrid models contain the best trade-off between efficiency and security. Nevertheless, the still prevalent use of DP as a leading strategy is opposite to the new research findings that indicate its vulnerability to inference attacks when auxiliary knowledge is present. The lack of uniformity points to a disparity between model validation in the experimental setting and a model of the threat in the real world.

The identified focus on trust and governance mechanisms can be explained by the sociotechnical views of cybersecurity studies, which claim that technical solutions are ineffective in guaranteeing safe cooperation. However, the fact that the empirical support for incentive models found in the current study is rather scant reflects the possibility of researchers exaggerating their application in practice in the existing literature. Additionally, the comparably poor incorporation of ZTA is in contrast with the new conceptual frameworks that discuss its need to secure distributed AI systems, which shows that the theoretical knowledge was uncovered quite a long time ago, and the practice is still undergoing.

5.3 Theoretical Implications

The research indicates that Federated cybersecurity systems require frameworks to integrate distributed learning, privacy protection, and resistance to attacks. Much of the existing work takes place in independent fields such as distributed optimization, information theory, and fault tolerance, and they are not easily balanced. This division presents a challenge to existing models that are unable to explain fully how risky collaboration can be managed.

Privacy versus utility is always a tradeoff that demonstrates the importance of multi-objective optimization. However, our findings indicate that fixed optimization models are inaccurate when changes in threats are involved. It also requires versatile models that are driven by risk checks using contextual scenarios and amendments to settings on-the-fly.

Another is that FL systems are not in the business of providing trust mechanisms, and there is a gap in the way to think of trust as something that can change and is measurable. Game theory only provides hints but disregards the actions of institutions and individuals. It should be able to apply mixed social-technical theories with enhancement to governance, rewards, and rules.

5.4 Implications in Practice and Industry

Practically, it requires firms relying on FL as the tool of threat intelligence to employ mixed privacy approaches that ensure that business processes remain flowing and the rules are observed. Due to the intensive level of computing capabilities, and communication, large-scale deployments demand improved facilities like edge computing network and effective transmission of information.

There are also trust issues, which are a great impediment and demonstrate that there must be rules and governance, as well as technology. To gain confidence in shared systems, companies should introduce accountability tools, such as logs and compliance checks.

Moreover, few organizations apply zero-trust architecture (ZTA) in FL, which means that they are potentially squandering the strength of identity-based security. Insider risks may be minimized by adding continuous verification and high-strength access control to prevent bad actors, particularly when the work of different companies is integrated.

The findings also indicate that specialists cannot use a single-size-fits-all environment but must choose privacy and resilience tools depending on the threats they face. This renders the system more adaptable and allows it to work in other environments.

5.5 Contradictions and Surprising Discoveries

There were contradictions in a few places in the manuscript. First, it is more evident that theory claims that decentralized systems are more resilient than centralized systems; however, decentralized systems are applied less often than centralized ones. This implies that theory is not equivalent to practice, which is likely due to the difficulty in coordinating decentralized systems and increased communication.

Second, numerous studies adopt Differential Privacy (DP) despite its limitations, which is practically selective favoring tools designed to be scalable and offer diminished protection. This contrasts with the belief that more powerful encryption would be applicable in situations with the greatest risk.

Third, many projects are being deployed, but they continue to complain about scalability and robustness. This implies that deployments may be critical in demonstrating ideas in the concept but not permanence in functioning. This demonstrates the distance between laboratory achievements and practice.

Finally, Zero-Trust Architecture is not commonly used, although it is well suited to FL. Mixed security models are cumbersome and do not allow adoption. More convenient

methods of adding them have to be available to reduce the cost of the effort.

5.6 Limitations of the Study

This research paper is subject to a number of limitations that typify systematic literature reviews. First, the study utilized only peer-reviewed journal articles indexed in Scopus and Web of Science; thus, it may have missed good conference papers, particularly in emerging disciplines such as machine learning and cybersecurity. Second, the majority of the articles used simulations; therefore, the findings might not be applicable to real-world cases where hackers are more diverse and attack systems are more complex. Third, the classification of studies in distinct themes may withhold the intricate complexity of the factors of federated learning, in which design, privacy, and security tend to intersect. Fourth, 2021-2026 only provides new work without including older underlying work that provides significant theory. Finally, the summary does not depend on the quality of the studies, data sets, or type of experiment, as it would alter the findings.

5.7 Future Research Directions

Further efforts in this area should involve the development of versatile federated learning infrastructures that can modify privacy and robustness configurations in real time based on more recent threat data. Such systems are capable of overcoming the optimization model constraints of fixed optimization models and have made the system resilient to evolving attacks. Another area that should be provided is security designs that incorporate privacy protection, attack detection, and trust management. To find complete solutions, machine learning, cryptography, and cybersecurity policy work will be required.

To better estimate the performance of the systems in the world, real-life tests with large numbers of participants and types of networks should be conducted. Future research needs to examine federated learning, which makes use of trust, such as redefining the reputations and regulations of organizations. This would help motivate individuals to participate and minimize the risk of bad actors. Finally, federated learning should be further developed in combination with zero-trust security protocols. This assists in the simplification of systems and collaboration. The creation of standard frameworks and rules will see more people utilize it and maintain the security level within the networks.

6 Conclusion

The review indicates that federated privacy-preserving machine learning is a feasible method for sharing threat data among groups that are not mutually trusted. 62 studies are consistent with the fact that models can be safely trained without central control. In the vast majority of studies (79%), federated learning is proven to be effective in identifying threats more effectively without violating the privacy of the data. In addition, 74% of the studies indicate that privacy tools, particularly mixed ones, are instrumental in maintaining the security level high and elasticity of the system. However, the results also demonstrate persistent issues: 66% of the studies report attacks that pass through, approximately 60% of the studies report a trade-off between usefulness and privacy, and 47% of the studies report trust issues. This translates to the fact that existing systems are not unified entirely.

To the industry members, the situation is that despite the

fact that federated learning systems are prepared to be employed more, they can only perform well when they are combined with privacy, strength as well as governance. As an investor, it is clear that there is bad news that 71% of the work consists of the implementation of the ideas, but still there exists the opportunity to develop something new in the field of the safe group data analysis. In practice, to maintain an acceptable balance between speed and safety companies ought to employ a combination of cryptography and differential privacy.

Overall, the research confirms that decentralized learning, adjustable privacy tools, and joined trust design are the three approaches that federated cyber security will have in the future. The fact that there is still low usage of zero-trusts, and attacks continue to occur signifies that the field is transitioning to a full system implementation. Therefore, both researchers and industries need to develop cohesive, scalable systems that translate these pieces into real-life scenarios. This is one of the reasons why making this integration better will be crucial to developing safe, strong and trustworthy collective threat-intelligence systems.

correlation analysis-based federated learning framework for defending against collusion-free-riding attacks,” *Cybersecurity*, vol. 8, no. 1, Sep. 2025, doi: 10.1186/s42400-025-00366-5.

References

- [1] M. H. Alsharif, R. Kannadasan, W. Wei, K. S. Nisar, and A.-H. Abdel-Aty, “A contemporary survey of recent advances in federated learning: Taxonomies, applications, and challenges,” *Internet of Things*, vol. 27, p. 101251, Oct. 2024, doi: 10.1016/j.iot.2024.101251.
- [2] R. Gosselin, L. Vieu, F. Loukil, and A. Benoit, “Privacy and Security in Federated Learning: A Survey,” *Applied Sciences*, vol. 12, no. 19, p. 9901, Oct. 2022, doi: 10.3390/app12199901.
- [3] H. U. Manzoor, A. Shabbir, A. Chen, D. Flynn, and A. Zoha, “A Survey of Security Strategies in Federated Learning: Defending Models, Data, and Privacy,” *Future Internet*, vol. 16, no. 10, p. 374, Oct. 2024, doi: 10.3390/fi16100374.
- [4] H. Li, L. Ge, and L. Tian, “Survey: federated learning data security and privacy-preserving in edge-Internet of Things,” *Artificial Intelligence Review*, vol. 57, no. 5, Apr. 2024, doi: 10.1007/s10462-024-10774-7.
- [5] E. M. Timofte et al., “Federated Learning for Cybersecurity: A Privacy-Preserving Approach,” *Applied Sciences*, vol. 15, no. 12, p. 6878, Jun. 2025, doi: 10.3390/app15126878.
- [6] R. Damaševičius, “Introductory Chapter: Recent Trends and Progress in Support Vector Machines,” *Federated Learning - A Systematic Review*, Apr. 2025, doi: 10.5772/intechopen.1005410.
- [7] “Federated Learning for Cybersecurity: Decentralized Threat Detection in Large Networks”, *WJFTCSE*, vol. 1, no. 2, pp. Apr (28–38), Nov. 2025, doi: [10.63345/wjftcse.v1.i2.103](https://doi.org/10.63345/wjftcse.v1.i2.103).
- [8] Y. Feng et al., “A survey of security threats in federated learning,” *Complex & Intelligent Systems*, vol. 11, no. 2, Jan. 2025, doi: 10.1007/s40747-024-01664-0.
- [9] N. Latif, W. Ma, and H. B. Ahmad, “Advancements in securing federated learning with IDS: a comprehensive review of neural networks and feature engineering techniques for malicious client detection,” *Artificial Intelligence Review*, vol. 58, no. 3, Jan. 2025, doi: 10.1007/s10462-024-11082-w.
- [10] P. Santos, R. Abreu, M. J. C. S. Reis, C. Serôdio, and F. Branco, “A Systematic Review of Cyber Threat Intelligence: The Effectiveness of Technologies, Strategies, and Collaborations in Combating Modern Threats,” *Sensors*, vol. 25, no. 14, p. 4272, Jul. 2025, doi: 10.3390/s25144272.
- [11] D. Sirohi, N. Kumar, P. S. Rana, S. Tanwar, R. Iqbal, and M. Hijjii, “Federated learning for 6G-enabled secure communication systems: a comprehensive survey,” *Artificial Intelligence Review*, vol. 56, no. 10, pp. 11297–11389, Mar. 2023, doi: 10.1007/s10462-023-10417-3.
- [12] M. Xue, H. Zhong, Y. Shi, Y. Zeng, J. Zhang, and N. Zhao, “A