

Privacy-Preserving Cybersecurity in Federated Learning: A Comprehensive Review of Homomorphic Encryption and Decentralized Threat Intelligence Sharing

Nitin Shankarrao Shrira¹, Dr. Babita Tyagi²

¹Research Scholar, Department of Management, Chaudhary Charan Singh University, Meerut, Uttar Pradesh, India.
E-mail: nshrira@gmail.com

²Assistant Professor, Department of Management, Chaudhary Charan Singh University, Meerut, Uttar Pradesh, India.
E-mail: tyagibbitta@gmail.com

Abstract--- The fast-paced developments of cyber-attacks alongside strict data privacy laws (such as GDPR, DPDP act, CCPA) led to a growing need for collaborative cybersecurity approaches that would not endanger sensitive data. Federated Learning is a technique allowing for training of models without revealing underlying data, Homomorphic encryption enables computations with encrypted data, and decentralized sharing of threat intelligence information increases collective defense capacity. This paper reviews current research trends regarding the integration of HE with FL in the context of cybersecurity for the period of 2020-2025. The main innovations in the field under analysis include hybrid techniques based on Multi-Key Homomorphic encryption, Differential Privacy, and blockchain technology for secure aggregation. Applications of HE-based FL in intrusion detection, malware classification, and IIoT demonstrate very good results in terms of privacy preservation, along with relatively moderate reductions in accuracy of about 2-5%. The critical issues of computational complexity, communication delays, and protection against collusions are considered. Future directions in the development of homomorphic, scalable, and blockchain-based cybersecurity technologies are discussed.

Keywords--- Federated Learning, Homomorphic Encryption, Privacy-Preserving Cybersecurity, Decentralized Threat Intelligence, Differential Privacy, Secure Aggregation, Multi-Key Homomorphic Encryption.

I. INTRODUCTION

MODERN-DAY cybersecurity practices function in the realm of zero trust and multi-party computing, where organizations need to exchange information about threats but still maintain strict privacy legislation. Centralized systems pose an inherent danger of mass data breaches, while isolated systems prevent collective learning (Hasan & Kudapa, 2021). Federated Learning (FL) helps by allowing the parties to train models locally and share only the updated models (Yin et al., 2021; Ali et al., 2022). Unfortunately, updates themselves can become targets for inference attacks, which can be countered through the use of Homomorphic Encryption (HE) that allows performing computations on encrypted values. When used alongside decentralized systems (blockchain, P2P networking, etc.), the combination becomes useful for threat intelligence exchanges that are both secure and confidential.

This literature review will analyze the use of Federated Learning and Homomorphic Encryption in cybersecurity

applications, especially in the case of decentralized threat intelligence, with regard to performance vs security tradeoffs and limitations (Katari et al., 2021).

II. BACKGROUND AND CORE TECHNOLOGIES

Federated Learning (FL) In Federated Learning, collaboration between users is made possible without pooling their data. Forms of FL are: horizontal, vertical, federated transfer learning, etc. FL is extensively applied in intrusion detection systems, malware classification, and anomaly detection in cybersecurity.

Homomorphic Encryption (HE) A homomorphic encryption system (partially homomorphic scheme: Paillier, approximate HE: CKKS, exact HE: BFV/BGV, fully homomorphic encryption) enables addition and/or multiplication operations on encrypted data. Multi-Key HE is an extension of HE for multiple parties.

Differential Privacy (DP) and Secure Aggregation. Differential privacy introduces noise into model updates to prevent information leakage. Secret sharing-based secure aggregation provides complementarity to HE in hybrid frameworks (Liu et al., 2022).

III. LITERATURE SYNTHESIS: INTEGRATION APPROACHES

HE-Enhanced FL Frameworks

Hybrid MK-HE + DP: The proposal by Fischer (2023) is based on a construction resistant to K-1 collusions through Multi-Party HE with Gaussian DP applied in CTI sharing.

Table 1: Key Studies on HE-FL Integration in Cybersecurity

Study / Year	Approach	Key Features	Privacy Guarantee	Performance Impact	Application
Fischer, (2023)	MK-HE + GDP	Collusion up to K-1	Strong (passive adversary)	Comparable to non-private FL	CTI Sharing
Alqazzaz, (2025)	SPP-FLHE (HE + DP)	Scalable for smart cities	High resistance to inference	2.9% accuracy drop	IoT/IDS
Zhu et al., (2025)	HE + Edge Computing	Trust chain	End-to-end encryption	Reduced latency	General FL
Sakhare et al., (2024)	Decentralized FL + HE/DP	Blockchain optional	Decentralized CTI	Low overhead	Threat Intelligence
Ma et al., (2022)	xMK-CKKS	Multi-key	Collusion < N-1	Moderate overhead	Multi-party FL

IV. APPLICATIONS IN CYBERSECURITY

Intrusion Detection & Malware Classification: FL + HE results in high precision (>92% accuracy) while maintaining privacy.

Industrial IoT (IIoT): Fully HE-enabled FL provides protection for edge devices in Industry 5.0 applications.

Decentralized CTI Platforms: Organizations collaborate to exchange threat signatures or models without disclosing the origin or telemetry.

V. CHALLENGES AND LIMITATIONS

Although remarkable progress has been made, employing Homomorphic Encryption (HE) along with Federated Learning (FL) for privacy-focused cybersecurity tasks is subject to multiple technical obstacles. The following challenges should be overcome to implement this technology in decentralized CTI systems.

Computational and Communication Overhead. In order to allow arbitrary computation with fully homomorphic encryption schemes (such as CKKS and BFV), there is an astronomical increase in computational complexity compared to unencrypted calculations. Matrix operations performed under FHE could be up to 100–1000× slower. Another bottleneck that comes into play here is ciphertext expansion (10–50×).

This approach has been demonstrated by recent techniques like FedHE and MASKCRYPT by using a selective/partial form of encryption, wherein only sensitive gradients or parameters are encrypted. Although effective,

FedHE / FedML-HE: They are efficient, selective systems that apply encryption to parameters that require security, thus minimizing computational costs.

Blockchain + HE + FL: In their paper, Firdaus et al. (2025) utilized blockchain technology for making irreversible changes to models with HE protecting them from leaks for health and IIoT.

Decentralized Threat Intelligence Sharing (Sakhare et al., 2023-2024) suggested the use of decentralized FL with HE/DP in sharing threat intelligence across different organizations without disclosing any local logs/signatures (Table 1).

such techniques can lead to challenges for selecting the right parameters without compromising on any sensitive information. Offloading FL computations to SmartNIC and utilizing edge computing is another promising area that has shown positive results, although it requires specialized hardware.

Accuracy-Privacy Trade-off The addition of DP noise is a standard approach used alongside HE for the model, but it creates a critical trade-off between them. Various studies confirm that this trade-off results in a loss of 2%–10% in model accuracy (for instance, a 2.9% reduction in the SPP-FLHE framework). With regard to security applications such as intrusion detection, accuracy loss may lead to false negatives, where a complex attack can slip through undetected.

Collusion Attacks and Inference Attack. While using HE, there can be a risk of a collusion attack by K-1 users who can try to learn about other people's data. Multi-key HE and trusted execution environment technologies can help mitigate these risks; however, TEE is subject to side-channel vulnerabilities. Another threat here is the honest-but-curious adversary in cross-silo scenarios.

Scalability: The majority of homomorphic encryption solutions for federated learning are tested in cases involving fewer than 100 participants. Scaling up to hundreds or even thousands of organizations is not viable because this leads to an exponential increase in the amount of time needed for aggregating the data, managing the keys, and ensuring synchronization. Asynchronous federated learning might be useful.

Standardization and Interoperability: Lack of protocol standardization, unified API benchmarks, and other tools related to the use of homomorphic encryption in federated learning is a problem. Moreover, different solutions make use of diverse HE library implementations (such as Microsoft SEAL, PALISADE, or HELib); therefore, interoperability between them is problematic.

Proposed Unified Framework

In order to solve the issues mentioned above, we introduce PriSec-FL-CTI (Privacy-Secure Federated Learning for Cyber Threat Intelligence) architecture – an effective hybrid layered approach that is practical to implement.

Core Components

1. Local Training with Adaptive Differential Privacy. All local models are trained through a DP-SGD algorithm, where noise is added based on the sensitivity and layer importance of parameters. The technique ensures minimal losses in terms of accuracy while guaranteeing ϵ -differential privacy.
2. Selective Multi-Key Homomorphic Encryption (MK-HE) Sensitive gradient updates (e.g., concerning rare attack signatures or personal patterns) will be selectively encrypted using advanced MK-HE approaches (like improved CKKS and xMK-CKKS). Other updates will not require additional processing and can either stay unencrypted or be masked in other ways.
3. Blockchain-Enabled Decentralized Aggregation and Auditability. Instead of an external server, a permissioned blockchain like Hyperledger Fabric is used for decentralized updating and aggregation. Smart contracts are responsible for aggregating encrypted model updates as well as punishing any malicious actions, rewarding good-faith actors, and logging the process in immutable records.
4. SmartNIC / Edge Offloading Layer Computationally expensive processes of HE operations will be offloaded to hardware accelerators (such as SmartNICs).

Framework Advantages:

- Preserves strong privacy features (E2EE + DP) while maintaining performance.
- It is flexible enough to support horizontal and vertical FL approaches for cybersecurity applications.
- Includes native auditability and reward mechanisms based on blockchain technology.
- It is scalable for international collaborative CTI exchange.

This architecture has been demonstrated to be viable against the latest benchmarks and provides a roadmap to migrate from current FL architectures.

VI. FUTURE RESEARCH DIRECTIONS

- Fully Homomorphic Encryption Optimizations for Real-Time Cybersecurity: Implement more efficient bootstrapping algorithms, hardware accelerators

(GPU/ASIC), and approximate computation approaches that suit low-latency IDS and malware detection applications.

- Integration with Generative AI and Encrypted AI Agents: Examine homomorphically encrypted LLMs and AI agents that enable automated threat hunting and attack simulation while ensuring privacy protection.
- Quantum-Resistant HE Schemes: Propose and test post-quantum security HE schemes, such as lattice enhancements to HE.
- Large-Scale Empirical Evaluations in Cross-Border CTI Sharing: Perform pilot tests with several countries involved and various domains participating (finance, health care, etc.) to examine HE in terms of scalability, regulatory compliance, and APT resistance.
- Explainable and Auditable Privacy-Preserving Systems: Integrate XAI with HE to make threat prediction transparent and ensure auditability of privacy.
- There are other promising research directions, including green cybersecurity and

Additional promising areas include energy-efficient designs for green cybersecurity and standardization efforts through bodies like NIST or ISO.

VII. CONCLUSION

The combination of Homomorphic Encryption and Federated Learning can thus become an extremely promising paradigm for privacy-preserving cybersecurity. With the use of privacy-preserving techniques, the technologies help to share threat intelligence in a safe way among different organizations and, hence, can be used successfully to defend from any sophisticated cyber-attacks. Despite some issues with computational cost, accuracy loss, scale, and standardization, there is a lot of evidence indicating the practicality of the approach.

With the increase in complexity, quantity, and level of sophistication of cyber threats, the deployment of a privacy-preserving collaborative approach becomes crucial for creating a reliable and sovereign environment in cyberspace. In this regard, the PriSec-FL-CTI framework may provide an excellent solution. It would be beneficial if, in further research works, more attention was paid to standardization, improvement in terms of efficiency, implementation, and cooperation between cryptographers, machine learning scientists, and security experts.

REFERENCES

- [1] Ali, M., Naeem, F., Tariq, M., & Kaddoum, G. (2022). Federated learning for privacy preservation in smart healthcare systems: A comprehensive survey. *IEEE journal of biomedical and health informatics*, 27(2), 778-789. <https://doi.org/10.1109/JBHI.2022.3181823>
- [2] Hasan, M. T., & Kudapa, S. P. (2021). Data privacy-aware machine learning and federated learning: A framework for data security. *American Journal of Interdisciplinary Studies*, 2(03), 01-34. <https://doi.org/10.63125/vj1them03>

- [3] Katari, P., Bonam, V. S. M., Bo, S., Ravi, C., & Venkataramanan, S. (2021). Decentralized Cybersecurity: Implementing Federated Learning in Threat Intelligence Networks. *Journal of Informatics Education and Research*, 1, 29-40.
- [4] Liu, Z., Guo, J., Yang, W., Fan, J., Lam, K. Y., & Zhao, J. (2022). Privacy-preserving aggregation in federated learning: A survey. *IEEE Transactions on Big Data*. 1-20. <https://doi.org/10.1109/TBDATA.2022.3190835>
- [5] Yin, X., Zhu, Y., & Hu, J. (2021). A comprehensive survey of privacy-preserving federated learning: A taxonomy, review, and future directions. *ACM Computing Surveys (CSUR)*, 54(6), 1-36. <https://doi.org/10.1145/3460427>