

Federated LLM-Based SIEM with Meta-Model Aggregation: Enabling Collaborative, Scalable, and Privacy-Preserving Threat Detection

Masoud GanjKhani*, Alireza Shameli-Sendi

*Ph.D. Student in Software Engineering, Shahid Beheshti University, Researcher in Cybersecurity and Computer Networks
Ph.D. in Computer Engineering, Shahid Beheshti University, Specialist in Cybersecurity and Computer Networks
m_ganjkhani@sbu.ac.ir, a_shameli@sbu.ac.ir

Abstract

This paper presents a Federated Large Language Model (LLM)-based Security Information and Event Management (SIEM) architecture for privacy-preserving distributed intrusion detection. The proposed framework combines federated learning with LLM reasoning capabilities, enabling multiple security nodes to collaboratively improve threat detection without sharing sensitive raw network traffic. Instead, nodes exchange abstracted attack patterns to maintain data sovereignty and confidentiality. All detected intrusions are classified using the MITRE ATT&CK framework, ensuring standardized threat intelligence and interoperability. Experimental results demonstrate 79% overall detection accuracy with a 16.1% improvement over five federated learning cycles. The system successfully detects 11 attack categories, including Port Scan, DDoS, Brute Force, SQL Injection, and Phishing, while generating 196 validated detection rules through collaborative learning.

Research method

The proposed method introduces a federated architecture consisting of three distributed nodes and a centralized knowledge base, where each node independently analyzes local network traffic without sharing raw data. Network features are first transformed into abstract semantic categories to ensure privacy preservation. A Large Language Model (LLM) performs classification using structured prompts that include abstracted traffic data, base attack signatures, and previously learned detection rules. The system operates in iterative learning cycles comprising classification, evaluation, rule generation, and knowledge sharing phases. Misclassified samples are used to generate new interpretable detection rules, which are stored and ranked in a shared knowledge base based on confidence scores. Over successive cycles, these rules enhance prompt context and improve detection performance. Attack types are standardized using the MITRE ATT&CK framework, ensuring consistent threat categorization and evaluation.

conclusion


This study presents a federated learning-based intrusion detection framework that enables multiple organizations to collaboratively detect cyberattacks while preserving data privacy. By leveraging Large Language Models for rule-based knowledge abstraction, the system shares high-level, human-interpretable detection patterns instead of raw traffic data. Experimental results show 79% overall accuracy with a 16.1% improvement over five learning cycles, successfully identifying 11 attack categories, including 100% detection for Port_Scan and 87.5% for Brute_Force. The framework collaboratively generated 196 validated rules across three distributed nodes without any leakage of sensitive information, with Node A achieving the highest accuracy of 81.8%. These results confirm the effectiveness of privacy-preserving feature abstraction, confidence-based rule validation, and incremental LLM-driven learning. The approach demonstrates strong applicability in multi-organizational environments, offering interpretable outputs and low computational overhead suitable for edge deployment, while highlighting future needs in scalability, class imbalance handling, and automated feature abstraction.

The purpose of the research

This research aims to address the limitations of traditional centralized and federated intrusion detection systems in handling evolving cyber threats while preserving data privacy. Existing approaches struggle with detecting novel and zero-day attacks and primarily rely on model parameter sharing, which limits interpretability and knowledge transfer. To overcome these challenges, this study proposes a federated LLM-based SIEM framework that enables collaborative threat intelligence through abstract, human-interpretable detection rules instead of raw data or model weights. The system enhances privacy preservation, improves detection of emerging attack patterns, and introduces semantic knowledge sharing across distributed security nodes. Furthermore, it integrates standardized threat modeling using the MITRE ATT&CK framework to improve consistency, interoperability, and interpretability in modern distributed cybersecurity environments.

Practical or simulation results

Results show that system accuracy improves from 68% to 79% (+16.1%), while a non-federated baseline degrades significantly, confirming the benefit of collaborative rule sharing. The knowledge base grows from 47 to 196 validated rules, enhancing detection capability without compromising privacy. Node-level analysis shows heterogeneous performance (81.8% best accuracy), while attack-wise results indicate perfect detection for Port Scan and Brute_Force. Overall, the system demonstrates scalable, privacy-preserving, and efficient intrusion detection suitable for distributed environments.



References

[1] TechTarget, "35 cybersecurity statistics to lose sleep over in 2025," <https://www.techtarget.com/whatis/34-Cybersecurity-Statistics-to-Lose-Sleep-Over-in-2025>, 2025, accessed: 2025-01-01.

[2] Gartner, "7 top trends in cybersecurity for 2022," <https://www.gartner.com/en/articles/7-top-trends-in-cybersecurity-for-2022,2022>, accessed: 2025-01-01.

[3] M. Viehberger, F. Böhmer, I. Fichtinger, and G. Pernul, "Security operationscenter: A systematic study and open challenges," *IEEE Access*, vol. 8, pp. 227 756–227 779, 2020.

[4] D. Shajee and N. Ware, "Integrated network and security operationscenter: A systematic analysis," *IEEE Access*, vol. 10, pp. 27 881–27 898, 2022.

[5] E. Agyepong, Y. Cherdantseva, P. Reinecke, and P. Burnap, "Challenges and performance metrics for security operations center analysts: A systematic review," *Journal of Cyber Security Technology*, vol. 4, no. 3, pp. 125–152, 2020.

[6] C. Onwubiko and K. Ouazzane, "Challenges towards building an effective cyber security operations centre," *arXiv preprint arXiv:2202.03691*, 2022.

[7] S. Y. Cho, J. Happa, and S. Creese, "Capturing tacit knowledge in security operation centers," *IEEE Access*, vol. 8, pp. 42 021–42 041, 2020.

[8] L. Y. Yeh, P. J. Lu, S. H. Huang, and J. L. Huang, "Sochain: A privacy-preserving ddos data exchange service over soc consortium blockchain," *IEEE Transactions on Engineering Management*, vol. 67, no. 4, pp. 1487–1500, 2020.

[9] A. U. Schmidt, S. Knudsen, T. Niehoff, and K. Schwietz, "Planning distributed security operations centers in multi-cloud landscapes: A case study," *arXiv preprint arXiv:2303.03141*, 2023.

[10] A. Shah, R. Ganesan, S. Jajodia, P. Samarati, and H. Cam, "Adaptive alert management for balancing optimal performance among distributed security operations centers using reinforcement learning," *IEEE Transactions on Parallel and Distributed Systems*, vol. 31, no. 1, pp. 16–33, 2019.

[11] A. A. Mughal, "Building and securing the modern security operations center (soc)," *International Journal of Business Intelligence and Big Data Analytics*, vol. 5, no. 1, pp. 1–15, 2022.

[12] U. Bartwal, S. Mukhopadhyay, R. Negi, and S. Shukla, "Security orchestration, automation, and response engine for deployment of behavioural honeypots," in *Proceedings of the 2022 IEEE Conference on Dependable and Secure Computing (DSC)*, IEEE, 2022, pp. 1–8.

[13] L. Liu, Y. Tian, C. Chakraborty, J. Feng, Q. Pei, L. Zhen, and K. Yu, "Multilevel federated learning based intelligent traffic flow forecasting for transportation network management," *IEEE Transactions on Network and Service Management*, vol. 20, no. 2, pp. 1052–1065, 2023.

[14] X. Jiang, J. Zhang, and L. Zhang, "Fedradar: Federated multi-task transfer learning for radar-based internet of medical things," *IEEE Transactions on Network and Service Management*, vol. 20, no. 3, pp. 2847–2861, 2023.

[15] V. Balasubramanian, M. Alogoskoufis, and M. Reisslein, "Fed-ssn: Joint failure probability based federated learning for fault-tolerant time-sensitive networks," *IEEE Transactions on Network and Service Management*, vol. 20, no. 4, pp. 4201–4216, 2023.

[16] Y. Sun, J. Shao, Y. Mao, J. H. Wang, and J. Zhang, "Semi-decentralized federated edge learning with data and device heterogeneity," *IEEE Transactions on Network and Service Management*, vol. 20, no. 1, pp. 9–22, 2023.

[17] A. Alsaacer, M. Abdallah, A. Al-Fuqaha, A. Mohammed, A. Erbad, and O. A. Dobre, "Fair selection of edge nodes to participate in clustered federated multi-task learning," *IEEE Transactions on Network and Service Management*, vol. 20, no. 2, pp. 1066–1080, 2023.

[18] D. M. S. Bharti and H. Nam, "Fedelis: Class-aware federated learning in heterogeneous environment," *IEEE Transactions on Network and Service Management*, vol. 20, no. 3, pp. 2818–2832, 2023.

[19] L. Teng, Y. Qiao, M. Shafiq, G. Srivastava, A. R. Javed, T. R. Gadekallu, and S. Yin, "Flpk-bisnet: Federated learning based on prior knowledge and bilateral segmentation network for image edge extraction," *IEEE Transactions on Network and Service Management*, vol. 20, no. 4, pp. 4321–4334, 2023.

[20] I. Shier and A. Shamir, "Correl: Correlation-based neural network architecture for unavailability concerns in a heterogeneous iot environment," *IEEE Transactions on Network and Service Management*, vol. 20, no. 2, pp. 1330–1344, 2023.

[21] A. P. Kalapaaking, I. Khalil, and M. Atiquzzaman, "Smart policy controller for securing federated learning management system," *IEEE Transactions on Network and Service Management*, vol. 20, no. 1, pp. 196–210, 2023.

[22] F. Sun, Z. Zhang, X. Chang, and K. Zhu, "Towards heterogeneous environment: Lyapunov-optimized implicit reinforcement learning for task offloading," *IEEE Transactions on Network and Service Management*, vol. 20, no. 3, pp. 2833–2846, 2023.

[23] A. Riahi, A. Mohamed, and A. Erbad, "Rl-based federated learning framework over blockchain (rl-fl-bo)," *IEEE Transactions on Network and Service Management*, vol. 20, no. 4, pp. 4187–4200, 2023.

[24] S. B. Saad, B. R. I. K. Bouziane, and A. Ksentini, "Toward securing federated learning against poisoning attacks in zero touch b5g networks," *IEEE Transactions on Network and Service Management*, vol. 20, no. 2, pp. 1315–1329, 2023.

[25] MITRE Corporation, "Mitre att&ck® matrix for enterprise," <https://attack.mitre.org/matrices/enterprise/>, 2023, accessed: 2025-01-01.

[26] —, "Mitre att&ck® techniques," <https://attack.mitre.org/techniques/>, accessed: 2025-01-01.

[27] —, "Mitre caldera™: Automated adversary emulation," <https://github.com/mitre/caldera>, 2023, accessed: 2025-01-01.