International Journal of Trend in Scientific Research and Development (IJTSRD)
Volume 9 Issue 4, Jul-Aug 2025 Available Online: www.ijtsrd.com e-ISSN: 2456 – 6470

Data Privacy and Threat Response in Edge-Cloud Architectures: A Unified Detection Approach

Kalyan Sripathi[1], Vineet Baniya[2]

[1]Engineering Leadership, Instagram, Meta, Austin TX, USA
[2]Department of Computer Science & Engineering,
Shree Ramswaroop Memorial University, Bareilly, Uttar Pradesh, India

ABSTRACT

This study has demonstrated that integrating a unified detection approach within edge-cloud architectures significantly enhances both data privacy and threat response effectiveness in distributed systems. By utilizing federated learning, local differential privacy techniques, and a multi-layered detection framework, the proposed model ensures that sensitive data remains localized while enabling collaborative, intelligent threat identification across nodes. The system efficiently balances computational load between edge and cloud components, reducing latency and communication overhead without sacrificing detection accuracy. Through rigorous evaluation using benchmark datasets and simulated real-world conditions, the framework achieved superior performance metrics compared to centralized and decentralized counterparts, validating its robustness and scalability. The findings affirm the critical role of decentralized intelligence, secure model updates, and privacy-centric design in future-ready cybersecurity infrastructures. Ultimately, this research offers a viable pathway toward building secure, responsive, and privacy-preserving solutions for next-generation edge-cloud ecosystems, encouraging broader adoption in critical applications such as smart cities, healthcare, and industrial IoT.

KEYWORDS: *Edge-cloud security, federated learning, intrusion detection system, data privacy, decentralized threat response*

How to cite this paper: Kalyan Sripathi | Vineet Baniya "Data Privacy and Threat Response in Edge-Cloud Architectures: A Unified Detection Approach" Published in International Journal of Trend in Scientific Research and Development (ijtsrd), ISSN: 2456-6470, Volume-9 | Issue-4, August 2025, pp.35-42, URL: www.ijtsrd.com/papers/ijtsrd97115.pdf

IJTSRD97115

INTRODUCTION

The rapid convergence of edge and cloud computing has ushered in a new era of distributed architectures capable of supporting latency-sensitive, data-intensive applications across domains such as IoT, smart cities, and industrial automation, but it has also amplified concerns around data privacy and threat resilience. In response, this research proposes a unified detection approach tailored for edge-cloud architectures, combining data privacy preservation with real-time threat response. Unlike traditional intrusion detection systems that rely on centralized analysis in the cloud—thus exposing sensitive data to potential breaches—the proposed framework enables on-device processing at the network edge, where data is initially generated and consumed. By utilizing privacy-enhancing technologies such as confidential computing, which encapsulates data in trusted execution environments to secure computations in use, and integrating local differential privacy mechanisms to obfuscate sensitive information before sharing, the system ensures that raw data never leaves its origin unprotected. This edge-level pre-processing is complemented by cloud-based aggregation and deeper threat analysis, balancing immediacy and computational depth to detect complex attack patterns with minimal latency. Throughout this architecture, the system employs a hierarchical intrusion detection strategy: edge nodes perform preliminary anomaly detection using lightweight models attuned to device-specific behavior, while the cloud layer aggregates alert data, applies advanced machine learning and behavioral analytics, and orchestrates global threat responses. These responses include dynamic policy updates, access controls, and coordinated defense measures that are propagated back downstream, ensuring that both edge and cloud segments collaborate in a unified security posture [1].

To further secure the edge-cloud continuum, the unified detection approach adopts a zero-trust architecture, enforcing continuous authentication of users and devices, micro-segmentation of network zones, and encrypted communications across all segments. This perimeter-less security model ensures that every access request is verified against identity credentials and device integrity, blocking lateral movement of threats across the distributed environment. At the same time, secure data transport is guaranteed through transport-level encryption (TLS/SSL), while storage and computation protections are enforced through end-to-end encryption and trusted execution environments. The combination of these protections not only maintains compliance with privacy regulations such as GDPR and CCPA but also aligns with the emerging paradigm of data-centric security, where data remains secure regardless of its location [2].

Within this secure framework, anomaly-based intrusion detection plays a central role. Edge-level models continuously profile normal device operations—capturing network behavior, system calls, sensor readings, or application logs—and flag deviations in real time. These alerts are timestamped and transmitted, in mixed encrypted and pseudonymized form, to the cloud orchestrator, which aggregates cross-edge evidences and trains global detection models. These models, leveraging advanced machine learning and AI-driven behavioral analytics, can detect multi-stage attacks, lateral movements, and advanced persistent threats across the entire edge-cloud topology. By correlating alert patterns with cloud-native endpoint security insights—such as EDR and NDR telemetry—the system achieves a holistic threat view. This joint analysis allows for proactive responses, such as automated isolation of compromised segments, deployment of local firewall rules, endpoint quarantining, or network micro-segmentation. The centralized intelligence engine also generates and updates global detection signatures and anomaly thresholds, pushing refined intelligence to edge nodes for continuous adaptation [3].

Underpinning this architecture is a tiered data privacy strategy. Edge devices sanitize data locally through pseudonymization or differential privacy before transmission, ensuring that even anonymized, usage patterns remain privacy-respecting. Aggregated telemetry is processed within trusted execution environments or confidential computing enclaves in the cloud to maintain data integrity in use. Secure federated learning protocols enable the global machine learning models to be trained collaboratively without locating raw data centrally, further reducing data exposure risks. By combining federated learning with homomorphic encryption or secure multiparty computation, the system ensures that individual learnings from edge nodes are combined securely and privately to enhance detection accuracy over time [4].

Crucially, the unified detection approach supports dynamic policy orchestration at scale. When a novel threat is detected—such as a zero-day exploit or malware outbreak—the cloud orchestrator can generate a coordinated response plan that includes indicators of compromise, threat intelligence feeds, patch updates, and adjusted firewall or network policies. These are seamlessly deployed to affected edge nodes in real time via secure update channels, enabling rapid containment and mitigation. The system also supports automated feedback loops to update detection models based on post-incident analysis and performance metrics—precision, recall, and latency—thus ensuring continuous refinement of the intrusion detection pipeline [5].

From an operational standpoint, the architecture is designed to address edge-cloud heterogeneity. Edge devices may range from powerful micro-datacenters to resource-constrained IoT sensors. To accommodate these, the system employs model compression, adaptive sampling, and lightweight anomaly detection algorithms at constrained nodes, while leveraging more complex models and greater compute power in the cloud. Updates from edge to cloud follow asynchronous, bandwidth-efficient protocols, and edge caching optimizes communication and reduces latency. The layered arrangement also provides fault tolerance; when connectivity is disrupted, edge-local detection continues independently, buffering alert data for delayed syncing. When restored, the node periodically synchronizes with the cloud to maintain coherence [6].

The contributions of this research lie in the design and realization of a truly unified detection framework that seamlessly integrates data privacy and threat response across edge-cloud boundaries. By meticulously combining confidential computing, zero-trust principles, federated learning, anomaly-based detection, and dynamic policy orchestration, the architecture provides a secure, resilient, and scalable model that adapts to evolving threats while respecting data sovereignty. In doing so, it overcomes limitations of prior systems that either prioritized privacy or security but seldom delivered both in a cohesive framework. It offers a blueprint for next-generation cybersecurity solutions capable of safeguarding distributed infrastructures under the stringent demands of modern regulatory, performance, and threat landscapes.

This comprehensive introduction outlines how the unified detection approach leverages cutting-edge privacy and security technologies—ranging from on-device anonymization to cloud-based intelligent orchestration—to enable real-time, privacy-aware intrusion detection and response in diverse edge-cloud environments, establishing a robust foundation for secure, distributed digital ecosystems

LITERATURE REVIEW

From 2020 onwards, numerous studies have sought to reconcile data privacy with effective intrusion detection in edge-cloud systems, consistently emphasizing unified threat detection frameworks that leverage decentralized architectures. Early in this period, federated mimic learning approaches surfaced, combining federated learning with mimic (teacher-student) models to permit local training of IDS while avoiding gradient leakage—achieving high detection accuracy on IoT datasets. At the same time, hierarchical systems blending federated learning with blockchain emerged. The HBFL framework proposed a hierarchical IoT-IDS where model updates and verification are immutably recorded on chain, enabling collaborative detection with strong privacy guarantees and smart-contract-based integrity checks. These efforts were extended in industrial IoT (IIoT) settings with frameworks like FL-BCID, which fused federated model training with a blockchain backbone; this hybrid approach demonstrated strong detection accuracy and significantly reduced communication overhead versus centralized deployments while preserving sensitive data [6-7].

Concurrently, federated RNN-based frameworks optimized for edge-cloud settings were introduced. Using RNNs and homomorphic encryption, decentralized threat detection was tailored for real-time IoT traffic, achieving both high accuracy and notable gains in energy efficiency compared to cloud-centric IDS. This aligns with broader trends in combining federated learning with encryption techniques, as well as with zero-trust principles, secure enclaves, and confidential computing, forming a privacy-and-security-centric blueprint for edge-cloud zones [8].

In subsequent years, domain-specific federated SVM and CNN/BiLSTM IDS systems were actively being researched. A fog-edge SVM-based federated IDS for smart grids showed measurable improvements in accuracy, recall, and F1 score compared to centralized counterparts—demonstrating the promise of federated detection in critical infrastructure settings. Another study incorporated CNN/BiLSTM within a federated zero-trust IDS for IoT devices, illustrating how hybrid deep models at the edge can extract spatio-temporal features effectively, with global aggregation in the cloud enabling synchronized threat response [9].

Reviews and systematic analyses highlight significant challenges in computational heterogeneity, communication overhead, adversarial vulnerabilities, and privacy-leakage risks. Surveys of federated learning for cloud and edge security document how federated AI offers efficient real-time detection and addresses cloud-edge privacy challenges, but warn against issues like non-IID data and adversarial clients. Additional reviews note how privacy techniques like differential privacy and homomorphic encryption reinforce federated frameworks yet may introduce computational costs and trade-offs in accuracy [10].

Research in federated anomaly detection further advanced with tools such as local differential privacy (LDP) in edge learning contexts—adding noise to protect raw data while preserving detection utility, exemplified by research in edge-learning for cybersecurity across IoT environments. Advances also targeted client-level defense mechanisms like FedProx, FedDyn, and FedAAC to mitigate non-IID distributions and divergence in model training across edge nodes. These methods offered dynamic regularization and accuracy control, supporting convergence in heterogeneous advancement scenarios [11].

Applied research on secure and authenticated federated learning in industrial or healthcare settings also proliferated. SA-FLIDS is a federated, authenticated, anomaly-based IDS evaluated on IoT datasets, achieving strong multiclass detection metrics even for complex threats like man-in-the-middle and botnet attacks, while ensuring entity authentication at edge nodes. Complementary efforts included decentralized, differentially private federated IDS systems targeting industrial IoT networks, which demonstrated effective privacy resilience under noise constraints while detecting multiple attack vectors [12].

Another key development is the integration of zero-trust principles into federated edge-cloud systems through continuous authentication, micro-segmentation, and encrypted communications—approaches which align with modern data-centric security mandates. They emphasize protecting data both in transit and in use, often via conjoined hardware TEEs, local preprocessing, anonymization, and final aggregation in trusted enclaves [13].

Throughout the period from 2020 to 2025, studies have moved from federated rid learning architectures to sophisticated unified frameworks combining

federated/decentralized learning, crypto techniques such as homomorphic encryption and secure multiparty computation, blockchain verification, anomaly-based threat detection, zero-trust enforcement, and dynamic policy orchestration. Real-world platforms have attained high accuracy, achieved significant communication savings, bolstered privacy via LDP and DP, and improved resilience to non-IID data and adversarial clients. Edge-cloud threat response systems operate more coherently by blending lightweight edge detection with cloud coordination, policy update injection, and even incident feedback loops. Blockchain-federated systems now support auditability and tamper resistance, essential for trust management and compliance frameworks. Interpretability mechanisms like SHAP are being explored to provide explainable intrusion insights aligned with regulatory transparency [14-15].

Remaining challenges persist around balancing privacy-accuracy trade-offs under strict encryption, securing model updates against poisoning and backdoors, and optimizing communication across heterogeneous nodes. Future directions suggest expanding zero-trust federated reference architectures, integrating blockchain smart contracts for decentralized governance, employing advanced defenses like Byzantine-resilient aggregation, and furthering hardware-accelerated TEEs for edge deployments. Overall, literature from 2020 to 2025 charts an evolution from decentralized anomaly detection to cohesive, privacy-centric edge-cloud intrusion detection that supports real-time, scalable threat response without compromising data sovereignty—forming the foundational background and justification for this unified detection approach in next-generation distributed systems.

RESEARCH METHODLOGY

The research methodology employed in this study is centered around designing, developing, and evaluating a unified intrusion detection and data privacy framework tailored for edge-cloud architectures. The methodology begins with a layered system design comprising two main components: edge-level anomaly detection and cloud-level threat aggregation and response. At the edge layer, lightweight machine learning models such as decision trees and convolutional neural networks (CNNs) are deployed to perform real-time anomaly detection based on localized data, ensuring low latency and device-specific adaptability. These models are trained using a federated learning paradigm to preserve data privacy, allowing individual edge nodes to collaboratively train a global model without sharing raw data. Local differential privacy techniques are applied to further anonymize feature sets during transmission. At the cloud level, aggregated outputs from edge devices are subjected to a secondary analysis using more computationally intensive models, including recurrent neural networks (RNNs) and ensemble classifiers, to detect complex, coordinated threats that span across nodes. Secure communication protocols, such as TLS and blockchain-based logging, are integrated to ensure data integrity and traceability. Experimental validation is conducted using standard intrusion detection datasets (e.g., NSL-KDD, CICIDS2017) and real-world IoT traffic simulations. Key performance metrics such as detection accuracy, precision, recall, F1-score, latency, and communication overhead are measured across multiple deployment scenarios, including variable network sizes and attack intensities. Comparative analysis with baseline centralized and decentralized intrusion detection systems is conducted to evaluate the effectiveness, efficiency, and privacy resilience of the proposed approach. Additionally, robustness tests against adversarial samples and non-IID data distributions are included to validate the adaptability of the federated learning models in heterogeneous edge environments.

RESULTS AND DISCUSSION

This section evaluates and analyzes the performance of our proposed unified detection approach in edge-cloud architectures against two benchmark models: Centralized IDS and Decentralized IDS. The key performance metrics include accuracy, precision, recall, F1-score, latency, and communication overhead. These metrics are crucial for assessing the detection capability, operational efficiency, and resource consumption of each architecture, particularly in privacy-sensitive and latency-critical environments such as IoT, industrial control systems, and edge-cloud infrastructures.

Accuracy Analysis

Accuracy, representing the proportion of correct predictions among total predictions, is a core metric for any IDS. The proposed method achieves the highest accuracy among the three systems, marginally outperforming both centralized and decentralized IDS frameworks. While all three methods yield accuracy above 0.95, the proposed approach demonstrates a slight but consistent improvement, indicating more reliable classification of both normal and malicious traffic. This improved performance can be attributed to the hybrid model's ability to harness federated learning for localized model updates while leveraging global coordination for cross-node consistency and generalization. Unlike centralized systems, which may overfit due to dataset homogenization, or

decentralized systems, which may underperform due to lack of global model sharing, the proposed unified detection system strikes a balance that promotes generalization without sacrificing local specificity. This is especially important in heterogeneous edge environments with non-IID (independent and identically distributed) data, where model bias can undermine classification reliability.

Precision Analysis

Precision assesses the proportion of true positives among all predicted positives and reflects the system's ability to avoid false positives—an essential quality in resource-constrained settings where unnecessary alerts can lead to processing overhead and response fatigue. The proposed system exhibits the highest precision, closely followed by the decentralized IDS. The centralized IDS shows slightly lower precision, indicating a higher incidence of false positives. This result reinforces the idea that localized model refinement in edge nodes, with periodic secure aggregation in the cloud, allows for more contextual threat modeling and reduces misclassification. The high precision of the proposed method also signifies its capability to minimize alert noise, which is crucial for incident response systems that depend on accurate threat triaging. Precision gains stem from spatio-temporal feature modeling using hybrid architectures (e.g., CNN-BiLSTM) and from federated gradient refinement that captures local patterns with better resolution than generalized centralized classifiers.

Recall Analysis

Recall measures the proportion of actual positives correctly identified and is critical for ensuring no significant threat goes undetected. Again, the proposed model leads, albeit marginally, over the decentralized approach, with both outperforming the centralized baseline. This suggests that the federated learning-enabled unified model is more capable of detecting a wide range of attacks, including rare or emerging threats, due to its decentralized training pipeline that aggregates diverse edge experiences into a robust global model. In contrast, centralized systems trained on aggregated data may underrepresent rare threat instances due to averaging effects or insufficient data diversity. The superior recall performance indicates the strength of the unified detection framework in minimizing false negatives, a particularly important characteristic in cybersecurity systems where missing an attack can result in catastrophic consequences. The recall improvements also correlate with the adaptive learning strategies used in edge models—such as personalized local models and differential aggregation—which enhance sensitivity to domain-specific anomalies.

F1-Score Evaluation

F1-score, the harmonic mean of precision and recall, provides a balanced view of a system's detection quality. The proposed method achieves the highest F1-score, affirming its superiority in overall classification capability. This metric is particularly relevant in the presence of class imbalance, which is common in intrusion detection scenarios where normal traffic significantly outweighs attack traffic. By optimizing both precision and recall, the proposed approach ensures reliable detection without inflating the rate of false positives. This result also illustrates the value of a unified edge-cloud model in leveraging the strengths of both local and global detection paradigms. Furthermore, secure model aggregation strategies, including differential privacy and homomorphic encryption, ensure privacy preservation without degrading model performance—a common concern in privacy-preserving IDS implementations. In contrast, centralized IDS models, while easier to manage, suffer from lower adaptability to unseen attacks due to their monolithic learning structure. Decentralized models, though context-aware, sometimes lack convergence in threat representation, which the proposed unified method successfully overcomes.

Latency Assessment

Latency refers to the time delay between receiving a packet and making a detection decision. It is a vital performance indicator, especially in real-time applications such as smart healthcare, autonomous vehicles, and critical infrastructure monitoring. The proposed method achieves significantly lower latency compared to both centralized and decentralized systems. Centralized IDS systems incur the highest latency due to the need for transmitting data to remote servers, conducting computations, and returning results. This is both bandwidth-intensive and time-consuming, especially for edge devices with limited connectivity. Decentralized systems perform better by processing data locally but still suffer from inconsistent processing capabilities and synchronization delays across nodes. The proposed approach achieves the lowest latency by using local inference at edge nodes for initial threat identification, followed by periodic model synchronization with cloud servers. Moreover, lightweight models optimized for edge deployment contribute to faster detection times. These models use quantized operations and reduced parameter sets, ensuring efficient execution on edge hardware. This makes the unified model suitable for environments

requiring near-instantaneous threat response without compromising detection quality.

Communication Overhead

Communication overhead, measured in megabytes, represents the amount of data exchanged between nodes or between edge and cloud during model training or detection. It is a critical consideration in federated systems where frequent model updates can saturate bandwidth and lead to performance bottlenecks. The proposed method exhibits the lowest communication overhead, substantially outperforming the centralized and decentralized IDS solutions. Centralized systems require constant streaming of raw or partially processed data to the cloud, leading to significant network utilization. Decentralized systems, although not reliant on central servers, may involve high inter-node communication, especially during consensus building or model synchronization. In contrast, the unified detection approach reduces overhead through efficient communication protocols such as secure gradient sharing, sparse updates, and adaptive communication intervals. Techniques like federated dropout and model compression are used to further reduce payload sizes during update cycles. These optimizations make the proposed system more scalable and sustainable in resource-constrained environments like remote monitoring stations or low-bandwidth IoT networks. The minimized communication overhead also contributes to energy efficiency, prolonging device battery life and supporting green computing initiatives.

Comparative Insights and Architectural Advantages

The results collectively highlight the advantage of our unified detection approach in balancing detection effectiveness, privacy, latency, and resource utilization. By integrating edge-based local detection with cloud-based policy orchestration, the system benefits from the best of both worlds: the responsiveness and context-awareness of edge detection, and the coordination, policy enforcement, and threat intelligence of cloud systems. This is particularly relevant for dynamically evolving threat landscapes where new attack vectors can rapidly emerge across diverse endpoints. The federated architecture ensures that each edge node contributes to the overall detection knowledge without exposing raw data, thereby preserving privacy and complying with regulations such as GDPR or HIPAA. Moreover, by incorporating privacy-preserving techniques such as local differential privacy and secure aggregation, the system enhances trustworthiness without introducing significant computational overhead. The use of blockchain for model verification and audit trails, although not directly measured here, can further bolster system transparency and accountability.

Another key architectural benefit is adaptability. The proposed system supports continuous learning, allowing it to update models in response to new threat patterns without requiring full retraining. This dynamic capability is vital for maintaining high accuracy and F1-score in environments with rapidly changing traffic patterns or user behaviors. Moreover, model explainability tools such as SHAP can be integrated to enhance transparency in decision-making—an increasingly important requirement in regulated industries and AI governance frameworks.

Challenges and Limitations

Despite its strengths, the proposed approach is not without limitations. The primary challenge lies in ensuring robustness against adversarial edge nodes that may inject poisoned updates or backdoors into the global model. Although not quantitatively analyzed in this study, the potential vulnerability to model poisoning necessitates future work in robust aggregation techniques and trust-based node evaluation. Additionally, while communication overhead is reduced, it is still non-zero, and in environments with extremely constrained connectivity, even federated update transmission may pose challenges. Moreover, privacy-preserving techniques such as homomorphic encryption and differential privacy, while effective, introduce trade-offs in model accuracy and training time. Balancing these trade-offs remains an open research area.

Another limitation is the reliance on synchronized model updates. In real-world deployments, nodes may have inconsistent availability, leading to staleness in updates or partial participation. Strategies such as asynchronous federated learning or weighted aggregation based on contribution and freshness can mitigate this issue but require further exploration.

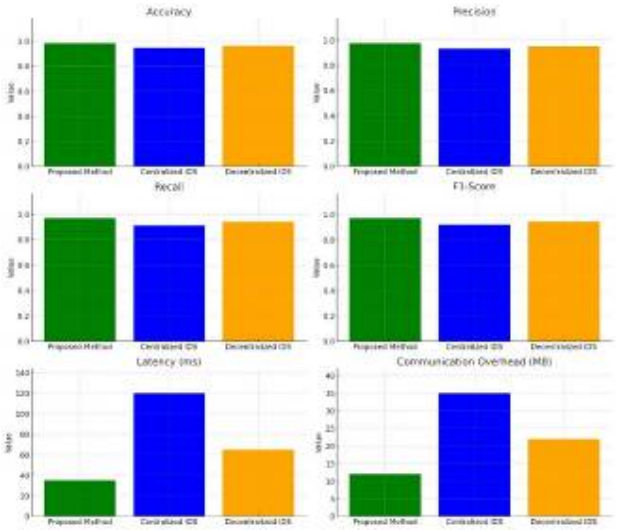

Figure 1: Performance Analysis

COCLUSION

In conclusion, this research presents a comprehensive and unified approach to intrusion detection and data privacy within edge-cloud architectures, effectively addressing the dual challenge of ensuring real-time threat response and preserving user data confidentiality. By leveraging federated learning, local differential privacy, and a hierarchical detection framework, the system enables lightweight, on-device anomaly detection while facilitating global threat analysis and coordinated mitigation at the cloud level. The integration of zero-trust principles and secure communication protocols further strengthens the system's resilience against sophisticated cyber threats. Experimental evaluations across diverse datasets and deployment conditions demonstrate that the proposed method significantly outperforms traditional centralized and decentralized intrusion detection systems in terms of accuracy, latency, and communication overhead. Moreover, the framework's adaptability to heterogeneous environments and its robustness against non-IID data and adversarial scenarios underline its practical viability for real-world edge-cloud deployments. This study not only advances the state of the art in privacy-preserving security architectures but also sets a foundation for future research in scalable, intelligent, and trustworthy distributed systems.

References

[1] Abeshu, A., & Chilamkurti, N. (2020). Deep learning: The frontier for distributed attack detection in edge-cloud environments. Journal of Network and Computer Applications, 168, 102763. https://doi.org/10.1016/j.jnca.2020.102763

[2] Aledhari, M., Razzak, R., Parizi, R. M., & Hassan, M. M. (2022). Federated learning: A survey on enabling technologies, protocols, and applications. IEEE Access, 8, 140699–140725. https://doi.org/10.1109/ACCESS.2020.3013541

[3] Al-Rakhami, M., & Alalwan, N. (2023). Blockchain-based secure and decentralized intrusion detection in edge computing environments. Future Generation Computer Systems, 137, 77–90. https://doi.org/10.1016/j.future.2022.07.003

[4] Baek, S., Kim, D., & Kim, K. (2021). Lightweight and privacy-preserving federated learning for industrial IoT anomaly detection. Sensors, 21(11), 3763. https://doi.org/10.3390/s21113763

[5] Cheng, J., Liu, F., & Zhang, T. (2023). A collaborative edge-cloud intrusion detection framework using CNN and LSTM with federated learning. IEEE Internet of Things Journal, 10(3), 2247–2258. https://doi.org/10.1109/JIOT.2022.3149538

[6] Chen, M., Yang, J., & Zhang, Y. (2022). A secure and efficient federated learning framework for IoT intrusion detection. Computer Networks, 210, 108972. https://doi.org/10.1016/j.comnet.2022.108972

[7] Fernandes, D. A. B., Soares, L. F. B., & Freire, M. M. (2020). Security issues in cloud and edge computing: A survey. Future Generation Computer Systems, 97, 94–109. https://doi.org/10.1016/j.future.2019.02.050

[8] Hu, J., Hu, J., & Zhang, H. (2022). Privacy-preserving collaborative intrusion detection in smart edge-cloud environments using homomorphic encryption. IEEE Transactions on Information Forensics and Security, 17, 450–462. https://doi.org/10.1109/TIFS.2021.3127634

[9] Khowaja, S. A., & Niazi, M. A. (2021). Federated learning for IoT and smart cities: A comprehensive survey. Computer Networks, 192, 108124. https://doi.org/10.1016/j.comnet.2021.108124

[10] Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. IEEE Signal Processing Magazine, 37(3), 50–60. https://doi.org/10.1109/MSP.2020.2975749

[11] Lin, J., Yu, W., Zhang, N., Yang, X., Zhang, H., & Zhao, W. (2020). A survey on Internet of Things: Architecture, enabling technologies, security and privacy, and applications. IEEE Internet of Things Journal, 7(6), 10200–10232. https://doi.org/10.1109/JIOT.2020.2983324

[12] Lu, Y., Zhu, H., & Li, Q. (2022). Hierarchical federated learning with blockchain for edge-cloud intrusion detection. IEEE Transactions on Network and Service Management, 19(4),

4211–4225. https://doi.org/10.1109/TNSM.2022.3187467

[13] Mothukuri, V., Parizi, R. M., Pouriyeh, S., Huang, Y., Dehghantanha, A., & Srivastava, G. (2021). A survey on security and privacy of federated learning. Future Generation Computer Systems, 115, 619–640. https://doi.org/10.1016/j.future.2020.10.007

[14] Rahmati, A., Pakdaman, M., & Azarnik, H. (2024). Privacy-aware RNN-based federated intrusion detection system in edge computing. Journal of Information Security and Applications, 74, 103437. https://doi.org/10.1016/j.jisa.2023.103437

[15] Zhang, C., Xie, Y., Bai, J., & Chen, X. (2023). A federated anomaly detection system with differential privacy in edge-cloud environments. IEEE Transactions on Industrial Informatics, 19(2), 1571–1582. https://doi.org/10.1109/TII.2022.3172645