International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Intelligent and Auditable: A Hybrid AI and Distributed Ledger Framework for Modern Cybersecurity

Naresh Kalimuthu

naresh.kalimuthu@gmail.com

Abstract:
Modern cyber threats, known for their complexity and constant change, surpass traditional intrusion detection systems (IDS). This paper explores a new security approach that combines Artificial Intelligence (AI) with decentralized architectures to develop IDS that are robust, scalable, and protect user privacy. It examines the core roles of Federated Learning (FL) and Blockchain, highlighting three main research challenges: The vulnerability of AI models to adversarial attacks, privacy and data integrity concerns in collaborative learning, and performance limitations in distributed systems. To address these issues, we suggest solutions such as adversarial training, differential privacy, and lightweight consensus mechanisms. Our analysis of case studies shows that hybrid FL-Blockchain systems outperform traditional methods in practical application environments.

Keywords: Intrusion Detection System (IDS), Artificial Intelligence, Federated Learning, Blockchain, Cybersecurity, Adversarial Attacks, Privacy-Preserving Machine Learning, Decentralized Systems, Internet of Things (IoT) Security.

I. INTRODUCTION

A. The Evolving Threat Landscape and the Limitations of Centralized IDS

The modern digital ecosystem features an ever-growing attack surface and increasingly complex cyber threats. Adversaries now run sophisticated campaigns involving zero-day exploits, Advanced Persistent Threats (APTs), and large-scale Distributed Denial of Service (DDoS) attacks, leveraging extensive networks of compromised devices. Traditional Intrusion Detection Systems (IDS), which are based on a centralized design, are ill-suited for this environment. They mainly depend on static, signature-based detection, which is effective against known threats but often ineffective against new, unknown attack methods.

Additionally, the centralized architecture itself poses a significant vulnerability. Collecting and analyzing traffic from multiple endpoints at a single point creates performance issues, a single point of failure, and raises serious concerns regarding data privacy. In the context of strict data protection laws, centralizing sensitive logs and user data poses a significant risk that many organizations seek to avoid. These structural and methodological issues necessitate a comprehensive reevaluation of network security strategies.

B. The AI Revolution in Cybersecurity: From Signature to Anomaly Detection

The rise of Artificial Intelligence (AI), especially through Machine Learning (ML) and Deep Learning (DL), has sparked a major shift in intrusion detection. It moves away from the reactive, signature-based approach to a proactive, anomaly-based one that adapts over time. Unlike older systems, AI-powered IDS do not depend on a fixed database of known threats; instead, they analyze large datasets to learn what normal network activity looks like. Any deviation from this baseline is flagged as a potential intrusion, allowing the system to detect new and zero-day threats accurately.

Deep Learning models such as Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), and Long Short-Term Memory (LSTM) networks excel at this. They can interpret complex, high-dimensional network data to spot subtle patterns over time and space that humans and simpler algorithms often miss. This enables AI-based systems to deliver higher detection accuracy and significantly reduce false positives compared to traditional methods.

C. The Decentralization Imperative: Motivations for Distributed Security Architectures

While AI offers the intelligence needed to tackle modern threats, its standard deployment in centralized systems remains a vulnerability. Decentralization has become a crucial architectural goal to address these issues, motivated by three major reasons. First, safeguarding privacy is critical; decentralized models prevent the need to send sensitive raw data to a central server, enabling analysis where the data already exists. Second, scalability is vital for contemporary networks, especially the vast Internet of Things (IoT), where billions of devices produce enormous data. A decentralized approach disperses computational and storage demands, avoiding central bottlenecks. Third, resilience is improved by eliminating single points of failure, resulting in a more robust and fault-tolerant security infrastructure.

Leading this decentralization shift are two key technologies: Federated Learning (FL) and Blockchain. FL allows collaborative training of models across distributed nodes without sharing the underlying data, while blockchain provides a distributed, tamper-proof ledger to ensure trust and integrity. Together, they underpin a new generation of Intrusion Detection Systems (IDS).

This shift from isolated, perimeter-focused security to a distributed system marks a fundamental change in cybersecurity philosophy. It moves from simple detection within a closed network to a collective defense approach. Traditional IDS operates independently, so an attack detected in one organization provides no immediate benefits to others. Even centralized AI models, while more advanced, are limited to their own data boundaries. Federated Learning helps overcome these barriers by enabling multiple entities to jointly train a powerful detection model, where insights from an attack on one participant quickly strengthen defenses across all, without risking data privacy. Blockchain technology then adds an essential trust layer: ensuring model updates and security alerts are tamper-proof and traceable. This creates a trustworthy framework for collective security efforts. By combining AI with decentralization, intrusion detection shifts from a solo pursuit to a collaborative effort, where the system's strength relies on the trusted, collective intelligence of the entire network.

D. Thesis Statement and Paper Structure

This paper argues that combining Artificial Intelligence with decentralized frameworks, especially hybrid models that integrate Federated Learning and Blockchain, is the most promising approach to develop next-generation Intrusion Detection Systems. These systems aim to be intelligent, private, scalable, and trustworthy. The paper supports this by examining key technologies, highlighting research challenges arising from their integration, suggesting effective mitigation strategies, and reviewing empirical evidence from cutting-edge implementations.

II. FOUNDATIONAL TECHNOLOGIES IN MODERN IDS

A. AI-Powered Threat Intelligence: A Deep Dive into Learning Models

1) Supervised and Unsupervised Learning Techniques (CNN, RNN, LSTM)

Modern IDS effectiveness primarily stems from the advanced pattern recognition capabilities of Deep Learning models. These models outperform traditional algorithms such as Support Vector Machines (SVMs) and Decision Trees, especially in handling the large volume and high dimensionality of current network traffic data.

Convolutional Neural Networks (CNNs), originally developed for image analysis, have been successfully adapted for intrusion detection. They excel at detecting spatial hierarchies and local patterns within network data, such as byte sequences in packet payloads or feature relationships in flow records. Using convolutional filters, CNNs can automatically identify features associated with malicious activity, thereby minimizing manual feature engineering.

Recurrent Neural Networks (RNNs) and their advanced forms, Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRU), are crucial for analyzing the time-based dynamics of network traffic. Designed for sequential data, these models are perfect for recognizing patterns across a series of network packets or events over time. This is especially important for identifying complex, multi-stage attacks like Advanced Persistent Threats (APTs), where malicious behavior is only evident when viewing the entire sequence of actions.

2) Reinforcement Learning for Adaptive Defense

While supervised and unsupervised models excel at learning from existing data, Reinforcement Learning (RL) adds a dynamic, adaptive layer to network defense. An RL-based IDS functions as an autonomous agent that learns to make optimal security decisions through direct interaction with its network environment. The RL agent is trained using a reward-and-penalty system: successful threat identification and mitigation are rewarded, while false alarms or missed detections are penalized.

This learning approach allows the IDS to adopt a proactive defense stance, automatically optimizing firewall rules, adjusting access policies, and executing responses in real time without needing explicit human input for each scenario. This ongoing learning process enables the system to adapt to new attack methods and changing network conditions, paving the way for self-healing, fully autonomous security systems.

B. Architectures of Decentralization

AI model intelligence is maximized when deployed within architectures that overcome the main limitations of centralization. Decentralized frameworks underpin privacy, scalability, and trust.

1) Federated Learning (FL): Collaborative Training with Data Privacy

Federated Learning is a distributed machine learning approach that separates model training from direct access to raw data. The process is iterative and collaborative:

- ***a)*** *Initialization:* A central server establishes a global model and shares it with selected client nodes (such as endpoints, IoT gateways, or entire organizational networks).
- ***b)*** *Local Training:* Each client trains the model locally using its own private data, which remains within its environment.
- ***c)*** *Update Transmission:* After local training, each client sends its model updates — such as weights or gradients — to the server without transmitting raw data.
- ***d)*** *Aggregation:* The server combines these updates (using algorithms such as Federated Averaging, or FedAvg) to build a better global model.
- ***e)*** *Iteration:* This improved model is then sent back to clients, and the cycle repeats.

This method enables diverse participants to collaboratively develop a robust AI model that leverages their collective data while safeguarding sensitive information, forming the foundation of privacy-preserving collaborative security.

2) Blockchain Technology: Establishing Trust and Immutability

Blockchain functions as the fundamental structure for establishing trust and integrity in a decentralized system without depending on a central authority. It is a distributed, cryptographically linked, and immutable ledger. The key components related to IDS include:

- ***a)*** *Distributed Ledger:* All network participants maintain a copy of the ledger, providing redundancy and resilience.
- ***b)*** *Immutability:* Once transactions are validated and added to the chain, they are secured with cryptographic hashes, making them nearly impossible to alter or delete. This provides a tamper-proof audit trail.
- ***c)*** *Consensus Mechanisms:* Algorithms like Proof-of-Work (PoW) or Proof-of-Stake (PoS) ensure agreement among all participants on transaction validity before they are recorded, preventing fraudulent entries.
- ***d)*** *Smart Contracts:* Self-executing contracts with terms embedded in code can automate tasks such as verifying security alerts or distributing updated IDS rules transparently and deterministically.

In an IDS setting, blockchain can be employed to generate a verifiable record of intrusion alerts, device identities, and shared threat intelligence, ensuring this vital information remains unaltered by malicious actors.

3) Hybrid Frameworks: The Synergy of FL and Blockchain

The most advanced decentralized IDS architectures combine Federated Learning and Blockchain, leveraging each to address the other's weaknesses. This combination aims to resolve the 'Trust-Privacy-Performance Trilemma,' where optimizing for two attributes often compromises the third.

For example, FL-only systems maintain strong privacy by keeping data local but rely on trust in a central aggregator. They lack mechanisms to verify the integrity of model updates, making them vulnerable to poisoning attacks and raising trust issues. On the other hand, traditional blockchain-based IDS provides high trust and data integrity through its immutable ledger but can compromise privacy due to its public nature. Additionally, its consensus mechanisms may cause high latency, affecting real-time performance.

Hybrid frameworks like Blockchain-enabled Federated Learning for Intrusion Detection Systems (BFLIDS) and "Federated Learning-Enhanced Blockchain Intrusion Detection" (FL-BCID) address this trilemma. They use FL primarily for privacy-preserving collaborative training, while a permissioned blockchain layer serves as a trusted, auditable platform. Model updates from FL clients are recorded as blockchain transactions, with smart contracts enforcing rules, validating contributions, and automating aggregation. This integrated approach makes the training process private, transparent, auditable, and tamper-resistant, providing a balanced, robust solution.

III. CORE RESEARCH CHALLENGES IN AI-POWERED, DECENTRALIZED IDS

The intersection of AI and decentralization shows promise but poses complex research challenges that must be addressed to unlock their full potential. These challenges include ensuring AI model robustness, securing collaborative learning processes, and optimizing the performance of the underlying distributed systems.

A. Challenge 1: Ensuring robustness Against Adversarial Attacks

1) Understanding Evasion and Poisoning Attacks on AI Models

The same learning mechanisms that empower AI models also make them vulnerable to adversarial threats. These attacks aim to exploit how models learn and predict, mainly falling into two types:

a) Evasion Attacks:

Predominant in deployed IDSs, attackers craft malicious inputs—such as network packets or traffic flows—by subtly altering them in ways invisible to signature-based systems. These changes can push the data across the model's decision boundary, leading to misclassification as benign. The goal is to evade detection during inference.

b) Poisoning Attacks:

Targeting the training phase, attackers inject small amounts of malicious samples into the training data. This "poisoned" data aims to corrupt the learning process, either reducing overall accuracy or creating a "backdoor". A backdoored model performs normally on most inputs but misclassifies specific, attacker-defined triggers.

2) Impact on IDS Efficacy: Misclassification and System Compromise

Errors and System Risks Adversarial attacks significantly threaten AI- based IDS. Successful evasion attacks can bypass defenses, allowing malicious activities to go unnoticed. These attacks exploit the high-dimensional linearity of neural networks, making it easier for attackers to find minimal perturbations for misclassification.

In decentralized systems like Federated Learning, poisoning attacks are especially dangerous. A single malicious participant can introduce poisoned updates, which are averaged into the global model. Such sabotage can undermine the entire network, turning a collective defense into a vulnerability point.

B. Challenge 2: Ensuring Privacy and Integrity in Collaborative Learning

1) Vulnerabilities in Federated Learning: Inference and Model Poisoning Attacks

Although Federated Learning aims to protect privacy, it is not foolproof. The updates shared by clients—such as gradients or weights—do not include raw data, but they still reveal subtle information about the private data used for training. Advanced inference attacks can examine these updates across multiple rounds to reconstruct or deduce sensitive details of a client's training set, resulting in potential privacy breaches.

Moreover, the accuracy of the global model relies on the honesty of all participants. As noted, a malicious client could deliberately carry out a model poisoning attack by sending specially designed, harmful updates. The aim is to distort the global model's functioning, either to reduce its overall effectiveness or to establish backdoors for future exploitation. Without a method to authenticate each contribution, the central aggregator remains susceptible to these manipulations.

2) The Paradox of Transparency: Blockchain and Data Exposure Risks

Blockchain's main advantage is its transparency and immutability, establishing a reliable basis for trust and audits. However, this transparency can pose risks if not properly controlled. Placing sensitive security details, like detailed intrusion alerts, network topology, or full model parameters, on a public blockchain would make them accessible to everyone, including malicious actors. This presents a paradox: the very system that fosters trust might also enable information leaks.

Effective architectures therefore require privacy-preserving designs. This typically entails using permissioned (private) blockchains, where access is limited to authorized nodes, and adopting an off-chain storage approach. In this setup, only cryptographic hashes or metadata are recorded on the blockchain, while most sensitive data are stored securely in separate locations such as the InterPlanetary File System (IPFS).

C. Challenge 3: Achieving Scalability and Real-Time Performance

1) The Overhead of Decentralization: Communication and Computational Costs

Decentralized systems inherently add overhead that centralized architectures do not. In Federated Learning, repeatedly distributing a global model and aggregating updates from many clients can result in substantial network traffic. This communication burden may become a key bottleneck, particularly in extensive IoT networks with limited bandwidth devices or when managing large, intricate deep learning models.

Blockchain systems face both computational and storage overheads. Every node must process transactions, run smart contracts, and store a copy of the entire ledger. As the blockchain expands, storage demands grow, and reaching consensus can become resource-heavy.

2) Consensus Bottlenecks and Transaction Latency in Blockchain-Based IDS

The primary performance challenge for blockchain in real-time security is latency from consensus mechanisms. Protocols like Proof-of-Work (PoW), used in public blockchains such as Bitcoin, are intentionally slow and require significant computation to guarantee security. Mining and confirming a transaction can take minutes to hours, which is too slow for an IDS that must detect and respond to threats within milliseconds.

This "consensus bottleneck" makes traditional public blockchains unsuitable for high-throughput, low-latency applications like intrusion detection. An alert recorded on the blockchain might only be confirmed long after the attack has happened. To solve this, research is focused on developing more efficient, lightweight consensus protocols designed for permissioned environments, where greater trust among participants enables faster validation.

These issues are interrelated, often creating a cycle where fixing one problem worsens another. For example, improving performance and scalability (Challenge 3) might mean using a simpler, faster consensus protocol. But these protocols are usually less secure, increasing the risk of manipulation and weakening integrity (Challenge 2). To address this, advanced AI models can vet transactions before they are added. However, complex AI models expand the attack surface and are more vulnerable to adversarial attacks (Challenge 1). The main defense is adversarial training, which is computationally demanding and adds load to each node, further impairing performance (Challenge 3). This cycle demonstrates that a resilient system must integrate security, privacy, and performance strategies rather than tackling each separately.

IV. RECOMMENDATIONS AND MITIGATION STRATEGIES

Tackling the intricate, interconnected issues of AI-driven, decentralized IDS requires a comprehensive approach. Strategies should focus on strengthening AI models through adversarial training and defensive structures, safeguarding privacy and integrity within the collaborative framework, and refining the distributed architecture for better performance.

A. Enhancing AI Models with Adversarial Training and Defensive Architectures

The main safeguard against adversarial threats is adversarial training. This involves enriching the training dataset with adversarial examples crafted to mislead the model. Training with both benign and malicious-looking, correctly labeled samples encourages the development of more robust, generalized features. This process helps smooth decision boundaries, thereby reducing sensitivity to small, targeted perturbations that are often exploited in evasion attacks. Though highly effective, adversarial training may slightly decrease natural accuracy on unaltered, benign data. Current research aims to strike a balance between robustness and maintaining high natural accuracy. Apart from training strategies, architectural defenses bolster resilience: -

a) *Ensemble Methods:*
Combining diverse AI models makes it harder for attackers to create adversarial examples that fool all of them simultaneously.
b) *Gradient Masking:*
Concealing or obfuscating gradient information impedes attacks that rely on gradient-based optimization to find optimal perturbations.

B. *Cutting-Edge Privacy Techniques for Federated Learning*
To mitigate remaining privacy threats, such as inference attacks, advanced cryptographic and statistical methods, collectively known as Privacy-Preserving Federated Learning (PPFL), can be used.

a) *Differential Privacy (DP):*
Adds statistically calibrated noise to model updates before they reach the aggregator, masking individual data contributions and making it infeasible to infer private data. A key parameter in DP is the privacy budget, epsilon (ϵ), which controls the trade-off between privacy and model accuracy (utility).
b) *Homomorphic Encryption (HE):*
Enables computations directly on encrypted data. In an FL setting, clients can encrypt their model updates before sending them to the server. The server can then aggregate these encrypted updates to create an encrypted global model update, all without ever decrypting individual contributions. This ensures that the server learns nothing about the personal updates, thereby maintaining a high level of privacy.
c) *Secure Multi-Party Computation (SMPC):*
Protocols that enable multiple parties to jointly compute functions without revealing their private inputs, thereby supporting decentralized aggregation without the need for a trusted server.

C. *Building Scalable and Efficient Decentralized Systems*
To address scalability and performance issues in real-time IDS, the system architecture must be optimized. For blockchain components, choosing lightweight consensus protocols over resource-heavy ones is crucial.
1) Delegated Proof-of-Stake (DPoS):
In this model, participants in the network vote for a limited, fixed number of "delegates" who are responsible for validating transactions and creating blocks. This approach significantly reduces the communication and computational effort required for consensus, enabling high throughput and low latency.
2) Proof-of-Authority (PoA):
This protocol is designed for private or consortium blockchains where participants are verified and trusted. Block creation is carried out by a designated group of authorized "authorities," ensuring rapid and energy-efficient operations. This makes it ideal for high-frequency activities such as logging security alerts in IoT environments.

For Federated Learning, hierarchical topologies enhance scalability by utilizing intermediate aggregators to reduce communication load, along with model compression to decrease update sizes and bandwidth utilization.

V. PERFORMANCE ANALYSIS AND CASE STUDIES

Empirical research is increasingly supporting the theoretical benefits of AI-powered, decentralized IDS. This section reviews key case studies illustrating their performance, security, and scalability in real-world or near-real-world scenarios. The table below offers a high-level comparison of the architectural paradigms discussed.

TABLE I. COMPARATIVE ANALYSIS OF DECENTRALIZED IDS ARCHITECTURES

Parameter	Centralized AI-IDS	Federated Learning (FL) IDS	Blockchain-only IDS	Hybrid FL-Blockchain IDS
Key Technology	Central Server, ML/DL	FL, Central Aggregator	DLT, Consensus, Smart Contracts	FL, DLT, Smart Contracts
Privacy Guarantees	Low (Data Centralized)	High (Data Localized)	Medium (Pseudonymous, on-chain data)	Very High (FL + Off-chain data)
Data Integrity	Medium (Server is trusted)	Low (Vulnerable to poisoning)	Very High (Immutable Ledger)	Very High (Blockchain-verified)
Scalability	Low (Server Bottleneck)	Medium (Communication Overhead)	Low (Consensus Overhead)	High (Hierarchical FL + Lightweight Consensus)
Latency	Low (Centralized)	Medium (Aggregation Rounds)	High (Block Confirmation Time)	Medium-Low (Optimized Consensus)
Adversarial Robustness	Dependent on model	Dependent on model + poisoning risk	High (Integrity of rules)	High (Model + Integrity)
Suitability for IoT	Poor	Good	Poor	Excellent

A. Case Study: The FL-BCID Framework for Industrial IoT Security

A clear example of hybrid architecture is the Federated Learning-Enhanced Blockchain Framework (FL-BCID), which is specifically designed to secure Industrial Internet of Things (IIoT) environments.

1) *Architectural Design:*

FL-BCID combines federated learning (FL) for privacy-focused, decentralized model training directly on IIoT edge devices, preventing the need to transmit sensitive industrial process data. It employs a permissioned blockchain as a trusted, auditable backend, with smart contracts managing the FL lifecycle. These contracts record model updates from edge devices and anomaly scores as unchangeable transactions, creating a verifiable record of the model's training history and performance, thereby ensuring accountability and integrity.

2) *Performance Gains:*

Tested on standard IIoT security datasets (ToN-IoT and N-BaIoT), the FL-BCID framework achieved impressive results. It attained a threat detection accuracy of 97.3% and reduced communication overhead by 41% compared to traditional centralized methods that depend on raw data collection.

This case study strongly indicates that hybrid architectures can effectively address the trust-privacy-performance trade-off, offering a secure, efficient, and practical solution for industrial applications.

B. Case Study: Scalable Blockchain IDS with Lightweight Consensus for IoT

Addressing the key performance bottlenecks of blockchain is crucial for its effective use in real-time intrusion detection systems, particularly within resource-constrained IoT networks. Several studies have focused on developing and evaluating lightweight consensus protocols.

1) *Consensus Analysis:*

Research indicates that protocols such as Delegated Proof-of-Stake (DPoS) and Proof-of-Authority (PoA) are more suitable for IoT security than Proof-of-Work (PoW). These protocols are tailored for permissioned environments where nodes have established identities, enabling much faster transaction validation.

2) *Performance Results:*

An evaluation of a DPoS-based framework for IoT data management revealed remarkable performance. The system achieved an extremely low transaction latency of under 0.976 ms and showed linear scalability, reaching 500 Transactions Per Second (TPS) with 500 nodes—a notable enhancement over PoS. Likewise, a PoA-based approach for securing smart grids achieved low latency suitable for real-time monitoring, while reducing energy use on IoT devices by shifting heavy computations to edge servers. These results demonstrate that with appropriate consensus mechanisms, blockchain can satisfy the high-performance requirements of real-time security applications.

C. 5.3 Case Study: Federated Learning for Anomaly Detection on the CICIDS2017 Dataset

To evaluate the overhead of adding blockchain to an FL workflow, a study used the well-known CICIDS2017 intrusion detection dataset.

1) *Model Performance and Implementation:*

The research involved a federated learning setup where multiple participants jointly trained an autoencoder for anomaly detection. A permissioned blockchain (MultiChain) was incorporated to ensure an unchangeable record of model updates during federated averaging.

2) *Impact of Blockchain Integration:*

The main finding was the measurable effect of blockchain on overall performance. Results indicated that the added complexity and processing for blockchain caused a modest overhead of only 5-15% to the federated learning process. This important conclusion demonstrates that blockchain's benefits, including trust, auditability, and integrity, can be achieved with only a minor performance hit, making the hybrid approach feasible in practice.

VI. CONCLUSION

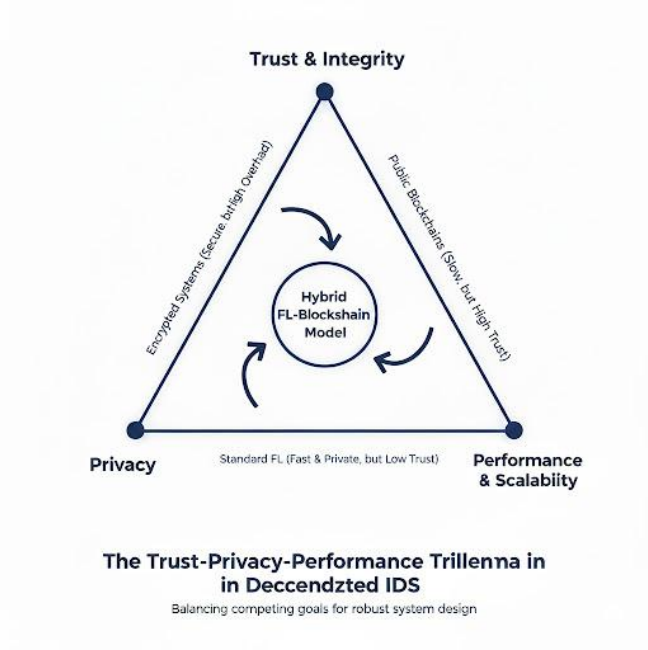

A. Synthesis of Findings and Answers to Research Questions

This paper discusses the shift in intrusion detection caused by the merging of Artificial Intelligence and decentralized systems. The findings show that combining these technologies provides a strong solution to the drawbacks of traditional, centralized security approaches. AI offers the intelligence needed to identify advanced, zero-day threats, while decentralization ensures privacy, scalability, and resilience. The study of three main research challenges uncovers a complex interaction of vulnerabilities and compromises.

1) *Adversarial Attacks:*

AI models are naturally susceptible to evasion and poisoning attacks, which can render an IDS ineffective. Adversarial training is recommended and proven to improve model robustness.

2) *Privacy and Integrity:*

Decentralized learning still faces risks of privacy leaks through inference attacks and integrity issues via model poisoning. Techniques such as differential privacy, Homomorphic Encryption, and blockchain-based verifiable integrity are crucial for creating a secure collaborative environment.

3) *Performance and Scalability:*

The extra overhead from decentralization can impact real-time performance. Case studies demonstrate that this can be managed effectively with lightweight consensus protocols and optimized design, making these architectures suitable for demanding settings like IoT.

Overall, empirical evidence suggests that hybrid frameworks combining Federated Learning and Blockchain are superior, effectively balancing trust, privacy, and performance to achieve high detection accuracy, strong privacy, and real-time scalability.

B. The Future Trajectory: Towards Autonomous, Resilient, and Trustworthy IDS

The field of AI-powered, decentralized IDS is advancing quickly, with several key research paths likely to influence its future. Incorporating Explainable AI (XAI) is vital for transitioning away from opaque "black box" models. XAI methods offer straightforward, human-understandable explanations for decisions made by sophisticated AI systems, increasing transparency, aiding audits, and fostering trust among security analysts.

As quantum computing advances, the cryptographic bases of existing systems will face significant threats. Research into post-quantum cryptography (PQC) is crucial to maintain the long-term security and viability of these decentralized systems. Moreover, ongoing progress in autonomous response and self-healing networks, driven by advanced Reinforcement Learning, aims to develop fully adaptive defense ecosystems capable of detecting, isolating, and recovering from security breaches in real-time with little human input.

C. Final Thoughts on the Emerging Security Framework

The merging of Artificial Intelligence and decentralization signals more than just an incremental tech breakthrough; it marks the emergence of a new security framework. This shift departs from the fragile, isolated defenses of earlier times and heads toward proactive, collaborative, and verifiably reliable security systems. By allowing entities to work together in defense while maintaining their privacy and control, these systems are well-suited to meet the challenges of today's complex threat environment. Although there are still many research hurdles, the principles and technologies presented here provide a strong foundation for future intrusion detection and network security.

REFERENCES:

1. Khraisat, Ansam & Alazab, Ammar & Singh, Sarabjot & Jan, Tony & Gopez, Alfredo. (2024). Survey on Federated Learning for Intrusion Detection System: Concept, Architectures, Aggregation Strategies, Challenges, and Future Directions. ACM Computing Surveys. 57. 10.1145/3687124.
2. Büyüktanır, Büşra & Altınkaya, Şahsene & Karatas Baydogmus, Gozde & Yildiz, Kazim. (Aug 2025). Federated learning in intrusion detection: advancements, applications, and future directions. Cluster Computing. 28. 10.1007/s10586-025-05325-w.
3. Wardana, Aulia & Sukarno, Parman. (2024). Taxonomy and Survey of Collaborative Intrusion Detection System using Federated Learning. ACM Computing Surveys. 57. 10.1145/3701724.
4. Al-Emari, Salam & Anbar, Mohammed & Sanjalawe, Yousef & Manickam, Selvakumar & Hasbullah, Iznan. (2021). Intrusion Detection Systems Using Blockchain Technology: A Review, Issues and Challenges. Computer Systems Science and Engineering. 40. 87-112. 10.32604/csse.2022.017941.
5. Al-Emari, Salam & Anbar, Mohammed & Sanjalawe, Yousef & Manickam, Selvakumar & Hasbullah, Iznan. (2021). Intrusion Detection Systems Using Blockchain Technology: A Review, Issues and Challenges. Computer Systems Science and Engineering. 40. 87-112. 10.32604/csse.2022.017941.
6. Vyas, Abhishek & Lin, Po-Ching & Hwang, Ren-Hung & Tripathi, Meenakshi. (2024). Privacy-Preserving Federated Learning for Intrusion Detection in IoT Environments: A Survey. IEEE Access. PP. 1-1. 10.1109/ACCESS.2024.3454211.
7. Tafreshian, Benyamin, and Shengzhi Zhang. "A Defensive Framework Against Adversarial Attacks on Machine Learning-Based Network Intrusion Detection Systems." *ArXiv*, 2025, https://arxiv.org/abs/2502.15561.
8. Preuveneers, Davy, et al. "Chained Anomaly Detection Models for Federated Learning: An Intrusion Detection Case Study." *Applied Sciences*, vol. 8, no. 12, 2018, p. 2663, https://doi.org/10.3390/app8122663.
9. Haque, Ehtisham & Shah, Adil & Iqbal, Jawaid & Ullah, Syed Sajid & Alroobaea, Roobaea & Hussain, Saddam. (2024). A scalable blockchain based framework for efficient IoT data management using lightweight consensus. Scientific Reports. 14. 10.1038/s41598-024-58578-7.
10. Dhieb, Najmeddine & Ghazzai, Hakim & Besbes, Hichem & Massoud, Yehia. (2020). Scalable and Secure Architecture for Distributed IoT Systems. 1-6. 10.1109/TEMSCON47658.2020.9140108.
11. Begum, Khadija, et al. "BFLIDS: Blockchain-Driven Federated Learning for Intrusion Detection in IoMT Networks." *Sensors*, vol. 24, no. 14, 2024, p. 4591, https://doi.org/10.3390/s24144591.

12. Ozturk, Ozer, et al. "Differential Privacy in Federated Learning: Mitigating Inference Attacks with Randomized Response." *ArXiv*, Sep 2025, https://arxiv.org/abs/2509.13987.
13. Adversarial Training for Raw-Binary Malware Classifiers. / Lucas, Keane; Bauer, Lujo; Pai, Samruddhi et al. 32nd USENIX Security Symposium, USENIX Security 2023. 2023. p. 1163-1180 (32nd USENIX Security Symposium, USENIX Security 2023; Vol. 2).
14. Levi, Matan, and Aryeh Kontorovich. "Splitting the Difference on Adversarial Training." *ArXiv*, 2023, https://arxiv.org/abs/2310.02480.
15. Feng, Zhengze, et al. "Identity Management Systems: A Comprehensive Review." *Information*, vol. 16, no. 9, July 2025, p. 778, https://doi.org/10.3390/info16090778.
16. S N, Prajwalasimha. (July 2025). AI-Powered Intrusion Detection and Privacy Preservation in 6G Networks. 10.1109/ICC-ROBINS64345.2025.11086288.
17. Vansiya, Jayveersinh & Chandi, Asha & Khan, Rashid. (May 2025). AI-Based Intrusion Detection & Prevention Models for Smart Home IoT Systems: A Literature Review. Journal of Computer Science and Technology Studies. 7. 982-996. 10.32996/jcsts . 2025.7.3.110.
18. Wenjuan LI, Yu WANG, Weizhi MENG, Jin LI, Chunhua SU, BlockCSDN: Towards Blockchain-Based Collaborative Intrusion Detection in Software Defined Networking, IEICE Transactions on Information and Systems, 2022, Volume E105.D, Issue 2, Pages 272-279, Released on J-STAGE February 01, 2022, Online ISSN 1745-1361, Print ISSN 0916-8532, https://doi.org/10.1587/transinf.2021BCP0013
19. Ali, A., Husain, M., & Hans, P. (2025). Federated Learning-Enhanced Blockchain Framework for Privacy-Preserving Intrusion Detection in Industrial IoT. *ArXiv*. https://arxiv.org/abs/2505.15376
20. Albanbay, N., Tursynbek, Y., Graffi, K., Uskenbayeva, R., Kalpeyeva, Z., Abilkaiyr, Z., & Ayapov, Y. (2025). Federated Learning-Based Intrusion Detection in IoT Networks: Performance Evaluation and Data Scaling Study. *Journal of Sensor and Actuator Networks*, *14*(4), 78. https://doi.org/10.3390/jsan14040078