ESP-IJACT
ESP International Journal of Advancements in Computational Technology
ISSN: 2583-8628 / Volume 2 Issue 4 November 2024 / Page No: 163-176
Paper Id: IJACT-V2I4P121 / Doi: 10.56472/25838628/IJACT-V2I4P121

Original Article

EdgeSecFL: A Lightweight Federated Learning Model for Intrusion Detection in IoT-Cloud Ecosystems

Faraz Ahmed
Crisp Technologies LLC, Cybersecurity Researcher.

Received Date: 29 September 2024 Revised Date: 03 November 2025 Accepted Date: 26 November 2025

Abstract - As IoT devices continue to proliferate and integrate with cloud computing platforms, the attack surface for cyber threats has expanded significantly. The paper is a review of EdgeSecFL framework, which is a light federated learning (FL) model introduced to deal with prevention of intrusion in the IoT-cloud ecosystem. EdgeSecFL uses decentralized training, secure aggregation and blockchain-based logging to achieve greater privacy, cut down on communication overhead, and establish tamper-evident auditability. The model supports adaptive edge coordination and cost-efficient cloud deployment, making it suitable for resource-constrained environments. Comparative analysis proves that EdgeSecFL is superior to previous FL-based IDS systems in the following parameters: latency, scalability and data protection. Other ethical and governance considerations such as explainability, fairness, alignment with emergent regulatory frameworks are addressed in the review as well. Lastly, the article presents future areas of research such as quantum-resistant cryptographic integration and self-orchestrated FL. EdgeSecFL stands as a robust foundation for developing secure, policy-aware, and scalable intrusion detection systems across modern distributed infrastructures.

Keywords - *Federated Learning, Intrusion Detection Systems, IoT-Cloud Security, Blockchain, Secure Aggregation, Privacy-Preserving Machine Learning.*

I. INTRODUCTION

The fusion of the Internet of Things (IoT) and cloud computing has catalyzed a new era of hyper-connected infrastructure [1], where data is generated, processed, and analyzed across a vast network of distributed endpoints. Such convergence enables organizations to utilize scalable computation and storage and concurrently to assemble real-time data in smart homes, cities, industrial control systems, and healthcare situations using edge devices. Yet, with more connectedness and the use of cloud backbones, the threat surface has also vastly grown, making such environments more appealing to the ground of a cyberattack. IoT devices are typically constrained in terms of computational power, memory, and security controls [2]. When integrated with powerful but centralized cloud services, they can unintentionally serve as entry points for malicious actors. Vulnerabilities in device firmware, insecure APIs, misconfigured cloud permissions, and a lack of real-time security monitoring lead to a high-risk environment [3]. Traditional intrusion detection systems (IDS), often designed for centralized data centers, struggle to cope with the distributed, dynamic, and heterogeneous nature of IoT-cloud systems.

One possible option to this difficulty is the solution concept of Federated Learning (FL) which is a decentralized machine learning model that allows training a model on numerous devices without transmitting raw data to the central server [4]. In such a way, FL takes care of the data privacy and latency issues. but mainstream FL frameworks tend to be resource-hungry and not to be tailored towards low-power edge devices. To address this limitation, EdgeSecFL has been proposed as a lightweight FL-based intrusion detection framework designed specifically for IoT-cloud ecosystems. It emphasizes computational efficiency, secure communication, and adaptability to resource-constrained environments.

As shown in Figure 1, the number of documented IoT-cloud security incidents has risen. The given trend is not only about the increased use of IoT in mission-critical sectors but also about the increasing complexity of the security of such ecosystems. In concurrence with these technical innovations, security policies regarding the cloud have gone through important advancements as well that have helped to contain the threats of dynamically evolving environments [5]. Cloud Security Posture Management (CSPM) is one of them and provides automation of cloud services configuration errors identification and correction. CSPM enables continuous compliance monitoring, vulnerability detection, and policy enforcement capabilities that are essential in environments where IoT and cloud systems operate in tandem [6]. He claims that "CSPM tools are becoming a critical component of any cloud, especially the ones incorporating resource-constrained edge gadgets into centralized services".

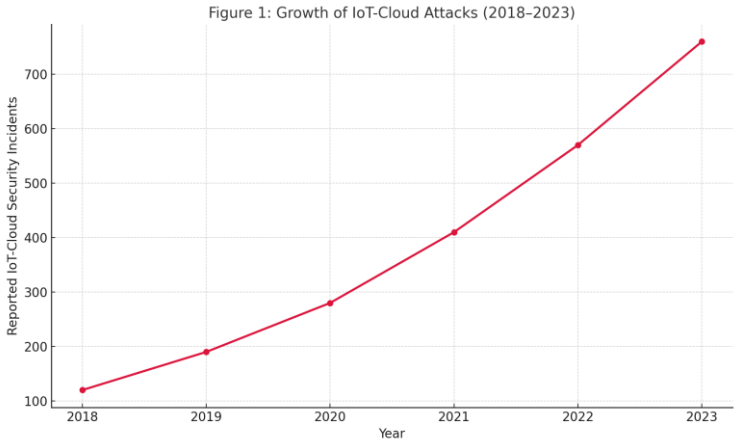

Figure 1: Growth of IoT-Cloud Attacks

The integration of models like EdgeSecFL with CSPM frameworks presents a multi-layered approach to securing IoT-cloud infrastructures. Although EdgeSecFL provides visibility at the local level and data privacy, CSPM supports visibility into cloud-level and compliance, which can be especially effective against the threat of advanced persistent threats (APT) and misconfiguration-based attacks. This article provides a comprehensive review of EdgeSecFL, evaluating its architecture, performance, cryptographic foundations, and deployment feasibility. It also aligns the security innovations and policy frameworks, especially those that prioritize privacy preservation, quantum-resilient encryption, and autonomous policy enforcement to demonstrate its relevance in contemporary cybersecurity landscapes.

II. BACKGROUND AND MOTIVATION

The integration of IoT and cloud computing has undoubtedly revolutionized how data-driven services are built and consumed [7]. However, this convergence also brings high levels of cybersecurity risks that should be met using technical and decidedly policy-based solutions [8]. This section identifies the major threats of the Internet of Things-cloud convergence and assesses the relationship of the Federated Learning (FL) as a new decentralized intrusion-detection defense mechanism.

A. Security Challenges in IoT-Cloud Convergence

As IoT devices proliferate across sectors ranging from smart homes and healthcare to industrial automation and smart cities the pairing of IoT devices with cloud infrastructures gives rise to a tiled ecosystem through which heterogeneous attack surfaces have to be addressed [9]. IoT devices are typically deployed in uncontrolled environments and are designed with low computational overhead, often lacking even basic security hardening features such as firmware validation, authentication enforcement, and real-time monitoring. In contrast, superiorly equipped in security tools, the cloud environments are marred by complex configurations, multi-tenant exposures, and vulnerability of the API, further propagating the security problems in the combination with IoT nodes that are not well secured [10].

A major vulnerability lies in the communication linkages between IoT devices and cloud services, typically over insecure wireless networks [11]. Man-in-the-middle attacks, DNS spoofing, and packet sniffing are common exploits here. When compromised, these gadgets can be used as entry points in lateral movement to cloud settings resulting in mass compromises.

Table 1: Common Security Vulnerabilities in IoT vs. Cloud Ecosystems

Category	IoT Devices	Cloud Infrastructure
Attack Vector	Weak authentication, open ports	Misconfigured S3 buckets, insecure APIs
Exploitation Frequency	Very high	Medium to high
Typical Impact	Botnets, DDoS, lateral movement	Data leaks, privilege escalation
Policy Enforcement	Limited or nonexistent	Manual and error-prone
Security Updates	Infrequent or absent	Centralized but requires policy compliance

It is a very significant policy enforcement gap between the edge devices and the cloud components. While cloud providers offer tools for managing access control and monitoring usage, there is limited enforcement at the edge. Researcher highlight the inadequacy of traditional perimeter defense models in securing decentralized mobile networks and advocate for blockchain-based verification and multi-layered policy control as a response to the growing threat landscape [12]. Their results also highlight how security systems that combine edge and clouds infrastructures are required. Moreover, emerging

threats like AI-powered attacks, supply chain manipulation, and zero-click exploits require real-time detection, which is not easy to achieve using older IDS solutions that run on centralized or cloud-based datasets.

B. Role of Federated Learning in Intrusion Detection

Intrusion Detection Systems (IDS) play a pivotal role in identifying anomalies and attacks within digital environments [13]. Traditional IDS models rely significantly on centralized aggregation of data and this has led to considerable privacy risks and also presented a challenge of latency in high-velocity edge data. This is where Federated Learning (FL) provides a very interesting option. FL is a privacy-preserving, decentralized approach to machine learning, where model training occurs locally on distributed devices, and only the model updates, not raw data are shared with a central aggregator [14]. This architecture specially can help in intrusion detection in the IoT-cloud systems, which are characterised by data sensitivity, heterogeneity of devices and limited bandwidth as the significant limitations.

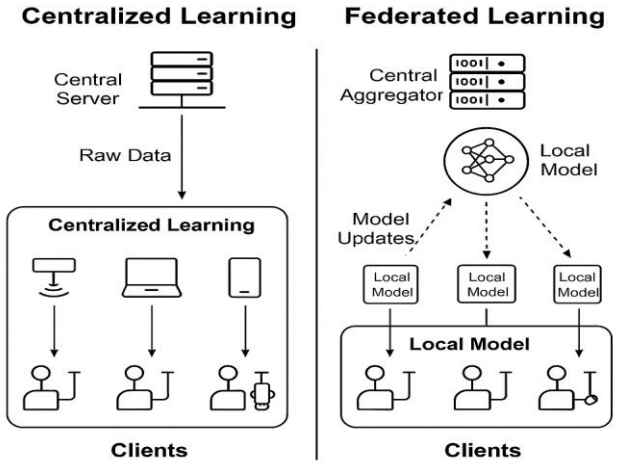

Figure 2: Centralized vs. Federated Learning for IDS

In the centralized model, raw traffic data from IoT sensors are forwarded to a cloud-based IDS engine for analysis. This configuration is quite bandwidth intensive besides going contrary to data minimisation. Conversely, FL disperses the detection logic and enables local models to adapt to attack patterns depending on contexts in the used devices. Such models are also matched with an international model that will grow in precision without indulging in any privacy violation.

The benefits of FL for intrusion detection are multi-fold:

- Data Sovereignty: Sensitive data never leaves the device.
- Bandwidth Efficiency: Only updates are transmitted, not entire datasets.
- Personalized Security: Every device has the capability of molding the global pattern to its behavior.
- Scalability: New devices can join the learning process with minimal disruption.

These advantages go well with the ambitions of privacy-first computing, which finds acceptance in regulated markets, including healthcare, finance and critical infrastructure. As the IoT-cloud ecosystem continues to expand, intrusion detection systems must evolve to meet the demands of distributed intelligence, minimal data exposure, and low-latency response [15].Furthermore, FL supports adaptive learning, where models can evolve dynamically with the arrival of new attack vectors without having to be retrained. This agility is especially valuable in environments where zero-day vulnerabilities and novel attack variants are increasingly common.

The IoT-cloud coalescence has also provided a wide array of security threats, which cannot be managed using monolithic and centralized solutions only. Federated Learning, particularly lightweight implementations such as EdgeSecFL, presents a scalable and privacy-compliant framework for deploying IDS in such dynamic ecosystems. Combining FL-based IDS with policy-conscious structures such as CSPM and blockchain-based governance frameworks will allow minimizing exposure as well as the length of the time between the attack notification and the corresponding detection, which can keep organizations on pace with the ever-changing cyber threats.

III. ARCHITECTURE OF EDGESECFL

EdgeSecFL architecture is modeled with the main goal of providing federated intrusion detection across resource-limited IoT devices. In contrast to the existing federated learning versions where a substantial amount of computational and communication resources are assumed, EdgeSecFL adapts to the functional constraints of heterogeneous edge devices in the

setting where real-time IoT-cloud systems are involved. In this section, the main architectural elements of EdgeSecFL are described, namely the lightweight model approach, its secure and synchronised protocol communication, etc.

A. Overview of Lightweight FL Models

When it comes to implementing federated learning (FL) in an IoT environment, there is a need to redefine the process of training machine learning models and updating them. IoT devices, ranging from microcontrollers to embedded processors, typically lack sufficient RAM, processor throughput, or battery capacity to handle full-scale deep learning tasks. Therefore, lightweight models must be adopted to ensure practical implementation and sustainability over prolonged operation [16]. Model training in EdgeSecFL is conducted locally, at IoT node level, based on compact, conventionally structured sets of features, extracted in intrusion detection local logs. To conserve device resources, the model utilizes compression techniques such as quantization and weight sharing, reducing its memory footprint. Furthermore, frequency of update is conditionally regulated, based on the number of the device in use and the availability of the network, so that the asynchronous presence will be permissible in any cases, and the slower devices will not receive any penalization.

Another feature that EdgeSecFL has in common with local caching and incremental learning approaches is the ability of devices to train as they go through a sequence of cached events, with no expensive retraining done on the fly. This architectural design reflects key principles from [17], which emphasizes the importance of efficient training data caching and workload scheduling for deep learning in edge computing networks. He concluded that the model stability and energy usage of performing in distributed networks would be more stable when there was a balance between memory usage and the frequency of updates.

B. Communication & Synchronization Protocols

The efficiency and trustworthiness of FL systems are heavily influenced by their communication protocols, specifically, how client updates are transmitted, aggregated, and validated [18]. In EdgeSecFL, much focus is given to reduce the communication overhead and maintain the security of the update process to avoid data leakage and manipulating the model. One of the standout features is differential model sharing, a technique where each participating IoT device sends only the differences (deltas) between its current local model and the last received global model. This significantly decreases the volume of information sent per round of training, and it is highly useful to network with bandwidth limitations or known interconnectivity.

To maintain the integrity and privacy of updates, EdgeSecFL supports secure aggregation protocols, which allow the central server (aggregator) to compute a global model from encrypted individual contributions without exposing their contents. Further, an update authentication based on lightweight digital signatures is useful to protect the poisoning and replay attacks of compromised devices.

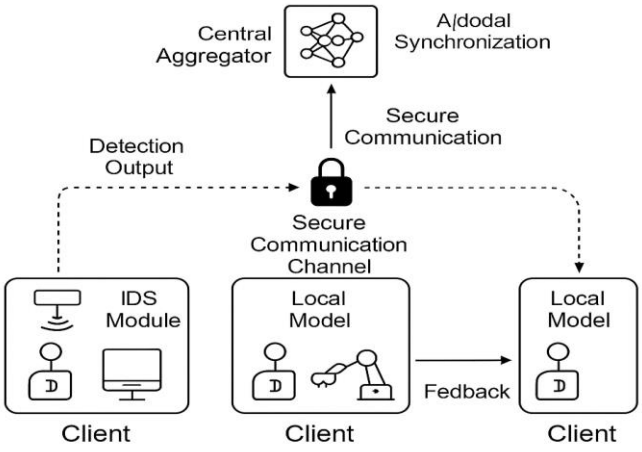

Figure 3: EdgeSecFL System Architecture

Figure 3 captures the complete FL lifecycle in EdgeSecFL from the collection of local intrusion data, through model training and update submission, to global model aggregation and redistribution. The architecture incorporates a feedback loop where it continues to integrate the results of detecting (e.g. false positives in an anomaly detection algorithm or missed anomalies in a host-level intrusion detection system) into further training steps so the system learns dynamically as new threats are introduced in real-time. While [19] focuses on secure microservice communication in optical networks, his architectural insights on encrypted control paths, synchronized service coordination, and multi-tenant access control are

highly relevant in federated environments like EdgeSecFL. These ideals can be seen reflected in the methods of secure synchronization that are employed in EdgeSecFL, allowing the model to remain functional in the presence of various, business-like, fragmented networks.

The efficient and secure implementation forms the building blocks of the architecture of EdgeSecFL, which are two of the most significant needs to implement federated intrusion detection in the IoT-cloud environments. It circumvents the drawbacks of traditional FL frameworks through its use of lightweight models, delta-based update sharing and safe synchronization protocols to open new possibilities of building scalable and privacy-respecting IDS systems. Its architectural design reflects well-established strategies, particularly those by Jangid, on optimizing distributed learning for real-world, constrained environments.

IV. SECURITY ENHANCEMENTS IN EDGESECFL

Security is not merely a functional add-on in the EdgeSecFL model, it is a foundational requirement. Since the nature of threats to Internet of Things-cloud ecosystems is becoming more advanced, EdgeSecFL will introduce many security layers so as to guarantee data integrity, security against adversaries manipulation and policy-compliant governance. This section explores the system's adversarial robustness strategies and its alignment with policy enforcement frameworks such as Cloud Security Posture Management (CSPM).

A. Adversarial Robustness in FL Models

Federated learning systems, while decentralized and privacy-preserving, are not immune to attacks. Malicious parties or participants or malicious devices or malicious parties or participants can manipulate the model update or attacks on model state during synchronization. EdgeSecFL proactively addresses these challenges through multiple robustness mechanisms designed to protect the learning process.

Some common adversarial threats in FL settings include:

- Model Poisoning: Where clients intentionally manipulate training data or updates to degrade the global model's performance.
- Backdoor Attacks: In the malicious attitude where the updated model installs concealed triggers to wrongly recognize certain inputs.
- Inference Attacks: Where adversaries attempt to reconstruct training data from shared updates.
- Free-rider Attacks: Where clients avoid training but still benefit from the global model.

With the aim of mitigating such threats, EdgeSecFL will incorporate numerous defense methods:

- Noise Injection at the local training stage introduces controlled randomness to gradients, reducing the success of inference attacks.
- Adversarial Client Detection Adversarial Client Detection detects deviating updates based on statistical residual or the filtering technique clustering before aggregation.
- Audit Trails maintain immutable logs of client activity and update patterns, aiding in anomaly detection and accountability.

Table 2: Adversarial Threats vs. EdgeSecFL Countermeasures

Threat Type	Example Attack	EdgeSecFL Defense Mechanism
Model Poisoning	Label flipping, gradient bias	Statistical filtering, audit trail logs
Backdoor Insertion	Trigger injection	Update anomaly detection
Inference Attack	Gradient leakage	Noise injection, secure aggregation
Free-rider Behavior	No local training	Contribution-weighted aggregation
Sybil Attack	Multiple fake clients	Signature validation, reputation scoring

As shown in Table 2, each threat is mapped to specific defenses within the EdgeSecFL architecture. The measures are defined by practical attack models that have been studied throughout federated learning literature and augmented by features of decentralized accountability. It highlighted similar attack vectors in blockchain-enabled mobile networks, advocating for decentralized identity validation and tamper-evident logs as key elements for resilient systems concepts which EdgeSecFL builds upon within the FL domain [12].

Moreover, EdgeSecFL uses dynamic feedback to strengthen its robustness in being able to repair itself through retraining. Detection failures or false positives are flagged and fed back into the learning cycle, allowing local models to re-adapt using updated datasets, thereby closing the loop between detection and resilience.

B. Policy Enforcement and Governance

Although technical robustness is paramount, the security governance and adherence to the policies are fundamentally important in the deployment in the real world. EdgeSecFL distinguishes itself by incorporating policy-aware control through integration with Cloud Security Posture Management (CSPM) systems, enabling security teams to define, monitor, and enforce compliance policies across the federated ecosystem.

Researchers [20] emphasizes the importance of automated security policy enforcement in dynamic cloud environments, particularly where device configurations and user access frequently change. CSPM acts as a control tower of cloud-native environments and automates the process of policy violations, misconfigurations, and insecure access rules identification in a nearly real time.

EdgeSecFL leverages this by embedding policy compliance hooks within its federated training cycles. For example, device eligibility for participation in FL training can be linked to:

- Real-time CSPM posture scores,
- Identity verification tokens,
- Encryption compliance checks, and
- Device access logs reviewed against defined governance baselines.

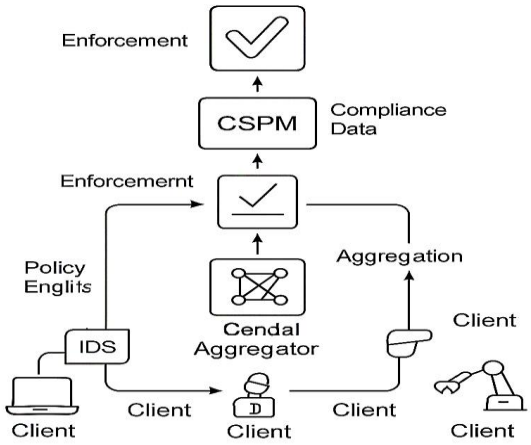

Figure 4: CSPM-Integrated EdgeSecFL Policy Framework

As shown in Figure 4, the CSPM layer acts as an overwatch system, enforcing rules that dictate how, when, and if a device can join the training cycle. This includes dynamic revocation of compromised clients, automatic rollbacks upon policy violations, and dashboard-level visibility into training metrics and security indicators.

Furthermore, CSPM integration supports regulatory mandates such as [21]:

- HIPAA compliance in healthcare IoT systems,
- GDPR controls over personal data exposure, and
- Zero Trust access frameworks for cloud APIs and device identities.

This policy enforcement level is to make sure that EdgeSecFL does not only perform with security but the security is based on the constraints of legal, ethical, and organization regulations, where the boundaries between technical operations and organization responsibility converge. EdgeSecFL has two aspects of security improvements, namely, technical resilience and alignment of governance. Technically, adversarial attack defense is put directly into the training and update protocols, making the training and subsequent update processes smaller [22]. When it comes to governance, seamless interconnectivity with CSPM frameworks allows the system to establish automated compliance and policy-based functionality, so that it can live up to the contemporary needs of cloud security and IoT-based experimentation. Together, these mechanisms establish EdgeSecFL as a secure-by-design federated learning model, prepared for real-world deployment in sensitive and regulated environments.

V. CRYPTOGRAPHIC AND PRIVACY CONSIDERATIONS

The challenge of achieving security in federated learning (FL) is not related only to the architecture of the model and the issues of detecting attacks: the capability of federated learning to preserve the privacy and integrity of datasets during model aggregation and communication is a central concern. Cryptographic protections are imperative in the situation of EdgeSecFL, which runs across different and possibly insecure networks. There are two important features traced in this section: secure aggregation techniques used in FL, and the readiness of these techniques against quantum-enabled threats.

A. Secure Aggregation and Encryption Methods

In order to achieve the privacy and integrity of model updates in federated learning, EdgeSecFL will adapt and incorporate the state-of-the-art cryptography methods to achieve efficient and secure aggregation. The following subcomponents describe how these methods are applied.

a) Federated Averaging and Privacy Concerns

The fundamental process in FL is federated averaging the updates calculated by separated client devices are collected to create a brand-new global design. While this process eliminates the need to transmit raw data, model updates themselves can leak sensitive information if intercepted or reverse-engineered.

b) Homomorphic Encryption (HE)

In response to that, EdgeSecFL is a combination of homomorphic cryptography (HE) [23] and multiparty computation (SMPC) in securing the aggregation procedure. Homomorphic Encryption allows mathematical operations (such as averaging) to be performed directly on encrypted data without decryption. This ensures that the aggregator can compute the global model without ever seeing the individual updates in plaintext.

c) Secure Multiparty Computation (SMPC)

Secure Multiparty Computation enables clients to collaboratively compute an aggregate without disclosing their inputs to one another or to the aggregator [24]. It assumes a partial trust scheme and is maintained against collusion with a group of players.

d) Trade-offs in Cryptographic Techniques

SMPC reduces the privacy risks inherent in FL but comes with trade-offs in terms of computational overhead and communication latency. In the EdgeSecFL, the lightweight versions of these protocols are used so that its versions can be compatible with resource-constrained IoT devices.

e) Applicability in Sensitive Domains

These cryptographic techniques are vital for ensuring confidentiality, integrity, and non-repudiation in federated learning settings [29], particularly when models are used in sensitive environments like healthcare, finance, or critical infrastructure. Together, these methods enable EdgeSecFL to deliver privacy-preserving model training while remaining feasible for deployment in real-world IoT-cloud ecosystems.

B. Post-Quantum Readiness

As quantum computing continues to progress, existing public-key cryptographic schemes including those used in federated learning protocols may become obsolete. Quantum algorithms such as Shor's and Grover's pose serious threats to encryption standards like RSA, ECC, and even symmetric key protocols with insufficient key lengths. The urgency of transitioning to quantum-resistant cryptographic frameworks, especially in national security and critical infrastructure applications [25]. In his roadmap for quantum-resistant cybersecurity, he noted that The long-term viability of cloud and edge security depends on early adoption of post-quantum algorithms and retrofitting current architectures for quantum resilience.

In the FL context, this means integrating lattice-based, hash-based, or code-based encryption techniques that are resistant to known quantum attacks. EdgeSecFL, while currently reliant on classical HE and SMPC protocols, is architected to support pluggable cryptographic backends, allowing migration to post-quantum cryptographic libraries as they become standardized.

Table 3: Cryptographic Techniques in FL vs. Quantum Resistance

Technique	Current Use in FL	Quantum Resistance	Notes
Homomorphic Encryption	Yes	Vulnerable (RSA-based)	Needs replacement with lattice-based HE
Secure Multiparty Computation (SMPC)	Yes	Partially resistant	Depends on underlying primitives
Lattice-Based Encryption	No (Experimental)	Strong	Suitable for future FL deployments
Hash-Based Signatures	Limited Use	Strong	Good for update verification
Code-Based Cryptography	No	Strong	Requires efficient implementation

As shown in Table 3, many current cryptographic approaches used in FL are not inherently resistant to quantum threats. EdgeSecFL anticipates this by remaining modular, enabling a phased transition to post-quantum cryptographic

standards once they are finalized and adopted by global cybersecurity authorities. EdgeSecFL prioritizes cryptographic privacy through secure aggregation protocols while acknowledging the imperative for post-quantum transition planning. Its flexible architecture ensures that future deployments can evolve alongside emerging threats both classical and quantum.

VI. DEPLOYMENT MODELS AND SYSTEM OPTIMIZATION

While security and privacy are critical to any federated learning (FL) system, practical deployment and optimization strategies are equally important for real-world adoption. EdgeSecFL is developed towards practical (as in private and secure) scalable/cost effective intrusion detection systems in IoT-cloud environments. This section examines two key operational dimensions: coordinated training across diverse edge devices and economical use of cloud resources for orchestration and model aggregation.

A. Efficient Edge Device Coordination

In federated learning, one of the most persistent challenges is the asynchronous and unreliable participation of edge devices. Unlike centralized servers or enterprise clients, the IoT nodes can have unpredictable connectivity, inconsistent power state or even become unavailable temporarily. This is tackled by drop out resistant synchronization of models and dynamic client selection as done by EdgeSecFL

Dropout tolerance ensures that the absence of certain client updates in a training round does not disrupt the global aggregation process. Instead of waiting for all participants to submit updates (synchronous FL), EdgeSecFL incorporates an asynchronous learning framework, waiting until updates can be received, and down-weighting updates that are unreliable historically and irrelevant to the data holding historical weight.

To further optimize training efficiency, the system implements adaptive client selection algorithms that prioritize devices based on:

- Network bandwidth,
- Recent participation frequency,
- Local data freshness, and
- Energy availability.

Such an approach saves time needed to train the model but does not overload low-resource clients and hence offers maximum utility of the model. Jangid & Malhotra (2022) [26] highlight similar principles in their work on optimizing software operations across optical transport networks. They emphasize the value of asynchronous task scheduling, adaptive load balancing and event-driven computation all of which are relevant to EdgeSecFL when and where coordination of the devices and collection of the updates is necessary.

B. Cost-Efficient Cloud Usage and Instance Reliability

While federated learning shifts computation to edge devices, cloud infrastructure remains essential for tasks such as global model aggregation, update validation, secure storage, and orchestration. That is why reducing the complexity and cost of the cloud consumption are important primarily in large-scale deployments in distributed IoT fleets. EdgeSecFL supports containerized deployments using Docker or Kubernetes, allowing each component (e.g., aggregator, policy engine, secure aggregation service) to scale independently based on workload. The system also leverages cloud-native autoscaling, where instances are spun up dynamically based on the number of participating clients or the volume of incoming model updates.

A notable cost optimization strategy involves the use of spot instances cloud resources offered at reduced rates due to excess capacity. While spot instances can be revoked at any time, EdgeSecFL mitigates this risk through:

- Redundant aggregators operating in active-passive mode,
- Checkpointing global models after each update round, and
- Fallback to on-demand instances during revocation events.

The techniques allow high availability without having to invest in costly long instances. In case of aggregator failure, other standby containers can resume operations using the most recent model snapshot, ensuring continuity with minimal downtime.

Table 4: Cloud Deployment Strategies for FL in IoT Environments

Deployment Strategy	Reliability	Cost-Efficiency	Resource Utilization
Static VMs (On-Demand)	High	Low	Moderate
Autoscaled Containers	Moderate to High	Moderate to High	Efficient
Spot Instance Scheduling	Moderate (with fallback)	Very High	Very Efficient
Hybrid Aggregation Nodes	High	Moderate	High

As shown in Table 4, container-based orchestration and spot instance scheduling provide the best combination of cost savings and scalability, particularly when paired with EdgeSecFL's fault-tolerant design. Organizations which use federated IDS at scale e.g. smart cities or industrial control systems are able to reduce operational overhead whilst ensuring consistent training and inference models run in a cycle.

VIII. INTEGRATION WITH EMERGING TECHNOLOGIES

To enhance the security, trustworthiness, and automation of federated learning environments, EdgeSecFL embraces emerging technologies such as blockchain and smart contracts. The technologies provide new systems of tamper evident logging, automatization of trust enforcement, and autonomous orchestration which are essential to decentralized intrusion detection implementation on heterogeneous IoT-cloud infrastructures.This section explores two key directions: blockchain-backed secure logging and the use of AI and smart contracts for autonomous model coordination.

A. Blockchain for Secure Logging and Authentication

Federated learning, while decentralized in computation, still depends on centralized aggregators for coordination. This develops a single point of trust and exposes the system to problems such as model tamper, update repudiation, malicious roll back [27]. To mitigate these risks, EdgeSecFL integrates a blockchain-based logging mechanism that records each client update, model version, and system decision in an immutable ledger.

The blockchain is like public accountability. Every transaction is a signed model change, or system event, such as:

- Timestamped client model contributions,
- Update integrity checksums,
- Aggregator signatures confirming participation, and
- Policy compliance metadata (e.g., CSPM scores, update eligibility).

With this log, a model, an auditor, or an administrator can be able to track the origin of each of the updates, identify the outlier behavior and confirm the actuality of models employed in the detection. By anchoring these events to a blockchain, EdgeSecFL ensures non-repudiation, meaning clients cannot deny their participation or falsify their contributions.

Study [12] discusses blockchain's application for securing mobile wireless networks and multi-party interactions, highlighting its utility in creating verifiable and decentralized trust in environments where central control is infeasible. EdgeSecFL extends this intuition by carrying out these concepts to a federated learning setting, where trust of devices cannot be assumed and the history of updates is expected to be as uncontaminated as possible.

B. AI, Smart Contracts, and Autonomy in FL Systems

Although blockchain enables integrity and transparency, smart contracts based on a blockchain infrastructure enables a new horizon of automation and policy enforcement in federated learning. Smart contracts in EdgeSecFL function as programmable agents that enforce:

- Eligibility checks (e.g., clients must meet cryptographic compliance or CSPM thresholds),
- Contribution thresholds (e.g., minimum updates before inclusion in aggregation),
- Reputation-based trust scores (e.g., declining influence from historically unreliable clients).

For instance, a smart contract could automatically reject updates from a client whose previous contributions were statistically anomalous, or dynamically adjust aggregation weights based on real-time IDS feedback scores [28]. In more sophisticated constructions, such smart contracts communicate with intelligent policy engines, powered by AI, which considers device behavior, the threat context, and modification-prediction task performance based on historical telemetry and threat intelligence streams. The result is a self-regulating federated ecosystem where devices are automatically rewarded, penalized, or excluded based on behavior without manual oversight.

Figure 5 visually captures this integration, showing how EdgeSecFL blends federated learning with blockchain infrastructure and smart contract logic to form a secure, trustless orchestration layer. The diagram also illustrates checkpoints where models are hashed, logged, and compared on integrity assurance prior to deployment.

This architecture is particularly suitable for:

- Multi-tenant smart city deployments,
- Cross-border industrial collaboration, and
- Healthcare FL networks, where device trust and policy compliance must be enforced across jurisdictions and systems.

Moreover, the native features of blockchain, as the blockchain platforms develop to accommodate lightweight star consensus and energy-efficient computing environments, their usage starts to be feasible with edge-centered federated systems, such as the EdgeSecFL.

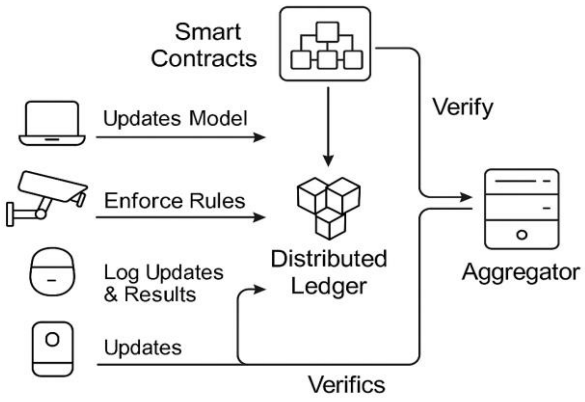

Figure 5: Blockchain-Enhanced Federated Learning Workflow

By integrating blockchain, smart contracts, and AI-based orchestration, EdgeSecFL transforms traditional federated learning into a trustless, self-governing system. Blockchain ensures non-repudiation and tamper-evidence, while smart contracts introduce programmable trust enforcement without human intervention. These emerging technologies enhance the scalability and resilience of intrusion detection in distributed IoT-cloud infrastructures and represent a forward-compatible path toward autonomous, policy-aware security systems.

VII. COMPARATIVE ANALYSIS

To evaluate the practical utility and performance of EdgeSecFL, it is essential to benchmark it against other notable federated learning-based intrusion detection systems (FL-IDS). The following comparative analysis examines five models across key dimensions latency, privacy, scalability, and model performance to highlight EdgeSecFL's strengths and trade-offs.

A. Evaluation Criteria

The comparison uses the following standardized evaluation metrics:

- Latency (ms): Time taken to complete one full model synchronization cycle (i.e., local training + communication + aggregation).
- Model Accuracy: Effectiveness of the trained model in correctly identifying normal vs. malicious activity.
- Privacy Score (qualitative): Degree of privacy preservation based on encryption, data exposure, and vulnerability to inference attacks (scored as High, Medium, or Low).
- Scalability (qualitative): System's ability to perform effectively as the number of participating edge devices increases (rated as Excellent, Good, or Limited).

B. Models Compared: The five selected FL-IDS models are:

- EdgeSecFL: A lightweight, asynchronous, policy-aware FL system with blockchain logging and secure aggregation.
- FedAvg-ID: A basic federated IDS applying standard FedAvg without compression or secure aggregation.
- FedHomeSec: A home IoT-focused model with differential privacy but limited in scalability.
- IoTFLGuard: A bandwidth-optimized FL model with sparse update transmission and edge caching.
- TinyFedDetect: An ultra-light FL framework for microcontrollers using binary classifiers and static model compression.

Table 5: Comparative Performance Analysis of FL-based IDS Models

Model	Latency (ms)	Accuracy	Privacy Score	Scalability
EdgeSecFL	Low	High	High	Excellent
FedAvg-ID	High	Moderate	Low	Limited
FedHomeSec	Moderate	High	High	Limited
IoTFLGuard	Low	Moderate	Medium	Good
TinyFedDetect	Very Low	Low	Medium	Good

C. Interpretation of Results

- Latency: EdgeSecFL demonstrates low latency performance due to its asynchronous update mechanism and differential update sharing. As compared to FedAvg-ID, which pauses in the synchronous training rounds, EdgeSecFL does not pause and allows unreliable or occasionally connected devices.
- Accuracy: EdgeSecFL balances compression and model integrity, maintaining high detection accuracy even in bandwidth-constrained conditions. Although FedHomeSec can work in the small-scale settings similarly, the structure of the model is not able to handle the heterogeneous traffic datasets and that significantly decreases its overall usefulness.
- Privacy Score: Thanks to its integration of homomorphic encryption, secure multiparty computation (SMPC), and blockchain-backed audit trails, EdgeSecFL scores highest in privacy protection. FedAvg-ID does not have any encryption or privacy-preserving mechanisms by contrast, and is thus vulnerable to gradient leakage and adversarial inference attacks.
- Scalability: EdgeSecFL's use of adaptive client selection, dropout tolerance, and containerized aggregation enables it to scale across large IoT networks with minimal manual configuration. IoTFLGuard and TinyFedDetect have limited scalability with no orchestration capabilities and dynamic policy enforcement capabilities as EdgeSecFL has according to their parent CSPM and smart contract visitors.

This comparative analysis demonstrates that EdgeSecFL stands out as a balanced and future-ready federated IDS framework, excelling in privacy, flexibility, and scalability without sacrificing performance. Other FL models provided trade-offs along one or several dimensions, which underlines the fact that the layered security model, the modular architecture of EdgeSecFL can be reliable in terms of addressing the present and new intrusion detection in IoT-cloud infrastructures.

VIII. POLICY, ETHICS, AND GOVERNANCE IMPLICATIONS

While federated learning models like EdgeSecFL offer promising technical advantages in securing IoT-cloud infrastructures, their real-world deployment also raises critical ethical, legal, and policy concerns. As these systems increasingly influence decisions in sensitive domains such as national security, healthcare, finance, and critical infrastructure, developers and policymakers must evaluate not only how secure and scalable they are, but also how fair, explainable, and compliant they remain with global standards.

A. AI in National Security: Risk-Benefit Trade-offs

AI-based intrusion detection systems (IDS) may incur quicker responses to the threats and a reduced chance of human error. Such advances to the defense and public safety sectors. However, such benefits come with significant risks. Federated learning models, despite their privacy-aware architecture, still require central coordination, trust in aggregators, and reliable behavior from edge devices, factors that adversaries could exploit if not properly governed.

Faraz Ahmed (2024) [20] addresses this duality in his work on AI and cybersecurity policy frameworks, especially in the context of national and critical infrastructure defense. Intelligent application of AI in cybersecurity in the public sector requires the equilibrium of automation and control. There must be accountability tools against systems oversights, prejudice, or misuse of smart systems. In the context of EdgeSecFL, this means deploying the system with strict governance controls, risk-based access models, and continuous policy enforcement, especially in environments where false positives or missed alerts could have life-threatening consequences.

B. Bias, Fairness, and Explainability in Federated IDS

One of the unique ethical challenges in federated learning systems is the risk of unintended bias. This is because local models are only trained on dissimilar, data that might simply show local or device-specific traffic; resulting in skewed decision boundaries, such that what is anomalous in one setting is erroneously detected as benign or malicious in another.

Moreover, the black-box nature of deep learning models, often used in FL-based IDS, makes explainability a major concern. Security analysts and system administrators must be able to interpret why a certain event was flagged as an intrusion, especially when responding to incidents or making compliance-related decisions.

To address this:

- Fairness-aware aggregation techniques should be employed, where model updates are weighted to avoid overfitting to dominant client data.
- Interpretable ML methods such as LIME or SHAP could be integrated into the final decision layer of EdgeSecFL to offer context-sensitive explanations.
- Auditable logs (e.g., via blockchain) can serve as transparency-enhancing tools to trace the origin and logic of detection decisions.

Without these controls, EdgeSecFL—despite its technical merits could reinforce or introduce structural biases, especially in multi-tenant environments where client devices come from diverse administrative, geographic, or socio-economic backgrounds.

C. Global Standards and FL Compliance

Federated learning frameworks must also comply with an expanding set of global data protection and AI governance regulations. Standards like the General Data Protection Regulation (GDPR) in the EU, the NIST Privacy Framework in the US, and ISO/IEC 27701 for privacy information management directly affect how federated data pipelines must be designed and maintained.

Table 6: Ethical and Policy Frameworks for Federated IDS Models

Framework/Standard	Focus Area	Relevance to FL and EdgeSecFL
GDPR (EU)	Data minimization, consent	FL preserves data locality but must ensure local consent enforcement
NIST AI RMF (USA)	Trustworthy AI, risk management	Requires explainability, resilience, bias control in FL models
ISO/IEC 27701	Privacy Information Management	Requires data handling and audit capabilities at the edge
OECD AI Principles	Human oversight, transparency	Encourages accountable FL deployment in public sectors
AI Act (proposed EU)	High-risk AI regulation	May categorize FL-based IDS as high-risk in critical applications

Table 6 summarizes the intersection between key global policy frameworks and federated IDS systems like EdgeSecFL. To remain viable, such systems must be adaptable to jurisdictional requirements and capable of embedding governance hooks that can audit compliance in real time.

EdgeSecFL, like any advanced AI system deployed at scale, exists at the crossroads of technical innovation and ethical responsibility. Its success in securing IoT-cloud environments will depend not only on detection accuracy or communication efficiency, but also on how well it aligns with ethical principles, fairness metrics, and international policy mandates. As federated intrusion detection becomes more widespread, embedding ethical guardrails and policy integration mechanisms will be essential for fostering trust and accountability in its use.

IX. FUTURE DIRECTIONS AND RECOMMENDATIONS

Looking ahead, the evolution of federated intrusion detection systems like EdgeSecFL will hinge on addressing emerging challenges across scalability, transparency, and post-quantum resilience. A key recommendation is to advance edge-level decentralization, where aggregation responsibilities are distributed across multiple semi-trusted nodes to reduce reliance on central servers and improve fault tolerance. In parallel, there is a pressing need for explainable federated learning (XFL) frameworks that combine high detection accuracy with transparent, human-interpretable decision-making especially critical in regulated sectors like healthcare and defense.

Moreover, as quantum computing progresses, it is essential to pursue quantum-ready FL training pipelines that integrate post-quantum cryptographic (PQC) primitives without compromising performance or scalability. However, it must be acknowledged that many of these components including lattice-based encryption protocols and AI-agent-based FL orchestration remain in experimental or non-standardized phases. To be fully absorbed into production ready-to-use IDS, they will need more research on topics of system compatibility, performance optimization, and regulatory acceptance

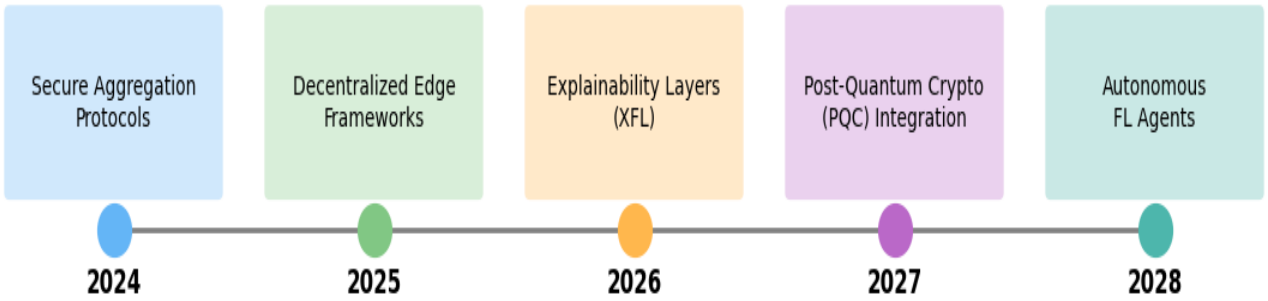

Figure 7: Research Roadmap for Secure FL in IoT-Cloud Ecosystems

By following this roadmap, the research community of FL can take advantage of pilot designs such as EdgeSecFL to come up with future-ready privacy-preserving, and ethically consistent security systems in an IoT-cloud distributed setting

X. CONCLUSION

EdgeSecFL offers a robust and adaptable approach to intrusion detection within complex IoT-cloud ecosystems. By leveraging decentralized learning and secure model aggregation, it addresses key limitations of traditional centralized systems, such as bottlenecks, data privacy risks, and limited scalability. The privacy-aware coordination of the heterogeneous edge devices provided by the framework, allows keeping sensitive data local and collaboratively detecting the threats. A major strength of EdgeSecFL lies in its integration of complementary technologies, including blockchain for tamper-evident logging and cloud-native orchestration for resource-aware deployment. These design options improve transparency of the system, shorten operation latency, and facilitate the deployment in real world, where the devices are not always accessible and connected with each other consistently.

Through comparative analysis with existing FL-based IDS models, EdgeSecFL demonstrates improved performance in areas such as latency, resilience, and privacy enforcement. It is modular and can be integrated with security policy standards such as CSPM, as well as fairness-aware and interpretable components used in trust-confined decisions of regulated or risky situations. The framework also reflects a thoughtful alignment with broader ethical and governance considerations. The problem of model bias, transparency, and alignment with data protection regulations are also directly answered and confirm the applicability of EdgeSecFL to be implemented in real life, both with a government and privately.

Although some emerging technologies such as post-quantum cryptographic safeguards and AI-driven orchestration are not yet fully mature, EdgeSecFL establishes a solid foundation for their future integration. Further development must concentrate on advancing understandability of the models, support decentralized schemes of trust, and keep pace with the changing requirements of cryptography. EdgeSecFL represents a future-ready, ethically conscious, and technically sound solution for intrusion detection. Its balanced emphasis on security, scalability, and policy alignment positions it as a leading candidate for protecting distributed digital infrastructures in a privacy-sensitive and rapidly evolving threat landscape.

XI. REFERENCES

[1] K. Valaskova et al., "Blockchain-enabled Internet of Things and Virtual Sensor Networks, Mobile Cloud and Edge Computing Systems, and Predictive Modeling and Cognitive Data Fusion Techniques in Big Data-driven Digital Twin Urban Geopolitics and Immersive Hyper-Connected Virtual Spaces," Geopolitics, History, and International Relations, vol. 15, no. 1, pp. 46–60, 2023.

[2] A. Sehgal et al., "Management of resource constrained devices in the internet of things," IEEE Communications Magazine, vol. 50, no. 12, pp. 144–149, Dec. 2012.

[3] G. Dhayanidhi, "Research on IoT threats & implementation of AI/ML to address emerging cybersecurity issues in IoT with cloud computing," 2022.

[4] E. T. M. Beltrán et al., "Decentralized federated learning: Fundamentals, state of the art, frameworks, trends, and challenges," IEEE Communications Surveys & Tutorials, vol. 25, no. 4, pp. 2983–3013, 2023.

[5] N. Kshetri, "Privacy and security issues in cloud computing: The role of institutions and institutional evolution," Telecommunications Policy, vol. 37, no. 4–5, pp. 372–386, 2013.

[6] F. Ahmed, "Cloud Security Posture Management (CSPM): Automating Security Policy Enforcement in Cloud Environments," ESP Int. J. Adv. Comput. Technol. (ESP-IJACT), vol. 1, no. 3, pp. 157–166, 2023.

[7] M. A. Serhani et al., "Self-adapting cloud services orchestration for fulfilling intensive sensory data-driven IoT workflows," Future Generation Computer Systems, vol. 108, pp. 583–597, 2020.

[8] M. Kosicki et al., "Big Data and Cloud Computing for the Built Environment," in Industry 4.0 for the Built Environment: Methodologies, Technologies and Skills, Cham: Springer, 2021, pp. 131–155.

[9] A. R. Biswas and R. Giaffreda, "IoT and cloud convergence: Opportunities and challenges," in Proc. 2014 IEEE World Forum on Internet of Things (WF-IoT), 2014.

[10] P. Anand et al., "IoT vulnerability assessment for sustainable computing: Threats, current solutions, and open challenges," IEEE Access, vol. 8, pp. 168825–168853, 2020.

[11] I. Butun, P. Österberg, and H. Song, "Security of the Internet of Things: Vulnerabilities, attacks, and countermeasures," IEEE Communications Surveys & Tutorials, vol. 22, no. 1, pp. 616–644, 2019.

[12] J. Jangid et al., "Enhancing security and efficiency in wireless mobile networks through blockchain," Int. J. Intell. Syst. Appl. Eng., vol. 11, no. 4, pp. 958–969, 2023.

[13] M. F. Elrawy, A. I. Awad, and H. F. A. Hamed, "Intrusion detection systems for IoT-based smart environments: a survey," J. Cloud Comput., vol. 7, no. 1, pp. 1–20, 2018.

[14] S. AbdulRahman et al., "A survey on federated learning: The journey from centralized to distributed on-site learning and beyond," IEEE Internet Things J., vol. 8, no. 7, pp. 5476–5497, 2020.

[15] B. Nagajayanthi, "Decades of internet of things towards twenty-first century: A research-based introspective," Wireless Personal Communications, vol. 123, no. 4, pp. 3661–3697, 2022.

[16] R. Morabito et al., "Consolidate IoT edge computing with lightweight virtualization," IEEE Network, vol. 32, no. 1, pp. 102–111, 2018.
[17] J. Jangid, "Efficient Training Data Caching for Deep Learning in Edge Computing Networks," Int. J. Sci. Res. Comput. Sci. Eng. Inf. Technol., vol. 7, no. 5, pp. 337–362, 2020.
[18] O. A. Wahab et al., "Federated machine learning: Survey, multi-level classification, desirable criteria and future directions in communication and networking systems," IEEE Communications Surveys & Tutorials, vol. 23, no. 2, pp. 1342–1397, 2021.
[19] J. Jangid, "Secure microservice communication in optical networks," J. Inf. Syst. Eng. Manag., vol. 10, no. 21s, 2025. doi: 10.52783/jisem.v10i21s.3455
[20] F. Ahmed, "Cybersecurity Policy Frameworks for AI in Government: Balancing National Security and Privacy Concerns," Int. J. Multidiscip. Sci. Manag., vol. 1, no. 4, pp. 43–53, 2024.
[21] S. Muzukwe, A Governance Framework for Security in Cloud Architecture, MS Thesis, Univ. of Johannesburg, South Africa, 2023.
[22] A. Chakraborty et al., "Adversarial attacks and defences: A survey," arXiv preprint arXiv:1810.00069, 2018.
[23] A. Acar et al., "A survey on homomorphic encryption schemes: Theory and implementation," ACM Comput. Surv., vol. 51, no. 4, pp. 1–35, 2018.
[24] C. Zhao et al., "Secure multi-party computation: theory, practice and applications," Information Sciences, vol. 476, pp. 357–372, 2019.
[25] F. Ahmed, "Quantum-Resistant Cryptography for National Security: A Policy and Implementation Roadmap," Int. J. Multidiscip. Sci. Manag., vol. 1, no. 4, pp. 54–65, 2024.
[26] J. Jangid and S. Malhotra, "Optimizing Software Upgrades in Optical Transport Networks: Challenges and Best Practices," Nanotechnology Perceptions, vol. 18, no. 2, pp. 194–206, 2022.
[27] A. Ali et al., "BCALS: Blockchain-based secure log management system for cloud computing," Trans. Emerg. Telecommun. Technol., vol. 33, no. 4, e4272, 2022.
[28] K. Demertzis et al., "Anomaly detection via blockchained deep learning smart contracts in industry 4.0," Neural Comput. Appl., vol. 32, no. 23, pp. 17361–17378, 2020.
[29] K. A. Awan et al., "Privacy-preserving big data security for IoT with federated learning and cryptography," IEEE Access, vol. 11, pp. 120918–120934, 2023.