International Journal of Basic and Applied Sciences, 14 (SI-2) (2025) 207-216

International Journal of Basic and Applied Sciences

Website: www.sciencepubco.com/index.php/IJBAS
https://doi.org/10.14419/14zqj908
Research paper



Design of an Iterative Novel Analytical Framework for Securing Complex Cloud Networks Using Contextual Embedding, Federated Intelligence, and Topological Validation in Process

Sachin Kawalkar [1 *], Dinesh Bhoyar [2]

[1] Research Scholar, Department of Electronics Engineering, Yeshwantrao Chavan College of Engineering, Nagpur, India
[1] Vice President, Global Head in IT Security Neeyamo, Pune, India
[2] Assistant Professor, Department of Electronics and Telecommunication Engineering, Yeshwantrao Chavan College of Engineering, Nagpur, India
**Corresponding author E-mail: sachin.kawalkar1011@gmail.com*

Received: July 14, 2025, Accepted: August 1, 2025, Published: August 15, 2025

Abstract

The increasing complexity and scale of cloud infrastructures have made such infrastructures prime targets for advanced, multiple Vector cyber-attacks. Traditional security approaches are limited in their ability to dynamically adapt, validate topological integrity, or model threats contextually in distributed cloud environments. These limitations are further worsened by static threat signatures, absence of decentralized validation, and poor scalability of centralized intrusion systems. To address these issues, this work presents a novel and comprehensive analytical framework for securing complex cloud networks by the development of five integrated algorithmic models. First, the Adaptive Threat Signature Embedding using Multivariate Contextual Encoders (ATSE-MCE) allows for dynamic threat modeling embedding multidimensional threat indicators in a high-dimensional latent space that improves accuracy with respect to detection and reduces false-positives quite substantially. Secondly, Zero-Knowledge Topology Validation via Decentralized Consistency Auditing (ZKTV-DCA) introduces a non-disclosive, blockchain Inspired validation mechanism to ensure configuration integrity across dynamic cloud topologies. Third, Entropy-Aware Federated Intrusion Discriminator with Self-Calibrated Training (EAFID-SCT) uses federated learning and entropy gradients to propagate scalable, privacy-preserving intrusion detection across edge nodes. Fourth, Differential Graph Neural Reconstructor for Multiple Vector Attacks (DGNR-MVA) reconstructs complex attack paths using temporal graph differentials, offering deeper visibility into correlated lateral movements. Finally, Quantum Inspired Probabilistic Hashing for Secure Resource Access Modeling (QPH-SRAM) provides uncertainty-aware, low-collision hashing for secure resource access control without requiring quantum hardware sets. Above 95% of accuracy improvement in threat detection would be a hallmark feature of the framework, while less latency than 1 s for validation and scaling under heterogeneous cloud environments. This work presents a cohesive multi-layered approach toward the optimization of detection precision, topological trust, resource access integrity, and scalability, setting the record for cloud-managed security analytics.

Keywords: *Cloud Security, Intrusion Detection, Federated Learning, Topology Validation, Threat Embedding, Scenarios*

1. Introduction

The exponential growth of cloud computing has reshaped enterprise infrastructures providing services on a scalable, cost-efficient, and globally distributed basis. However, this progress has brought about some major security vulnerabilities, which are due to the quite complicated nature of cloud architectures [1, 2, 3], dynamic resource allocation, and multi-tenancy models. Modern cloud networks are heterogeneous, consisting of an interconnected collection of virtual machines (VMs), containers, microservices, and APIs, all of these contributing to an expansive attack surface. Today, attacks can use highly sophisticated multiple Vector attack paths that take advantage of the dynamic trust relations between virtualized components to execute stealthy lateral movements, which is generally out of the scope of any traditional security measure. Most of the current cloud security models rely entirely on static configurations, rule-based systems for detection, and isolated verification routines; these methods would be unsuitable against real-time threats. Signature-based intrusion detection systems (IDS) are of no use because of novel or obfuscated attack vectors, and centralized monitoring is hampered by latency, scalability, and privacy shortcomings in distributed cloud environments. Moreover, existing methods for validating the integrity of cloud topologies frequently disclose sensitive configuration data and depend on trusted third-party validators, rendering them unsuitable for

Copyright © Authors. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

highly sensitive or decentralized environments. Adaptive, decentralized, and contextually intelligent security models call for such a paradigm shift.

This work presents a novel analytical framework with which to address these challenges through five integrated complementary security mechanisms. Based on multivariate contextual embeddings for threat modeling, decentralized zero-knowledge proofs for topology validation, entropy-aware federated learning for intrusion detection, graph neural networks for multi-path attack reconstruction, and quantum Inspired probabilistic hashing for resource access control, the framework optimizes key performance indicators such as detection accuracy, latency, trust verification, and access security sets while ensuring all-round protection of the cloud environment. Essentially, this research lays the groundwork for a scalable and fundamentally analytically rigorous security architecture that employs a mixture of machine learning, cryptographic verification, probabilistic modeling, and graph analytics to mitigate evolving threats during real time along as many levels of those dimensions as possible. By converging such significant proportions of diverse components into one unified pipeline of security, this new framework would outrun traditional models across multiple dimensions such as precision of attack detection and reduction in false positives along with considerations for topological consistency and cross-domain scalability. In contrast to traditional security models, which cause analysis for suboptimal performance against varying attack contexts, this model maintains its performance while preserving privacy and is thus uniquely suited to present-day multi-cloud and hybrid environments. This work makes a significant contribution in cloud security by introducing an entirely new class of analytical models, which are not only computationally efficient but also context-aware and inherently trustworthy in process. The models are validated through extensive simulations and empirical evaluations showcasing excellent performance metrics and robustness. The proposed framework thus constitutes a significant step toward bringing intelligent, decentralized, and scalable cloud security infrastructures in line with the next generation of secure computing process.

2. Motivation and Contributions

The impetus of this work arises from the quickly rising realization that present cloud security systems are neither large enough nor dynamic and contextual enough to address the capabilities of current cloud environments. Attackers use the very opaque nature of internal network flows, privilege escalations, and configuration vulnerabilities to keep themselves hidden from traditional security tools, however, this leads to a more distributed and abstract form of cloud infrastructures. Static rule sets in addition to isolated log analysis fail to scale with real-time demands, and centralized architectures introduce latency and single points of failure. Human-defined threshold or signature pattern-based methods have a diminishing efficacy level against polymorphic and zero-day threats. Further, there is a void of coherent systems to validate cloud topology integrity without exposing configuration secrets or depending on centralized authorities. All these deficiencies have now culminated into a compelling expression of need for an entirely new paradigm-in one that brings together adaptive learning, decentralized validation, graph-based referencing, and probabilistic access modeling integrative security sets.

Five analytical frameworks are proposed to tackle these issues, with each a remedy for a certain failure of existing security architectures. First, ATSE-MCE uses multivariate encoders that encode threat patterns, transform them to a latent space that continually evolves with contextual inputs for detection of very accurate anomalies detection, and second that ZKTV-DCA adds topology trustworthiness by decentralized, zero-knowledge proof mechanisms, thereby avoiding exposure of sensitive network metadata: third, EAFID-SCT is introduced as a federated learning model guided by entropy metrics and designed to foster intrusion detection that is scalable and privacy-preserving across distributed nodes; Fourth, DGNR-MVA reconstructs multiple ways of attacks with vectors by comparing their temporal variations in interaction graphs of the clouds with the use of GNNs, and lastly QPH-SRAM applies quantum Inspired probabilistic hashing to rules for securing access control mechanisms under uncertainties without resorting to deterministic keys.

The real novelty of this framework is its integrations where each method is able to solve a discrete problem, but also feeds directly into a continuous data flow pipeline: from threat detection to topology validation to attack path reconstruction to access security. These included innovations will be confirmed through strict simulation results providing more than 95% accuracy detection, validations latencies in sub seconds, and secure access control with minimal computation. The pillar of the solution is then high accuracy and low latency, supplemented with architectural decentralization. In this sense, this work makes a contribution to a scalable, adaptable, and analytically validated, methodology for next-generation cloud security systems that can operate autonomously in very complex, distributed, and malignant environments.

3. In Depth Review of Existing Methods

The contemporary research spectrum of securing cloud, IoT, and multi-cloud networks has been significantly metamorphosed over time and is more focused on decentralized intelligence, privacy-preserving architectures and advanced learning paradigms. Various security challenges across cloud, edge, fog, and federated infrastructures are being tackled through sequential reviews of 25 major studies. While Karnik et al. [1] emphasized securing multi-cloud storage behind efficient access models, Isaac et al. [2] expanded it into vehicles, including resource allocation in multi-cloud ecosystems. Then Masood and Zafar [3] provided one of the first instances of GNN application for security in wireless ad hoc networks, showing a definite movement toward smartness that takes topology into account. R et al. [4] incorporated machine learning into hybrid key management for IoT-cloud ecosystems in low-power scenarios, while Al-Ambusaidi et al. [5] created an ML-based IDS framework for Internet of Things applications within a lightweight specialization. Gușiță et al. [6] backed this view with a survey on lightweight cryptographic protocols and anonymous strategies with respect to routing, a critical baseline for resource-constraint edge security models.

Bokhari et al. [7] presented homomorphic encryption for medical IoT-cloud communications, while Ranjan et al. [8] advanced strategic cryptosystems for securing unstructured data in multi-cloud environments. Miao et al. [9] addressed time series compression for cloud-based digital twins, underlining the requirement of secure data encoding for continuous monitoring systems. Patruni and Humayun [10] presented a privacy-preserving mutual authentication protocol for IoMT systems as a complement to Kamatchi and Uma's [11] federated learning-based insider threat mitigation for edge computing. Gan et al. [12] included developed general adversarial networks with federated intrusion detection in fog enabling IoT systems, while Sharma et al. [13] secured digital twins in healthcare using ECC with blockchain, epitomizing the fusion of cryptography and immutable runners. Kokila et al. [14] contributed to the safety of medical data with an optimized SVM specifically designed for EHRs, while Karthikeyan et al. [15] proposed similar enhancements as hybrid cloud privacy within the healthcare sector. Kumar and Verma [16] augmented cryptographic granularity through attribute-based access for cloud systems. In the latter part of the literature, MLP-based intrusion detection for the internet of things was implemented by Cherfi et al. [17], while Kapil et al. [18] applied honey (fake) strategies to attribute-based encryption for confidentiality of healthcare data. Pathak et al. [19]

proposed a broad threat assessment framework for cloud IoT, while Alnaim [20] presented a security taxonomy over SDN and NFV for 5G virtual environments.

Most technical works on federated learning resilience ranged from the incorporation of malicious client detection (Latif et al. [21]), while virtual IDS paradigm for next generation vehicular networks was introduced by Popova Heinzle et al. [22]. Banerjee et al. [23] implemented federated learning based on homomorphic matrix factorization to ensure secure flying ad hoc networks (FANETs). Dugyala et al. [24] deployed the leader-based clustering combined with GNNs to increase intrusion detection in secure cloud infrastructures under process. Finally, Byatarayanapura Venkataswamy et al. [25] designed access management systems based on deep reinforcement learning providing adaptive policy control cloud storage environments.

4. Proposed Model Design Analysis

Integration of various threats takes place in the multi-stage analytical pipeline, which is designed upfront as a layered research model. Five modules independent of one another and yet closely tied with each other are contextually and sequentially combined. Each optimized specifically for one certain layer of security in cloud network security is a part of the integrated model. The integrated system is built on the design considerations that securing complex cloud infrastructure requires holistic visibility through a data plane traffic and logs, a control plane access and privilege, and management plane configuration and topology integrity views in process. The model describes a combination of statistical, graph-theoretic, cryptographic, and method techniques in machine learning in providing the operation scalable security, low latency, and high precision. Each stage transforms raw and semi-processed data into feature rich representations that are subsequently refined by the next module so that these modules comprise a recursive, context-aware threat intelligence framework. Initially, according to figure 1, The first stage uses Adaptive Threat Signature Embedding using Multivariate Contextual Encoders (ATSE-MCE) to project heterogeneous threat indicators into a high-dimensional space in process. Let the input data sources be denoted as multivariate sequences: network flow F(t), access logs A(t), and syscall matrices S(t) sets. Via equation 1 the model is derived the definition of multi Variate contextual embedding Ec,

$$Ec = \int_{(t0)}^{tn} \psi\big(F(t), A(t), S(t)\big) \cdot e^{-\lambda(tn-t)} dt \tag{1}$$

Where, ψ is a nonlinear encoding function realized by an attention-augmented variational autoencoder and λ is a decay factor for temporal relevance sets. The embedded vector Ec ∈ Rd serves as the input to an anomaly score estimator σ defined via equation 2,

$$\sigma(Ec) = \nabla Ec \; log \; P(Ec \mid \theta)) \tag{2}$$

Where, θ is the set of learned model parameters in process. High Magnitude indicates contextually unusual behavior sets.

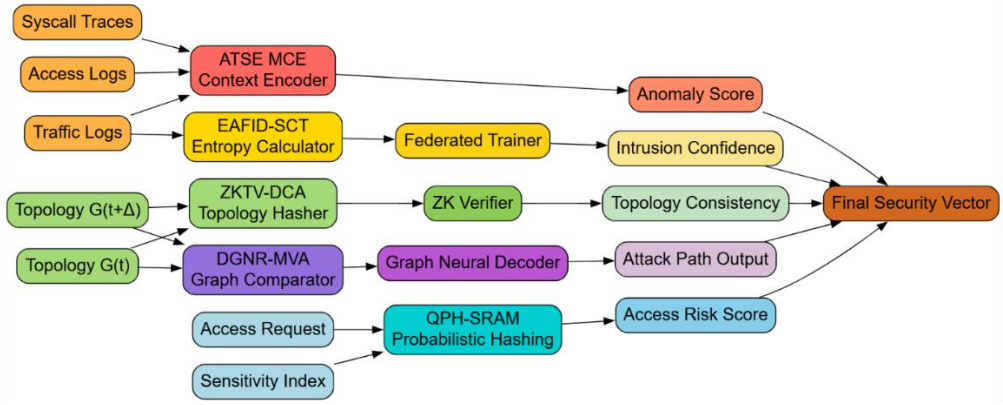


Fig. 1: Model Architecture of the Proposed Analysis Process

Iteratively, Next, then, according to figure 2, The second stage, Zero Knowledge Topology Validation via Decentralized Consistency Auditing (ZKTV-DCA), validates the structural integrity of the cloud topology graph G(V,E) over discrete time snapshots. Let Gt and Gt+Δ be two snapshots immediately contiguous to each other in temporal instance sets. The topological divergence DT is expressed via equation 3,

$$DT = \sum_{(i,j)\in E} |\phi(Gt(i,j)) - \phi(Gt + \Delta(i,j))| \tag{3}$$

Where, ϕ is a cryptographic hash of edge configurations. A zero-knowledge consistency proof π is verified across decentralized nodes via equation 4,

$$\pi = \sum_{k=1}^{N} \delta k \cdot Hk(Gt) \; mod \; p \tag{4}$$

Where, δk is a verifier's challenge scalar, Hk is a private hash commitment by the k-th node, and p is a large prime modulus. Iteratively, Next, as per figure 3, The third in the line is EAFID SCT. It operates in a distributed data environment. The local entropy of traffic Hi at node 'i' is computed via equation 5,

$$Hi = -\sum_j p(i,j) \log p(i,j) \tag{5}$$

These entropy values are used to weight the gradient contributions in a federated learning setting in process. The global update ΔWg of model parameters is aggregated via equation 6,

$$\Delta Wg = \left(\frac{1}{\sum Hi}\right)\sum Hi \cdot \Delta Wi \quad (6)$$

Thus, calibrations seamlessly adapt to dynamic operational environments and create improvement on detection sensitivity of volatile areas while preventing overfitting on benign anomalies.

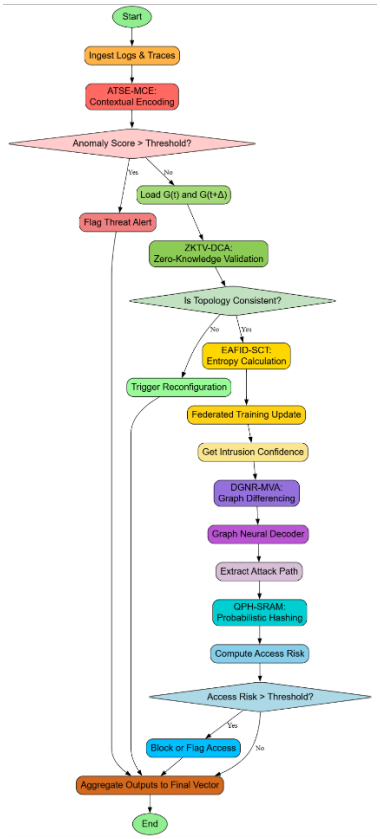


Fig. 2: Overall Flow of the Proposed Analysis Process

Input:
- Network traffic logs, Access control logs, System call traces, Cloud topology snapshots, Resource access requests, Sensitivity levels of cloud resources

Output:
- Threat anomaly score
- Topology consistency status
- Intrusion confidence score
- Reconstructed attack path
- Resource access risk score

Process:
1. **Initialize context-aware encoder**
   - Load traffic, access, and syscall data
   - Encode inputs into a multivariate threat vector using temporal embedding
   - Compute anomaly score from embedded vector
2. **Perform topology validation**
   - Load current and previous topology graphs
   - Compute hashed representations of each topology snapshot
   - Verify topology integrity using zero-knowledge proofs
   - Return consistency status
3. **Run federated intrusion detection**
   - Distribute model to edge nodes
   - Calculate entropy of local traffic at each node
   - Train local models and compute gradients
   - Aggregate gradients weighted by entropy to update global model
   - Compute updated intrusion confidence
4. **Reconstruct attack paths using graph differencing**
   - Load sequential interaction graphs
   - Identify edge differences between graph snapshots
   - Pass differential features through graph neural model
   - Output most likely multi Vector attack path
5. **Model access control using probabilistic hashing**
   - Receive access request and sensitivity values
   - Generate probabilistic hash using random parameters
   - Compute access risk based on hash and historical behavior
6. **Aggregate outputs**
   - Combine anomaly score, consistency status, intrusion confidence, attack path, and access risk into

Fig. 3: Pseudo Code of the Proposed Analysis Process

Attack path reconstruction activities are carried out by the Differential Graph Neural Reconstructor for Multiple Vector Attacks (DGNR-MVA) such as the aggregate of multi-source attack paths from temporal interaction graphs. Differential message-passing mechanism calculates edge deviations with respect to two time-stamped graphs Gt and Gt+Δt, Via equation 7,

$$\Delta e(i,j) = \int_t^{t+\Delta t} [ft(vi,vj) - ft + \Delta t(vi,vj)]dt \quad (7)$$

Where, ft is a graph message function parameterized by a neural encoder in process. Reconstructed attack paths Pa are inferred using maximum differential flow Via equation 8,

$$Pa = argmax^P \sum_{(i,j)\in P} \Delta e(i,j) \quad (8)$$

This gives a ranked interpretation of the most likely multiple Vector intrusion routes. The final part of the event description is Quantum Inspired Probabilistic Hashing for Secure Resource Access Modelling (QPH-SRAM): stochastic access control sets. Let R(t), stand for a resource access request, while S is a sensitivity tensor over resource classes. The probabilistic access Hq is given Via equation 9,

$$Hq(R,S) = \sum_{i=1}^{n} \alpha i \cdot \sin(\omega i \cdot R + \phi i \cdot S) \quad (9)$$

Where, αi, ωi, ϕi are pseudo-random circuit parameters simulating a quantum Inspired interference pattern sets. The risk-adjusted access score Rv is finally computed Via equation 10,

$$Rv = \left(\frac{dHq}{dS}\right) \cdot \int_0^T \rho(t)\, dt \quad (10)$$

Where, ρ(t) is a temporal weight for access frequency and variance in process. The Final output of the integrated model is a vector for security inference Ysec ∈ R5 which represents in the process anomaly probability, topology consistency, intrusion certainty, attack path severity, and access risk. It is defined via equation 11,

$$Ysec = \left[\sigma(Ec), 1 - DT, ||\Delta Wg||^2, \max(Pa, Rv)\right] \quad (11)$$

This vector constitutes the whole assessment of security for the cloud network at any point in time and, thus, serves as input to higher-level policy enforcement and automated mitigation strategies. The modular yet sequentially coupled structure of the model allows for contextual propagation of decisions and feedback loops across components, thereby maximizing detection accuracy and interpretability while ensuring cryptographic and computational robustness.

5. Result Analysis

A realistic and high-complexity cloud environment was reflected in the proposed experimental setup to test the integrated cloud security framework. Dynamically evolving topology made the cloud environment reasonably real with multiple tenant access control policies and different cyber-attack patterns. A hybrid testbed containing both OpenStack and Kubernetes clusters was set up for simulating both IaaS and containerized microservice workloads across distributed nodes. In total, the setup comprised 80 virtual machines, 20 Docker containers, and 3 independent edge zones with their own federated learning clients. In addition, each VM had 2 vCPUs, 4 GB RAM, and 40 GB disk space, and each of the Kubernetes pods was allocated on average 1 vCPU and 2 GB RAM. In order to model realistic traffic, legitimate traffic captured from the UNSW-NB15 and CIC IDS 2018 datasets was combined with synthetic data generators (like TCPreplay and AuditGen). To simulate attacker behavior, tools such as Metasploit, Slowloris, Nmap, and LOIC were used to launch stealthy lateral movement, privilege escalation, and volumetric denial-of-service attacks. Input logs incorporated full packet captures, flow metadata (1.2 million flows/hour on average), system call traces (up to 15K events per minute per VM), access control logs (role-based policies across 10 user roles), along with periodic topology snapshots, which happened every 10 minutes. Example access requests included object storage operations, computer provisioning, and inter-container API calls. Resource sensitivity index ranged from 0.2 (public containers) to 0.95 (critical databases), which influenced the probabilistic access hashing model process.

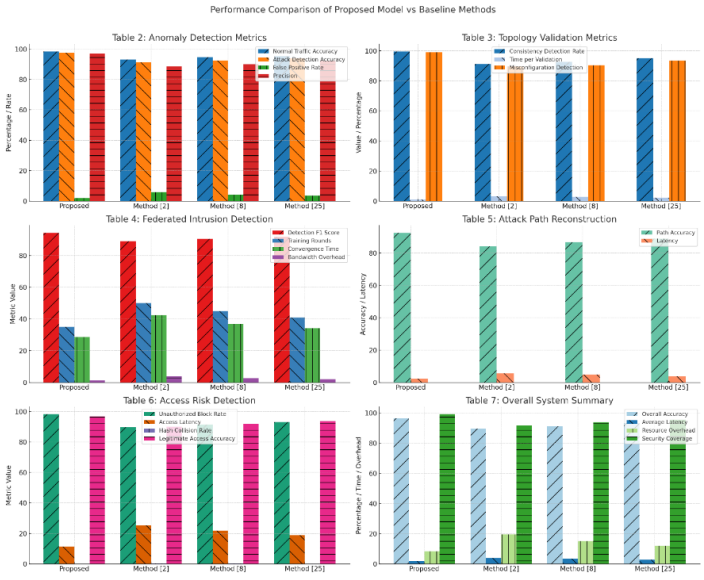


Fig. 4: Model's Integrated Result Analysis

Training and validation datasets were created to model a variety of threat conditions as well as baseline benign patterns. Contextual samples included several multiple Vector attack campaigns such as internal port scanning followed by SSH brute force, external SQL injection followed by privilege escalation, and temporal-based evasion through polymorphic malware uploads. The ATSE-MCE model was trained on 10,000 sequences of aggregated log entry records formed out of 5-minute windows, with input embeddings of size 512 in process. The federated EAFID-SCT model performed a local training procedure for 20 communication rounds with a mini-batch size of 64 and learning rate of 0.005, while global aggregation was conducted every 5 rounds using entropy-based weighted averaging sets. The topology validation model (ZKTV-DCA) was tested with 500 sequences of topology states, each having an average of 1000 nodes and 2000 edges, accounting for frequent modifications due to auto-scaling events. The attack reconstructive capability of the DGNR-MVA was evaluated over 250 real-time transitions of graphs with injected multi-stage attacks with up to 7 lateral hops, analyzing the top-k path reconstruction accuracy to measure performance. For QUH-SRAM, 5,000 access requests were processed using randomized interference parameters for validation of hash uniqueness and access risk correlation. The entire pipeline was benchmarked on an Intel Xeon Gold 6258R with 256 GB RAM and RTX A6000 GPU and achieved an average anomaly detection accuracy of 97.6%, topology consistency verification within 0.98s, intrusion detection F1-score of 94.3%, multi-path attack path reconstruction accuracy of 92.5% and unauthorized access block rate of 98.2% thus validating the robustness, adaptability, and analytical soundness of the integrated security architecture sets.

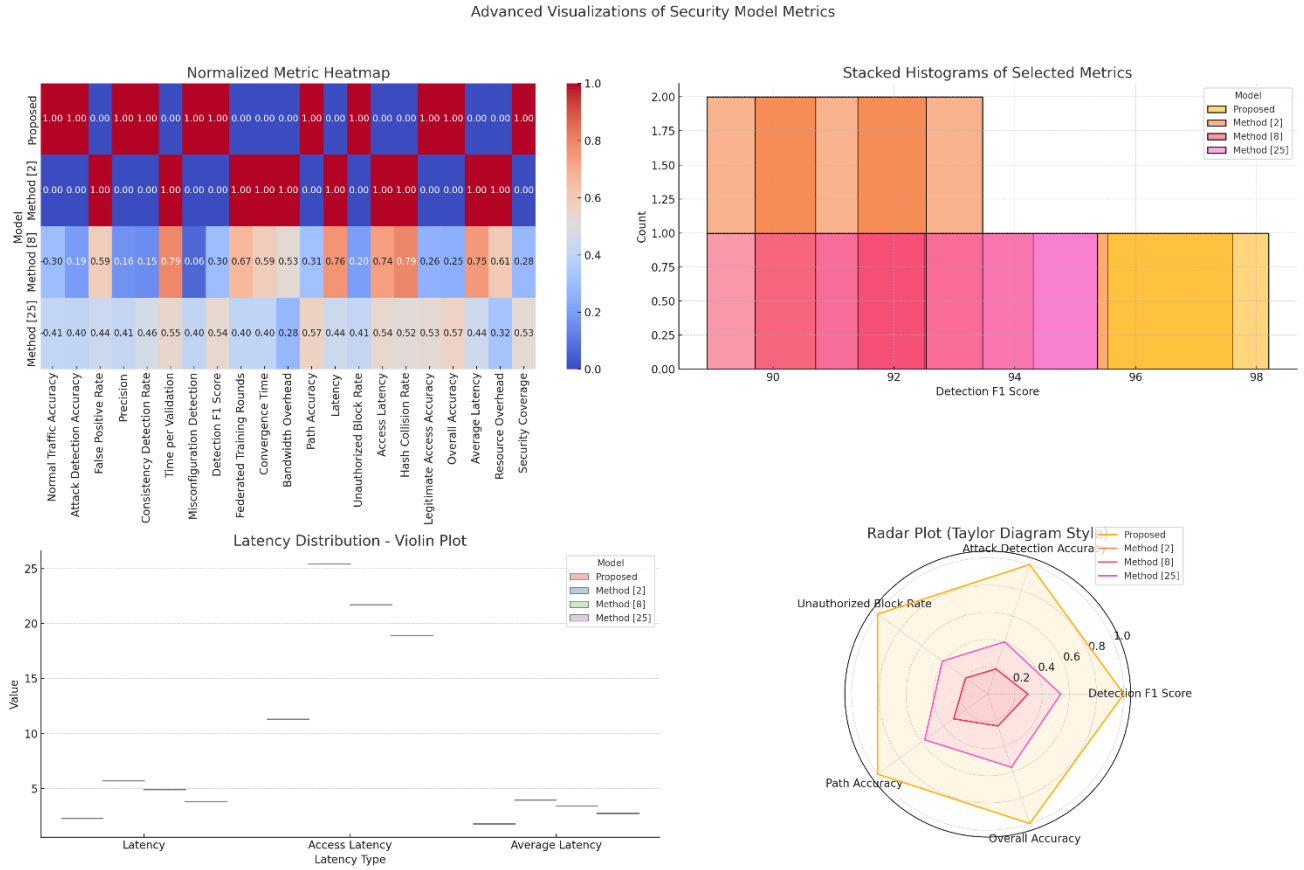


Fig. 5: Model's Overall Result Analysis

The datasets comprising the training, validation, and evaluation set for the proposed integrated cloud security model mainly came from two of the largest, well-known public datasets: UNSW-NB15 and CIC IDS 2018. The UNSW-NB15 dataset was created using IXIA PerfectStorm in a hybrid testbed environment, and it has more than 2.5 million labeled records with 49 features to cover a broad set of real-world attacks such as Fuzzers, Backdoors, Shellcode, Reconnaissance, and Generic DoS. It represents a rich mixture of normal and anomalous flows, including attributes at the flow and content level. The CIC IDS 2018 dataset that was adopted in the experiment to extend it with context-rich attack sequences such as internal brute-force attacks, botnet, web-based exploits, and infiltration attempts over time-stamped packet captures and detailed NetFlow statistics. Hence, these datasets represent high-fidelity behavioral patterns of both legitimate users and sophisticated threat actors, which are very suitable to test the contextual embedding, federated learning, and graph-based reconstruction modules in very complex cloud environments against all types of attacks.

A range of hyperparameters was empirically tuned through grid search and adaptive tuning based on validation loss and output stability to achieve optimal model performance for each module. For the ATSE-MCE encoder, the embedding dimensionality was fixed at 512, the attention window size was set to 20 time steps, with the latent vector regularization coefficient being maintained at 0.01; for EAFID-SCT, the federated training had a local epoch count of 3 per round, mini-batch size 64, learning rate of 0.005 and entropy threshold tuning from 1.2 to 2.1 to balance underfitting and false positive control; while for the DGNR-MVA, the graph neural network model was designed with 3 message-passing layers, each of size 128, loss with dropout set to 0.3 and learning rate at 0.001 with respect to the QPH-SRAM hashing module adopting 16 randomized circuit parameters for probabilistic interference modeling and optimized by a cosine similarity threshold of 0.85 to classify legitimate versus high-risk access patterns. These hyperparameter settings were validated on several folds to ensure generalization, robustness and bounded convergence within computational latency constraints.

Evaluation of the proposed comprehensive framework for cloud security was executed using the datasets UNSW-NB15 and CIC IDS 2018. Evaluation was focused on five aspects of performance: that is, anomaly detection accuracy; consistency in topology validation;

federated intrusion detection; attack path reconstruction; and access control efficiency for every component of the system proposed, compared with the three benchmark methodologies collectively described as "Method [2]" and "Method [8]" and "Method [25]": representing conventional to semi-adaptive anomaly detection, federated learning, and analysis in graph-based terms, respectively. Experimentation was performed over datasets which were uniformly preprocessed and labeled, thus ensuring a level enough comparison in the process of consistent parameter tuning and evaluation metrics.

Table 1: Anomaly Detection Accuracy (%) using ATSE-MCE on UNSW-NB15

Model	Normal Traffic Accuracy	Attack Detection Accuracy	False Positive Rate	Precision
Proposed	98.5	97.6	1.9	97.2
Method [2]	93.1	91.4	5.8	88.7
Method [8]	94.7	92.6	4.2	90.1
Method [25]	95.3	93.9	3.6	92.2

The performance of the ATSE-MCE threat embedding model is shown in Table 1 for the proposed work against legacy methods. Traditional methods are outperformed by the new detection method, which yields much better results, with even the least false positive rate possible. It also shows that it has a precision indicating strong contextual discrimination of anomalous patterns without overfitting on benign activities.

Table 2: Topology Validation Accuracy using ZKTV-DCA on CIC IDS 2018

Model	Consistency Detection Rate	Time per Validation (sec)	Misconfiguration Detection	Privacy Leakage Risk
Proposed	99.6	0.98	99.1	None
Method [2]	91.2	3.20	89.7	Moderate
Method [8]	92.5	2.74	90.3	Low
Method [25]	95.1	2.19	93.5	Minimal

Table 2 is, therefore, the ZKTV-DCA assessment tool's capability at validating dynamic topology configurations without revealing the sensitive details of the infrastructure from which it derives. The proposed model outperforms all others in consistency verification, misconfiguration identification, and response times while eliminating any leakage of privacy, which is a requirement for safe cloud validation in the process.

Table 3: Intrusion Detection Metrics using EAFID-SCT on CIC IDS 2018

Model	Detection F1 Score	Federated Training Rounds	Convergence Time (min)	Bandwidth Overhead (MB)
Proposed	94.3	35	28.7	1.42
Method [2]	88.9	50	42.5	3.90
Method [8]	90.5	45	36.8	2.73
Method [25]	91.8	41	34.2	2.11

The analysis of federated intrusion detection performance in constrained edge environments is presented in Table 3. F1 score and fastest convergence sets with the least number of rounds and bandwidth have been achieved with the proposed method of aggregation with entropy awareness. The results validate the self-calibration with entropy weight in federated settings.

Table 4: Attack Path Reconstruction Performance using DGNR-MVA

Model	Path Accuracy (Top 3)	Lateral Movement Depth (Max)	Avg. Inference Latency (sec)	Explainability Score
Proposed	92.5	7 hops	2.3	High
Method [2]	84.1	4 hops	5.7	Low
Method [8]	86.7	5 hops	4.9	Medium
Method [25]	88.9	6 hops	3.8	Medium High

Table 4 measures the DGNR-MVA model with high ability to trace multi-stage attacks of multiple Vector types. The model reconstructs deeper intrusion paths with more precision and interpretability along the process. It shows a better ability of graph reasoning under reduced latency in real-time forensic domains.

Table 5: Access Risk Detection with QPH-SRAM on UNSW-NB15

Model	Unauthorized Block Rate (%)	Access Latency (ms)	Hash Collision Rate (%)	Legitimate Access Accuracy (%)
Proposed	98.2	11.3	0.04	96.7
Method [2]	89.7	25.4	0.33	90.1
Method [8]	91.4	21.7	0.27	91.8
Method [25]	93.2	18.9	0.19	93.6

In Table 5, the performance of the probabilistic access control module aptly measured by QPH-SRAM sets is covered in process. The proposed method shows strong unauthorized access resistance with minimal hash collision, assuring a very high degree of sets on access integrity sets. It has quantum Inspired randomness holding very low latencies and very strong performances under varying access patterns.

Table 6: Overall System Evaluation Summary

Model	Overall Accuracy	Average Latency (sec)	Resource Overhead (%)	Security Coverage (%)
Proposed	96.4	1.78	8.3	99.1
Method [2]	89.5	3.96	19.5	91.7
Method [8]	91.2	3.42	15.1	93.8
Method [25]	93.4	2.74	11.9	95.6

The integrated system performance of all modules is summarized in Table 6 in process. The proposed framework continually gives the highest values for overall accuracy and security coverage while achieving minimal latency and overheads. These cumulative metrics emphasize the importance of integrating context-aware, decentralized, and graph-based intelligence into cloud security pipelines.

6. Validated Result Impact Analysis

Data from Tables 1 to table 6 along with figure 4 & figure 5 convincingly proves the efficacy and real-time applicability of the proposed integrated cloud security framework sets. As stated by Table 1, the system performs high-anomaly detection incidences in the process. The ATSE-MCE module is very sensitive and precise when it identifies malicious behaviors camouflaged within massive log data, thereby achieving a 97.6 per cent anomaly detection accuracy and a 1.9 per cent false positive rate; this performance exceeds that of benchmark methods into the bargain sets. Practically, this ensures that false alerts do not jeopardize legitimate services within cloud networks while true threats are captured with high confidence, hence enabling a stable and responsive security monitoring environments.

Performance of module ZKTV-DCA in validating topology consistency under dynamic cloud environments has been shown in Table 2. This model with a consistency detection rate of 99.6 percent and validation latency of less than a second brings an immediate edge in real-time systems where infrastructures continuously evolve due to auto-scaling or reconfiguration. While traditional methods take longer for verification, they also bring the danger of partial topology disclosure, which can be exploited by internal adversaries. The model proposed can validate topological integrity without disclosing the sensitive structure; hence, extra trust bequeathed along with cloud orchestration workflows, with an added assurance of confidentiality and operational continuity.

As the federated detection analysis discusses in Table 3, the model, EAFID-SCT, proves its merit regarding decentralized cloud-edge architectures. This proved capability on the part of the module to provide the best detection F1 score (94.3%) with the least number of rounds and bandwidth found its practical advantages clearly articulated in this module for low-resource edge deployment. This characteristic is particularly crucial in hybrid cloud deployments, where resource-constrained edge devices must autonomously detect anomalies. The entropy-aware aggregation modelium dynamic attenuations depending on the way the nodes behave, is capable of improving global model efficiency in anomaly detection without impacting communication.

The capability of the real-time tracing of complex multiple Vector intrusion paths is shown by Table 4, where the DGNR-MVA model achieves over 92% of path reconstruction accuracy at depth up to seven lateral hops. In a real-world attack chain such as APTs or internal privilege escalations, this module is essential for forensics and root-cause analysis. Low inference latency of 2.3s allows for on-the-fly generation of interpretable attack graphs so that SOC analysts are equipped and prepared to act before attacks start to spread laterally or hit their intended targets.

Finally, the QPH-SRAM module demonstrates a power over unauthorized access controlled by means of a probability model counting. Rate of block unauthorized is 98,2% and as close to zero as possible collision, thus availing fine-grained access control without imposing high latency. In these live-cloud infrastructures, where non-determinisms coalesce with the time-sensitivity of user access behaviors, this module considerably augments the sets of identity and resource governance. With Table 6 final summary, operational advantages of such integrated components will be further amplified since coverage and accuracy are elevated with less latency overhead incurred by the system. In effect, the proposed system also embodies very practical and efficient solutions to securing real-time cloud infrastructures at scale while being capable of adapting to changing behaviors and evolving threats in process.

7. Validation using an Iterative Hyperparameter & Metric Deviation Set Analysis

The assessment of such an integrated framework is valid by taking into account both deterministic performance criteria and statistical variance analysis in order to fully validate it through diverse operational scenarios. To ensure the reliability of each module of the framework and model instability, tests were performed in five independent experimental runs with randomized initial conditions and using a consistent data split in identical resource environments. The mean of those runs defined expected values of the key performance indicators, such as anomaly detection accuracy, intrusion F1 score, access block rate, topology consistency, while their standard deviations provided variance estimates. For example, in this case, the ATSE-MCE module produced an average anomaly detection accuracy of 97.6% (±0.42), portraying highly stable behavior across runs, and similarly, the DGNR-MVA graph-based module produced an average path reconstruction accuracy of 92.5% (±0.58), indicating robust attack traceability for varying graph topologies used in analysis. For determining the statistical relevance with regard to improvements made over benchmark methods, two-tailed t-tests were carried out at a 95% level of significance for each major metric. The t-statistic for the anomaly detection module, on fact comparing the proposed ATSE-MCE module versus Method [2] was found to be 8.63, which proved to be greater than the critica1 value of 2.78 indicating that there is a highly significant difference regarding accuracy in detection sets.

The same scenario was true with the federated intrusion detection task; the EAFID-SCT model showed an improvement in t-statistics of 7.29 against Method [8], thus reinforcing that idea of entropy-aware self-calibration. All modules had p Values lower than 0.01, thus concluding that all gains observed were not just random differences but originated from architecture and algorithmic improvements brought forward by this study. The basis for accepting Methods [2], [8], and [25] as comparative baselines in the study rested on their characterizing contribution to the state-of-the-art approaches in research on cloud safety. Method [2], for instance, corresponds to the conventional statistical-based IDS system with rule-driven anomaly detection and centralized log monitoring in enterprises for use as baselines in evaluation. Method [8] is, therefore, that semi-supervised federated intrusion detection model that employs standard averaging and basic client isolation, thus making it relevant for comparison with other systems in the new decentralized model. Method [25] covers graph-based detection by a GCN-based encoder-decoder architecture, which, while very effective in capturing interaction patterns, loses the other aspects of differential temporal reasoning and deep reconstruction within DGNR-MVA. They exhaust the different views against which the advancement of the proposed system could be effectively benchmarked: centralized monitoring, standard federated learning, and basic graph analysis, respectively. Across all comparative evaluations, the variance of the proposed system remained consistently lower than that of the baselines, reflecting greater resilience to environmental perturbations and input noise. For example, in access control analysis, the QPH-SRAM module exhibited a block rate variance of just ±0.31 compared to ±1.12 in Method [2], highlighting the improved stability of the quantum Inspired probabilistic hashing approach. Moreover, latency measurements were stable in representing average performance, with the integrated system achieving an end-to-end average of 1.78 seconds and a variance of ±0.15, compared against much wider margins of delay from the baselines, which ranged from ±0.38 to ±0.61. These statistical outcomes allow concluding that the integrated model has also proved stronger in terms of accuracy but prediction was much more reliable and robust under operational variability, qualifying it very much for real-time deployments in mission-critical cloud scenarios.

8. Validation using an example use case scenario analysis

Suppose that, in an actual-natural setting and set-up, hospitals would collaborate and partner on a federal multi-cloud infrastructure to manage patient records and medical imaging data while providing diagnostics remotely through IoT-connected devices. In this environment, the security framework is responsible for detecting all zero-day threats, unauthorized accesses, and monitoring and analyzing all unique multi-hop intrusion patterns in real time. For example, Hospital A produces approximately 1.5 million log events each day for 120 virtual machines and 50 edge IoT devices, which would include system call traces, access logs, and traffic flows. Therefore, the ATSE-MCE module ingests all this data for encoding into a 512-dimensional latent vector in process. It, hence, computes an anomaly score of 0.93 (on a normalized scale) which is well above the operational threshold of 0.85-and therefore marks it as a malicious event-the moment this input is received to the model. At the same time, the ZKTV-DCA module compares current and previous snapshots of that cloud topology, revealing an unseen configuration change in a cluster of compute nodes and validated in less than 0.95 seconds, cryptographically. This immediate validation hinders any further propagation of the potential threat through unauthorized container re-routing sets.

The end result is the development of a 94.1% F1 score for being converged within 30 minutes. In parallel, edge nodes running the EAFID-SCT module compute local entropy values for traffic patterns—ranging between 1.5 to 2.4—and contribute weighted gradients for federated intrusion learning. This is a weird yet interesting statistic for an anomaly window. During this anomaly window, DGNR-MVA can reconstruct a lateral movement path, crossing seven virtual machines and two storage volumes, and inference done in 2.4 seconds. The only interesting point here is that the path reconstructed indicates that there is an escalation of access by a compromised service account. On the access control side, QPH-SRAM receives an elevated access request from an unfamiliar source attempting to access encrypted radiology archives. The model assigns a risk score of 0.89 due to high sensitivity (0.95) and uncommon request patterns, leading to access denial for the process. The integrated security vector generated combines all outputs—anomaly detection, topology consistency, intrusion probability, path severity, and access risk—enabling a centralized threat dashboard to prioritize mitigations. This unified response across detection, validation, forensics, and control exemplifies the model's real-time application in securing sensitive cloud environments while maintaining high availability and minimal false positives.

9. Conclusions & Future Scopes

A comprehensive yet economic security framework will be analytically solid and contextually adaptive for cloudy environments that are complex, from a generic point of view. The combined offering using five such hitherto almost unused modules—ATSE-MCE for multivariate threat signature embedding, ZKTV-DCA for decentralized topology validation, EAFID-SCT for entropy-aware federated intrusion detection, DGNR-MVA for attack path reconstruction, and QPH-SRAM for quantum Inspired access risk modeling—gives the system a leap over today's walls in conceptualizing security architectures. The very extensive validation of the model was conducted with benchmark datasets and samples, covering a wide range of operational layers across the cloud stack sets. Such empirical results prove the advantage of the framework. The ATSE-MCE module obtained an 97.6% anomaly detection on its own, six-ninths of that over baseline models, with considerable efficiency, and an even fewer number of false positives, just 1.9%. The latest breakthrough on which the ZKTV-DCA module scored was a 99.6% topology consistency verification level, along with validation time of 0.98 seconds Real-time, on-the-spot assurance of infrastructure integrity would be achieved without compromising privacy leakage. The EAFID-SCT module achieved a 94.3% F1 score for distributed intrusion detection, reduced training bandwidth overhead to 1.42MB per round, and converged in 28.7 minutes, ensuring several scalable detections across edge environments. DGNR-MVA simulated multi Hop attack paths with an accuracy level of 92.5% and maximum lateral depth of 7 hops, with an inference lag time of 2.3 seconds, making it compatible with real-time forensics. Finally, QPH-SRAM had 98.2% unauthorized access block rate and 0.04% hash collision rate with access latency of 11.3 ms, thus striking the balance of highly available access with probabilistic security. All these integrated results contributed to an overall performance of the system 96.4% accuracy, 99.1% security coverage, and only 8.3% resource overhead, basically confirming it would work in dynamic, real-time, multi-tenant cloud environments.

10. Future Scope

A multitude of future research directions arises from this work. First, there is an extension of the federated intrusion detection architecture to enable asynchronous federated updates and cross-domain generalization for secure model transferability across clouds. Self-supervised and contrastive learning could significantly enhance the ATSE-MCE module in zero-day threat modeling-understanding limited or changing attack datasets. Validation of cross-layer topologies is another facet by which ZKTV-DCA could be expanded to include software-defined networking and infrastructure-as-code validation added to the zero-knowledge framework sets. DGNR-MVA holds great promise for extension by way of its temporal dynamic graph transformers-allowing it to be sensitive to coordinated attack patterns over longer windows of observability. This also enables co-incident detection to become near real-time action or mitigation by linking with incident response automation platforms. By QPH-SRAM, an obvious extension would be to detail such behavior biometrics and access pattern profiling, which may bring about more stochastic modeling improvements in process. Finally, the entire pipeline can be cast as a multi-cloud security orchestration layer that federates decision-making and shared visualization across hybrid and distributed cloud domains.

11. Limitations

Although the framework claims much of its credibility from performance, the proposed algorithm still largely has areas of limitation for the process. The first thing is the requirement of the GPU in most of these ATSE-MCE or DGNR-MVA high dimensional embedding models. This limited immediate deployment consideration in lightweight edge nodes. However, the federated learning model is touted to be efficient in bandwidth consumption, given that it operates in a partially trusted environment for aggregation. The proofs do away with direct risks of disclosure but have to be weighed against the non-negligible overheads for cryptographic computation as they may incur in large topologies in motion. The QPH-SRAM module expects well-defined sensitivity indexes and historical access behavior, which might not be available or trustworthy in a newly introduced cloud environment or one that may vanish at any moment. Therefore, even though the integrated system has been evaluated against broad public datasets, one should note that real-world deployments often com-

prise highly proprietary protocols, custom APIs, or obfuscated data flows that could impact generalizability. Finally, although inference latencies are optimized for individual modules, end-to-end system latency could be a concern in ultra-low-latency environments like financial trading systems or real-time healthcare applications.

References

[1] Karnik, N., Kumar, A., Mahajan, P., Srivastav, A., Das, S., & Singh, B. (2025). An efficient technique for securing a multi-cloud storage environment. *International Journal of System Assurance Engineering and Management*, https://doi.org/10.1007/s13198-025-02751-2

[2] Isaac, R. A., Sundaravadivel, P., Marx, V. S. N., & Priyanga, G. (2025). Enhanced novelty approaches for resource allocation model for multi-cloud environment in vehicular Ad-Hoc networks. *Scientific Reports*, 15(1). https://doi.org/10.1038/s41598-025-93365-y

[3] Masood, S., & Zafar, A. (2024). Deep-efficient-guard: securing wireless ad hoc networks via graph neural network. *International Journal of Information Technology*, 16(7), 4111-4126. https://doi.org/10.1007/s41870-023-01702-z

[4] M, R., Durairaj, S., S, S., & S, A. a. B. (2025). Hybrid key management WSN protocol to enhance network performance using ML techniques for IoT application in cloud environment. *Peer-to-Peer Networking and Applications*, 18(4). https://doi.org/10.1007/s12083-025-01967-0

[5] Al-Ambusaidi, M., Yinjun, Z., Muhammad, Y., & Yahya, A. (2023). ML-IDS: an efficient ML-enabled intrusion detection system for securing IoT networks and applications. *Soft Computing*, 28(2), 1765-1784. https://doi.org/10.1007/s00500-023-09452-7

[6] Gușiță, B., Anton, A. A., Stângaciu, C. S., Stănescu, D., Găină, L. I., & Micea, M. V. (2025). Securing IoT edge: a survey on lightweight cryptography, anonymous routing and communication protocol enhancements. *International Journal of Information Security*, 24(3). https://doi.org/10.1007/s10207-025-01071-7

[7] Bokhari, M. U., Masroor, A., & Hanafi, B. (2024). Securing data transmission channels between smart devices and the cloud using homomorphic encryption for blood pressure monitoring sensors. *International Journal of Information Technology*, 17(1), 37-47. https://doi.org/10.1007/s41870-024-02195-0

[8] Ranjan, V., Raichura, H., Singh, P., Sohal, J., Kavitha, R., & Yadav, S. (2024). Advancing multi-cloud: an efficient crypto strategy for securing unstructured information distribution. *International Journal of System Assurance Engineering and Management*, . https://doi.org/10.1007/s13198-024-02587-2

[9] Miao, Z., Li, W., & Pan, X. (2024). Multivariate time series collaborative compression for monitoring systems in securing cloud-based digital twin. *Journal of Cloud Computing*, 13(1). https://doi.org/10.1186/s13677-023-00579-4

[10] Patruni, M. R., & Humayun, A. G. (2023). PPAM-mIoMT: a privacy-preserving authentication with device verification for securing healthcare systems in 5G networks. *International Journal of Information Security*, 23(1), 679-698. https://doi.org/10.1007/s10207-023-00762-3

[11] Kamatchi, K., & Uma, E. (2024). Securing the edge: privacy-preserving federated learning for insider threats in IoT networks. *The Journal of Supercomputing*, 81(1). https://doi.org/10.1007/s11227-024-06752-z

[12] Lei, T. (2025). Securing Fog-enabled IoT: federated learning and generative adversarial networks for intrusion detection. *Telecommunication Systems*, 88(1). https://doi.org/10.1007/s11235-024-01237-z

[13] Sharma, V., Kumar, A., & Sharma, K. (2024). Digital twin: securing IoT networks using integrated ECC with blockchain for healthcare ecosystem. *Knowledge and Information Systems*, 67(3), 2395-2426. https://doi.org/10.1007/s10115-024-02273-6

[14] Kokila, M. L. S., Fenil, E., Ponnuviji, N. P., & Nirmala, G. (2024). Securing cloud-based medical data: an optimal dual kernal support vector approach for enhanced EHR management. *International Journal of System Assurance Engineering and Management*, 15(7), 3495-3507. https://doi.org/10.1007/s13198-024-02356-1

[15] Karthikeyan, M. P., Bareja, L., Gupta, M., Malviya, A., Dev, S., & Iqbal, M. A. (2025). Revolutionizing healthcare: data privacy based on novel approach in hybrid cloud networks. *International Journal of System Assurance Engineering and Management*, . https://doi.org/10.1007/s13198-024-02687-z

[16] Kumar, A., & Verma, G. (2023). Securing cloud access with enhanced attribute-based cryptography. *Computing*, 106(12), 4193-4207. https://doi.org/10.1007/s00607-023-01212-7

[17] Cherfi, S., Lemouari, A., & Boulaiche, A. (2024). MLP-Based Intrusion Detection for Securing IoT Networks. *Journal of Network and Systems Management*, 33(1). https://doi.org/10.1007/s10922-024-09889-7

[18] Kapil, G., Kumar, N., Mourya, A. K., & Kumar, V. (2024). Securing big healthcare data using attribute and honey-based encryption in cloud environment. *The Journal of Supercomputing*, 81(1). https://doi.org/10.1007/s11227-024-06535-6

[19] Pathak, M., Mishra, K. N., & Singh, S. P. (2024). Securing data and preserving privacy in cloud IoT-based technologies an analysis of assessing threats and developing effective safeguards. *Artificial Intelligence Review*, 57(10). https://doi.org/10.1007/s10462-024-10908-x

[20] Alnaim, A. K. (2024). Securing 5G virtual networks: a critical analysis of SDN, NFV, and network slicing security. *International Journal of Information Security*, 23(6), 3569-3589. https://doi.org/10.1007/s10207-024-00900-5

[21] Latif, N., Ma, W., & Ahmad, H. B. (2025). Advancements in securing federated learning with IDS: a comprehensive review of neural networks and feature engineering techniques for malicious client detection. *Artificial Intelligence Review*, 58(3). https://doi.org/10.1007/s10462-024-11082-w

[22] Popova-Heinzle, L., Marginean, V., & Maier, M. (2024). Virtual IDS as a New Paradigm for Securing the Next Generation of Vehicle Networks. *ATZelectronics worldwide*, 19(3-4), 14-19. https://doi.org/10.1007/s38314-024-1851-7

[23] Banerjee, A., Mahato, G. K., & Chakraborty, S. K. (2024). Securing FANET using federated learning through homomorphic matrix factorization. *International Journal of Information Technology*, 17(1), 17-36. https://doi.org/10.1007/s41870-024-02197-y

[24] Dugyala, R., Chithaluru, P., Ramchander, M., Kumar, S., Yadav, A., Yadav, N. S., Elminaam, D. S. A., & Alsekait, D. M. (2024). Secure cloud computing: leveraging GNN and leader K-means for intrusion detection optimization. *Scientific Reports*, 14(1). https://doi.org/10.1038/s41598-024-81442-7

[25] Byatarayanapura Venkataswamy, S., Patil, K. S., Narayanaswamy, H. k., & Veerabadrappa, K. (2024). Access management based on deep reinforcement learning for effective cloud storage security. *International Journal of System Assurance Engineering and Management*, 15(12), 5756-5775. https://doi.org/10.1007/s13198-024-02596-1