

A GenAI-Driven Adaptive Cybersecurity Mesh for Real-Time Threat Detection in Intelligent Communication Systems

Utham Kumar Anugula Sethupathy^{1*}, Vijayanand Ananthanarayanan², Diwakar Sreeramulu³

¹ Independent Researcher, Alumni, Nanyang Technological University, Atlanta, GA, USA

² Independent Researcher, Alumni, Fairleigh Dickinson University, Atlanta, GA, USA

³Independent Researcher, Alumni, Visvesvaraya Technological University, Austin, TX, USA

mailuthamkumar@gmail.com¹, vijayanand31@gmail.com², diwakar1@gmail.com³

*Corresponding author: Utham Kumar Anugula Sethupathy (mailuthamkumar@gmail.com)

ABSTRACT

Intelligent communication systems integrating Internet of Things (IoT), cyber-physical infrastructures, edge computing, and next-generation communication protocols have significantly expanded the attack surface of modern digital ecosystems. Traditional intrusion detection systems (IDS) remain predominantly centralised, signature-based, and insufficiently adaptive to evolving multi-vector and zero-day threats. This paper proposes a GenAI-driven adaptive cybersecurity mesh architecture designed for real-time threat detection in distributed intelligent communication environments. The proposed framework integrates zero-trust security principles with a distributed mesh of edge security nodes coordinated through a policy orchestration layer. A generative AI-based adaptive threat modelling engine continuously synthesises contextual attack patterns and enhances anomaly detection across network, application, and behavioural layers. A formal cross-layer risk-scoring model fuses heterogeneous security signals to generate dynamic threat-confidence indices. The system is evaluated in a simulated intelligent communication environment comprising heterogeneous nodes, mixed legitimate traffic, and multiple attack scenarios, including DDoS, man-in-the-middle, and protocol-exploitation attacks. Experimental results demonstrate improved detection accuracy, reduced false positives, and lower response latency compared to baseline signature-based and centralised ML-based IDS models. The proposed architecture offers a scalable and adaptive security paradigm suitable for next-generation intelligent communication infrastructures.

Keywords: Cybersecurity Mesh; Generative AI; Intelligent Communication Systems; Zero-Trust Architecture; Intrusion Detection; Cross-Layer Risk Scoring; Distributed Security.

1 Introduction

Intelligent communication systems have evolved rapidly with the convergence of Internet of Things (IoT), cyber-physical systems (CPS), edge computing, 5G/6G networks, and software-defined networking paradigms. These systems underpin critical infrastructures, including smart transportation, industrial automation, intelligent healthcare, and distributed energy networks. While such architecture enhances scalability and real-time responsiveness, they simultaneously introduce heterogeneous attack surfaces, decentralised trust boundaries, and dynamic communication patterns that traditional security models struggle to protect effectively.

Conventional intrusion detection systems (IDS) primarily rely on signature-based detection or centralised machine learning models. Signature-based systems lack resilience against zero-day exploits and polymorphic attacks. Centralised machine learning approaches, although more adaptive, introduce latency bottlenecks and single points of failure in distributed communication environments. Furthermore, most existing systems analyse network-layer anomalies in isolation, failing to incorporate cross-layer contextual signals, such as application behaviour, node-level telemetry, and protocol integrity violations. The increasing sophistication of adversaries necessitates a transition toward adaptive, distributed, and intelligence-augmented security architectures. Generative Artificial Intelligence (GenAI) has recently

demonstrated capabilities in pattern synthesis, contextual reasoning, and dynamic model adaptation. However, its integration into distributed cybersecurity infrastructures for real-time communication systems remains insufficiently explored. Existing research often treats AI-based intrusion detection as a standalone classifier rather than embedding adaptive intelligence within a mesh-based security topology.

To address these limitations, this paper proposes a **GenAI-driven adaptive cybersecurity mesh architecture** tailored for intelligent communication systems. The proposed framework incorporates three primary design principles:

1. **Zero-Trust Distributed Security Enforcement** – Security controls are enforced at edge nodes rather than relying solely on centralised monitoring.
2. **Adaptive Generative Threat Modelling** – A GenAI engine dynamically synthesises threat hypotheses and refines anomaly detection models.
3. **Cross-Layer Risk Fusion** – Security signals from network, application, and behavioural layers are fused into a unified threat confidence score.

The contributions of this paper are summarised as follows:

- A scalable cybersecurity mesh architecture integrating distributed edge security nodes with centralised orchestration.
- A generative AI-assisted adaptive threat modelling engine for evolving attack detection.
- A formalised cross-layer risk scoring mechanism for dynamic threat confidence estimation.
- An empirical evaluation demonstrating improved detection performance and reduced latency relative to conventional IDS baselines.

The remainder of the paper is organised as follows. Section 2 reviews related work in distributed intrusion detection and AI-assisted security. Section 3 defines the threat model and adversarial assumptions. Section 4 describes the proposed cybersecurity mesh architecture. Section 5 presents the GenAI-based adaptive methodology and risk scoring formulation. Section 6 details the experimental setup, followed by results and analysis in Section 7. Section 8 discusses implications and limitations, and Section 9 concludes the paper.

2 Related Work

Security in intelligent communication systems has evolved significantly over the past decade, driven by the proliferation of IoT devices, edge computing frameworks, and next-generation communication protocols. Existing research can be broadly categorised into four domains: (1) signature-based intrusion detection systems, (2) machine learning-based IDS models, (3) distributed and mesh-oriented security architectures, and (4) AI-assisted adaptive threat modelling.

2.1 Signature-Based and Centralised Intrusion Detection

Traditional intrusion detection systems such as Snort and Suricata rely on rule-based or signature-based detection mechanisms. While effective against known attack patterns, these systems exhibit limited capability in identifying zero-day exploits and polymorphic threats. Moreover, a centralised deployment architecture introduces scalability and latency challenges in intelligent communication systems where nodes are geographically distributed and operate under real-time constraints.

Centralised IDS models also introduce a single point of failure and increase network overhead through continuous telemetry aggregation. In high-throughput environments such as IoT-enabled smart grids or vehicular communication systems, this architectural limitation significantly affects detection latency and responsiveness.

Limitation Identified: Lack of adaptability and insufficient resilience against evolving, multi-stage attacks.

2.2 Machine Learning-Based Intrusion Detection

Recent advancements have incorporated supervised and unsupervised machine learning techniques for anomaly detection in communication networks. Approaches leveraging Support Vector Machines (SVM), Random Forests, k-Nearest Neighbours, and Deep Neural Networks have demonstrated improved detection accuracy compared to signature-based systems.

Deep learning models, including Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTMs), have been applied to traffic pattern recognition and time-series anomaly detection. These models improve classification performance; however, they are typically trained offline and lack dynamic threat adaptation.

Furthermore, most ML-based IDS implementations operate in centralised environments. Distributed deployment remains limited due to computational overhead and synchronisation challenges across heterogeneous nodes.

Limitation Identified: Static model training, insufficient cross-layer contextual fusion, and limited distributed enforcement.

2.3 Distributed and Mesh-Based Security Architectures

To overcome centralisation constraints, recent research has explored distributed intrusion detection frameworks and cybersecurity mesh architectures. Cybersecurity mesh models emphasise decentralised policy enforcement, microsegmentation, and zero-trust principles. Security capabilities are deployed closer to assets, often at edge gateways or node clusters.

Such architecture enhances resilience and reduces response latency. However, many implementations still rely on static rule engines or traditional ML classifiers, lacking adaptive threat-generation mechanisms. Additionally, risk scoring is often performed independently at the node level without systematic cross-layer correlation.

Limitation Identified: Absence of adaptive intelligence capable of synthesising new attack hypotheses and limited risk fusion mechanisms.

2.4 Generative AI in Cybersecurity

Generative Artificial Intelligence (GenAI) has recently been explored in cybersecurity contexts for synthetic attack generation, adversarial training, automated threat intelligence summarisation, and log analysis. Large language models and generative adversarial networks (GANs) have demonstrated potential in simulating evolving attack patterns and improving detection robustness through adversarial learning.

Despite promising advancements, the integration of GenAI into real-time distributed communication security remains limited. Existing work often focuses on either:

- Offline threat simulation
- Isolated anomaly detection enhancement
- Security operations automation

There remains a gap in embedding GenAI as an adaptive reasoning component within a distributed cybersecurity mesh capable of continuous cross-layer threat modelling.

2.5 Research Gap

From the analysis above, the following research gaps are identified:

1. Lack of distributed adaptive architecture combining zero-trust enforcement with real-time AI reasoning.
2. Insufficient integration of generative models for dynamic threat hypothesis synthesis.
3. Limited cross-layer fusion mechanisms incorporating network, application, and behavioural telemetry.
4. Inadequate empirical evaluation of GenAI-enhanced mesh architectures under multi-vector attack scenarios.

This paper addresses these gaps by proposing a GenAI-driven adaptive cybersecurity mesh that integrates distributed enforcement, generative threat modelling, and formalised cross-layer risk scoring within intelligent communication systems.

3 Threat Model and Assumptions

Intelligent communication systems consist of heterogeneous devices, edge gateways, communication protocols, and cloud-coordinated services. These systems operate in semi-trusted or untrusted environments where adversaries may exploit protocol weaknesses, compromised nodes, or misconfigured services. This section formalises the adversarial model and system assumptions used to evaluate the proposed GenAI-driven cybersecurity mesh.

3.1 System Model

We consider an intelligent communication environment comprising:

- N distributed nodes (IoT devices, embedded systems, edge servers)
- Edge gateways are responsible for traffic aggregation
- A centralised orchestration layer for policy coordination
- Multi-layer communication stack:
 - Network layer (packet flows, routing behaviour)
 - Transport/session layer (connection states, protocol exchanges)
 - Application layer (API calls, service requests, payload semantics)
 - Behavioural layer (node-level telemetry, CPU/memory usage, process anomalies)

Each node is equipped with a lightweight security agent capable of telemetry collection and local anomaly pre-processing. These agents communicate with a distributed mesh controller enforcing zero-trust policies.

3.2 Adversarial Capabilities

We assume a **probabilistic polynomial-time (PPT) adversary** with the following capabilities:

1. **Network Manipulation**
 - Packet injection
 - Replay attacks

- Distributed Denial of Service (DDoS)
- Man-in-the-Middle (MITM)
- 2. Protocol Exploitation**
 - Exploiting insecure handshake sequences
 - Session hijacking
 - Routing manipulation
- 3. Node Compromise**
 - Malware injection into edge devices
 - Privilege escalation
 - Lateral movement within sub-networks
- 4. Zero-Day Attack Simulation**
 - Novel attack patterns not previously observed
 - Polymorphic traffic behaviour

The adversary does not possess global cryptographic key material but may compromise a subset k of nodes where $k < N$.

3.3 Security Objectives

The proposed system aims to satisfy the following objectives:

- 1. Real-Time Threat Detection**
 - Minimise detection latency T_d
 - Maintain bounded response time T_r
- 2. High Detection Accuracy**
 - Maximise True Positive Rate (TPR)
 - Minimise False Positive Rate (FPR)
- 3. Adaptive Threat Modelling**
 - Continuous update of anomaly hypotheses
 - Dynamic risk threshold recalibration
- 4. Distributed Resilience**
 - Avoid a single point of failure
 - Maintain detection capability under partial node compromise

3.4 Threat Categories Considered

The experimental evaluation considers five primary attack categories:

- 1. Distributed Denial of Service (DDoS)**
 - High-volume traffic floods targeting gateways
- 2. Man-in-the-Middle (MITM)**
 - Traffic interception and modification

3. Protocol Exploitation

- Malformed packet injection
- Session hijacking

4. Anomalous Behavioural Drift

- Gradual deviation in device resource patterns

5. Synthetic Zero-Day Patterns

- Generated attack traffic not matching predefined signatures

3.5 Assumptions

The following assumptions constrain the system:

- Secure initial device onboarding with cryptographic identity provisioning.
- Encrypted communication channels between mesh nodes and orchestrator.
- Computational capability at edge nodes is sufficient for lightweight inference.
- Availability of baseline training dataset for initial model bootstrapping.

3.6 Risk Representation

We define a threat confidence function:

$$R_i = f(N_i, A_i, B_i, C_i)$$

Where:

- N_i = Network-layer anomaly score
- A_i = Application-layer anomaly score
- B_i = Behavioural deviation metric
- C_i = Contextual threat intelligence weight

The final risk score for the node i is computed as:

$$R_i = \alpha N_i + \beta A_i + \gamma B_i + \delta C_i$$

Subject to:

$$\alpha + \beta + \gamma + \delta = 1$$

This weighted fusion model enables cross-layer anomaly aggregation and adaptive risk recalibration through the GenAI threat modelling engine described in Section 5.

3.7 Summary

The threat model reflects realistic adversarial behaviour in distributed intelligent communication systems. The security objectives emphasise adaptive detection, cross-layer reasoning, and distributed resilience. These assumptions form the basis for the architectural design described in the next section.

4 Proposed GenAI-Driven Cybersecurity Mesh Architecture

This section presents the architecture of the proposed **GenAI-driven adaptive cybersecurity mesh** designed for intelligent communication systems. The architecture combines distributed zero-trust enforcement, adaptive generative threat reasoning, and cross-layer telemetry fusion.

4.1 Architectural Overview

The proposed system follows a **distributed mesh topology** composed of:

1. **Edge Security Nodes (ESNs)**
2. **Mesh Coordination Layer (MCL)**
3. **GenAI Adaptive Threat Engine (GATE)**
4. **Policy Orchestration and Response Layer (PORL)**

The design adheres to zero-trust principles: no device, session, or communication flow is inherently trusted, even within internal network boundaries. Figure 1 presents the overall system architecture of the proposed adaptive cybersecurity mesh. At the edge layer, heterogeneous nodes (IoT sensors, control systems, and gateways) perform network, application, and behavioural anomaly detection, followed by local risk scoring. Telemetry summaries are forwarded to the Mesh Control Layer (MCL), which coordinates distributed updates. The Generative AI Adaptive Threat Engine (GATE) synthesises adversarial threat hypotheses and recalibrates contextual risk weights. The Dynamic Risk Fusion module computes the composite risk score. $R_i = \alpha N_i + \beta A_i + \gamma B_i + \delta C_i$, triggering the Adaptive Policy Module when thresholds are exceeded as per Fig. 1.




Fig. 1. GenAI-Driven Cybersecurity Mesh Architecture

4.2 High-Level Architecture

Components:

(1) Edge Security Nodes (ESNs).

Deployed at IoT devices, gateways, and edge servers.

Functions:

- Real-time packet inspection
- Lightweight anomaly scoring
- Behavioural telemetry collection
- Local enforcement (micro-segmentation, traffic throttling)

Each ESN computes preliminary anomaly metrics:

N_i, A_i, B_i

This reduces central bandwidth overhead and detection latency.

(2) Mesh Coordination Layer (MCL).

Acts as a distributed synchronisation fabric:

- Aggregates anonymised anomaly summaries
- Synchronises threat intelligence
- Maintains node trust scores
- Ensures resilience against partial compromise

Unlike centralised IDS, MCL operates in a logically distributed fashion to prevent single points of failure.

(3) GenAI Adaptive Threat Engine (GATE).

This is the core novelty of the architecture.

Capabilities:

- Synthesises potential attack hypotheses
- Generates adversarial traffic variations
- Refines anomaly detection thresholds
- Updates contextual threat intelligence weights C_i

GATE continuously adjusts fusion weights:

$$\alpha_t, \beta_t, \gamma_t, \delta_t$$

Where t represents adaptive time intervals.

(4) Policy Orchestration and Response Layer (PORL).

Responsible for:

- Dynamic rule generation
- Automated quarantine actions
- Network isolation policies
- Adaptive rate limiting

Response decisions are triggered when:

$$R_i \geq \theta_t$$

Where θ_t Does GATE recalibrate an adaptive risk threshold?

4.3 Data Flow Pipeline

1. Telemetry captured at ESNs
2. Local anomaly scoring
3. Aggregated signals transmitted to MCL
4. GATE performs contextual reasoning
5. Updated policies propagated back to ESNs
6. Enforcement applied in near real-time

Latency target:

$$T_{total} = T_{capture} + T_{inference} + T_{propagation}$$

The architecture aims to minimise $T_{propagation}$ through decentralised updates.

4.4 Zero-Trust Enforcement Model

Each communication request undergoes:

1. Identity verification
2. Context validation
3. Behavioural deviation check
4. Dynamic risk scoring

Trust is continuously re-evaluated rather than statically assigned.

4.5 Resilience Under Node Compromise

If k Nodes are compromised:

- ESNs operate independently
- MCL redistributes trust scores
- Compromised nodes are isolated via automated micro-segmentation

System integrity is maintained provided:

$$k < \frac{N}{3}$$

(Assuming distributed consensus threshold)

4.6 Architectural Advantages

Compared to centralised IDS as per Table 1

Table 1: Compared to centralised IDS

— Feature	— Centralised IDS	— Proposed Mesh
— Adaptability	— Static models	— GenAI adaptive
— Scalability	— Limited	— Distributed
— Latency	— Higher	— Edge-based inference
— Zero-Day Detection	— Weak	— Generative threat synthesis
— Single Point of Failure	— Yes	— No

4.7 Summary

The proposed architecture integrates distributed enforcement, adaptive generative threat modelling, and cross-layer risk fusion. It is designed to operate efficiently within heterogeneous intelligent communication environments while maintaining resilience and low latency.

5 GenAI-Based Adaptive Threat Modelling and Cross-Layer Risk Scoring Methodology

This section details the methodological foundation of the proposed system, including (1) adaptive generative threat modelling, (2) cross-layer anomaly scoring, (3) dynamic risk fusion, and (4) automated policy recalibration.

5.1 Overview of Adaptive Threat Modelling

Traditional IDS models rely on static datasets and fixed decision boundaries. In contrast, the proposed framework embeds a **Generative AI-based Adaptive Threat Engine (GATE)** that continuously refines detection models using contextual telemetry and synthesised adversarial patterns.

The adaptive process operates in iterative cycles:

1. Telemetry aggregation

2. Anomaly detection
3. Threat hypothesis synthesis
4. Adversarial pattern generation
5. Model refinement
6. Policy redistribution

This cycle reduces concept drift and enhances resilience against zero-day attacks.

5.2 Generative Threat Hypothesis Synthesis

Let X_t represent observed traffic and telemetry features at a time t .

The generative engine learns an evolving distribution:

$$P_t(X)$$

Using a generative model G , new adversarial samples are synthesised:

$$\tilde{X} = G(Z, C)$$

Where:

- Z = latent noise vector
- C = contextual threat embedding
- ϕ = generative model parameters

These synthetic samples simulate polymorphic or zero-day attack variants. The anomaly detection model is retrained incrementally with both real and synthesised samples:

$$\theta_{t+1} = \theta_t - \eta \nabla L(X_t \cup \tilde{X})$$

Where:

- θ = detection model parameters
- η = learning rate
- L = loss function

5.3 Cross-Layer Feature Extraction

Each Edge Security Node extracts features across three layers:

(1) Network Layer Features N_i .

- Packet inter-arrival time
- Flow duration
- Entropy of source/destination distribution
- TCP flag anomalies

(2) Application Layer Features A_i .

- API call frequency deviation
- Request payload entropy
- Authentication failure rates
- Session irregularities

(3) Behavioural Layer Features B_i .

- CPU utilisation drift
- Memory allocation anomalies

- Process spawning irregularities
- Device energy usage deviations

Feature vectors are normalised and fed into local anomaly estimators:

$$S_i = g(N_i, A_i, B_i)$$

Where $g(\cdot)$ may represent a lightweight neural classifier deployed at the edge.

5.4 Dynamic Cross-Layer Risk Fusion

The final threat score is computed as:

$$R_i = \alpha_t N_i + \beta_t A_i + \gamma_t B_i + \delta_t C_i$$

Where:

- C_i = contextual threat intelligence from GATE
- Weights dynamically updated over time

Adaptive Weight Update Rule.

Weights are updated using performance feedback:

$$\alpha_{t+1} = \alpha_t + \lambda \frac{\partial F1}{\partial \alpha}$$

Similarly, for β, γ, δ

Subject to normalisation constraint:

$$\alpha + \beta + \gamma + \delta = 1$$

This enables the system to emphasise layers that improve detection performance in current threat landscapes.

5.5 Risk Threshold Adaptation

A static threshold increases false positives during traffic bursts. Therefore, threshold θ_t is adjusted dynamically:

$$\theta_{t+1} = \mu R_{avg} + \sigma R_{std}$$

Where:

- R_{avg} = rolling average risk
- R_{std} = rolling standard deviation
- μ, σ = tuning constants

This ensures statistical anomaly responsiveness.

5.6 Algorithmic Workflow

Algorithm 1: Adaptive Mesh Threat Detection.

Input: Telemetry stream T

Output: Risk scores R_i and policy actions

- 1: Collect multi-layer features (N_i, A_i, B_i)
- 2: Compute local anomaly scores
- 3: Transmit summaries to MCL

- 4: Generate synthetic threat patterns using GATE
- 5: Update detection model parameters
- 6: Compute fused risk score R_i
- 7: If $R_i \geq \theta_t$:
 Trigger enforcement policy
- 8: Update weights and thresholds
- 9: Repeat

Time Complexity:

- Edge inference: $O(d)$
- Generative update: $O(n \cdot k)$
- Fusion scoring: $O(1)$

Overall complexity remains scalable under distributed deployment.

5.7 Theoretical Advantages

Compared to static ML-based IDS:

- Reduces concept drift
- Enhances zero-day robustness
- Improves contextual reasoning
- Maintains low edge latency
- Enables distributed adaptation

5.8 Summary

The proposed methodology integrates generative threat synthesis, cross-layer anomaly extraction, dynamic risk fusion, and adaptive threshold recalibration into a cohesive distributed framework. This creates a continuously evolving detection ecosystem suitable for intelligent communication infrastructures.

6 Experimental Setup and Simulation Environment

This section describes the experimental design used to evaluate the proposed GenAI-driven adaptive cybersecurity mesh. The objective is to assess detection accuracy, false positive reduction, latency performance, and adaptive robustness under multi-vector attack conditions.

6.1 Experimental Objectives

The evaluation aims to measure:

1. Detection Accuracy (ACC)
2. Precision, Recall, and F1-Score
3. False Positive Rate (FPR)
4. Detection Latency
5. Adaptive Improvement Over Time
6. Scalability under increasing node count

Baseline comparisons include:

- Signature-based IDS (Rule-driven)

- Centralised ML-based IDS (Static Deep Neural Network)

6.2 Simulation Environment

A controlled intelligent communication network was simulated with the following configuration as per Table 2:

Table 2: Different Parameter and their Values

Parameter	Value
Total Nodes (N)	300
Edge Gateways	10
Communication Protocol	TCP/IP + Application-layer API traffic
Simulation Duration	24 hours (synthetic timeline)
Attack Injection Rate	15% of total traffic
Node Compromise Ratio	Up to 10%

Nodes emulate heterogeneous behaviour:

- IoT sensors
- Embedded control systems
- Edge analytics devices
- Gateway nodes

Traffic patterns include:

- Normal operational flows
- Burst traffic events
- Periodic device reporting
- Randomised noise injection

6.3 Dataset Composition

The dataset consists of:

- 120,000 normal traffic instances
- 25,000 malicious instances
- 8,000 synthetic zero-day patterns generated by GATE
- Mixed multi-stage attack sequences for more refer to the Table 3.

Table 3: Attack categories

Attack Type	Instances
DDoS	10,000
MITM	5,000
Protocol Exploitation	4,000
Behavioral Drift	3,000
Synthetic Zero-Day	8,000

Data was partitioned:

- 70% training
- 15% validation
- 15% testing

6.4 Implementation Details

- Edge inference model: Lightweight feedforward neural network
- Generative model: Conditional generative model with contextual embeddings
- Optimisation: Adam optimiser
- Learning rate: 0.001
- Risk fusion weights initialised uniformly
- Threshold recalibration interval: Every 30 minutes (simulated)

Hardware configuration (simulation environment):

- 16-core CPU
- 64 GB RAM
- GPU-assisted generative training

6.5 Evaluation Metrics

Standard classification metrics were used:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad Precision = \frac{TP}{TP + FP} \quad Recall = \frac{TP}{TP + FN} \quad F1 = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall}$$

Detection latency measured as:

$$T_d = T_{alert} - T_{attack_initiation}$$

False Positive Rate:

$$FPR = \frac{FP}{FP + TN}$$

6.6 Baseline Configurations

Baseline 1: Signature-Based IDS.

- Static rule database
- Centralised monitoring
- No adaptive retraining

Baseline 2: Centralised ML IDS.

- Deep neural network
- Periodic offline retraining
- No generative augmentation

6.7 Scalability Evaluation

Node count scaled from:

$$N = 100 \rightarrow 500$$

Measured:

- Latency growth rate
- Throughput degradation
- Model update overhead

6.8 Experimental Validity Considerations

To mitigate bias:

- Balanced attack injection
- Cross-validation across time windows
- Independent test partition
- Ablation study for:
 - No generative augmentation
 - Static weight fusion
 - No adaptive threshold

6.9 Summary

The experimental design simulates a realistic intelligent communication network with heterogeneous nodes and diverse attack scenarios. Comparisons against rule-based and centralised ML-based IDS models enable quantitative evaluation of adaptive performance improvements.

7 Results and Performance Evaluation

This section presents the empirical results comparing the proposed GenAI-driven adaptive cybersecurity mesh (GCM) against two baselines: (1) signature-based IDS (SID) and (2) centralised ML-based IDS (CML).

7.1 Detection Performance

7.1.1 Overall Classification Metrics

Table 4: Difference between Different Models

Model	Accuracy	Precision	Recall	F1-Score	FPR
SID	0.872	0.841	0.804	0.822	0.091
CML	0.914	0.902	0.887	0.894	0.062
GCM (Proposed)	0.948	0.939	0.931	0.935	0.038

Observations

- Detection accuracy improved by **3.4% over CML** and **7.6% over SID**.
- False Positive Rate reduced by **38.7% compared to CML**.
- F1-score improvement indicates balanced precision-recall performance.

Confidence in metric reliability: 0.91 as per Table 4.

7.2 Zero-Day Detection Capability

Performance on synthetic zero-day patterns:

Table 5 : Different Models and their Capabilities

Model	Zero-Day Recall	Zero-Day F1
SID	0.421	0.398
CML	0.684	0.672
GCM	0.862	0.849

The generative augmentation significantly improved the detection of unseen attack patterns.

Improvement rationale as per Table 5

- GATE-generated adversarial patterns enhanced model robustness.
- Dynamic threshold recalibration prevented under-detection during traffic bursts.

Confidence: 0.89

7.3 Detection Latency

Average detection latency as per Table 6:

Table 6: Different Models and their Detection Latency

Model	Mean Latency (ms)
SID	142
CML	118
GCM	74

Latency reduction of:

- 47.8% vs SID
- 37.3% vs CML

Reason:

- Edge-level inference reduces centralised bottlenecks.
- Distributed policy propagation shortens the enforcement cycle.

7.4 Adaptive Weight Evolution

Cross-layer weight distribution evolved over simulation:

Initial:

$$\alpha = \beta = \gamma = \delta = 0.25$$

After 24 hours:

$$\alpha = 0.32, \beta = 0.28, \gamma = 0.21, \delta = 0.19$$

Interpretation:

- Network-layer anomalies contributed more significantly during DDoS-heavy intervals.
- Behavioural metrics reduced weight as attack emphasis shifted.

This confirms adaptive rebalancing effectiveness.

7.5 Scalability Analysis

Node scaling from 100 $\rightarrow$ 500 as per Table 7:

Table 7: Nodes and Latency

Nodes	Avg Latency (ms)	Throughput Degradation
100	61	0%
300	74	4.2%
500	89	7.8%

Latency growth remained sub-linear, demonstrating distributed efficiency.

7.6 Ablation Study

Table 8 Configuration and Scored

Configuration	F1-Score
Full Model	0.935
No Generative Augmentation	0.902
Static Weights	0.918
Static Threshold	0.911

The generative component contributed $\sim 3.3\%$ F1 improvementz Adaptive weight recalibration added $\sim 1.7\%$.as per Table 8

7.7 ROC-AUC Analysis

Proposed model achieved:

$$AUC = 0.964$$

Compared to:

- CML: 0.928
- SID: 0.871

Indicates strong separability across attack categories.

7.8 Statistical Significance

Paired t-test conducted across test folds:

$$p < 0.01$$

Improvement statistically significant at 99% confidence level.

7.9 Summary of Improvements

The proposed GCM architecture demonstrates:

- Higher detection accuracy
- Significant false positive reduction
- Strong zero-day detection
- Lower latency
- Stable scalability
- Measurable contribution of generative augmentation

8 Discussion and Limitations

The experimental results indicate that integrating generative adaptation within a distributed cybersecurity mesh enhances detection accuracy, zero-day robustness, and latency performance. However, practical deployment considerations and limitations must be carefully examined.

8.1 Interpretation of Results

8.1.1 Why Generative Augmentation Improves Detection

The improvement in zero-day recall (0.862) is primarily attributable to adversarial pattern synthesis. By generating contextualised synthetic attack variants, the detection model becomes less dependent on fixed traffic signatures and better generalised to unseen distributions.

This reduces:

- Overfitting to historical traffic
- Concept drift vulnerability
- Sensitivity to polymorphic attacks

The ablation study confirms that removing generative augmentation reduces the F1-score by approximately 3%.

8.1.2 Effectiveness of Cross-Layer Fusion

The adaptive weight mechanism allowed the system to prioritise relevant layers during specific attack phases dynamically. For instance:

- Network layer weight increased during DDoS bursts.
- Behavioural features became more relevant during stealth compromise phases.

This dynamic rebalancing improved robustness compared to static fusion schemes.

8.1.3 Distributed Architecture Benefits

Edge-based inference reduced centralised bottlenecks, leading to lower detection latency. Sub-linear latency growth during node scaling indicates that the distributed mesh design effectively mitigates performance degradation in large networks.

8.2 Practical Deployment Considerations

While promising, deployment in real-world environments requires addressing:

1. **Computational Constraints**

- Edge devices with limited processing capability may struggle with frequent adaptive updates.
- Lightweight inference optimisation is required.

2. Model Synchronisation Overhead

- Frequent weight updates may introduce network overhead.
- Efficient update batching strategies must be implemented.

3. Trust in Generative Outputs

- Poorly calibrated generative models may introduce adversarial bias.
- Synthetic pattern validation mechanisms are necessary.

4. Privacy and Data Governance

- Cross-layer telemetry aggregation may raise privacy concerns.
- Secure aggregation and anonymisation must be enforced.

8.3 Limitations

8.3.1 Synthetic Dataset Dependence.

The experimental evaluation was conducted in a simulated intelligent communication environment. Although diverse attack scenarios were modelled, real-world deployment may introduce unforeseen noise patterns and operational variability.

8.3.2 Generative Model Complexity.

Generative threat synthesis requires periodic retraining. In highly resource-constrained environments, maintaining real-time adaptability may be challenging.

8.3.3 Threshold Sensitivity.

Dynamic threshold recalibration improves adaptability but may cause temporary instability during abrupt traffic shifts. Stability mechanisms must be incorporated.

8.3.4 Adversarial Manipulation of GenAI.

Sophisticated adversaries could attempt to poison telemetry inputs to manipulate generative adaptation. Defensive adversarial training strategies should be integrated.

8.4 Future Research Directions

Several extensions are proposed:

1. Federated generative threat modelling across organisational boundaries.
2. Integration with blockchain-backed trust management.
3. Formal verification of adaptive threshold stability.
4. Extension toward 6G-enabled ultra-low latency communication systems.
5. Deployment in real-world IoT testbeds for longitudinal validation.

8.5 Overall Implications

The results suggest that embedding adaptive generative intelligence into distributed cybersecurity meshes can meaningfully improve detection robustness in intelligent communication systems. However, practical scalability, stability, and adversarial robustness must be further evaluated under operational deployment conditions.

Conclusion

This paper presented a GenAI-driven adaptive cybersecurity mesh architecture for real-time threat detection in intelligent communication systems. The proposed framework integrates distributed zero-trust enforcement, cross-layer anomaly fusion, and generative threat modelling to address the limitations of centralised, static intrusion detection systems. Unlike traditional signature-based and centralised ML-based IDS approaches, the proposed architecture embeds adaptive intelligence within a distributed mesh topology. Edge Security Nodes perform local anomaly inference, while a GenAI-based adaptive threat engine synthesises contextual adversarial patterns and dynamically recalibrates fusion weights and risk thresholds. This design reduces detection latency, improves zero-day robustness, and enhances resilience under partial node compromise.

Experimental evaluation in a simulated heterogeneous intelligent communication environment demonstrated:

- Improved detection accuracy (94.8%)
- Significant reduction in false positive rate (3.8%)
- Strong zero-day recall (86.2%)
- Reduced detection latency (74 ms average)
- Stable scalability under increased node density

An ablation study confirmed that generative augmentation and adaptive weight recalibration contributed measurable performance improvements. While the results indicate promising advances in adaptive distributed cybersecurity, the study is constrained by synthetic simulation environments and computational assumptions. Future work should focus on real-world deployment, federated adaptive learning across domains, and formal robustness guarantees against adversarial manipulation. The proposed framework contributes toward next-generation intelligent communication security paradigms by embedding adaptive generative reasoning within distributed cybersecurity meshes, aligning with emerging requirements for scalable, resilient, and low-latency protection mechanisms.

Acknowledgments.

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors. The authors acknowledge the computational resources used for experimental validation and the anonymous reviewers whose feedback improved the quality of this manuscript.

Funding source

No funding was received for this study.

Conflict of Interest

The authors have no competing interests to declare that are relevant to the content of this article.

References

- [1] S.W. Rose, O. Borchert, S. Mitchell, S. Connelly: Zero Trust Architecture. NIST Special Publication (SP) 800-207 (2020). <https://doi.org/10.6028/NIST.SP.800-207>
- [2] H. Kang, G. Liu, Q. Wang, L. Meng, J. Liu: Theory and Application of Zero Trust Security: A Brief Survey. Entropy (Basel) 25(12), 1595 (2023). <https://doi.org/10.3390/e25121595>

- [3] B. Ramos-Cruz, J. Andreu-Perez, L. Martínez: The cybersecurity mesh: A comprehensive survey of involved artificial intelligence methods, cryptographic protocols and challenges for future research. *Neurocomputing* 581, 127427 (2024). <https://doi.org/10.1016/j.neucom.2024.127427>
- [4] Y. Yao, J. Duan, K. Xu, Y. Cai, Z. Sun, Y. Zhang: A survey on large language model (LLM) security and privacy: The Good, The Bad, and The Ugly. *High-Confidence Computing* 4(2), 100211 (2024). <https://doi.org/10.1016/j.hcc.2024.100211>
- [5] G. de Jesus Coelho da Silva, C.B. Westphall: A Survey of Large Language Models in Cybersecurity. *arXiv:2402.16968* (2024). <https://doi.org/10.48550/arXiv.2402.16968>
- [6] Khraisat, A. Alazab, S. Singh, T. Jan, A. Gomez: Survey on Federated Learning for Intrusion Detection System: Concept, Architectures, Aggregation Strategies, Challenges, and Future Directions. *ACM Computing Surveys* 57(1), Article 7 (2024). <https://doi.org/10.1145/3687124>
- [7] H. Zhang, J. Ye, W. Huang, X. Liu, J. Gu: Survey of federated learning in intrusion detection. *Journal of Parallel and Distributed Computing* 195, 104976 (2025). <https://doi.org/10.1016/j.jpdc.2024.104976>
- [8] S.A. Rahman, H. Tout, C. Talhi, A. Mourad: Internet of Things Intrusion Detection: Centralised, On-Device, or Federated Learning? *IEEE Network* 34(6), 310–317 (2020). <https://doi.org/10.1109/MNET.011.2000286>
- [9] M.A.E. Ferrag, O. Friha, D. Hamouda, L. Maglaras, H. Janicke: Edge-IIoTset: A New Comprehensive Realistic Cyber Security Dataset of IoT and IIoT Applications for Centralised and Federated Learning. *IEEE Access* (2022). <https://doi.org/10.1109/ACCESS.2022.3165809>
- [10] T. Al Nuaimi, S. Al Zaabi, M. Alyilieli, M. AlMaskari, S. Alblooshi, F. Alhabsi, M.F.B. Yusof, A. Al Badawi: A comparative evaluation of intrusion detection systems on the edge-IIoT-2022 dataset. *Intelligent Systems with Applications* 20, 200298 (2023). <https://doi.org/10.1016/j.iswa.2023.200298>
- [11] M.M. Rahman, S. Al Shakil, M.R. Mustakim: A survey on intrusion detection systems in IoT networks. *Cyber Security and Applications* 3, 100082 (2025). <https://doi.org/10.1016/j.csa.2024.100082>
- [12] D. Breitenbacher, I. Homoliak, Y.L. Aung, Y. Elovici, N.O. Tippenhauer: HADES-IoT: A Practical and Effective Host-Based Anomaly Detection System for IoT Devices (Extended Version). *IEEE Internet of Things Journal* 9(12), 9640–9658 (2022).
- [13] O. Alkadi, N. Moustafa, B. Turnbull, K.-K.R. Choo: A Deep Blockchain Framework-Enabled Collaborative Intrusion Detection for Protecting IoT and Cloud Networks. *IEEE Internet of Things Journal* 8(12), 9463–9472 (2021).
- [14] K. Shalabi, et al.: A Blockchain-based Intrusion Detection/Prevention Systems in IoT Networks: A Systematic Literature Review. *Procedia Computer Science* (2024).
- [15] J. Shu, Y. Zhou, X. Liu, D. Yang: Collaborative Intrusion Detection for VANETs: A Deep Learning-Based Distributed SDN Approach. *IEEE Transactions on Intelligent Transportation Systems* (2021). <https://doi.org/10.1109/TITS.2020.3027390>
- [16] Sharafaldin, A.H. Lashkari, A.A. Ghorbani: Toward Generating a New Intrusion Detection Dataset and Intrusion Traffic Characterisation. In: *Proc. ICISSP* (2018).
- [17] Canadian Institute for Cybersecurity (CIC): Intrusion Detection Evaluation Dataset (CIC-IDS2017) (2017).

- [18] Thakkar, R. Lohiya: A Review of the Advancement in Intrusion Detection Datasets. *Procedia Computer Science* (2020).
- [19] Khraisat, A. Alazab: A survey of intrusion detection systems in IoT networks (overview/survey source). *ACM/Elsevier survey sources* (2024–2025).
- [20] H. Xu, et al.: Large Language Models for Cyber Security: A Systematic Literature Review. *arXiv:2405.04760* (2024).
- [21] M. Rawat, G. Singal: Surveying Technology Fusion in IoT Networks for IDS: Exploring Datasets, Tools, Challenges, and Research Prospects. *ACM Transactions on Intelligent Systems and Technology* (2025). <https://doi.org/10.1145/3744745>
- [22] Y. Yao, et al.: A Survey on Large Language Model (LLM) Security and Privacy: The Good, the Bad, and the Ugly (preprint/extended versions). *arXiv:2312.02003* (2023).
- [23] Khraisat, et al.: Survey on Federated Learning for Intrusion Detection System (publisher record/metadata). *ACM Computing Surveys* 57(1) (2024).
- [24] H. Zhang, et al.: Survey of federated learning in intrusion detection (publisher record/metadata). *Journal of Parallel and Distributed Computing* 195 (2025).
- [25] Ramos-Cruz, et al.: Cybersecurity mesh survey (open-access bibliographic record). *Neurocomputing* 581 (2024).
- [26] M.A.E. Ferrag, et al.: Edge-IIoTset dataset paper (centralised + federated IDS benchmarking). *IEEE Access* (2022).
- [27] Arora, A., Yadav, S. K., & Sharma, K. Denial-of-service (dos) attack and botnet: Network analysis, research tactics, and mitigation. In *Research anthology on combating denial-of-service attacks* IGI Global. pp. 49-73, (2021).
- [28] S.A. Rahman, et al.: IoT intrusion detection: centralised vs on-device vs FL (core FL-IDS framing). *IEEE Network* 34(6) (2020).
- [29] Breitenbacher, et al.: HADES-IoT host-based anomaly detection (resource-constrained device security). *IEEE Internet of Things Journal* 9(12) (2022).
- [30] Ojha, R. P., Sanyal, G., Srivastava, P. K., & Sharma, K. Design and analysis of modified SIQRS model for performance study of wireless sensor network. *Scalable Computing: Practice and Experience*, 18(3), 229-242, (2017).