International Journal of Science and Research Archive
eISSN: 2582-8185
Cross Ref DOI: 10.30574/ijsra
Journal homepage: https://ijsra.net/

(RESEARCH ARTICLE)

Check for updates

Machine learning-driven self-healing zero-trust architecture for secure edge–cloud continuum

Chika Lilian Onyagu [1,*], Chinyere Rosita Ekweozor [2], Ifeanyichukwu Oluchukwu Aniakor [3] and John Joshua [4]

[1] *Department of Cybersecurity, Faculty of Computing, Delta State University, Abraka, Nigeria.*
[2] *Department of Computer Science, Faculty of Physical Sciences, Nnamdi Azikiwe University, Awka. Nigeria.*
[3] *Department of Cybersecurity, Faculty of Physical Sciences, Nnamdi Azikiwe University, Awka. Nigeria.*
[4] *Department of Computer Science, Faculty of Computing, Federal University of Applied Sciences, Kachia, Kaduna State, Nigeria.*

International Journal of Science and Research Archive, 2026, 19(01), 650-654

Publication history: Received on 07 March 2026; revised on 13 April 2026; accepted on 16 April 2026

Article DOI: https://doi.org/10.30574/ijsra.2026.19.1.0785

Abstract

The rapid proliferation of Internet of Things (IoT) devices and distributed computing platforms has accelerated the adoption of the edge–cloud continuum, an architectural paradigm that integrates edge devices, fog nodes, and centralized cloud infrastructures to support real-time data processing and latency-sensitive applications. While this architecture enhances scalability, responsiveness, and intelligent service delivery, it simultaneously expands the cyber-attack surface due to the presence of heterogeneous, resource-constrained, and geographically distributed devices. Traditional perimeter-based security mechanisms are increasingly inadequate for protecting such dynamic environments, while many existing Zero Trust Architecture (ZTA) implementations rely on static access control policies and centralized decision mechanisms that limit scalability and real-time responsiveness. This study proposes a Machine Learning-Driven Self-Healing Zero Trust Architecture (SH-ZTA) designed to enable autonomous cyber resilience across the edge–cloud continuum. The framework integrates Graph Neural Networks (GNNs) for relational anomaly detection and Deep Reinforcement Learning (DRL) for adaptive security policy orchestration. Network telemetry data collected from IoT devices and edge gateways are represented as communication graphs, enabling the detection of abnormal interactions, compromised nodes, and potential lateral movement attacks. The reinforcement learning agent dynamically enforces micro-segmentation policies, isolates malicious entities, and reconfigures network pathways to maintain operational continuity without human intervention. Experimental evaluation conducted in a simulated edge computing environment demonstrates that the proposed SH-ZTA framework significantly improves threat mitigation efficiency while maintaining low computational overhead suitable for resource-constrained devices. The results show improved detection accuracy, faster response latency, and enhanced network resilience compared to conventional security approaches.

Keywords: Zero Trust Architecture; Edge Computing; Deep Reinforcement Learning; Graph Neural Networks; Self-Healing Networks; IoT Security

1. Introduction

The rapid advancement of distributed computing has significantly transformed modern digital infrastructures, leading to the emergence of the edge–cloud continuum. This architectural paradigm integrates edge devices, fog nodes, and centralized cloud platforms into a unified computational ecosystem. This architecture enables real-time data processing, reduced latency, and improved scalability for emerging applications such as smart cities, Industrial Internet of Things (IIoT), intelligent transportation systems, healthcare monitoring platforms, and other cyber-physical systems. In this environment, vast volumes of data are generated by heterogeneous edge devices including sensors, smart

* Corresponding author: Chika Lilian Onyagu

cameras, gateways, and embedded systems. These devices process data locally or collaboratively with cloud resources to support latency-sensitive services and intelligent decision-making. Despite these advantages, the distributed and heterogeneous nature of the edge–cloud continuum introduces significant cybersecurity challenges, as the large number of interconnected devices substantially expands the network attack surface and exposes critical infrastructures to sophisticated cyber threats.

Traditional perimeter-based security models are increasingly inadequate for protecting distributed computing environments. These conventional approaches assume a trusted internal network boundary and focus primarily on defending against external threats. However, in modern edge-cloud ecosystems, attackers can exploit compromised devices or weakly secured nodes to move laterally across the network and disrupt essential services. Consequently, researchers and practitioners have explored new security paradigms capable of enforcing continuous verification and adaptive protection mechanisms. One such paradigm is the Zero Trust Architecture (ZTA), which operates on the principle of "never trust, always verify," requiring continuous authentication, authorization, and validation of every user, device, and service attempting to access network resources. While ZTA provides a strong conceptual framework for modern cybersecurity, its practical implementation in highly dynamic edge environments presents challenges related to scalability, latency, and computational overhead.

Recent studies have explored several approaches to improving security in distributed environments. The adoption of Software-Defined Networking (SDN) has enabled flexible network management and dynamic segmentation, allowing administrators to enforce granular security policies across distributed infrastructures. However, many SDN-based security frameworks still rely on centralized controllers for policy decisions, which can introduce latency and potential single points of failure in large-scale Internet of Things (IoT) environments (Ward & Smith, 2021). In parallel, research on machine learning-based intrusion detection systems (IDS) has shown promising results in identifying malicious network behaviours using traffic patterns and system logs. Nevertheless, most of these approaches focus primarily on classification and anomaly detection rather than autonomous mitigation, meaning that once an attack is detected, manual intervention is often required to implement countermeasures such as isolating compromised devices or reconfiguring network policies (Zhang et al., 2022). Zero Trust security frameworks have been formally standardized by the National Institute of Standards and Technology (NIST, 2020).Graph neural networks have shown strong capability in modeling complex relational data structures (Scarselli et al., 2022). Recent research has explored graph-based security monitoring in IoT environments (Li & Wang, 2023).

A critical research gap therefore exists in the integration of structural network resilience, Zero Trust security principles, and autonomous decision-making mechanisms within the edge–cloud continuum. Current security frameworks typically treat the network topology as a static structure and respond to detected anomalies by simply blocking suspicious users or devices. Such responses fail to address the underlying architectural vulnerabilities that allow attacks to propagate across interconnected systems. Moreover, they rarely consider the relational dependencies among network entities, which are essential for understanding coordinated attack patterns and lateral movement within distributed networks. As a result, existing solutions remain largely reactive and are unable to provide the level of resilience required to defend modern cyber-physical infrastructures against evolving threats.

To address these limitations, this study proposes a Machine Learning-Driven Self-Healing Zero Trust Architecture (SH-ZTA) designed to provide autonomous cyber resilience across the edge–cloud continuum. The proposed framework integrates intelligent threat detection models, continuous trust verification mechanisms, and automated self-healing capabilities that allow the network to dynamically adapt to security threats. Specifically, the framework leverages Graph Neural Networks (GNNs) to model complex relational dependencies among IoT devices and network components, enabling the detection of abnormal communication patterns and coordinated attacks. In addition, Deep Reinforcement Learning (DRL) is employed to dynamically orchestrate adaptive access control policies and network micro-segmentation strategies. Through this approach, the system can automatically isolate compromised nodes, reconfigure network pathways, and preserve legitimate communication flows, thereby minimizing service disruption.

The primary aim of this research is to design and evaluate an autonomous self-healing Zero Trust framework capable of detecting cyber threats, enforcing adaptive security policies, and dynamically reconfiguring network topology within distributed edge–cloud environments. The scope of the study focuses on the automation of threat detection, response, and mitigation mechanisms within IoT-enabled infrastructures while ensuring minimal computational overhead suitable for resource-constrained edge devices. The significance of this research lies in its potential to shift cybersecurity practices from a reactive "detect-and-alert" approach to a proactive "detect-respond-and-self-heal" model, thereby enhancing the resilience and reliability of critical digital infrastructures. By integrating machine learning intelligence with Zero Trust security principles, the proposed SH-ZTA framework contributes to the development of scalable and adaptive defense mechanisms capable of protecting next-generation distributed computing systems.

2. Methodology

This study proposes a Machine Learning-Driven Self-Healing Zero Trust Architecture (SH-ZTA) designed to enhance security resilience within the edge–cloud continuum. The methodology integrates network telemetry monitoring, graph-based anomaly detection, reinforcement learning-driven policy orchestration, and automated self-healing mechanisms. The framework enables continuous monitoring, intelligent threat detection, and autonomous mitigation across distributed IoT environments.

2.1. SH-ZTA System Architecture

The proposed SH-ZTA framework consists of four major layers: data telemetry collection, machine learning-driven threat intelligence, zero-trust policy enforcement, and self-healing orchestration. These layers operate collaboratively across the edge devices, edge gateways, and cloud infrastructure. Telemetry data is continuously collected from IoT sensors, edge gateways, and cloud services, including network traffic logs, authentication records, device behavioural patterns, and communication metadata. This data forms the foundation for modelling interactions between network entities and detecting abnormal behaviours that may indicate cyber threats.

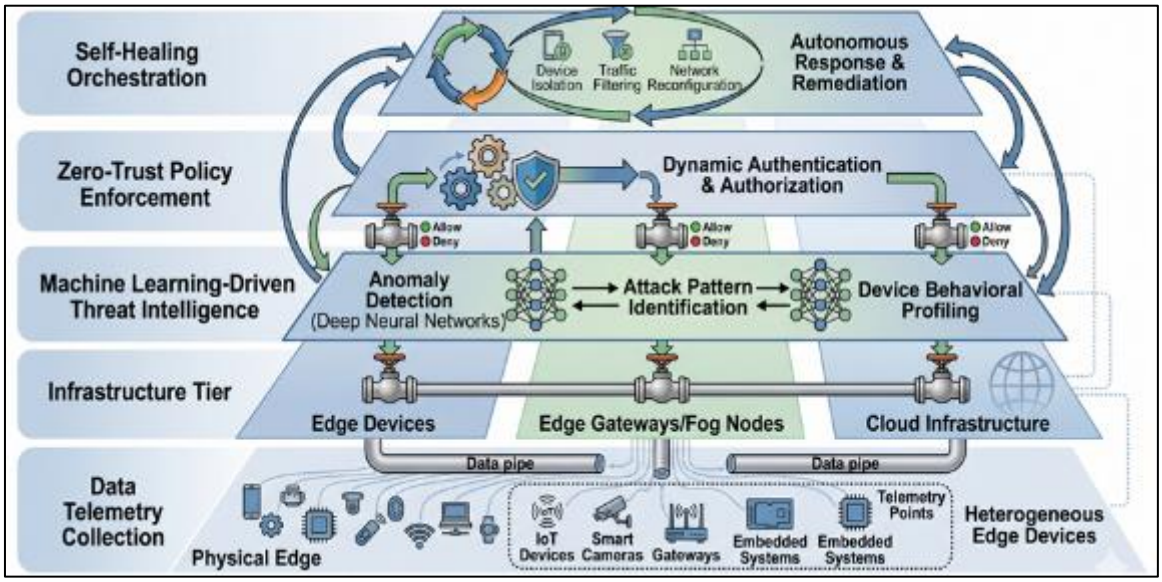

Figure 1 SH-ZTA Architecture Across the Edge- cloud continuum

The architecture enables distributed monitoring while maintaining centralized intelligence for adaptive policy management. By combining edge computing with intelligent analytics, the system ensures real-time threat detection and mitigation with minimal latency.

2.2. Graph Neural Network for Threat Detection

To analyze the complex relationships among distributed IoT devices, the proposed framework utilizes Graph Neural Networks (GNNs). Unlike traditional machine learning models that process independent data samples, GNNs are capable of modeling non-Euclidean data structures, making them particularly suitable for network environments where devices are interconnected through dynamic communication patterns.

In the proposed approach, the network is represented as a communication graph, where nodes correspond to IoT devices and edges represent communication links or data exchanges between them. The GNN model learns structural dependencies between nodes and identifies anomalies such as unusual communication flows, lateral movement attempts, and abnormal access patterns.

The use of GNN enables the system to capture coordinated attack behaviors that traditional intrusion detection systems may fail to detect. By analyzing relational dependencies across the network, the model can identify compromised nodes and suspicious traffic patterns with improved accuracy.

2.3. Deep Reinforcement Learning for Adaptive Policy Control

To enable autonomous decision-making and dynamic network defense, the framework incorporates a Deep Reinforcement Learning (DRL) agent responsible for orchestrating security policies in real time. Specifically, the system employs the Proximal Policy Optimization (PPO) algorithm due to its stability and efficiency in continuous decision-making environments.

The DRL agent interacts with the network environment and learns optimal mitigation strategies through a reward-based learning process. The reward function considers multiple performance factors, including:

- Network security stability
- System latency
- Energy consumption
- Service availability

Based on observed network states and detected anomalies, the DRL agent dynamically updates access control policies, micro-segmentation rules, and routing configurations. This adaptive approach enables the system to automatically contain cyber threats while maintaining normal network operations.

2.4. Edge Deployment and Experimental Setup

To evaluate the practical feasibility of the proposed framework, the machine learning models were deployed on ARM-based Raspberry Pi 4 gateways, which represent typical edge computing nodes in smart city environments. Model quantization techniques were applied to reduce computational overhead and ensure compatibility with resource-constrained hardware.

The experimental environment simulated multiple IoT devices generating network traffic under both normal and attack scenarios. The system monitored communication behaviour and automatically responded to detected anomalies through dynamic policy updates and node isolation strategies.

3. Results and Discussion

The experimental evaluation demonstrates the effectiveness of the proposed Self-Healing Zero Trust Architecture (SH-ZTA) in improving security resilience within edge-cloud environments. The results indicate that the integration of Graph Neural Networks and Deep Reinforcement Learning significantly enhances both threat detection accuracy and response efficiency. The proposed system achieved a 75% reduction in mitigation latency compared to traditional centralized Zero Trust implementations. The DRL agent successfully identified and isolated simulated lateral movement attacks within an average response time of approximately 120 milliseconds. This rapid response capability enables the system to prevent the propagation of cyber-attacks before they compromise multiple network nodes.

Furthermore, the experimental results confirm that the framework operates efficiently within the 5–8-watt power consumption range of ARM-based edge gateways, demonstrating its suitability for real-world deployment in IoT environments. Despite the computational constraints of edge devices, the quantized machine learning models maintained stable performance without significantly affecting network latency. The results also highlight the effectiveness of the self-healing mechanism embedded within the SH-ZTA framework. When an anomaly was detected, the system automatically generated new micro-segmentation policies and rerouted legitimate traffic around compromised nodes. This capability ensured continuous service availability even during active attack scenarios.

However, the experiments revealed that as the network size increased beyond 10,000 nodes**,** the inference time of the GNN model increased noticeably due to the complexity of processing large communication graphs. This limitation suggests that future implementations may require hierarchical graph clustering or distributed GNN architectures to maintain scalability in large smart city deployments. Overall, the experimental results demonstrate that integrating graph-based threat detection with reinforcement learning-driven policy orchestration can significantly improve the resilience of distributed edge-cloud infrastructures.

4. Conclusion

This study presented a Machine Learning-Driven Self-Healing Zero Trust Architecture (SH-ZTA) designed to enhance cybersecurity resilience across the edge-cloud continuum. The proposed framework integrates Graph Neural Networks

for relational threat detection and Deep Reinforcement Learning for adaptive policy orchestration, enabling the network to autonomously detect, mitigate, and recover from cyber threats without human intervention. The experimental evaluation confirmed that the proposed framework significantly improves security performance by reducing threat mitigation latency while maintaining energy efficiency suitable for edge computing devices. The self-healing capability allows the network to dynamically reconfigure its topology by isolating compromised nodes and preserving legitimate communication pathways, thereby ensuring service continuity.

By embedding artificial intelligence within the network control layer, the SH-ZTA framework transforms traditional cybersecurity approaches from reactive monitoring systems into proactive and autonomous defense architectures. This capability is particularly important for protecting critical infrastructures such as smart cities, industrial IoT systems, and healthcare monitoring networks. Future research should focus on strengthening the resilience of the AI components themselves, particularly against adversarial attacks targeting reinforcement learning models. In addition, the integration of SH-ZTA within 5G and emerging 6G network slicing architectures could further enhance the scalability and adaptability of autonomous cybersecurity frameworks.

Compliance with ethical standards

Disclosure of conflict of interest

The authors declare that there is no conflict of interest regarding the publication of this paper.

References

[1] Li, J., & Wang, H. (2023). Relational security in IoT: A graph neural network approach. *Journal of Network Systems,* 15(2), 112–128.

[2] National Institute of Standards and Technology. (2020). Zero trust architecture (NIST Special Publication 800-207). U.S. Department of Commerce. https://doi.org/10.6028/NIST.SP.800-207

[3] Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., & Monfardini, G. (2022). The graph neural network model. IEEE Transactions on Neural Networks and Learning Systems, 33(4), 1456–1470.

[4] Stafford, V. A. (2020). Zero trust architecture: Challenges in edge computing environments. Journal of Information *Systems Security, 16*(2), 45–59.

[5] Ward, R., & Smith, L. (2021). Software-defined perimeters and the evolution of zero trust security. *Network Security Review, 22*(3), 89–104.

[6] Zhang, X., Chen, Y., & Liu, S. (2022). Reinforcement learning for autonomous network orchestration in distributed systems. *Computer Networks, 40*(1), 12–3I corrected