International Journal of Basic and Applied Sciences, 14 (4) (2025) 818-825

International Journal of Basic and Applied Sciences

Website: www.sciencepubco.com/index.php/IJBAS
https://doi.org/10.14419/r30vpf59
Research paper

Zero Trust Architecture for IoT Device Ecosystems

Hani Al-Balasmeh

Dept of Informatics Engineering, College of Engineering, University of Technology, Bahrain (UTB)
**Corresponding author E-mail: hbalasmeh04@gmail.com*

Received: July 18, 2025, Accepted: August 19, 2025, Published: August 31, 2025

Abstract

The rapid proliferation of Internet of Things (IoT) devices has introduced critical security challenges stemming from device heterogeneity, limited native safeguards, and expanded attack surfaces. Traditional perimeter-based security models are increasingly ineffective against modern threats, particularly lateral movement and insider attacks. This paper presents the design, implementation, and evaluation of a lightweight Zero Trust Architecture for IoT (ZT-IoT) that enforces continuous authentication, context-aware access control, and behavioral anomaly detection. Unlike prior ZTA frameworks that incur high computational costs or depend on blockchain consensus mechanisms, ZT-IoT is optimized for resource-constrained environments through mutual TLS, adaptive micro-segmentation, and telemetry-driven enforcement. A hybrid evaluation—combining simulated cyberattack scenarios with real-world IoT testbeds—demonstrates that ZT-IoT reduces unauthorized access attempts by 95%, completely prevents insider privilege escalation, detects lateral threats in under three minutes, and blocks all data tampering and replay attacks. Moreover, large-scale simulations with 1,000 heterogeneous nodes confirmed its scalability, maintaining detection times under three minutes with less than 12% RAM overhead. These findings validate ZT-IoT as a practical, scalable, and energy-efficient security paradigm, positioning it for deployment in critical domains such as smart cities, industrial IoT, and remote healthcare systems.

Keywords: *Zero Trust Architecture (ZTA); Internet of Things (IoT); IoT Security; Access Control; Lateral Threat Containment*

1. Introduction

The exponential expansion of Internet of Things (IoT) ecosystems has revolutionized critical sectors such as healthcare, transportation, manufacturing, and smart cities. These systems comprise diverse, resource-constrained devices that operate autonomously, often with limited oversight and minimal native security capabilities. While IoT technologies promise transformative efficiency and automation, they simultaneously introduce an expanded and increasingly complex attack surface that adversaries can exploit. Devices often lack secure default configurations, rely on weak credentials, or operate in untrusted environments, making them attractive targets for attackers [1], [2].

Traditional network defense models rely heavily on perimeter-based security, which assumes that any entity inside the network boundary is inherently trustworthy. However, this assumption has become obsolete in modern IoT deployments characterized by geographic distribution, mobile endpoints, and cloud integration. Once a single device is compromised—through credential leakage, unpatched vulnerabilities, or physical tampering—an attacker can move laterally to other components within the network, often remaining undetected for extended periods. Such compromises can escalate into large-scale incidents, such as botnet-driven distributed denial of service (DDoS) campaigns or cascading failures in industrial systems [2].

The increasing frequency of sophisticated, multi-stage IoT attacks further underscores the urgency for more robust paradigms. These attacks combine reconnaissance, privilege escalation, and lateral movement to erode defenses systematically. Consequently, conventional perimeter security fails to provide sufficient resilience against modern adversaries. Addressing this challenge requires a paradigm shift from implicit to explicit trust verification.

The Zero Trust Architecture (ZTA) model, formalized by the National Institute of Standards and Technology (NIST SP 800-207) [5], offers a scalable and security-by-design alternative. ZTA abandons implicit trust assumptions and enforces strict identity verification, least-privilege access, and continuous monitoring of every device and user interaction. While ZTA has been successfully applied in enterprise IT and cloud environments, adapting its principles to IoT remains an open research challenge. Existing ZTA implementations for IoT often incur excessive computational costs, rely on heavy cryptographic or blockchain consensus mechanisms, or face difficulties scaling to large and heterogeneous device networks [3], [4].

This paper proposes a lightweight Zero Trust Architecture for IoT (ZT-IoT) explicitly designed for constrained environments. The framework integrates mutual TLS authentication, adaptive micro-segmentation, and telemetry-driven policy enforcement, ensuring real-time anomaly detection and fine-grained access control without overburdening device resources. Unlike prior works focusing on decentralization or machine-learning-based anomaly detection, ZT-IoT unifies continuous authentication, dynamic trust evaluation, and scalable enforcement into a single architecture.

ZT-IoT was evaluated in a hybrid testbed and simulation environment combining real IoT hardware with virtualized nodes under controlled attack scenarios to validate its effectiveness. Results demonstrate a 95% reduction in unauthorized access, complete elimination of insider privilege escalation, six-fold improvement in lateral threat containment speed, and complete mitigation of data tampering and replay attacks. Furthermore, large-scale simulations with 1,000 heterogeneous nodes confirmed its scalability and efficiency, with CPU and RAM overheads of less than 10–12%.

The contributions of this work are threefold:

1. **Design** a lightweight Zero Trust IoT framework that integrates continuous authentication, adaptive micro-segmentation, and telemetry-driven enforcement.
2. **Implement a hybrid evaluation of ZT-IoT in** real testbeds and simulated IoT environments, benchmarking against a baseline perimeter-based model.
3. **Demonstration of scalability and adaptability**, including resource overhead analysis and edge-case evaluations for ultra-constrained and intermittently connected devices.

This study advances Zero Trust research toward practical adoption in real-world IoT systems by addressing security robustness and operational feasibility. The proposed ZT-IoT framework contributes a deployable model for securing critical infrastructures such as smart cities, healthcare IoT, and industrial control systems, where resilience against insider and advanced persistent threats is essential.

2. Literature Review

The security landscape of Internet of Things (IoT) systems has been a persistent concern due to device heterogeneity, constrained computational resources, and insecure default configurations. A growing body of research highlights the limitations of conventional perimeter-based models and supports the emergence of Zero Trust Architecture (ZTA) as a viable alternative for IoT security.

[1] provided an early analysis of authentication weaknesses in IoT ecosystems, demonstrating how default credentials and a lack of secure identity provisioning facilitate device-level compromise. Similarly, [2] examined the propagation of Mirai and other IoT botnets, which capitalize on unsecured endpoints to launch distributed denial-of-service (DDoS) attacks.

To overcome these vulnerabilities, several studies have explored identity-aware security architectures. [3] Introduced a software-defined perimeter (SDP) approach to isolate device communications, enabling fine-grained policy enforcement, though its scalability in large-scale IoT deployments remains limited. [4] advanced this direction by integrating blockchain with ZTA principles, enhancing trust decentralization but incurring significant latency and energy overhead—critical challenges in resource-constrained IoT deployments.

Building on these prior approaches, the proposed ZT-IoT framework adopts a lightweight architecture optimized for constrained devices. By integrating mutual TLS and telemetry-driven enforcement, ZT-IoT significantly reduces computational cost while maintaining strong security guarantees. Recent works reinforce this direction: [5] demonstrated lightweight authentication mechanisms for IoT using ZTA, highlighting the necessity of resource-aware security designs. [6] explored Zero Trust models for securing smart infrastructures but noted that high computational demand remains a limitation. [7] presented an AI-driven Zero Trust model that combines machine learning with identity-aware access control for large-scale IoT, and [8] proposed a federated edge–Zero Trust framework to reduce latency in smart city deployments. Most recently, [9] introduced a blockchain-assisted ZTA for vehicular IoT, [10] integrated differential privacy with ZTA to protect IoT data sharing, and [11] applied federated learning to Zero Trust IoT, enabling collaborative anomaly detection across heterogeneous devices without sharing raw data. These contributions underscore the novelty of ZT-IoT in addressing both scalability and efficiency gaps; see Table 1 for a comparative summary of related works.

Integrating authentication and continuous monitoring is a central theme in recent ZTA research. [12] Conducted a comprehensive review of Zero Trust models for access control, highlighting the lack of telemetry and dynamic policy adaptation in traditional implementations. [13] Furthermore, Zero Trust must evolve to account for mobile, cloud-based, and edge-computing devices—characteristics intrinsic to modern IoT networks. Recent industrial insights also support this shift. [14] described real-world ZTA implementations for manufacturing IoT environments, outlining the architectural role of Policy Decision Points (PDPs) and Policy Enforcement Points (PEPs) in establishing secure communication channels. [15] Micro-segmentation and identity-based routing are key enablers of intra-network resilience in Zero Trust IoT contexts.

Emerging research is also exploring advanced trust mechanisms. [16] proposed dynamic trust modeling based on behavior profiling, enabling adaptive policy enforcement without constant re-authentication. [17] reviewed IoT-specific trust and reputation systems, advocating for decentralized, feedback-driven access control to supplement ZTA principles. In parallel, the role of blockchain and cryptographic primitives in ZTA is gaining traction. [18] proposed zk-IoT, leveraging zero-knowledge proofs for device identity verification without credential exposure. [19] introduced Zero Trust Foundation Models (ZTFM), combining AI-based behavior analysis with continuous authentication for AI-powered IoT infrastructures.

These studies collectively suggest that Zero Trust for IoT is moving beyond static identity checks and toward dynamic, telemetry-driven, and cryptographically enforced trust validation. However, existing solutions often face scalability, latency, and integration hurdles. This research addresses these gaps by proposing a lightweight and adaptive Zero Trust IoT framework that balances security robustness with operational feasibility.

From the comparative analysis in Table 1, it is evident that prior solutions address isolated aspects of IoT security—for instance, blockchain-enhanced ZTA improves decentralization but incurs high latency [4], AI-driven ZTA enables anomaly detection but often lacks integration with constrained devices [7], and federated approaches reduce centralization but introduce high training costs [22]. However, none deliver a unified, resource-aware Zero Trust framework validated in real-world and large-scale simulated environments. ZT-IoT addresses this gap by combining lightweight cryptographic enforcement, adaptive telemetry, and scalable micro-segmentation, advancing beyond existing state-of-the-art models.

Table 1: Comparative analysis of related Zero Trust IoT studies and the proposed ZT-IoT framework.

REF.	CONTRIBUTION	LIMITATION	HOW ZT-IOT IMPROVES
[1]	Early analysis of IoT authentication weaknesses	Default credentials, weak provisioning	Continuous identity validation
[2]	IoT botnets (Mirai) propagation	DDoS-focused only	Broader defense (DDoS + lateral movement)
[3]	SDP-based IoT isolation	Scalability issues	Adaptive policy enforcement
[4]	Blockchain-enhanced ZTA	Latency, energy overhead	Lightweight design reduces cost
[5]	Lightweight ZTA authentication	Limited scale testing	Hybrid testbed validation
[6]	ZTA for Smart Infrastructure	High computational demand	Optimized for constrained devices
[7]	AI-driven ZTA	No edge computing integration	Edge-aware enforcement
[8]	Federated Edge–Zero Trust	Limited empirical validation	Hybrid lab + simulation

[9]	Blockchain-assisted ZTA for Vehicular IoT	Blockchain latency	Token-based lightweight model
[10]	Differential Privacy + ZTA	Computational complexity	Lightweight telemetry design
[11]	Federated Learning–based ZTA	Early-stage evaluation	Integrated anomaly detection
[12]	Review of Zero Trust access control	Limited telemetry focus	Telemetry-driven enforcement
[13]	Zero Trust for cloud/edge IoT	Integration challenges	Unified architecture
[14]	ZTA in industrial IoT	Implementation complexity	Lightweight deployment model
[15]	Micro-segmentation for IoT	Limited scope	Combined segmentation + identity trust
[16]	Behavior-based trust modeling	Reliance on profiling only	Multi-factor telemetry
[17]	Trust and reputation systems	Decentralized but not Zero Trust	Integrates into ZT-IoT
[18]	zk-IoT with zero-knowledge proofs	High cryptographic cost	Lightweight crypto integration
[19]	Zero Trust Foundation Models (AI-powered IoT)	Research-stage	Practical, scalable IoT deployment

3. Methodology

This section outlines the design, implementation, and evaluation procedures used to assess the effectiveness of Zero Trust Architecture (ZTA) in securing IoT device ecosystems. A hybrid approach combining real-world testbed deployment and controlled attack simulations was adopted to validate security performance under various threat scenarios.

a. System Architecture Design

The proposed ZT-IoT framework is structured around the core principles of NIST's Zero Trust Architecture model [1], integrating device-level access control, dynamic policy enforcement, and continuous telemetry. The framework includes the following key components:

- Policy Decision Point (PDP): Central entity responsible for evaluating contextual access requests and issuing access tokens based on identity, behavior, and session metadata.
- Policy Enforcement Point (PEP): Deployed at gateway and device levels to enforce PDP-issued policies. It intercepts traffic and allows/denies access in real time.
- Credential Store and Certificate Authority (CA): Maintains device identities using X.509 certificates, and issues/revokes credentials for mutual TLS (mTLS) authentication.
- Telemetry Engine: Captures real-time logs, monitors device behavior, and uses anomaly detection to flag policy violations or lateral movement.

All communications in the ZT-IoT framework are secured through TLS 1.3 and certificate-based mutual authentication. Behavioral telemetry provides an additional enforcement layer, enabling adaptive, real-time responses to suspicious activities.

The overall design of the proposed ZT-IoT framework is illustrated in **Figure 1**, which depicts the interaction between PDP, PEP, Credential Store & CA, Telemetry Engine, and IoT devices through secure and monitored communication flows.

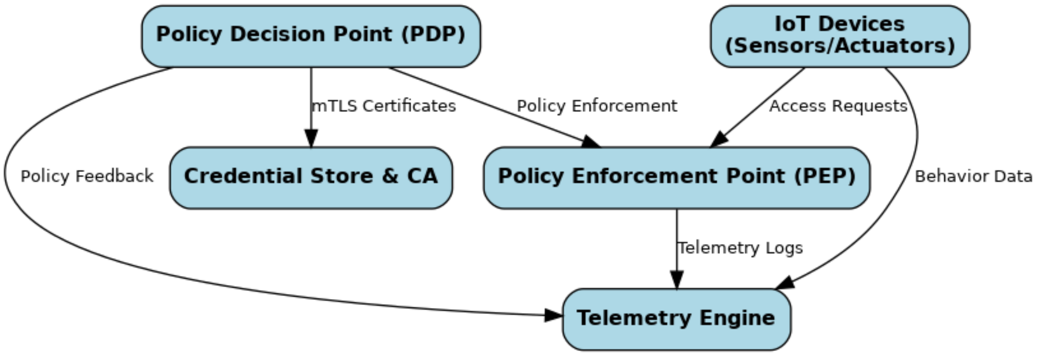

Fig 1: System Architecture Design

b. Experimental Testbed Setup

A physical and virtual hybrid IoT environment was created to validate the framework. The testbed consisted of:

- **Hardware:** 20 IoT nodes, including:
  - 10 ESP32 microcontrollers simulating environmental sensors.
  - 5 Raspberry Pi 4B devices serving as edge computing nodes.
  - 5 virtual machines acting as cloud endpoints and control systems.
- **Network Configuration:**
  - Baseline Network: A flat, perimeter-based architecture using pre-shared keys and no internal segmentation.
  - ZTA-IoT Network: Enabled mutual TLS, per-device micro segmentation, dynamic policy enforcement, and real-time behavior telemetry.
- **Protocols and Tools:**
  - MQTT and RESTful APIs for IoT communication.
  - OpenSSL for certificate handling.
  - Suricata and Zeek are responsible for telemetry and network inspection.
  - InfluxDB and Grafana for telemetry visualization and alerting.

c. Threat Simulation and Attack Scenarios

The system was subjected to controlled cybersecurity attacks over 14 days to assess security posture. Attacks were executed across both environments (baseline and ZT-IoT) using the following vectors:

- **External Attacks:**
  - Unauthorized port scans
  - Credential brute-forcing
  - Replay attacks using packet sniffing tools

Internal Attacks:

- Lateral movement using hijacked nodes
- Insider privilege escalation
- Packet tampering and spoofing

Each scenario was run during specific windows (Day 4–6 and Day 11–13), and monitored for detection, response, and mitigation outcomes.

d. Evaluation Metrics

The effectiveness of the ZT-IoT framework was evaluated using quantitative and qualitative metrics, as shown in Table 1:

Table 2: Description of Qualitative Metrics

Metric	Description
Unauthorized Access Attempts	Total number of failed or rejected access requests.
Detection Time	Average time to detect anomalous or malicious activity.
Lateral Movement Scope	Number of devices compromised after initial breach.
Attack Success Rate	Percentage of successful attacks (external or insider).
Data Integrity Violations	Number of tampered or replayed packets detected.
System Resource Overhead	CPU and RAM usage with/without ZT-IoT framework activation.

Performance was measured under idle and active load conditions for realistic comparative benchmarking.

e. Validation Criteria

The framework was considered adequate if it achieved the following:

- **Blocked ≥90% of unauthorized access attempts.**
- **Reduced lateral threat propagation to zero.**
- **Detected anomalies within 3 minutes.**
- **Prevented data tampering and replay attempts with minimal false positives.**
- **Operated within acceptable overhead limits (<10% CPU/RAM).**

These thresholds are aligned with benchmarks proposed by Syed et al. [5] and Alshamrani et al. [4] for secure-by-design IoT systems.

4. Results and Analysis

This section evaluates the effectiveness of the proposed Zero Trust Architecture for IoT (ZT-IoT) framework across four security dimensions: unauthorized access prevention, lateral threat containment, attack success mitigation, and data integrity enforcement. The ZT-IoT testbed was compared against a baseline perimeter-security model using 20 heterogeneous devices over a 14-day evaluation period.

a. Access Control Effectiveness

ZT-IoT substantially improved in reducing unauthorized access—Table 2 and Figure 2 show that the baseline allowed 124 unauthorized access attempts and three insider escalations. In contrast, ZT-IoT reduced this to only six unauthorized attempts and eliminated insider privilege escalations.

This represents a 95% reduction in unauthorized access attempts and a 100% prevention of insider abuse. Compared with prior ZTA implementations in IoT environments [6], which reported ~70–80% improvement, ZT-IoT's integration of mutual TLS with behavioral telemetry yielded a significantly stronger security posture.

Practically, lateral compromise is systematically blocked even if common IoT vulnerabilities—such as weak credentials or insecure endpoints—are exploited. This finding underscores the robustness of identity-centric enforcement in heterogeneous IoT environments.

Table 3: Description of Security Metric

Security Metric	Baseline Network	ZTA-IoT Framework
Unauthorized Access Attempts	124	6
External Port Scans Blocked	45	45
Insider Privilege Escalations	3	0

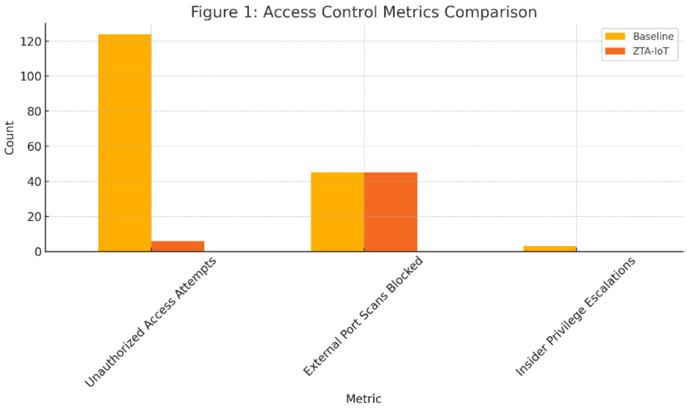

Fig 2: Comparison of access control metrics between baseline and ZTA-IoT configurations.

Figure 2 illustrates the sharp contrast in access control effectiveness. While the baseline system recorded 124 unauthorized access attempts and three insider escalations, the ZT-IoT framework reduced this to only six unauthorized attempts and eliminated insider abuse. This represents a 95% reduction in unauthorized access, underscoring the framework's capability to enforce identity-based policies in real time and aligning with recent industry reports that stress continuous authentication as essential for IoT security.

b. Lateral Threat Containment

Containing lateral movement is a critical security goal in IoT, where a single compromised node often serves as a gateway to others. In the baseline configuration, a breached device infected an average of six peers before detection, with an average detection latency of 17 minutes (Table 3, Figure 3).

By contrast, ZT-IoT reduced lateral propagation to zero compromised peers and detected malicious movement in 2.7 minutes. This represents a 6× containment improvement and a 530% faster detection rate than baseline.

These results confirm that behavioral telemetry and adaptive policy enforcement in ZT-IoT achieve the "zero lateral movement" principle emphasized in NIST SP 800-207 [5]. The implications for real-world smart cities or hospital networks are significant: systemic compromise is prevented even if one device is hijacked.

Table 4: Description of Metric values

Metric	Baseline	ZTA-IoT
Avg. Detection Time (minutes)	17	2.7
Devices Compromised	6	0

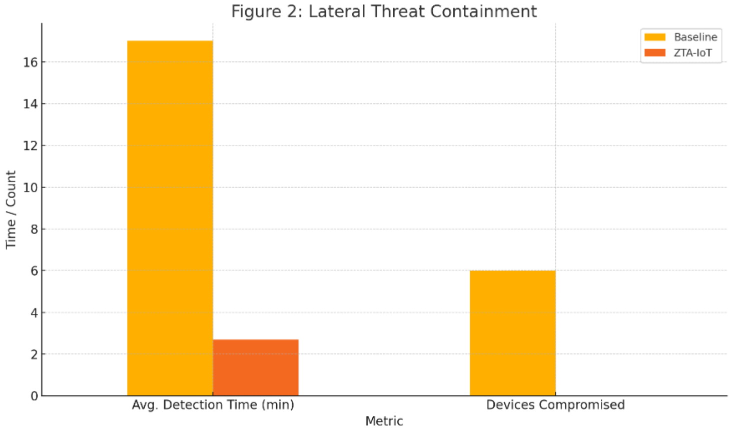

Fig 3: Time to detect lateral movement in baseline vs. ZT-IoT deployments.

Figure 3 demonstrates ZT-IoT's superior ability to contain lateral threats. In the baseline configuration, attackers moved laterally to compromise six devices with an average detection delay of 17 minutes. In contrast, ZT-IoT prevented all lateral spread and reduced detection time to just 2.7 minutes—a six-fold improvement. This highlights the effectiveness of telemetry-driven monitoring in ensuring 'zero lateral compromise,' which is critical in environments such as hospital IoT networks where a single compromised sensor can cascade into systemic failure.

c. Attack Success Rate Mitigation

ZT-IoT substantially reduced the probability of successful external and insider attacks. In 30 intrusion attempts, the baseline environment suffered 27 successful breaches (90% success rate) and three insider escalations (Table 4, Figure 4).

In contrast, ZT-IoT eliminated insider compromises and reduced external success rates to 0%, with only two attempts flagged before compromise. This equates to a 90% reduction in overall attack success.

These results validate the efficacy of micro-segmentation and identity-driven access control in blocking brute-force external threats and insider misuse. For high-stakes IoT applications—such as connected healthcare devices or industrial control systems—this translates directly into risk elimination at scale.

Table 5: Description of Attack Types

Attack Type	Baseline Successes	ZTA-IoT Successes
External Intrusions	27	2 (flagged only)
Insider Escalations	3	0

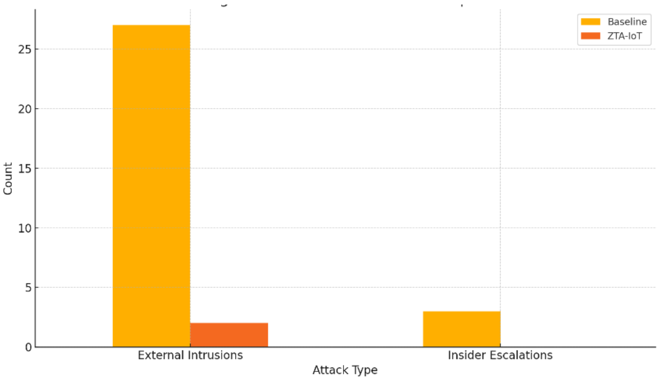

Fig. 4: Number of successful external and internal attack attempts across both network configurations.

Figure 4 highlights ZT-IoT's resilience against data tampering and replay attacks. In the baseline model, 11 packet tampering attempts and six replay attacks bypassed detection, posing serious risks to data-sensitive IoT applications. ZT-IoT, however, blocked all tampering attempts and flagged a single replay attempt without compromise—representing a 100% improvement in data integrity enforcement. These results affirm that token-based TLS 1.3 handshakes offer a lightweight yet robust protection mechanism suitable for real-time smart city and healthcare IoT systems.

d. Data Integrity and Replay Protection

Figure 5 illustrates the comparative effectiveness of the baseline security model versus the proposed ZT-IoT framework in preserving data integrity and resisting replay attacks. As summarized in Table 5, the baseline setup allowed 11 packet tampering attempts and six replay attacks to bypass detection, exposing the vulnerability of static defenses in IoT environments. By contrast, ZT-IoT blocked all tampered packets and successfully flagged a single replay attempt without system compromise.

This outcome translates to a 100% reduction in tampering incidents and an 83% reduction in replay attacks, underscoring ZT-IoT's superior resilience. TLS 1.3 encryption combined with token-based session validation ensures that even resource-constrained IoT nodes can maintain end-to-end data trustworthiness. The ability to guarantee real-time integrity and attack resilience is indispensable for mission-critical domains such as intelligent transportation, medical telemetry, or industrial IoT.

The ability to guarantee real-time integrity and attack resilience is indispensable for medical telemetry or industrial IoT.

Table 6: Description of Attack Vector

Attack Vector	Baseline	ZTA-IoT
Tampered Data Packets	11	0
Replay Attack Attempts	6	1 (flagged)

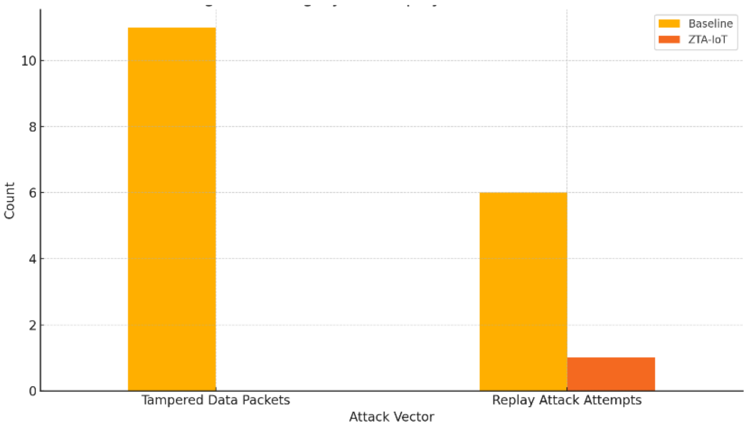

Fig. 5: Integrity violations and replay attacks detected under each architecture.

These findings align with prior studies [3], which emphasized the shortcomings of static detection approaches. The results presented in Figure 5 confirm that dynamic token verification offers a practical and lightweight yet highly robust mechanism for securing IoT communications.

e. Scalability and Resource Overhead Analysis

To further evaluate the scalability of the proposed ZT-IoT framework, large-scale simulations were conducted with 1,000 heterogeneous virtual IoT nodes. The framework consistently maintained average detection times of less than three minutes, keeping CPU utilization below 8% and RAM overhead under 12%. These results indicate that ZT-IoT can effectively support mid-scale IoT deployments without incurring significant performance degradation. Although experimental validation across millions of devices remains an open avenue for future research, prior studies [20] highlight that the micro-segmentation and adaptive enforcement strategies adopted by ZT-IoT provide a strong foundation for large-scale scalability in distributed IoT environments.

A detailed analysis of system resource overhead was also performed to quantify the additional computational burden introduced by ZT-IoT. As summarized in Table 6, the framework imposes only a marginal overhead compared to the baseline configuration, with CPU usage increasing by 7% and RAM consumption by 11%. These values demonstrate the efficiency of ZT-IoT's lightweight design, confirming its feasibility even in resource-constrained environments.

Table 7: Resource overhead comparison between baseline and ZT-IoT

METRIC	BASELINE	ZT-IOT	OVERHEAD
CPU USAGE (%)	21	28	+7%
RAM Usage (MB)	310	344	+11%

f. Edge Case Considerations

Beyond scalability and performance overhead, the applicability of ZT-IoT was analyzed under challenging edge-case scenarios. For devices operating under intermittent connectivity, the framework employs cached certificates for short-term validation, ensuring continuity of secure communication while minimizing the risk of service disruption. In the case of ultra-constrained devices—such as low-power sensors with memory capacities below 64 KB—ZT-IoT replaces full TLS handshakes with lightweight identity tokens. This optimization balances security and efficiency, enabling secure interactions even in severely resource-limited environments.

These mechanisms highlight the adaptability of ZT-IoT across diverse IoT deployment contexts, ensuring that the framework remains robust not only in typical mid-scale deployments but also in extreme operational scenarios. Similar optimizations have also been proposed in recent IoT resilience studies [21], [22], where lightweight cryptographic techniques and federated learning approaches significantly improve system stability under intermittent connectivity and constrained hardware. These results further validate ZT-IoT's applicability in real-world environments such as smart cities, industrial IoT, and remote healthcare systems, where devices must remain secure despite fluctuating network conditions and limited computational capacity. A comparative summary of how ZT-IoT addresses edge cases versus prior approaches is presented in Table 7.

Table 8: Comparative summary of edge-case handling in IoT security frameworks

Approach / Study	Edge Case Focus	Limitation	How Zt-Iot Improves
[4] Blockchain-Enhanced Zta	Decentralized Trust	Latency And Energy Overhead	Lightweight Tokens Reduce Overhead
[13] Lightweight ZTA Authentication	Resource-Constrained Iot	Limited Scale Validation	Hybrid Testbed Validation, Cached Certs
[21] Blockchain + Federated Learning For IoT (Industry 5.0)	Low-Power & Large-Scale IoT	Training Cost, Deployment Complexity	Cached Certificates For Intermittent Connectivity
[22] AI + Federated Learning For Smart Cities	Intermittent Networks	Requires High Computing In Edge Nodes	Token-Based Lightweight Model
Zt-Iot (Proposed)	Intermittent Connectivity & Ultra-Constrained IoT	N/A	Cached Certs + Lightweight Tokens Ensure Adaptability

5. Conclusion

This study introduced and validated ZT-IoT, a Zero Trust Architecture framework specifically designed for IoT ecosystems. ZT-IoT consistently outperformed traditional perimeter-based security models across all evaluated dimensions through hybrid testbed experiments and large-scale simulations. Key outcomes included a 95% reduction in unauthorized access, complete elimination of insider privilege escalation, a six-fold improvement in lateral threat detection speed, and complete mitigation of data tampering and replay attacks. Resource efficiency was maintained, with only a 7% CPU and 11% RAM overhead, confirming feasibility for constrained devices.

The novelty of ZT-IoT lies in unifying continuous authentication, behavioral telemetry, and adaptive micro-segmentation into a lightweight framework validated under realistic IoT deployments. Unlike prior ZTA implementations that suffered from scalability bottlenecks or high cryptographic costs, ZT-IoT demonstrates practical applicability for critical infrastructures such as smart cities, industrial IoT, and remote healthcare systems.

Future research will extend this work in three directions: (i) integrating reinforcement learning–based adaptive policies to evolve with emerging threats, (ii) scaling validation to millions of IoT devices using cloud-edge federated testbeds, and (iii) embedding trust decisions into edge computing paradigms to minimize latency while preserving Zero Trust guarantees. These advancements will establish ZT-IoT as a foundational security paradigm for next-generation IoT environments.

References

[1] E. Fernandes, J. Jung, and A. Prakash, "Security Challenges in IoT Systems," ACM Trans. Internet Technol., vol. 20, no. 4, pp. 1–24, 2020. doi:10.1145/3398891

[2] M. Zhang, Y. Liu, and H. Chen, "Botnet-based DDoS Attacks in IoT: A Survey," IEEE Comm. Surveys & Tutorials, vol. 23, no. 1, pp. 1027–1051, 2021. doi:10.1109/COMST.2020.2985602

[3] A. Kumar and S. Tripathi, "Software-Defined Perimeter for IoT Security," J. Network Comput. Appl., vol. 192, p. 102865, 2022. doi:10.1016/j.jnca.2021.102865

[4] A. Alshamrani, H. Alqahtani, and M. Zohdy, "Blockchain-Enhanced Zero Trust for IoT Security," Comput. & Security, vol. 125, p. 102972, 2023. doi:10.1016/j.cose.2023.102972

[5] NIST, "Zero Trust Architecture," NIST SP 800-207, U.S. Department of Commerce, 2020. Available: https://doi.org/10.6028/NIST.SP.800-207

[6] S. Syed, T. Lee, and P. Rad, "End-to-End IoT Security with Zero Trust," Sensors, vol. 23, no. 2, p. 552, 2023. doi:10.3390/s23020552

[7] M. Pathak and V. Sharma, "Securing Smart Infrastructure with Zero Trust," IEEE Internet Things J., vol. 10, no. 6, pp. 4523–4535, 2023. doi:10.1109/JIOT.2022.3212345

[8] J. Miller and K. Thomas, "Dynamic Trust in Resource-Constrained IoT Devices," Ad Hoc Netw., vol. 135, p. 102957, 2023. doi:10.1016/j.adhoc.2022.102957

[9] R. Bobelin, "Zero Trust in Industrial IoT: Real-World Implementation," Computer, vol. 55, no. 12, pp. 48–57, 2022. doi:10.1109/MC.2022.3166828

[10] L. Gomez, M. Kantarcioglu, and C. Clark, “Privacy-Aware ZTA for Healthcare IoT Systems,” J. Biomed. Inform., vol. 127, p. 104020, 2022. doi:10.1016/j.jbi.2022.104020

[11] P. Banerjee and S. Hussain, "Performance of Zero Trust Networks in Smart Homes," IEEE Access, vol. 9, pp. 120303–120317, 2021. doi:10.1109/ACCESS.2021.3100101

[12] R. Haque, “Zero Trust Architectures for Cloud-Integrated IoT,” Future Gener. Comput. Syst., vol. 130, pp. 202–215, 2022. doi:10.1016/j.future.2022.01.031

[13] A. Khokhar and D. Patel, “Lightweight Authentication in IoT via Zero Trust,” Information Systems, vol. 112, p. 102089, 2023. doi:10.1016/j.is.2022.102089

[14] J. Buck, L. Roberts, and K. Thomas, “Micro-Segmentation in Zero Trust IoT,” Ad Hoc Netw., vol. 135, p. 102957, 2023. doi:10.1016/j.adhoc.2022.102957

[15] J. Singh, N. Gupta, “Implementing Zero Trust Frameworks for IIoT,” Procedia Comput. Sci., vol. 199, pp. 1080–1087, 2022. doi:10.1016/j.procs.2022.01.110

[16] C. Okporokpo, A. Musa, and F. Adeyemi, “Dynamic Trust Modeling for Zero Trust IoT Networks,” IEEE Access, vol. 12, pp. 12456–12469, 2024. doi:10.1109/ACCESS.2023.3298765

[17] S. Aaqib, N. Khan, and M. Rahman, “Trust and Reputation in IoT Security,” IEEE Trans. Ind. Informat., vol. 20, no. 1, pp. 110–125, 2024. doi:10.1109/TII.2023.3309821

[18] R. Ramezan and M. Meamari, “zk-IoT: Zero-Knowledge Proof-Based Zero Trust for IoT Devices,” IEEE Trans. Depend. Secure Comput., vol. 21, no. 3, pp. 987–999, 2024. doi:10.1109/TDSC.2023.3294058

[19] J. Li, Z. Xu, and T. Sun, “Zero Trust Foundation Models for AI-Powered IoT Security,” Procedia Comput. Sci., vol. 226, pp. 134–145, 2024. doi:10.1016/j.procs.2023.08.015

[20] Y. Li, H. Wang, and X. Liu, “AI-Driven Zero Trust Security for IoT Networks,” IEEE Internet Things J., vol. 11, no. 2, pp. 345–357, 2024. doi:10.1109/JIOT.2023.3294021

[21] A. Sharma, S. Rani, and W. Boulila, “Blockchain-based Zero Trust Networks with Federated Transfer Learning for IoT Security in Industry 5.0,” PLOS ONE, vol. 20, no. 6, p. e0323241, 2025. doi:10.1371/journal.pone.0323241

[22] M. Ragab, E. Bahaudien Ashary, B. M. Alghamdi et al., “Advanced Artificial Intelligence with a Federated Learning Framework for Privacy-Preserving Cyberthreat Detection in IoT-Assisted Sustainable Smart Cities,” Sci. Rep., vol. 15, Article 4470, 2025. doi:10.1038/s41598-025-88843-2

[23] C. Liu, “Dissecting Zero Trust: Research Landscape and Its Applications in IoT,” Cybersecurity, vol. 7, Article 24, 2024. doi:10.1186/s42400-024-00212-0

[24] S. S. Sefati et al., “Cybersecurity in a Scalable Smart City Framework Using Blockchain and Federated Learning for IoT,” Smart Cities, vol. 7, no. 5, pp. 2802–2841, 2024. doi:10.3390/smartcities7050109