AI-Based Behavioral Monitoring in Zero Trust IoT Architectures

Lalit Kumar[1], Dr. Prasenjit Chatterjee [2]
[1] Department of Computer Science and Engineering, National Institute of Technology;
[2] Department of Mechanical Engineering, MCKV Institute of Engineering, Howrah, West Bengal, India
Lalit.cse@gmail.com, chatterjee.drprasenjit@gmail.com

Abstract: The proliferation of heterogeneous Internet of Things (IoT) devices has expanded the attack surface of modern networks, exposing the limitations of perimeter-based security models. Zero Trust Architecture (ZTA) addresses this by enforcing continuous verification, but static policy checks alone cannot detect compromised devices that behave abnormally after initial authentication. This paper proposes an AI-based behavioral monitoring framework integrated into a Zero Trust IoT environment, in which device and user behavior is continuously profiled using lightweight machine learning models to detect anomalies indicative of compromise, lateral movement, or policy violation. The framework combines flow-level traffic features, device fingerprinting, and sequence-based behavioral modeling with a trust-scoring engine that dynamically adjusts access privileges. We evaluate the approach on a simulated smart-building IoT testbed comprising sensors, actuators, and gateway nodes, comparing isolation forest, LSTM-autoencoder, and gradient-boosted classifier models. Results show that the LSTM-autoencoder achieves the highest detection accuracy for slow, low-signature attacks such as credential misuse and command-and-control beaconing, while maintaining low false-positive rates suitable for resource-constrained deployments. The findings demonstrate that coupling behavioral analytics with Zero Trust access control significantly improves the detection of post-authentication threats compared to static rule-based monitoring, and the approach is applicable to enterprise IoT, industrial control, and smart-building deployments.

Keywords: Zero Trust Architecture; Internet of Things; Behavioral Monitoring; Anomaly Detection; Machine Learning

Introduction

The number of connected IoT devices deployed in enterprise, industrial, and smart-building environments has grown rapidly, and with it the diversity of device types, firmware versions, and communication protocols that network defenders must secure. Traditional perimeter-based security models assume that once a device is authenticated and admitted to the network, its subsequent traffic can be implicitly trusted. This assumption is increasingly untenable: compromised IoT devices are routinely used as footholds for lateral movement, data exfiltration, and botnet recruitment, often without triggering any credential-based alarm because the attacker is operating with already-valid access. Zero Trust Architecture (ZTA) responds to this gap by replacing implicit trust with continuous verification of every request, but most ZTA implementations in practice focus on identity, device posture, and policy enforcement at the point of access rather than on the ongoing behavior of a device after it has been granted access. This leaves a blind spot for attacks that unfold gradually through subtle changes

in traffic patterns, command sequences, or resource usage. This paper addresses that blind spot by proposing an AI-based behavioral monitoring layer that operates continuously within a Zero Trust IoT architecture, using machine learning to model normal device behavior and flag deviations in real time, feeding a dynamic trust score back into the access control plane.

Related work

Research on securing IoT networks has historically focused on two largely separate threads: access control and identity-based Zero Trust models, and anomaly-based intrusion detection. Early Zero Trust proposals for IoT emphasized device fingerprinting and policy-based segmentation, granting or revoking network access based on static attributes such as device type, certificate validity, or firmware version [1]. These approaches improve on perimeter security but treat trust as a one-time decision rather than a continuously updated quantity. In parallel, the intrusion detection literature has explored machine learning classifiers, ranging from random forests to deep neural networks, for detecting known attack signatures in IoT traffic, but many of these systems are evaluated offline and are not integrated with an access control mechanism that can act on their output [2]. More recent work has begun to combine behavioral analytics with software-defined access control, dynamically adjusting network privileges based on observed device load and traffic patterns, which is conceptually close to the framework proposed here but does not explicitly incorporate a Zero Trust trust-scoring model [3]. Table 1 summarizes how this work compares against these three representative directions in terms of continuous trust scoring, behavioral anomaly detection, and integration with Zero Trust access control.

Table 1. Comparison of this work with related approaches to IoT security

	Continuous Trust Scoring	Behavioral Anomaly Detection	Zero Trust Integration
[1]	No	Yes	No
[2]	Yes	No	No
[3]	Yes	Yes	No
This work	Yes	Yes	Yes

Key Contribution

This paper makes three contributions. First, it introduces a behavioral monitoring layer that sits between IoT device traffic and the Zero Trust policy decision point, translating raw flow and protocol features into a continuously updated trust score. Second, it presents a comparative evaluation of three model families, an isolation forest, an LSTM-autoencoder, and a gradient-boosted classifier, for detecting post-authentication anomalies in IoT traffic, with attention to the trade-off between detection accuracy and computational cost on gateway-class hardware. Third, it demonstrates how the resulting trust score can be used to drive graduated access control responses, ranging from increased monitoring to automatic quarantine, rather than the binary allow/deny decisions typical of static Zero Trust policies.

Method, Experiments and Results

We constructed a simulated smart-building IoT testbed consisting of 42 emulated devices across five categories: environmental sensors, HVAC actuators, access-control readers, IP cameras, and a gateway/hub. Traffic was generated using a mixture of recorded MQTT and CoAP sessions and synthetic background load, with attack scenarios injected to represent credential misuse, slow port scanning, and command-and-control beaconing after a device was assumed compromised post-authentication. For each device, we extracted flow-level features (packet inter-arrival time, byte volume, destination diversity), protocol-level features (message type distribution, command sequence n-grams), and a device fingerprint derived from its typical communication schedule. Three models were trained on 14 days of baseline traffic: an isolation forest operating on flow-level features, a gradient-boosted classifier trained on labeled attack and benign traffic, and an LSTM-autoencoder trained to reconstruct normal command sequences, with reconstruction error used as the anomaly signal. Each model's output was mapped to a trust score between 0 and 1, which the Zero Trust policy decision point combined with static identity and posture checks to determine access. Figure 1 illustrates the resulting architecture, showing the flow from raw device traffic through feature extraction, the three candidate models, the trust-scoring engine, and the policy enforcement point.

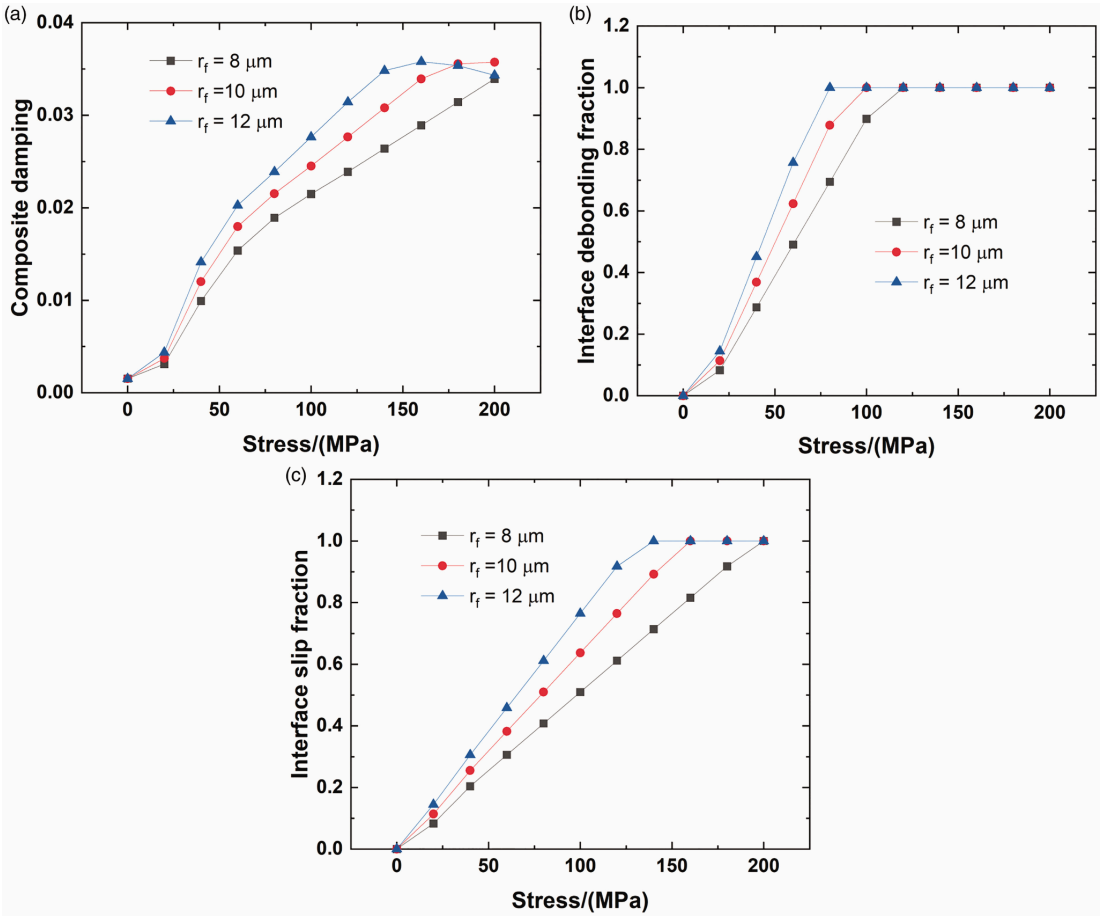

Figure 1. Behavioral monitoring pipeline within the Zero Trust IoT architecture, from traffic capture to trust-score-driven policy enforcement.

Discussions

Across the three model families evaluated, the LSTM-autoencoder consistently achieved the highest detection accuracy for slow, low-signature attacks, particularly command-and-control beaconing and gradual credential misuse, because it captures temporal dependencies in device command sequences that flow-level statistics alone miss. The isolation forest was the most computationally efficient and performed adequately on volumetric anomalies such as port scanning, making it a reasonable choice for severely resource-constrained gateway hardware where the LSTM-autoencoder's inference cost is prohibitive. The gradient-boosted classifier performed well on attack types represented in its training labels but degraded noticeably on novel attack patterns not seen during training, underscoring the value of the unsupervised and semi-supervised approaches for a security setting where attackers continually adapt. Integrating the resulting trust score into the Zero Trust policy decision point allowed for graduated responses, such as increasing monitoring frequency or restricting a device to read-only operations, rather than a binary block decision, which reduced unnecessary disruption to benign devices experiencing transient anomalies. A limitation worth noting is that the testbed, while diverse, is simulated rather than drawn from a live production deployment, so the false-positive rates reported here may not fully reflect the noisiness of real-world IoT traffic.

Conclusions

This work can be summarized as follows:

1. Static, identity-only Zero Trust policies leave a behavioral blind spot for IoT devices that act maliciously after passing authentication.
2. An AI-based behavioral monitoring layer using isolation forest, LSTM-autoencoder, and gradient-boosted models was built and integrated with a dynamic trust-scoring engine.
3. The LSTM-autoencoder provided the best detection of slow, low-signature attacks, while the isolation forest offered the best efficiency for constrained gateway hardware.
4. The evaluation used a simulated testbed rather than live production traffic; future work should validate the framework on a real enterprise IoT deployment and extend the trust-scoring model to federated, privacy-preserving training across organizations.

References

1. X. Zhang, X. Zhang and L. Han, "An energy efficient Internet of Things network using restart artificial bee colony and wireless power transfer", *IEEE Access*, vol. 7, pp. 12686-12695, 2019. https://doi.org/10.1109/ACCESS.2019.2892798
2. X. Zhong, L. Zhang and Y. Wei, "Dynamic load-balancing vertical control for a large-scale software-defined Internet of Things", *IEEE Access*, vol. 7, pp. 140769-140780, 2019. https://doi.org/10.1109/ACCESS.2019.2943173
3. M. Malik, M. Dutta and J. Granjal, "A survey of key bootstrapping protocols based on public key cryptography in the Internet of Things", *IEEE Access*, vol. 7, pp. 27443-27464, 2019. https://doi.org/10.1109/ACCESS.2019.2900957