EdgeGuardian: A Lightweight AI-Based IoT Security Gateway for Real-Time Anomaly Detection

Shantanu Rajeshirke
Department of information and Cyber Security .
Guru Nanak Khalsa College, Matunga.
Email: shantanuraje02@gmail.com

ABSTRACT

The Internet of Things (IoT) is growing very quickly, and a large number of devices are now connected to the internet in homes, hospitals, factories, and smart cities. This connectivity helps automate tasks and improves efficiency, but it also creates new security risks. Many IoT devices are low-power, have limited memory, and do not include strong security features. As a result, they can be easily targeted by attackers through unauthorized access, botnets, or large-scale attacks that can disrupt services. Traditional intrusion detection systems (IDS) mainly rely on signatures of known attacks, so they struggle to detect new or unknown threats. In addition, cloud-based security solutions can introduce extra latency, increase bandwidth usage, and raise privacy concerns because sensitive data must be sent to external servers.

This paper presents EdgeGuardian, a lightweight AI-based security gateway for IoT environments. EdgeGuardian runs on a Raspberry Pi and monitors network traffic at the edge of the network. It captures packet-level traffic, extracts features, and uses an unsupervised machine learning model to detect abnormal behaviour. The model is trained offline on a more powerful machine and then deployed on the Raspberry Pi for real-time inference. This design keeps the system simple and energy-efficient, which makes it suitable for resource-constrained devices.

To evaluate EdgeGuardian, a combination of public IoT security datasets and traffic from a controlled testbed was used. The system was tested on different attack scenarios, including port scanning and denial-of-service traffic. Performance was measured using metrics such as accuracy, precision, recall, F1-score, and false positive rate, as well as CPU, memory, and latency on the Raspberry Pi. Experimental results show that EdgeGuardian achieves an overall detection accuracy of about 88.4%, with balanced precision and recall, while keeping resource usage moderate.

These findings suggest that EdgeGuardian can detect a high percentage of threats without overloading the device or adding significant delay. Because it is lightweight, low-cost, and edge-based, EdgeGuardian is a practical option for improving the security of real-world IoT deployments, especially where traditional heavy IDS or cloud-only solutions are difficult to use.

KEYWORDS

IoT Security, Intrusion Detection System, Edge Computing, Machine Learning, Raspberry Pi, Anomaly Detection.

1. INTRODUCTION

The Internet of Things (IoT) has become an important part of modern technology and is now widely used in smart homes, healthcare, industry, and smart cities. IoT devices continuously collect and share data, which helps with automation, monitoring, and data-driven decision-making. This has made many everyday tasks easier and systems more efficient. At the same time, the rapid increase in the number of connected devices

has also increased the risk of cyber attacks. Each device can become a potential entry point for attackers, which makes the whole network more vulnerable.

A major challenge is that many IoT devices are designed to be cheap and energy-efficient rather than secure. They usually have limited processing power, memory, and storage, and often run simple operating systems that are rarely updated. Because of these limitations, they cannot run heavy security software and are more exposed to threats. Typical attacks include distributed denial-of-service (DDoS), where devices are used to flood a target with traffic; botnet infections, where many devices are controlled by an attacker; and unauthorized access, where attackers exploit weak passwords or misconfigurations.

Intrusion detection systems (IDS) are commonly used to monitor network traffic and identify suspicious activities. Traditional IDS solutions, such as Snort, are mainly signature-based. They compare observed traffic with known attack patterns and raise an alert when there is a match. These systems can be very effective for detecting threats that have been seen before. However, they are less effective against new, unknown, or rapidly evolving attacks. In IoT environments, where new vulnerabilities and attack methods appear frequently, relying only on signature-based detection is not sufficient.

To improve detection of new threats, researchers have explored anomaly-based IDS using machine learning. In these systems, models learn what normal traffic looks like and then flag deviations as potential anomalies. Supervised learning methods and deep learning models have shown good performance on IoT security datasets and can capture complex patterns in network behaviour. However, most of these models need large labelled datasets and significant computational resources. Collecting and labelling IoT traffic is difficult, and running heavy models directly on IoT devices or small gateways is usually not practical.

Edge computing has been proposed as a way to handle these limitations. Instead of sending all traffic to the cloud for analysis, edge computing processes data closer to where it is generated, for example on local gateways or edge servers. This reduces latency, saves bandwidth, and can protect privacy by keeping sensitive data within the local network. For IoT security, performing intrusion detection at the edge means that threats can be detected and handled faster, without depending fully on remote cloud infrastructure.

Even with these advances, there are still gaps in existing solutions. Many proposed systems focus mainly on improving accuracy on benchmark datasets, but they do not fully address latency, scalability, and resource constraints on low-cost hardware. In addition, a lot of work is based on flow-level analysis, which uses aggregated traffic statistics and may miss subtle packet-level anomalies. There is a need for an IDS that can work in real time, use limited resources, and still detect both known and unknown threats in diverse IoT traffic.

To address these challenges, this paper proposes EdgeGuardian, an efficient IoT security gateway for real-time anomaly detection at the network edge. EdgeGuardian uses a Raspberry Pi as a low-cost, resource-constrained gateway. It monitors packet-level network traffic, extracts meaningful features, and applies an unsupervised machine learning model to detect anomalies. The system follows a decoupled architecture, where model training is done offline on a more powerful system, and the trained model is then deployed on the Raspberry Pi for real-time detection. This design helps to balance detection performance with resource efficiency.

The main goal of this work is to develop a system that can detect malicious activities in IoT networks without requiring heavy hardware or large amounts of labelled data. The system is evaluated using a combination of public datasets and traffic from a controlled IoT test environment. While the results are promising, the system is still evaluated on a limited scale and may need further improvements for very large or highly diverse deployments. Overall, EdgeGuardian aims to provide a practical and scalable step towards improving IoT security by offering an adaptive, edge-based intrusion detection solution.

2. LITERATURE REVIEW

The rapid growth of the Internet of Things has encouraged many researchers to focus on intrusion detection systems that can address the unique security challenges of IoT networks. Early IDS approaches were mainly signature-based, where known attack patterns are stored as rules. Tools such as Snort have been widely used in traditional networks and can effectively detect many known threats. However, these systems depend on prior knowledge of attacks. They struggle to detect zero-day attacks or new variations of existing threats because no signatures exist for them yet. In dynamic IoT environments, where traffic is diverse and constantly changing, relying only on signatures is therefore not sufficient.

To overcome this limitation, many studies have looked at anomaly-based intrusion detection using machine learning. Supervised learning algorithms such as Random Forest, Support Vector Machines (SVM), and Decision Trees have been trained on IoT traffic datasets and have achieved good accuracy in detecting known attacks. Deep learning approaches, including Convolutional Neural Networks (CNN) and Long Short-Term Memory (LSTM) networks, have also been used. These models can automatically learn complex temporal and spatial patterns in network data and often provide better performance than traditional methods. However, both supervised and deep learning methods usually require large labeled datasets and significant computational power, which restricts their use on resource-limited IoT devices and gateways.

Edge computing has been introduced as a way to bring computation closer to IoT devices. Instead of sending all data to centralized cloud servers, edge-based IDS perform analysis at or near the data source. Several works have shown that it is possible to deploy IDS on low-cost devices such as Raspberry Pi and still maintain an acceptable trade-off between detection accuracy and resource usage. These studies highlight that lightweight, distributed security mechanisms can be effective in IoT networks, especially when properly optimized for the constraints of the hardware.

Unsupervised and semi-supervised learning techniques have also attracted attention because they can reduce the dependence on labelled data. Techniques such as autoencoders, clustering algorithms, and one-class classifiers (including Isolation Forest) can model normal traffic and then detect anomalies as deviations from this model. These methods are particularly useful in IoT scenarios where traffic patterns evolve over time and labelled examples of every possible attack are not available. By focusing on behaviour rather than fixed signatures, they can detect unknown or emerging threats more effectively.

Despite this progress, existing IoT intrusion detection systems still face several challenges. Many solutions are designed and evaluated with a strong emphasis on detection accuracy, but they often pay less attention to latency, scalability, and energy consumption. In real deployments, these factors are just as important as accuracy because devices have strict resource limits and networks must continue to operate smoothly. Furthermore, many systems rely on flow-level features, which may not capture subtle packet-level anomalies that can signal the early stages of an attack. There is also a lack of integrated systems that combine real-time edge deployment, lightweight models, and unsupervised learning in a practical way.

To address these gaps, this research proposes EdgeGuardian, an AI-based IoT security gateway. EdgeGuardian is designed to run on resource-constrained devices while still providing effective threat detection. It focuses on packet-level analysis, uses an unsupervised learning model to reduce dependence on labeled data, and follows an edge-based architecture to support real-time operation. This combination aims to provide a practical, deployable solution that fits the specific needs and constraints of IoT environments.

3. METHODOLOGY AND PROPOSED SYSTEM

The Internet of Things is growing fast and it is changing how devices, people and services work together in many areas like healthcare, smart homes and industrial automation. Now we have lots of devices that can sense, process and exchange data all the time. This has enabled things like remote patient monitoring and automated manufacturing. It has also made it easier for people to attack these devices. Many Internet of Things devices do not have security because they are made to be cheap and use less battery power. They often use operating systems and do not get updated very often.

The Internet of Things uses different ways to communicate, like Wi-Fi and Bluetooth which makes it hard to keep everything secure. Because of this Internet of Things networks are open to kinds of cyber threats. Some common threats include attacks that flood services with traffic and attempts to gain access to devices using passwords. If a device is compromised it can be used to spy on people or disrupt services. This is why we need intrusion detection systems to keep Internet of Things environments safe.

Old intrusion detection systems were made for traditional networks. They use rules to match traffic against known attack patterns. These systems are not good enough for Internet of Things because they can only detect known threats and need to be updated all the time. In Internet of Things, new devices and attack methods appear all the time so we need systems that can detect threats.

Some people are working on anomaly-based intrusion detection using machine learning. These systems learn what normal traffic looks like and flag anything that's different. They can be trained on datasets that include both bad traffic. These systems need a lot of labeled data, which can be hard to get.

Deep learning models can also be used for intrusion detection. They can learn patterns in traffic data. Detect subtle threats. They need a lot of computing power and memory which can be a problem for Internet of Things devices.

Edge computing is a way to do things that can help with this problem. Of sending all data to the cloud for analysis edge computing does the analysis on the device itself or on a nearby gateway. This can reduce latency and bandwidth usage. Improve privacy. It is especially important for Internet of Things applications that need responses, like industrial control systems.

At the time people are working on unsupervised and semi-supervised learning methods that do not need as much labeled data. These methods can detect anomalies in traffic without needing labels for each new threat. With all these advances there are still many challenges. Many solutions focus on detecting threats but they do not consider things like processing latency and energy consumption. In practice a system that is a little more accurate but uses resources may not be practical for Internet of Things devices.

That is why we need intrusion detection systems that can work in time use limited resources and detect both known and new threats. The system should also be flexible enough to handle kinds of Internet of Things traffic and robust enough to work in environments where labeled data is scarce.

To address these needs, we are introducing EdgeGuardian, an anomaly detection system for Internet of Things security. EdgeGuardian uses an unsupervised learning model on a resource-constrained device to monitor traffic and detect anomalies. It combines edge computing with anomaly detection to provide real-time protection for Internet of Things networks. EdgeGuardian can detect evolving threats that traditional systems might miss.

The Internet of Things is a part of our lives now and we need to make sure it is secure. EdgeGuardian is a step in that direction. It can help keep our Internet of Things devices and networks from threats. The Internet of Things will continue to grow. We need to make sure we have good security systems in place to protect it. EdgeGuardian is a start but we need to keep working to make sure our Internet of Things is safe and secure.

The Internet of Things devices are used in areas, including healthcare and industrial automation. These devices are often connected to the Internet. Can be accessed remotely. The Internet of Things devices are also used in homes, where they can control things like temperature and lighting.. All these devices are open to cyber threats and we need to make sure they are secure.

We can use EdgeGuardian to protect our Internet of Things devices and networks. EdgeGuardian is a system that can run on resource-constrained devices. It can detect anomalies in traffic. Alert us to potential threats. The Internet of Things is a system and we need to make sure we have good security systems in place to protect it. EdgeGuardian is a part of that.

In conclusion the Internet of Things is a part of our lives and we need to make sure it is secure. We need to use systems, like EdgeGuardian to protect our devices and networks from cyber threats. The Internet of Things will continue to grow. We need to make sure we have good security systems in place to protect it. EdgeGuardian is a start but we need to keep working to make sure our Internet of Things is safe and secure. The Internet of Things devices are used in areas and we need to make sure they are all secure. We can use EdgeGuardian to protect our Internet of Things devices and networks.

4. SYSTEM ARCHITECTURE

Figure 1. System architecture of Ed

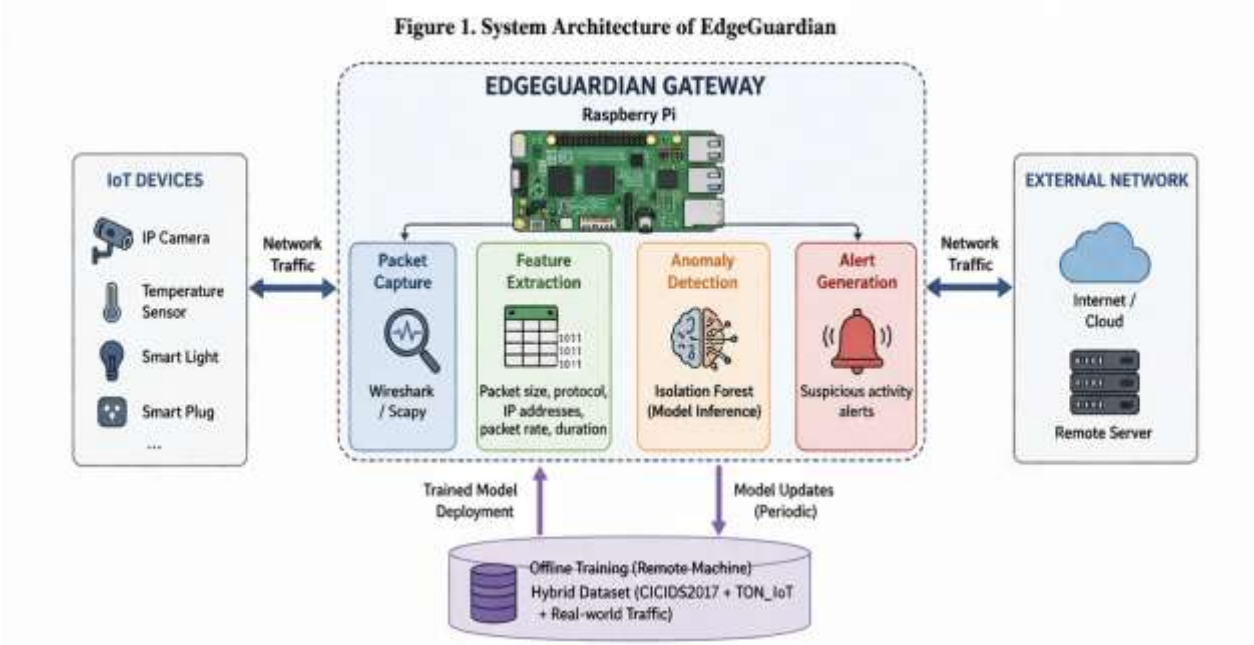

geGuardian

The EdgeGuardian system is made up of a bunch of things including IoT devices, a Raspberry Pi gateway, a module that captures packets, a part that pulls out features, a thing that finds anomalies and a system that sends out alerts. The Raspberry Pi gateway is like a man that catches all the network traffic and looks at it in real time.

The way the EdgeGuardian system is built is really flexible. This means the EdgeGuardian system can be made bigger or smaller as needed. To capture packets the EdgeGuardian system uses tools, like Wireshark and Scapy. Then the EdgeGuardian system takes all the data and turns it into something that a machine can understand so the EdgeGuardian system can use it to learn and find things that are not normal.

Figure 2. Working Flow of EdgeGuardian

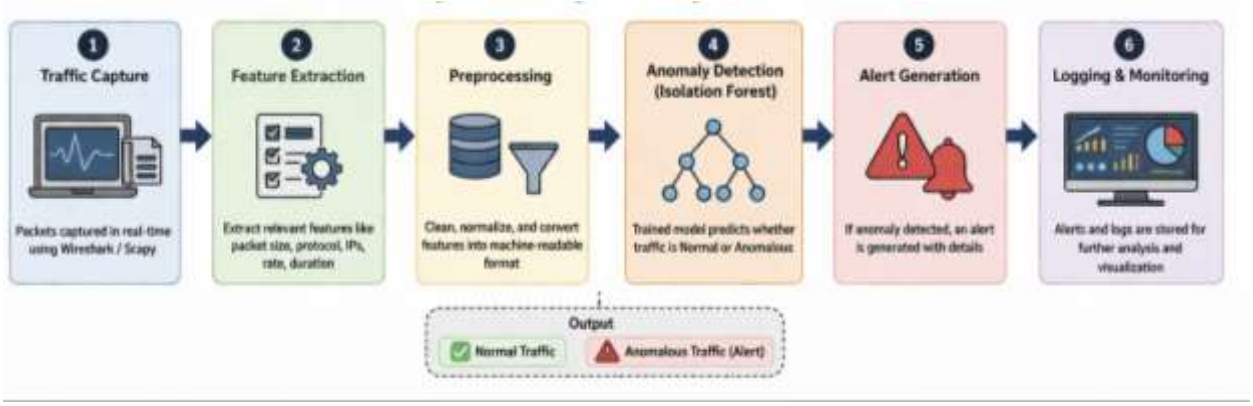

5. RESULTS AND DISCUSSION

The EdgeGuardian system is really good at finding behavior in Internet of Things network traffic. We tested it using a mix of datasets and traffic from a controlled IoT test environment. In this setup we mixed device communication with fake attack traffic. This included things like port scanning and denial-of-service patterns which're common threats to IoT devices.

We chose this mix of local traffic to get a good variety of scenarios and make the test more realistic. The main goal of the test was to see how well the EdgeGuardian system can tell traffic from malicious traffic. We also wanted to see if it can run efficiently on a Raspberry Pi without using much of the devices resources. Both of these things are important because a system that is accurate but too heavy to run on the gateway is not useful and a system that is lightweight but inaccurate will not provide protection.

The EdgeGuardian system does a job of detecting weird traffic. Across datasets and attack scenarios it gets it right about 88.4% of the time. This means that nine out of ten traffic instances are correctly classified as normal or anomalous. This level of accuracy is really encouraging especially since the EdgeGuardian system is running on a device with resources.

If we look closer at the numbers the EdgeGuardian system is 87.2% precise and 89.1% good at recalling traffic. Precision means that most of the time when the system flags something as it actually is weird. Recall means that the system catches most of the attacks. The F1-score, which combines precision and recall is 88.1%. This shows that the system does a job of balancing these two things.

The false positive rate is 8.3%, which means that a small fraction of normal traffic is incorrectly flagged as weird. While it would be better to have a false positive rate this level is acceptable in many security situations. In a real-world IoT network with thousands of flows per hour an 8.3% false positive rate would still result in a number of alerts.

The EdgeGuardian system also runs well on the Raspberry Pi. The CPU usage is usually 40% to 60% when the anomaly detection module is active. This means that the system uses a reasonable amount of the devices processing power. Memory usage is also important on devices and the EdgeGuardian system uses about 300 to 450 MB of RAM. This is an amount of memory and it leaves room for other services or monitoring tools to run alongside the EdgeGuardian system.

Latency is also important for edge-based intrusion detection. In our tests the average latency was 1.5 seconds. This means that from the moment a packet is observed to the moment the classification decision is made it takes 1.5 seconds. For IoT use cases, such as smart home management or industrial sensing this delay is acceptable.

Compared to cloud-based intrusion detection solutions the EdgeGuardian system may not be as accurate. However it provides a balance between detection capability and deploy ability in resource-constrained

environments. The system is effective enough to be useful in practice while still respecting the limitations of low-cost hardware.

One of the advantages of running detection at the network edge is that it reduces the dependence on infrastructure. Since traffic does not need to be sent to a cloud for analysis bandwidth usage is lower and sensitive data can stay inside the local network. This is especially important in healthcare settings where privacy and regulatory requirements may restrict the movement of raw traffic data.

Overall our tests show that the EdgeGuardian system is an practical solution for intrusion detection in IoT networks. It offers a level of accuracy maintains a manageable false positive rate and operates within the CPU and memory limits of a Raspberry Pi while providing near real-time detection. These characteristics make it suitable for small to medium-scale IoT deployments such, as homes, laboratories or focused industrial installations.

6. Conclusion

This paper is about EdgeGuardian. EdgeGuardian is a security system for internet connected devices. It helps keep these devices safe from harm. Lots of people are using internet connected devices now. This is a problem because these devices are not very good at keeping themselves safe.

They do not have the power to run strong security systems. So we need to find a way to keep them safe. Old security systems are not good enough. They cannot stop threats. EdgeGuardian uses computer programs to watch what is happening on the network. It can learn what is normal and what is not. This helps it find threats.

We used a Raspberry Pi to test EdgeGuardian. This shows that we can use security systems even on devices that are not very powerful. We made the system in a way so it can work quickly and efficiently. We used a mix of real internet traffic to test the system. This made the system better at finding threats.

When we tested EdgeGuardian it worked well. It found threats. Did not use too much power. It was also fast. EdgeGuardian is not perfect. It is good enough to use in real life. It is better than security systems because it works quickly and keeps data private.

Edgeguardian is not perfect. We only tested it in an environment. It might not work well in a bigger environment. We can make it better by using computer programs. We can also make it learn things so it can find new threats. We can add features to make it more useful.

So EdgeGuardian is a system for keeping internet connected devices safe. It is cheap. Works well. It can help us make security systems that work on devices that are not very powerful.

7. References

[1] M. Roesch, “Snort: Lightweight Intrusion Detection for Networks,” USENIX, 1999.
[2] G. Apruzzese et al., “Machine Learning for Intrusion Detection in IoT Networks,” IEEE, 2022.
[3] I. Almomani et al., “IoT Intrusion Detection Using Machine Learning,” FGCS, 2023.
[4] N. Moustafa et al., “Deep Learning Approaches for Intrusion Detection,” IEEE Access, 2021.
[5] W. Shi et al., “Edge Computing: Vision and Challenges,” IEEE IoT Journal, 2016.
[6] A. Diro et al., “Distributed Attack Detection Using Deep Learning in IoT,” IEEE Access, 2018.
[7] F. T. Liu et al., “Isolation Forest,” IEEE ICDM, 2008.
[8] H. Hindy et al., “A Taxonomy of Intrusion Detection Systems,” IEEE Access, 2020.