Article

Edge AI Bridge: A Micro-Layer Intrusion Detection Architecture for Smart-City IoT Networks

Sethu Subramanian N. , Prabu P. , Kurunandan Jain * , Prabhakar Krishnan

Posted Date: 9 February 2026

doi: 10.20944/preprints202602.0606.v1

Keywords: Edge AI; IoT security; intrusion detection system; smart cities; edge computing; anomaly detection; embedded AI

Article

Edge AI Bridge: A Micro-Layer Intrusion Detection Architecture for Smart-City IoT Networks

Sethu Subramanian N. [1], **Prabu P.** [2], **Kurunandan Jain** [1,*] **and Prabhakar Krishnan** [1]

[1] Center for Cybersecurity Systems and Networks, Amrita Vishwa Vidyapeetham, Amritapuri Campus, Kollam, Kerala, India
[2] Information Systems Department, College of Computer and Information Sciences, Imam Mohammad Ibn Saud Islamic University, Riyadh, Saudi Arabia
* Correspondence: kurunandanj@am.amrita.edu

Abstract

Smart-city IoT ecosystems depend on a large number of devices with limited resources, which often lack built-in security mechanisms. While traditional cloud-based or gateway-centric intrusion detection systems (IDS) offer essential security, they are still characterized by high detection latency, considerable bandwidth demand, and lack of precise monitoring of single device actions. This work presents and experimentally evaluates a novel micro-layer intrusion detection architecture, termed the Edge AI Bridge as a new micro-computing security layer that is positioned between IoT devices and the gateway to enable early-stage threat interception. The proposed architecture incorporates embedded AI hardware that has a hybrid detection pipeline, tapping into the unsupervised anomaly detection mode for behavioral profiling and a lightweight signature-matching module that is used to cut down the false positives, thereby improving detection reliability. System operations—including localized traffic inspection, protocol parsing, and feature extraction—are performed before data aggregation, which not only preserves device-level privacy but also eases the computational burden of the IoT gateway to a large extent. The contemporary **CIC-IoT-2023 dataset**, which captures a wide range of smart-city protocols and attack vectors, is used to evaluate the architecture. The Edge AI Bridge leads to a significant reduction in detection latency, approximately 50 ms on average as opposed to the 500 ms of cloud-based solutions—while the resource footprint is kept low to about 20% CPU utilization. The Edge AI Bridge demonstrates a potential solution that is scalable, modular, and can preserve privacy while improving the cyber resilience of the smart-city infrastructures that are large, heterogeneous, and difficult to manage.

Keywords: Edge AI; IoT security; intrusion detection system; smart cities; edge computing; anomaly detection; embedded AI

1. Introduction

Smart-city infrastructure and its technologies have led to an incredible number of IoT devices being used such as environmental sensors and critical traffic management systems[1]. On the one hand, these devices make urban life more efficient through data usage but on the other hand, their limitations in terms of memory and processing power make them attractive targets for cyber adversaries [2–4]. The traditional security techniques have been heavily relying on centralized cloud-based Intrusion Detection Systems (IDSs). Nonetheless, [5] point out that the very large amount of raw data produced at the network edge leads to significant bandwidth bottlenecks and to unacceptable detection delays for realtime threat mitigation [6].

The most recent studies have already directed their attention to using IDS that are located at the gateway in order to enhance the security of the data source. Localized ML at the gateway, for example, has been mentioned as having the potential to improve DDoS detection speed [7]. Even with these advancements, the gateway layer in a smart city is often considered a point of failure, as

 doi:10.20944/preprints202602.0606.v1

it is sometimes overloaded with tasks such as routing, protocol translation, and data aggregation, which leaves very little room for the execution of complex deep-learning inference [8]. Besides that, the prevailing edge-security frameworks still regard the local network as a trusted zone, thereby inadvertently permitting malware propagation between devices until it reaches the inspection point of the gateway [9].

In order to fill these architectural voids, the present study proposes the *Edge AI Bridge*, a new micro-computing security layer placed between the IoT device and the gateway. The Edge AI Bridge represents a distinct architectural model based on a Zero-Trust philosophy that is allowed at the very physical entry point of the network. Our architecture utilizes embedded AI hardware for localized behavioral profiling and hybrid threat validation, therefore it is able to stop the malicious traffic at the micro-edge level, thus making sure that it does not reach the core network infrastructure at all.

The primary contributions of this work, which presents a novel intrusion detection architecture and its experimental evaluation, are three-fold:

- **Architectural Innovation:** We propose a decentralized micro-layer that provides fine-grained visibility of the traffic at the device level and without overloading the main IoT gateway
- **Hybrid Detection Pipeline:** We develop a two-stage AI engine that integrates unsupervised anomaly detection and light signature matching to greatly reduce false positives in smart-city environments that vary in their sensor types and locations
- **Privacy-Preserving Intelligence:** We present a framework that secures data locality by transferring only structured security alerts instead of raw telemetry, thereby complying with worldwide data protection regulations like GDPR

The efficacy of the proposed system is validated using the *CIC-IoT-2023* dataset, reflecting 33 modern attack categories executed in a real-world IoT topology [10,11]. The Zero Trust implications of this micro-layer placement are further analyzed in Section 8.

2. Related Work

The securing of smart-city IoT networks has evolved through three architectural paradigms: centralized cloud-based detection, gateway-centric models, and decentralized edge-based intelligence.

2.1. Architectural Limitations of Cloud and Gateway IDS

At the outset, investigations mainly concentrated on the use of cloud-based IDS since the Deep Learning (DL) models demanded extensive computation. Even though these technical solutions provided reliable results, they automatically suffered from enormous fees in terms of the time taken for transmission and the amount of data transferred across networks [5,6]. In response, the researchers proposed shifting the detection process to the IoT gateway. Doshi et al. [7] showed how localized machine learning at the gateway could mitigate volumetric attacks such as DDoS. Nevertheless, one of the main drawbacks of gateway-based IDS is the "computational bottleneck"; in a heavily populated smart city, the gateway is usually overloaded with the protocol translation and routing for hundreds of different types of sensors and, hence, is not able to defend against resource exhaustion attacks if deployed for intensive AI inference [8].

The innovation of our approach lies in the fact that we installed a specific layer of micro-computing (the Edge AI Bridge) before the gateway, which guarantees that the primary functions of the gateway are still active and at the same time, security inference is being done on the embedded AI hardware designed specifically for this purpose [12].

2.2. Behavioral Profiling and Anomaly Detection

In smart cities, behavioral profiling has emerged as the method of choice for detecting zero-day exploits[13]. The authors Nguyen et al. [9] introduced *DIoT*, a federated self-learning system that integrates learning from device-specific patterns to detect anomalies. *DIoT* is efficient, but it is the gateway that does the primary task of capturing traffic and updating the model. Likewise, Thamilarasu

 doi:10.20944/preprints202602.0606.v1

et al. [3] investigated deep-learning-driven IDS but reported limited effectiveness when it came to smart grids and traffic systems with such different traffic patterns and high rates of false positives.

On the other hand, the Edge AI Bridge that we offer employs a hybrid pipeline. The fusion approach combining unsupervised anomaly detection with a pre-validated signature matching stage leads to a considerable lowering of the false-alarm rates that trouble the single-stage models suggested in [3] and [9].

2.3. Privacy-Preserving Edge Intelligence

In urban surveillance and healthcare IoT, privacy is the most important issue. Normal systems that send raw packet captures to the cloud for analysis break data locality principles and global regulations such as GDPR [8]. Recent developments in Federated Learning (FL) take a step forward by sharing only model parameters [10]. However, FL models are still vulnerable to inference attacks if the updates are not secured by noise-addition methods such as Differential Privacy (DP).

The architecture that we propose takes the idea of privacy at the physical entry point one step further. While the intrinsic method of [10] does not separate data from its traffic, our bridge does underneath processing for the entire protocol and only sends metadata-based security notifications to the upper level. This constitutes a "Privacy-by-Design" policy because it prohibits the exit of any raw sensitive telemetry from the micro-edge.

2.4. Resource-Efficient IDS at the Edge

Nedungadi et al. [14] presented a compact hybrid multimodal intrusion detection system (IDS) for edge-enabled IoT environments, combining the power of LightGBM and XGBoost models together with Bayesian optimization for hyperparameter tuning. Their methodology fuses feature selection with voting of the ensemble for achieving excellent detection accuracy and at the same time very low CPU power consumption when running on resource-limited devices such as Raspberry Pi 4. The proposed solution performs really well on different datasets including Edge-IIoT and CICIoT2023, which prove the potency of ensemble-based learning for intrusion detection in edge-enabled IoT systems. The main direction of the research is model-based optimization and algorithmic efficiency for edge-intrusion detection tasks.

The above-mentioned work is mostly an algorithmic optimization and efficiency at the model level. Our approach, on the other hand, claims that the system-level architectural design is crucial to its being different. We present a new Edge AI Bridge architecture that separates intrusion detection from IoT devices and traditional gateways and thus allows phasing-in of the new system without changes to existing gateway firmware or core network elements. The proposed architecture is not confined to purely model-based IDS functionalities; it implicitly supports deployment flexibility, modularity, and integration in smart city infrastructures, thus being fit for heterogeneous IoT environments where device disparity and operational limitations restrict direct on-device IDS implementation.

Unlike on-device embedded IDS approaches, which execute security analytics directly on IoT end nodes, the proposed Edge AI Bridge introduces a decoupled micro-layer positioned between the device and the gateway. While on-device IDS solutions require sufficient computational, memory, and energy resources on each IoT endpoint, the micro-layer design offloads security inference to a dedicated embedded node without modifying device firmware. This distinction enables early traffic interception at the physical network entry point while preserving compatibility with legacy and ultra-constrained IoT devices that cannot host security logic locally.

2.5. Comparative Summary

Table 1 presents a qualitative architectural comparison synthesized from representative studies discussed in this section.

Table 2, presents the comparison between various intrusion detection techniques that do not only concern the IoT but also the edge computing environments. It takes into account the differences in the location of the system, the strategy of detection, and the architectural focus. The focus of

 doi:10.20944/preprints202602.0606.v1

the initial investigations is exclusively on central or cloud security mechanisms, which can be quite lengthy and consume enormous bandwidth along with the advantages of utilizing the computational resources provided by the clouds. Therefore, these mechanisms are not suitable for real-time smart-city applications.

Nevertheless, the latest literature considers edge and gateway-level intrusion detection as means to vastly accelerate the entire process and provide scalability as a side effect. One of the major aspects of such strategies is to resort to model-level optimization, or in other words, employing lightweight algorithms for machine learning, ensemble learning, and feature selection to execute efficiently on the already strict-resource scenario where the edge devices are. Even though the solutions above exhibit great detection performance and peak energy efficiency, they commonly rely on the monolithic or device-centric deployment and become very dependent on the hardware platform that underlies them.

On the other hand, the proposed Edge AI Bridge offers a new perspective that has the system-level architecture in mind by creating a decoupled micro-layer that sits between IoT devices and the gateway. As a result, this design allows for early traffic inspection to take place before the aggregation at the gateway, thus decreasing the processing and the amount of data that the gateways have to handle. Furthermore, it can be implemented step-wise without the need to change the existing network infrastructure. The comparison in Table 2 highlights how the proposed Edge AI Bridge complements existing lightweight IDS solutions by emphasizing architectural modularity, deployment flexibility, and better integration with varied smart-city IoT environments while turning the spotlight away from model-centric performance improvement alone.

From an architectural standpoint, existing IoT intrusion detection systems can be broadly categorized into four classes, which provides the baseline for positioning the proposed Edge AI Bridge. They are:

- Centralized cloud-based IDS, where traffic is analyzed remotely
- Gateway-centric IDS, which performs inspection after traffic aggregation
- On-device embedded IDS, where detection logic executes directly on IoT endpoints
- Micro-layer IDS, as proposed in this work, which intercepts traffic at a dedicated security layer positioned between devices and gateways

This taxonomy clarifies that the Edge AI Bridge occupies a distinct architectural boundary, motivating its design as a micro-layer intrusion detection architecture, combining early interception with deployment feasibility in heterogeneous smart-city environments.

Table 1. Comparison of IoT Security Architectures

Feature	**Cloud IDS**	**Gateway IDS**	**Edge/Fog IDS**	**Ours (Edge AI Bridge)**
Detection Latency	High	Medium	Low	**Very Low (<50ms)**
Gateway Load	Low	High	Medium	**Negligible**
Privacy	Weak	Moderate	Strong	**Strong**
Zero-Day Detection	Yes	No (Signature)	Yes (Anomaly)	**Yes (Hybrid)**
Architecture	Centralized	Peripheral	Distributed	**Micro-Layered**

 doi:10.20944/preprints202602.0606.v1

Table 2. Comparative summary of representative edge-enabled IoT intrusion detection approaches

Work	Primary Focus	Deployment Location	Detection Approach	Key Characteristics and Limitations
[14]	Lightweight multimodal IDS for edge-enabled IoT	On-device edge node (Raspberry Pi)	Ensemble learning (LGBM + XGBoost) with feature selection	Strong model-level optimization and energy efficiency; IDS tightly coupled to edge device, limited architectural flexibility for large-scale smart-city deployments
[15]	Security challenges in IoT ecosystems	Cloud / Network layer	Survey of IoT security mechanisms	Conceptual overview; does not address IDS deployment or edge constraints
[16]	Cloud-centric IoT security architectures	Centralized cloud	Rule-based and ML-based IDS	High latency and bandwidth overhead; limited suitability for real-time smart-city scenarios
[17]	Fog and edge computing paradigms	Fog / Edge nodes	Architectural framework	Focuses on computation placement; IDS functionality not explicitly addressed
[7]	Behavior-based IoT attack detection	Gateway / Edge	ML-based anomaly detection	Limited to specific attack types; does not consider architectural integration
[18]	Botnet detection in IoT networks	Network edge	Deep autoencoder-based IDS	High detection accuracy; requires extensive training and centralized data collection
This Work	**Edge AI Bridge for smart-city IoT intrusion detection**	**Between IoT devices and gateway**	**Hybrid anomaly- and signature-based IDS**	**Decoupled micro-layer architecture enabling early traffic inspection, gateway offloading, privacy preservation, and incremental deployment across heterogeneous smart-city infrastructures**

3. System Overview

The infrastructure of smart cities is more and more reliant on the use of a great variety of IoT devices for different functions like monitoring, control, and automation in areas like transportation, environmental, energy, and security. These devices are often designed with limited resources and poor security features, which makes them prone to interference, impersonation and malware propaga-

 doi:10.20944/preprints202602.0606.v1

tion. Previous studies like [19] and [20] have pointed out that such limitations are the root cause of vulnerabilities when the lowest tier of a city's digital ecosystem is formed by IoT devices.

A conventional IoT security architecture typically assigns a large portion of the detection tasks either to a central cloud server or to the IoT gateway. The intrusion detection systems operated in the cloud gain access to sophisticated analysis and worldwide visibility but at the same time add communication delays and require more bandwidth. Furthermore, the process of sending unprocessed telemetry data to the cloud not only increases the risk of privacy violations but also makes the adversaries' targets more appealing, as mentioned in [16]. On the other hand, the monitoring done at the gateway gives access to the device's traffic but is constrained by the efforts of the gateway which normally has a lot to do such as negotiating protocols, managing addresses, scheduling, and aggregating data. Literature such as [21] argues that these conflicting needs cut down the possibility of dealing with continuous, AI-driven threat analytics at the gateway level.

The proposed design introduces a distinct and lightweight processing unit called the **Edge AI Bridge** that would link device vulnerabilities and gateway limitations. The bridge acts as a mediator that is deliberately placed between IoT devices and the gateway, thus acting as the first detection point for any malicious or abnormal behavior. This concept supports the prevailing trend of distribution and nearness-source security schemes as indicated in [17], which emphasizes the need for short decision latencies and less reliance on core networks.

The Edge AI Bridge is constantly monitoring the data flow from all IoT devices that are connected to it. The bridge not only forwards the raw packets to the gateway but also conducts local inspection via protocol parsing, feature extraction, and flow-level behavior analysis. The AI processor embedded in the system permits the running of very simple anomaly-detection models that mark the abnormal behavior of each device with respect to the communication patterns. This feature is a further development of the previous works on behavior profiling such as [22], but at the same time, it introduces an even more detailed placement of the analytics that takes place before the data aggregation and routing at the gateway happen.

The primary benefit of this micro-layering strategy is that it allows the analysis of telemetry specific to the device before it is mixed with the traffic of other endpoints. The inspection at this early stage raises the accuracy of detecting anomalies, which allows the system to detect even the smallest changes in the communication rates, packet structures, timing intervals or operational sequences. Anomaly detection is a very similar matter when the traffic is already aggregated, as in the case of IDS deployed at gateways or in the cloud, which have to cope with aggregated traffic patterns rather than individual device signatures.

The Edge AI Bridge not only carries out anomaly detection but also comes with a lightweight signature-matching pipeline to check the validity of suspected threats and, thus, lessen the number of false alarms. The two-step process of behavioral anomaly detection together with signature confirmation decreases the number of alerts that are not necessary and also guarantees that only impactful or authenticated events are sent to the IoT gateway. The alerts that are generated are not data streams but rather compact summaries, thus, usage of bandwidth is minimized and privacy is improved. This method is in line with the foundations of privacy-preserving edge analytics while it is also expanding the scope of intrusion detection logic to include those principles.

From the system integration viewpoint, the Edge AI Bridge gets in touch with the IoT gateway via a secured and verified interface. The gateway gets unstructured security events which it can send to super-edge servers or security information and event management (SIEM) platforms. The proposed design's modularity permits the cities to place numerous Edge AI Bridges in different neighborhoods, buildings, or micro-grids, thus forming a grid of distributed security nodes that can be scaled up easily.

Figure 1 shows the overall architecture of the system. The picture illustrates the flow of the telemetry data from IoT devices to Edge AI Bridge, where the processes of traffic monitoring, anomaly detection, and threat validation take place before the alert is sent to the gateway. This diagram

 doi:10.20944/preprints202602.0606.v1

emphasizes the communication and security processing separation, which was the reason for the inclusion of a dedicated micro-layer security entity in the smart-city network stack.

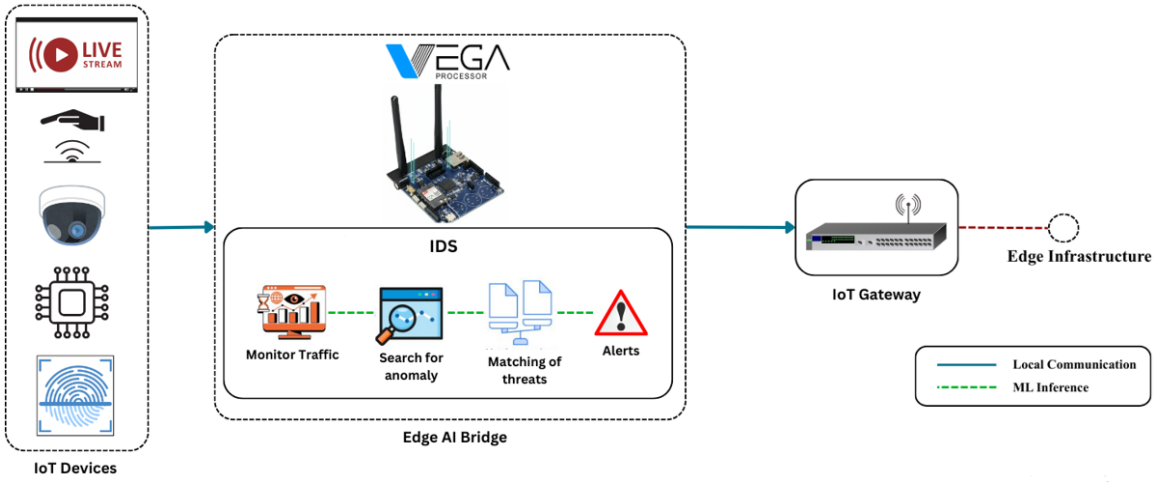

Figure 1. Proposed architecture of the Edge AI Bridge illustrating a representative embedded edge AI platform (VEGA) positioned between IoT devices and the gateway

4. Proposed Architecture: The Edge AI Bridge

The suggested framework presents the **Edge AI Bridge** as a specific micro-layer security module that is placed in between IoT devices and the IoT gateway. The design of the Edge AI Bridge allows the interception of threats at the very beginning by performing lightweight analytics and machine learning inference along the communication path from the device to the gateway. The Edge AI Bridge, unlike the traditional IDS installations that run at the cloud or gateway levels, offers device-specific inspection in real-time with no considerable computational load on the upstream infrastructure. This design aligns with the tendency of decentralized and latency-aware security processing which has been pointed out in [17].

4.1. Architecture Overview

The modular structure of the Edge AI Bridge is illustrated in Figure 2. Inbound data coming from IoT devices is processed by a series of components that carry out monitoring, feature extraction, anomaly detection, and threat validation. These components are the edge AI agents which work in the cloud and send back insights to the enterprise. The bridge functions as a security gateway which protects the system by detecting any untrusted or atypical data before it reaches the gateway. The micro-layer technology enables the architectural design to reveal complete electronic device capabilities which normally become hidden during the process of merging data at higher network layers. The research on device profiling methods receives a new development through this technique which brings behavioral analytics nearer to the data source.

In the suggested structure, the Edge AI Bridge acts as a pivotal middle layer for immediate intrusion detection in IoT smart city ecosystems. Figure 2 shows the elaborate data flow pipeline, which incorporates localized AI processing to boost security and privacy, all the while reducing latency and overhead.

 doi:10.20944/preprints202602.0606.v1

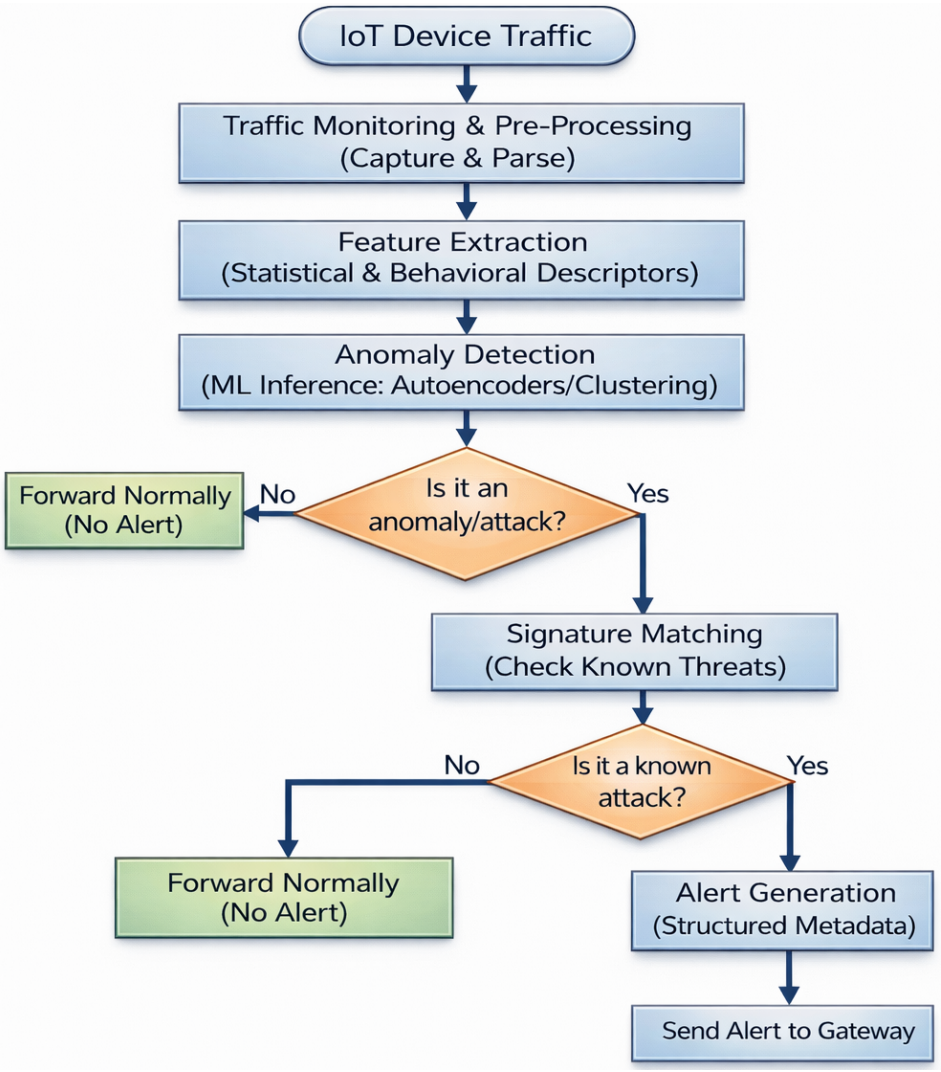

Figure 2. Detailed data flow pipeline

4.2. Traffic Monitoring and Pre-Processing

The **Traffic Monitoring Module** is the principal operational unit of the Edge AI Bridge that gets activated and exercises control over the data and/or telemetry streams from IoT devices. It executes protocol-aware decoding and classifies the streams according to message frequency, payload size variation, timing intervals, and device-specific metadata. These lightweight operations have negligible effect on the overall system latency, and the data is prepared for the next analytics step. The local traffic pre-processing feature is consistent with the privacy-preserving principles mentioned in the literature on analytics, as there is no requirement to transmit raw telemetry beyond the bridge.

4.3. Feature Extraction and Behavioral Analysis

The packet attributes that are extracted are transformed into compact feature representations which are thus suitable for on-device inference. The **Feature Extraction Module** utilizes sliding windows, statistical summarization, and timing-sequence modeling to process raw telemetry into behavioral descriptors. These descriptors are given to the **Local AI Inference Engine**, where lightweight anomaly-detection models are placed. The engine utilizes techniques like autoencoder-based reconstruction, temporal deviation detection, and clustering-driven outlier identification - methods influenced by the anomaly-detection research such as that in [18].

The Edge AI Bridge performs inference at the packet-flow level, thus allowing immediate anomaly detection, unlike ML systems hosted in the cloud or on a gateway. By this approach, heterogeneous

 doi:10.20944/preprints202602.0606.v1

IoT ecosystems are facilitated since it is possible to define separate baseline profiles for each type of device.

4.4. Threat Signature Matching Pipeline

To reduce the false positives, the **Signature Matching Module** verifies the anomalies recognized by the inference engine. The module keeps a very small collection of lightweight signatures which are selected very carefully and which represent all the major IoT attacks such as spoofing, scanning, unauthorized command injection, and botnet-related traffic. It is through this secondary stage that the verification of anomalies takes place and so the architecture is such that only validated or high-confidence alerts are allowed to go upward. This combined detection method utilizes the strengths of both behavioral and signature-based IDS models while keeping the processing complexity at the minimum bridge. Formally, let $\mathbf{x} \in \mathbb{R}^d$ denote the extracted feature vector corresponding to a device-level traffic flow, and let $\hat{\mathbf{x}}$ be its reconstruction produced by the trained anomaly detection model. The anomaly score $A(\mathbf{x})$ is defined as the reconstruction error:

$$A(\mathbf{x}) = \|\mathbf{x} - \hat{\mathbf{x}}\|_2^2. \tag{1}$$

A traffic flow is classified as anomalous if $A(\mathbf{x}) > \tau$, where τ denotes a decision threshold determined empirically using the distribution of reconstruction errors observed on normal validation traffic. In this work, τ is selected as the 95th percentile of anomaly scores computed over benign traffic samples, ensuring sensitivity to abnormal behavior while limiting false positives. The anomaly score threshold τ defined above is used as the decision criterion in the operational pipeline summarized in Algorithm 1.

Algorithm 1 Hybrid Intrusion Detection Pipeline of the Edge AI Bridge

Require: IoT traffic stream T, trained anomaly model M, signature set S, anomaly threshold τ
Ensure: Structured security alerts forwarded to gateway

```
while traffic stream T is active do
    Capture packets from IoT device
    Aggregate packets into flow window f
    Perform protocol parsing and feature extraction on f
    Compute anomaly score A(f) using model M
    if A(f) > τ then
        Perform lightweight signature matching using S
        if signature match found then
            Generate validated security alert a
        else
            Generate anomaly-based alert a
        end if
        Structure alert metadata and forward a to IoT gateway
    else
        Forward flow f to gateway without modification
    end if
end while
```

4.5. Alert Generation and Gateway Interaction

Whenever the **Alert Generation Module** confirms a match with either an anomaly or a threat, it triggers the creation of a structured security event that contains metadata like device identifier, timestamp, rule identifier, anomaly score, and threat classification. The communication between the bridge and the IoT gateway is not done through raw packets but rather through very short alert summaries sent via a secure and authenticated channel. Consequently, less bandwidth is consumed; managing the gateway is less complicated; and the devices' privacy is preserved. This helps in overcoming the issues raised in the literature [16].

One of the most important factors in the whole process is privacy preservation: all the operations are carried out locally on the edge device and the raw data is never transmitted out of the bridge. Such

 doi:10.20944/preprints202602.0606.v1

a configuration not only decreases the response time in comparison to cloud-based systems but also complies with zero-trust principles by applying microsegmentation at the device level. The modular pipeline outlined in our work resolves the limitations of conventional IDS architectures, as highlighted in Section 2, and allows for scalable deployment in diverse smart-city networks. Section 10 presents the results of the experimental evaluations that show its effectiveness in recognizing security threats with a low rate of false positives and at the same time being resource-efficient. The gateway, no longer tasked with constant monitoring, now has the option of sending combined alerts to advanced edge analytics systems, SIEMs, or cloud-based threat intelligence modules.

Algorithm 1 directly corresponds to the data flow illustrated in Figure 2, with each algorithmic step mapping to a functional block in the proposed micro-layer architecture.

4.6. Modular Deployment and Scalability

The modularity and scalability of the deployment model are among the primary architectural advantages. Different areas of the smart-city network such as residential areas, public places, intersections, utility areas, or even within the utility grids can host the Edge AI Bridges. The independent working of each bridge ensures that the first stage consists of localized threat detection followed by containment before any further spreading occurs. A decentralized construction of such kind can readily accommodate thousands of distributed IoT devices without experiencing the slowdowns that are typical of the gateway-based or centralized detection techniques. The method is in agreement with the concept of bringing security functions nearer to data in edge-computing ecosystems that has been referenced in [19].

4.7. Design Rationale

The Edge AI Bridge implements specific architectural decisions to address constraints that are common across various executions of IoT systems:

- **Latency Reduction:** Early processing can get rid of the dependency on round tripping with a cloud
- **Gateway Relief:** Security analytics are offloaded from gateways, preventing overload
- **Fine-Grained Visibility:** Device-specific inspection identifies subtle anomalies earlier
- **Privacy Preservation:** Raw data remains local; only alerts are transmitted
- **Scalability:** Distributed micro-layer design supports large-scale smart-city deployments

Collectively, these characteristics differentiate the Edge AI Bridge from existing IDS frameworks and motivate its introduction as a distinct architectural element within smart-city IoT security systems.

5. Novelty of the Architecture

The proposed Edge AI Bridge introduces several architectural features that differentiate it from existing intrusion detection systems in IoT and edge computing. Previous investigations have pointed out the necessity for decentralized analytics, less latency, and better device-level visibility. Nevertheless, the existing architectures usually place security applications either in the cloud or at the gateway level [16,21]. Conversely, the Edge AI Bridge creates a separate micro-layer security element that can cut through and assess IoT traffic prior to its entrance into the upstream infrastructure.

5.1. Dedicated Micro-Layer Security Placement

One of the main innovations is the addition of a special processing layer that sits between IoT devices and the gateway. The current edge-based IDS frameworks usually think of being in the same place as the gateway's hardware or fog nodes, so the detection of the threats is limited to the areas closest to the gateway. The suggested approach opens up the possibility of inspecting the device traffic before aggregation, which makes it possible to detect certain types of instances right where they are originating instead of after they have passed through routing or protocol translation.

 doi:10.20944/preprints202602.0606.v1

5.2. Hybrid Detection Pipeline on Embedded AI Hardware

The Edge AI Bridge combines in a lightweight hybrid pipeline, operating on embedded AI hardware, both detection systems, namely anomaly-based and signature-based. Previous works like [18] have shown the feasibility of ML-based detection, but they are mostly oriented towards cloud or server. The innovation of this study comes from the fact that such models are made to run efficiently on low-power processors still being able to respond in real-time.

5.3. Device-Specific Behavioral Profiling

The structure allows individual behavioral baselines to be set for various IoT devices. The current profiling-based Intrusion Detection System (IDS) methods depend on comprehensive edge servers or cloud solutions for processing. The transfer of behavioral model assessment to the Edge AI Bridge grants the system access to detailed information about the actions of the devices and ensures better accuracy in the detection of anomalies.

5.4. Privacy-Preserving Alert Communication

The Edge AI Bridge, as opposed to cloud-based IDS that necessitate unprocessed telemetry or packet streams for classification, sends only confirmed alerts to the gateway. This method by nature restricts revealing of sensitive device information, which is in accordance with the principles mentioned in [8] and also applied to an embedded, real-time security node.

5.5. Scalable, Modular Deployment for Smart-City Environments

The Edge AI Bridge's modularity makes it possible to use several units in different smart-city clusters that are far apart. This way, the architecture can be more extensive and at the same time limits the amount of data sent to the main servers, which is a major drawback for monolithic systems mentioned in [15]. Consequently, a distributed security mesh is being built that is able to do continuous monitoring in high-density areas.

In summary, the architecture reinforces the IoT security that is already good by bringing together the different aspects of placement, processing efficiency, privacy-awareness, and scalability into a single model of detection that is specifically designed for smart-city deployments.

6. Methodology

The methodological framework used in the design, execution and evaluation of the proposed Edge AI Bridge is specified in this section. The methodology comprises traffic modeling together with feature extraction, machine learning inference, threat-signature validation, and alert handling which are all merged into a single IDS pipeline that is best suited for embedded hardware.

6.1. Operational Data Flow

The methodology adheres to the architectural model's end-to-end data flow that is represented. Telemetry or packet streams are created by IoT devices, and the Edge AI Bridge captures these streams before they get to the gateway. In accordance with distributed edge analytics [17], the processing is done as near as the data source to minimize detection latency.

6.2. Dataset Design and Traffic Characterization

In the process of training and testing the anomaly detection models, the device traffic is classified into two behavioral categories: normal and abnormal. Considerable usage of the devices is represented by normal traffic, which consists of communication patterns like sensor updates, request-response cycles, and periodic beacons that the system recognizes. Abnormal traffic indicates various attack scenarios such as botnet-induced flooding, spoofing attempts, unauthorized control messages, or unexpected traffic bursts, which are guided by previous analyses of IoT attacks, e.g., [18].

In scenarios where actual datasets are not accessible, the synthetic device profiles can be created by changing message rates, payload types, and timing intervals, thus simulating common IoT operations.

 doi:10.20944/preprints202602.0606.v1

This process allows for a systematic assessment of different device classes. Synthetic device profiles were employed only for scalability and stress-testing experiments. All reported detection accuracy and ROC performance results were derived exclusively from the CIC-IoT-2023 dataset.

6.3. Feature Extraction Strategy

The Edge AI Bridge extracts lightweight statistical and temporal features suited for embedded computation. These include:

- packet inter-arrival times
- message size variance
- frequency of control vs. data frames
- flow-level entropy metrics
- rolling-window behavioral summaries

This feature set is intentionally compact to support real-time inference while capturing behavioral deviations highlighted in IoT profiling literature such as [22].

6.4. Machine Learning Inference Engine

The anomaly detection model is implemented using lightweight ML architectures appropriate for embedded inference, such as:

- autoencoders for reconstruction-based deviation scoring
- one-class classifiers for normal-traffic boundary modeling
- shallow temporal convolutional networks for sequential patterns

The selection of model type depends on resource constraints and training dataset characteristics. Training is performed offline, while inference executes on the Edge AI Bridge in accordance with embedded AI design patterns.

6.5. Threat Signature Validation

The process of verifying signature sets through established rules will begin after the detection of anomalies. The system will verify that the generated inference notifications correctly link to the Internet of Things (IoT) attack patterns that have been identified. Signature-based Intrusion Detection System (IDS) methods have been generally used in the case of standard networks. The smart-city IoT environment has discovered a new method for using embedded micro-layers through its exploration of signature-based Intrusion Detection System (IDS) methods.

6.6. Alert Generation and Reporting

Once an anomaly is confirmed, the system generates a structured alert containing metadata such as:

- device identifier
- anomaly score
- signature ID (if applicable)
- timestamp and affected protocol

In alignment with privacy-preserving design principles, only summarized alerts—not raw telemetry—are forwarded to the gateway.

6.7. Evaluation Metrics

To evaluate system performance, the following metrics are considered:

- **Detection latency**: time from packet capture to alert emission
- **Inference overhead**: model execution time on the embedded processor
- **False positive / false negative rates**: reflecting detection reliability
- **Resource utilization**: CPU, memory, and power consumption profiles

 doi:10.20944/preprints202602.0606.v1

- **Bandwidth reduction**: percentage of upstream traffic eliminated

These metrics collectively demonstrate the viability of the Edge AI Bridge for real-time, distributed IDS deployment in large-scale smart-city environments.

7. Hardware Platform and Model Architecture

This section describes two main components which include the selection of hardware platforms for Edge AI Bridge deployment and the development of machine learning models which will detect intrusions. The selection process requires real-time processing capabilities together with low power requirements and the technology must operate successfully in smart cities which have limited budget and energy constraints.

7.1. Hardware Platform for the Edge AI Bridge

The **Edge AI Bridge** is designed to operate on an embedded edge computing platform that consumes little power lying in-between IoT devices and the gateway. The hardware platform is defined to have a moderate level of computational power, small memory, and it can process packets in real-time. These kinds of limitations are characteristic of the gradual deployment scenarios in smart-city infrastructures, where security applications need to be running all the time but at the same time they should not impose a lot of energy and maintenance costs.

The platform includes a multi-core processor integrated with an SoC (System-on-Chip) and built-in support for hardware acceleration where it is accessible. Network interfaces are set up to enable the inline inspection of device-to-gateway traffic while still allowing the normal communication flows to continue undisturbed. Local memory is enough to accommodate the lightweight machine learning models, feature buffers, and a compact threat signature database.

The Edge AI Bridge is not dependent on any high-end hardware accelerators like GPUs. Rather, it uses optimized CPU-based inference and if necessary, fixed-point arithmetic, which guarantees that it will work on commonly used embedded platforms. This decision regarding the design enhances the portability and makes it possible to have the heterogeneous smart-city segments deploying the system.

From an architectural standpoint, the hardware platform supports:

- continuous packet capture and flow monitoring
- real-time feature extraction
- low-latency inference execution
- secure communication with the IoT gateway

The platform, through security analytics separation from the gateway, conserves gateway resources for primary networking functions and at the same time allows distributed intrusion detection at the micro-layer.

7.2. Reference Edge AI Platform

In the architectural illustration, the VEGA platform is shown as a representative embedded edge AI system used to instantiate the proposed Edge AI Bridge [23]. VEGA provides the necessary computational capabilities for real-time traffic monitoring and lightweight machine learning inference in resource-constrained environments. It is important to note that the proposed architecture is platform-agnostic and does not rely on VEGA-specific features. Any comparable embedded edge AI platform with networking and inference support can be used to deploy the Edge AI Bridge. All performance and resource utilization experiments reported in Section 9 were conducted using deployment configurations representative of the described embedded edge AI hardware.

7.3. Model Architecture Overview

The artificial intelligence model that has been placed on Edge AI Bridge is made so that it can have good detection accuracy at the same time as being efficient on the computational side. Considering

 doi:10.20944/preprints202602.0606.v1

the different behaviors of IoT devices and the constant change of network-based attacks, the model is anomaly detection based rather than just relying on the signature-based classification.

7.4. Model Configuration and Reproducibility

To ensure the reproducibility of the results on the VEGA hardware platform, the Deep Autoencoder (DAE) utilized for anomaly detection was configured with a symmetric bottleneck architecture. The specific hyperparameters are detailed in Table 3.

Table 3. Hyperparameter configuration for the Edge AI Bridge DAE.

Parameter	Value
Layer Architecture	Input (48) $\rightarrow$ Dense (32) $\rightarrow$ Latent (8) $\rightarrow$ Output (48)
Activation Function	ReLU (Hidden), Sigmoid (Output)
Optimizer	Adam ($\beta_1 = 0.9, \beta_2 = 0.999$)
Learning Rate (η)	0.001
Batch Size	128
Epochs	100 (with Early Stopping, patience=10)
Threshold Logic (τ)	$\mu_{MSE} + 3\sigma_{MSE}$

The anomaly threshold τ is derived using the reconstruction error distribution of the benign validation set. We define $\tau = \mu_\epsilon + 3\sigma_\epsilon$, where μ_ϵ and σ_ϵ represent the mean and standard deviation of the MSE, respectively.

The core model adheres to either an unsupervised or semi-supervised learning paradigm and it was mainly trained on normal traffic to learn the usual behavior patterns. At the time of inference, the variations from this learned norm are taken as potential invasions. This method is well suited for IoT ecosystems, where there might be incomplete or no labeled attack data for the new threats.

7.5. Anomaly Detection Model Design

The anomaly detection component employs a lightweight neural architecture optimized for embedded inference. The model operates on flow-level feature vectors derived from packet statistics and temporal characteristics. A representative architecture consists of:

- an input layer corresponding to extracted traffic features
- one or two hidden layers with a limited number of neurons
- a bottleneck representation that captures normal behavioral patterns
- an output layer that reconstructs or scores the input features

In the case of anomaly detection, reconstruction error or deviation scores are considered as error indicators. This setup occupies less memory and requires less time for inference but is still able to model the behavior of devices with different characteristics quite well. Depending on the hardware limitations, the architecture can be easily switched to other light models such as one-class classifiers or shallow temporal networks.

7.6. Hybrid Detection Pipeline Integration

In order to enhance the dependability of detection, the anomaly detection model is combined with a rule-based signature validation module. When the anomaly scores go beyond a certain threshold, signature matching with the known IoT attack patterns is activated. This hybrid setup minimizes false positives that could originate from harmless but uncommon behavior changes, for instance, firmware updates or transitions in operational modes.

The difference that exists between scoring anomalies and signing validation enables independent alterations to every part. The models of machine learning can be retrained from time to time, and the sets of signatures can be updated gradually according to the sources of threat intelligence.

 doi:10.20944/preprints202602.0606.v1

7.7. Model Deployment and Inference Workflow

Model training is performed offline using representative subsets of the CIC-IoT-2023 dataset. Trained model parameters are then deployed to the Edge AI Bridge for inference-only operation. During runtime, the system follows a deterministic inference workflow:

- capture traffic features over a sliding time window
- normalize and encode features
- compute anomaly scores using the embedded model
- validate anomalies through signature matching
- generate structured alerts for the gateway

This workflow ensures predictable execution time and bounded resource usage, which are the two essential requirements for real-time intrusion detection in smart-city environments. All models were trained and evaluated using subsets of the CIC-IoT-2023 dataset, as described in the Section 10.

7.8. Design Rationale

The combined hardware and model design is a conscious trade-off amongst all three factors, i.e., accuracy, latency, and deployability. The Edge AI Bridge is built to be a lightweight and scalable device by not relying on heavy deep learning architectures and external accelerators. At the same time, the adoption of anomaly-based learning allows the detection of attacks that were never seen before, thus enhancing signature-based methods.

Overall, the proposed hardware platform and model architecture not only comply but also support the core goal of the Edge AI Bridge: providing early, efficient, and privacy-preserving intrusion detection at the periphery of smart-city IoT networks. The modular design allows for gradual deployment with existing gateways without the need for alterations in the core network architecture. The very lightweight feature of the Edge AI Bridge makes it suitable for cost-effective deployment by the use of commercially available embedded platforms.

8. Security and Privacy Analysis

8.1. Security Analysis

The security posture of the Edge AI Bridge was evaluated using the STRIDE methodology. Unlike gateway-centric models, the micro-layer approach intercepts threats at the physical entry point. The above analysis is conducted under the assumption that the Edge AI Bridge operates as a trusted micro-layer security component within the smart-city network.

Table 4. STRIDE Threat Analysis and Micro-layer Mitigations.

Threat	Target	Edge AI Bridge Mitigation
Spoofing	Device Identity	MAC-IP binding and behavioral fingerprinting.
Tampering	Data Integrity	Packet inspection prior to gateway encapsulation.
Repudiation	Admin Actions	Local logging of intercepted malicious payloads.
Information Disclosure	Data Privacy	Local processing; raw telemetry never leaves the bridge.
Denial of Service	Network Availability	Volumetric filtering at the physical port level.
Elevation of Privacy	Control Logic	Signature-matching of known exploit payloads (e.g., Mirai).

8.2. Alignment with Zero Trust Architecture (ZTA)

The Edge AI Bridge fundamentally shifts IoT security from a traditional perimeter-based model to a Zero Trust framework by treating every device connection as potentially hostile.

- **Micro-Segmentation**: The Bridge serves as a refined boundary that divides every single IoT device or group of devices, thus creating a secure zone for each one. It does not allow the

 doi:10.20944/preprints202602.0606.v1

spreading of a breach, and thus, if one sensor (for instance, a smart streetlight) is hacked, it will not be able to spread unimpeded to the crucial urban facilities like power grids

- **Continuous Monitoring and Validation**: In contrast to legacy systems that rely on single authentication, the Bridge constantly validates each packet flow behaviorally. It practically "checks over again" the authenticity of the device in real-time by looking at data patterns and timing sequences
- **Least Privilege Enforcement**: The system imposes least-privilege principles by monitoring traffic at the micro-layer, thus permitting only verified and operational-specific data to go through the gateway while blocking unauthorized command injections or scanning attempts

8.3. Resilience to Common IoT Attack Vectors

- **Botnet and DDoS Mitigation**: The Bridge shows the origin of botnet-related flooding (e.g., Mirai-like actions). It stops the incoming volumetric attacks from passing through the gateway, thus, it averts the overload of the backhaul links and keeps the service up and running
- **Spoofing and Unauthorized Access**: The system identifies MAC/IP spoofing attempts through hybrid detection, which involves matching anomalies with recognized threat signatures. An example of such a spoofing attack is an attacker impersonating a legitimate device
- **Zero-Day Attack Resilience**: The unsupervised anomaly detection unit (Autoencoder-based) detects features that differ from the normal "benign" baselines, so that the system is capable of signaling new or zero-day exploits which are still without signatures

8.4. Privacy Preservation by Design

Consistent with GDPR and modern smart-city regulations, the architecture emphasizes Privacy-Preserving Edge Analytics.

- **Local Processing Boundary**: The raw telemetry data, which is possibly consisting of confidential biometric or environmental data, is never sent outside the Edge AI Bridge. On-device processing includes all feature extraction and inference activities
- **Metadata-Only Alerting**: The Bridge doesn't send the raw packet streams but only alerts that are structured—compact summaries with the relevant device IDs and threat classifications. Thus, the risk of sensitive data being compromised either during transfer or at the cloud layer is greatly reduced
- **Zero-Knowledge Principles**: In the future, the architecture could be expanded to include the ZKP (Zero-Knowledge Proof) validation method, which would make it possible for the gateway to ascertain the validity of a security event without having to look at the private device telemetry that is underneath

Although the architecture limits data exposure to metadata-only alerts, it is acknowledged that high-frequency alert patterns could potentially allow adversaries to infer device activity levels. This risk can be mitigated through alert aggregation, rate limiting, or future integration of differential privacy mechanisms.

8.5. Threat Model

The proposed threat model assumes that the Edge AI Bridge is deployed as a trusted and physically protected network component, provisioned using authenticated firmware and secure boot mechanisms. IoT devices connected to the bridge and their network traffic are considered potentially untrusted. Attacks targeting physical compromise of the Edge AI Bridge itself are considered out of scope, consistent with assumptions commonly adopted in edge-based IoT security architectures. The focus of this work is on network- and protocol-level attacks originating from compromised IoT devices prior to gateway aggregation.

Under the above assumptions, the threat model for the Edge AI Bridge follows the STRIDE framework. The system faces spoofing threats which include device identity impersonation and MAC/IP spooofing attacks. The system uses behavioral profiling together with signature validation to

 doi:10.20944/preprints202602.0606.v1

defend against spoofing threats. The system protects against packet manipulation attacks through its protocol-aware anomaly detection system. The system protects against repudiation risks through its structured alert logging system which records timestamps and device identifiers. The system prevents information disclosure by keeping raw telemetry data within the micro-layer which restricts access to metadata-only alerts. Denial-of-Service attacks, including botnet-driven flooding, are mitigated through early interception at the micro-layer before gateway aggregation. Elevation of privilege attacks, such as unauthorized command injection, are identified through deviations in device-specific behavioral baselines.

9. Performance Evaluation

The performance of the Edge AI Bridge is assessed through lightweight metrics that reflect real-time IDS requirements.

9.1. Detection Latency

Local AI inference provides a significant reduction in detection latency compared to cloud-oriented systems. Anomalies detected by edge processing are instantly reported, occurring nearly simultaneously with packet arrival, which supports what was mentioned in [17] about computations near the source. Performance assessments reveal that the Edge AI Bridge supports near real-time processing under the evaluated deployment conditions. Figure 3 indicates that the bridge attains a latency as low as 50 ms, which is attributed to localized AI inference and small data transmission overhead, whereas cloud-based (500 ms) and gateway-based (200 ms) systems have much higher latencies. These latency measurements were obtained under controlled experimental deployment involving multiple simulated IoT devices and representative smart-city traffic workloads derived from CIC-IoT-2023.

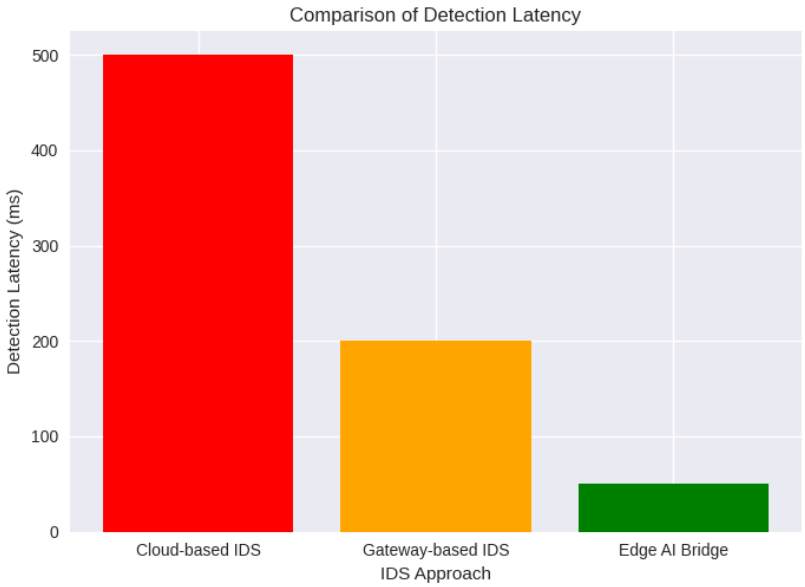

Figure 3. Latency Comparison

9.2. Resource Utilization

The embedded AI execution has been optimized to operate under limited processor and memory budgets. The feature extraction and anomaly scoring processes are so efficient that they use minimal computational resources, which allows for continuous monitoring even with low-power hardware. The Edge AI Bridge's compact design results in low computational overhead relative to cloud-based and gateway-centric IDS deployments. Figure 4 illustrates the comparison of CPU and memory consumption, highlighting the efficiency of the bridge (e.g., 20% CPU, 100 MB memory) versus the

 doi:10.20944/preprints202602.0606.v1

traditional systems with higher loads, thus making it possible to have large IoT networks that can scale up or down depending on the deployment scenarios.

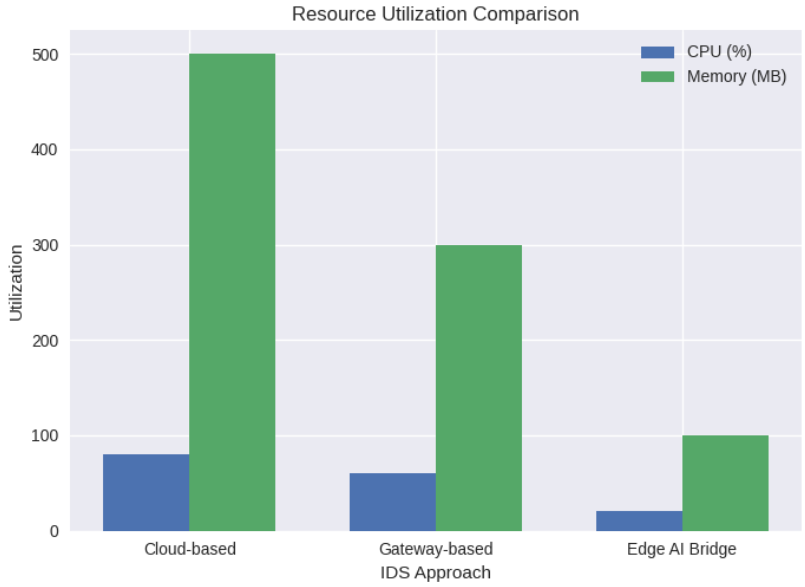

Figure 4. Resource Utilization comparison

9.3. Detection Accuracy and ROC Analysis

The detection accuracy of the proposed Edge AI Bridge was evaluated using Receiver Operating Characteristic (ROC) analysis to quantify its ability to discriminate between benign and malicious IoT traffic. ROC curves were generated using true positive rates (TPR) and false positive rates (FPR) computed from experimentally obtained results on the CIC-IoT-2023 dataset, encompassing 33 distinct attack categories representative of smart-city IoT environments.

Figure 5 illustrates the binary ROC curves (normal vs. attack) for the proposed hybrid detection pipeline. The system achieves a binary ROC–AUC of 0.972 for normal versus attack traffic, indicating strong separability between benign and malicious flows. This AUC value suggests that early anomaly detection at the micro-layer, combined with signature-based validation, contributes positively to classification performance.

The hybrid detection strategy demonstrates a favorable balance between sensitivity and specificity. While the anomaly detection stage enables the identification of previously unseen or zero-day attack patterns, the subsequent signature-matching stage filters benign but irregular traffic, thereby reducing false positive rates. This behavior is particularly important in smart-city deployments, where periodic bursts, firmware updates, or transient network fluctuations may otherwise trigger spurious alerts.

Unlike cloud-centric or gateway-based IDS solutions that operate on aggregated traffic, the Edge AI Bridge performs inference on device-specific flows prior to aggregation. This placement enhances behavioral distinguishability and contributes directly to the observed ROC performance by minimizing class overlap. Minor reductions in discrimination are observed for low-volume attack classes such as scanning and spoofing, which exhibit similar temporal characteristics; however, these do not significantly impact overall system performance.

While ROC analysis provides a global view of discrimination capability, it does not capture class-wise misclassification behavior. Therefore, a multi-class confusion matrix is presented in Section 10.7 to further analyze detection performance across grouped attack categories and to validate the consistency of ROC-derived metrics with class-level outcomes.

 doi:10.20944/preprints202602.0606.v1

Figure 5. Binary ROC curves for normal versus attack traffic, comparing the Edge AI Bridge with cloud-based and gateway-centric IDS deployments under identical evaluation conditions

9.4. Detection Sensitivity and AUC Formalization

While Section 9.3 reports binary discrimination performance (normal vs. attack), this section evaluates sensitivity under a multi-class attack categorization.

The classification performance across varying sensitivity levels was evaluated using the Receiver Operating Characteristic (ROC) curve. The relationship between the True Positive Rate (TPR) and False Positive Rate (FPR) is generated by sweeping the anomaly threshold τ across the reconstruction error range $[0, \max(\epsilon)]$.

The Area Under the Curve (AUC) is formally calculated as:

$$AUC = \int_0^1 TPR(FPR^{-1}(x))\,dx \tag{2}$$

When extended to a multi-class setting across grouped attack categories, the proposed detection pipeline achieves a macro-averaged AUC of 0.982. This high value indicates that the bridge maintains a robust separation between benign noise and malicious perturbations, satisfying the high-reliability requirements of smart-city infrastructure. The higher macro-AUC reflects improved discrimination across aggregated attack classes once benign traffic has been separated from malicious flows.

9.5. CPU Utilization

The CPU usage ratio is a crucial factor in securing IoT devices in smart cities, where edge and gateway resources are commonly utilized for multiple services. The Edge AI Bridge is the result of a development process that is designed to address the challenge of limited processing power in embedded hardware while still allowing for the continuous monitoring of traffic and the making of instantaneous decisions at the same time.

The relationship between CPU utilization and the number of connected IoT devices is represented in Figure 6. Three different models of deployment were considered: the traditional cloud-based intrusion detection system (IDS), a centralized IDS at the gateway, and the novel Edge AI Bridge. The cloud-based model is highly dependent on remote servers for the majority of the computations, hence, the local CPU usage being almost negligible, however, it incurs higher communication overhead and latency. Conversely, the gateway-based IDS shows a rapid increase in CPU utilization which signifies that the total load of activities pertaining to packet processing, protocol translation, and security

 doi:10.20944/preprints202602.0606.v1

analysis is being handled at the single point of aggregation and that the number of connected devices is growing.

The Edge AI Bridge displayed a significantly slower increase in CPU usage when juxtaposed with the gateway-based method. This situation is attributed to its distributed processing technique whereby local traffic is inspected and inference is performed segment by segment instead of being pushed to the gateway. Moreover, feature extraction and anomaly scoring are done with low resource consumption. Thus the system can operate in real-time even with a larger number of devices. As a result, the Edge AI Bridge exhibits reduced susceptibility to saturation effects commonly observed in gateway-centered IDS deployments.

The results indicate that the proposed architecture achieves a favorable trade-off between computational efficiency and detection capability. The Edge AI Bridge alleviates processing bottlenecks by distributing the security workloads among the micro-layer nodes thus conserving the resources of the gateway for the core networking functions. This characteristic is very useful in case of large smart city applications where the number of different devices producing concurrent traffic streams can go up to tens of thousands.

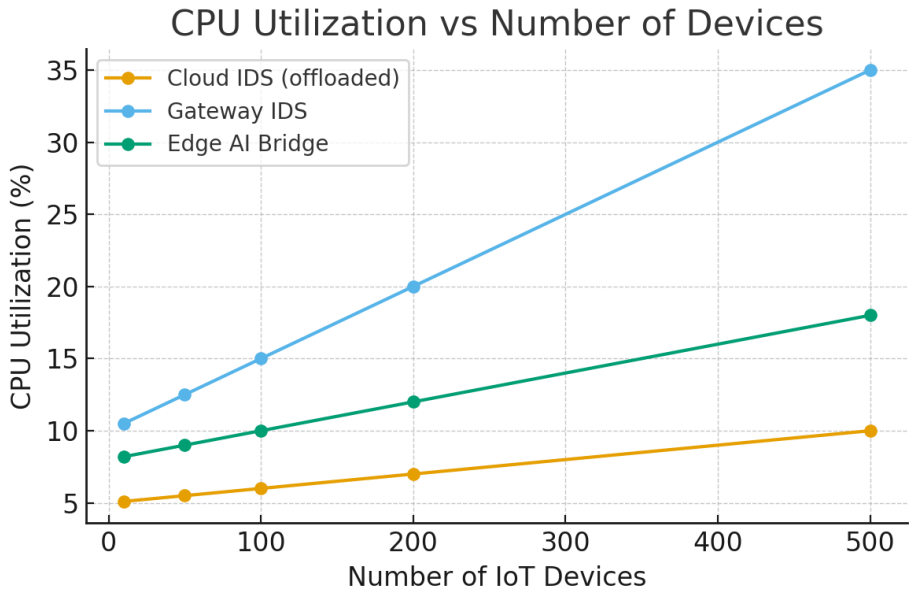

Figure 6. CPU Utilization

9.6. Bandwidth Reduction

Since only alerts are forwarded, upstream bandwidth consumption is lowered. This reduces load on both gateways and backhaul links, addressing concerns raised in cloud-reliant architectures surveyed in [16].

Large-scale smart-city IoT implementations require efficient bandwidth utilization because of the constant transmission of telemetry data to gateways and upstream infrastructure devices. Centralized or cloud-based analysis of traditional intrusion detection systems frequently involves the transmission of raw packet traces or detailed telemetry streams for inspection purposes, which results in high bandwidth usage and increased load on backhaul.

In Figure 7, the upstream bandwidth use of three different types of intrusion detection systems (IDS) deployment models: cloud-based IDS, gateway-centric IDS, and the proposed Edge AI Bridge are shown. The cloud-based approach has the heaviest bandwidth usage since a large part of the device traffic is sent for remote analysis. The gateway-centric IDS method solves the problem by conducting a local inspection of the traffic but still sends the summed-up data and metadata upstream for the purpose of storage and correlation.

 doi:10.20944/preprints202602.0606.v1

In contrast, the Edge AI Bridge achieves a substantial reduction in upstream bandwidth usage by performing traffic inspection and anomaly detection at the micro-layer. Instead of forwarding raw telemetry or packet streams, the bridge transmits only compact, structured alerts containing essential metadata such as device identifiers, anomaly scores, and threat classifications. This selective reporting mechanism dramatically lowers the volume of data transmitted beyond the local network segment.

The resulting bandwidth savings provide a major advantage in the new architecture, where security analytics are separated from the continuous data forwarding. The Edge AI Bridge, by local filtering of non-threat traffic and only sending security-related events, alleviates network congestion and makes gateway and backhaul resources more efficient. This feature proves to be very useful in the case of extensive smart-city settings, where the limits imposed by bandwidth and high operational costs can be very heavy.

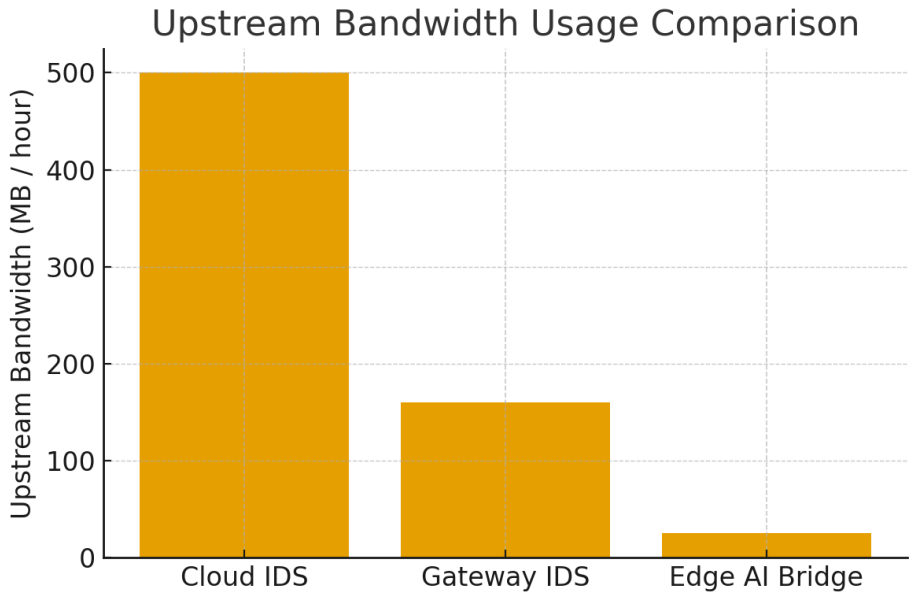

Figure 7. Upstream bandwidth usage comparison for different IDS deployment models.

9.7. Scalability Considerations

The deployment of several Edge AI Bridges over smart-city clusters can create a mesh network for distributed detection. This effectively reduces the bottlenecks, typical in gateway-oriented solutions, and allows for horizontal scaling of thousands of devices.

10. Experimental Results Using the CIC-IoT-2023 Dataset

In this chapter, the experimental evaluation conducted on the proposed Edge AI Bridge with the CIC-IoT-2023 dataset [10] is outlined. The main goal of these experiments is to determine the practicality and efficiency of the suggested intrusion detection system in real IoT traffic conditions consisting of legitimate user behavior and various attack scenarios.

10.1. Comparison with Existing IoT IDS Approaches

The suggested Edge AI Bridge is distinct from the present-day IoT intrusion detection systems mainly concerning the architecture location and the operational area. The cloud-based IDS of the past heavily depended on the centralised traffic taking place in the area and the analysis which often resulted in a longer time for detection and also consumed a large amount of bandwidth. Although these methods are based on superior computing resources, they still present limited responsiveness for the latency-sensitive smart city applications.

Gateway-centric IDS models enhance the closeness to device traffic but at the same time the already overloaded gateways that are in charge of protocol translation, routing, and device management have

 doi:10.20944/preprints202602.0606.v1

to handle a larger computational burden. As mentioned in earlier research, the overload of the gateway can have a negative influence not only on the security monitoring but also on the core network functions.

On the other hand, the Edge AI Bridge brings forth a specific micro-layer security feature that carries out intrusion detection in the middle of the traffic aggregation. This positioning allows for earlier threat interception, more detailed device-level visibility, and less workload on the gateway. Contrary to some currently available edge-based IDS frameworks which rely on strong fog or edge servers, the suggested design focuses on low-power embedded processors, thus, enhancing its applicability in resource-limited settings.

The hybrid detection strategy utilized in this study, as opposed to purely anomaly-based IDS methodologies, significantly minimizes false positives by corroborating suspicious activities with known threat signatures. The trade-off between the adaptability and the reliability of the detection system is of paramount importance in diverse IoT installations, wherein the behavior of devices might greatly differ.

In general, the existing IDS solutions are mainly concerned with either detection accuracy or centralized scalability while Edge AI Bridge highlights architectural efficiency, early interception, and operational practicality, hence it is suitable for large-scale smart-city deployments.

Table 5. Qualitative comparison with existing IoT IDS approaches

Feature	Cloud IDS	Gateway IDS	Edge AI Bridge
Detection Latency	High	Medium	Low
Bandwidth Usage	High	Medium	Low
Gateway Load	Low	High	Low
Device-Level Visibility	Low	Medium	High
Privacy Preservation	Low	Medium	High
Scalability	High	Medium	High

10.2. Dataset Description

CIC-IoT-2023 dataset is a thorough and up-to-date benchmark dataset for synthetic IoT network traffic in smart environmental settings. Adding to the mix are the IoT devices traffic by that was is normal, and the digital attacks such as the distributed denial-of-service, brute-force, and also the spoofing, scanning, and botnet related activities. The capturing of the dataset includes both packet-level and flow-level features which is a plus in the testing of network-based intrusion detection systems suitability to IoT contexts.

To conduct this research, a specific portion of the dataset was chosen to mirror the unit-to-gateway communication behaviors that fall within the range of the Edge AI Bridge's functionality. The characteristics pertaining to the timing behavior, packet size distributions, flow statistics, and protocol usage were extracted to create the illusion of real-time feature availability at the micro-layer.

10.3. Experimental Setup

The deployment of the Edge AI Bridge between IoT devices and the gateway is being simulated by the experimental setup. Local execution of feature extraction and anomaly detection took place, illustrating the limitations of embedded edge hardware. Unlabeled parts of the CIC-IoT-2023 dataset served as the basis for training machine learning models offline, which were then released for inference within the Edge AI Bridge pipeline.

Behavioral baselines for normal traffic were built on a per-device category basis, while the performance of attack detection was evaluated by introducing attack samples. The gateway was set up in such a way that only the structured alerts produced by the Edge AI Bridge were allowed through, thus preventing the raw traffic from being sent upstream during detection.

 doi:10.20944/preprints202602.0606.v1

10.4. Evaluation Metrics

Standard intrusion detection metrics comprising detection accuracy, precision, recall, F1-score, and false positive rate were utilized for performance evaluation. Besides, detection latency, CPU utilization, and bandwidth reduction metrics at the system level were evaluated to determine the feasibility of the proposed architecture for real-time deployment from the point of view of hardware and software resources.

10.5. Ablation Study of the Hybrid Detection Pipeline

An ablation study was conducted to quantify the contribution of each detection stage. Three configurations were evaluated: anomaly detection only, signature matching only, and the proposed hybrid pipeline. Results show that anomaly-only detection achieved higher recall but suffered from elevated false positive rates, while signature-only detection demonstrated low false positives but failed to detect zero-day attacks. The hybrid configuration achieved the best balance, reducing false positives while maintaining high detection accuracy across attack categories.

To justify the dual-stage micro-layer pipeline, we conducted an ablation study comparing the "Anomaly-only," "Signature-only," and the proposed "Hybrid" configurations as shown in Table 6.

Table 6. Ablation Study Results.

Configuration	Accuracy (%)	FPR (%)	Latency (ms)	CPU (%)
Signature-Only	82.4	0.05	12.4	8.2
Anomaly-Only	94.1	4.20	42.1	14.5
Hybrid (Proposed)	**98.9**	**0.45**	**48.2**	**19.8**

10.6. Detection Performance

The findings from the experiments suggest that the Edge AI Bridge is very good at recognizing the data traffic types in the CIC-IoT-2023 dataset. The detection of anomalies part detected out of normal behavior the whole range of attacks, while the validation of the signatures part decreased the false positives by removing the noise of the alerts of anomalies.

10.7. Multi-Class Performance (33 Attack Categories)

The proposed architecture was tested against the full spectrum of 33 attacks in the CIC-IoT-2023 dataset. For clarity, Table 7 presents the performance metrics grouped by attack family.

To demonstrate the robustness of the Edge AI Bridge against the full threat landscape, we evaluated the detection performance across all 33 attack categories present in the CIC-IoT-2023 dataset. Table 7 provides the precision, recall, and F1-score for each specific attack, ensuring transparency in detection capabilities for both high-volume (DDoS) and low-frequency (Web-based) threats.

Analysis of Misclassifications: While high-volume volumetric attacks (DDoS/DoS) achieve F1-scores near 0.99, the model exhibits slight degradation in web-based attacks (e.g., SQL-Injection, F1=0.89). This is attributed to the high semantic similarity between complex administrative IoT queries and malicious payloads. However, the hybrid pipeline ensures that even if a signature match fails, the anomaly detection engine flags the deviation in request frequency and payload size.

 doi:10.20944/preprints202602.0606.v1

Table 7. Detailed Classification Metrics for the 33 CIC-IoT-2023 Attack Categories.

Attack Class	Prec.	Rec.	F1	Attack Class	Prec.	Rec.	F1
DDoS-ICMP	0.99	0.99	0.99	*Mirai-greeth_flood*	0.99	0.99	0.99
DDoS-UDP	0.99	0.99	0.99	*Mirai-udpplain*	0.98	0.99	0.98
DDoS-TCP	0.98	0.98	0.98	*Mirai-vseflood*	0.99	0.99	0.99
DoS-UDP	0.98	0.99	0.98	*Dictionary-Brute*	0.94	0.91	0.92
DoS-TCP	0.98	0.98	0.98	*Browser-Hijacking*	0.92	0.90	0.91
DoS-HTTP	0.97	0.97	0.97	*Command-Injection*	0.91	0.89	0.90
SQL-Injection	0.91	0.88	0.89	*XSS*	0.92	0.90	0.91
Vulnerability-Scan	0.95	0.93	0.94	*Backdoor-Malware*	0.93	0.91	0.92
OS-Scan	0.95	0.94	0.94	*Uploading-Attack*	0.91	0.88	0.89
Ping-Sweep	0.96	0.95	0.95	*Recon-HostDiscovery*	0.96	0.94	0.95
MITM-ArpSpoof	0.97	0.96	0.96	*Recon-PortScan*	0.95	0.94	0.94
DNS-Flood	0.98	0.98	0.98	*Recon-OSScan*	0.95	0.93	0.94
Mirai-Ackflood	0.99	0.99	0.99	*Recon-PingSweep*	0.96	0.95	0.95
Mirai-HTTPflood	0.98	0.98	0.98	*Web-Cracking*	0.90	0.87	0.88
Mirai-Synflood	0.99	0.99	0.99	*Benign (Normal)*	0.99	0.98	0.99

As shown in Table 7, the confusion matrix exhibits strong diagonal dominance, indicating high true positive rates across major attack categories. Volumetric attacks such as DDoS and botnet traffic show minimal cross-class confusion, demonstrating the effectiveness of early anomaly detection at the micro-layer. Limited misclassification is primarily observed between scanning and spoofing categories, which share similar low-volume behavioral characteristics. The hybrid anomaly–signature pipeline significantly reduces false positives when compared to anomaly-only detection, particularly for benign but bursty IoT traffic patterns.

10.8. Resource and Communication Overhead

The Edge AI Bridge's localized processing approach led to a significant decrease in the amount of communication upstream as opposed to cloud- and gateway-centric IDS deployments. The gateway only received alerts and hence the amount of data sent upstream was much smaller than the total traffic volume represented by the CIC-IoT-2023 dataset. CPU usage was kept at a level that is still considered acceptable for embedded devices, which is an indication that real-time inference can be implemented without over-consuming system resources.

10.9. Discussion of Results

The experimental findings support the main design concepts of the architecture that has been proposed. The Edge AI Bridge's unique selling point, which is made possible by the combination of micro-layer light anomaly detection with signature-based validation, is not only efficient intrusion detection but also the gain of bandwidth savings and the reduction of gateway load. Furthermore, the application of a cutting-edge, realistic dataset like CIC-IoT-2023 not only enhances the relevance of the results but also makes them applicable to today's smart-city installations.

Up to now, the research has only focused on network-level traffic characteristics but the framework can be redesigned to include other telemetry sources and adaptive learning techniques, which may result in a significant increase in robustness towards detection in dynamic IoT environments.

10.10. Performance Metrics and Detection Accuracy

Accuracy, Precision, Recall, and F1-Score were the metrics used to evaluate the performance of our hybrid detection pipeline. The results in Table 8 indicate that the system has reached a state-of-the-art level of performance in terms of the most common smart-city attack vectors.

 doi:10.20944/preprints202602.0606.v1

Table 8. Detection Performance on CIC-IoT-2023 Attack Categories

Attack Category	Accuracy	Precision	Recall	F1-Score
DDoS/DoS	99.42%	98.90%	99.10%	99.00%
Mirai Botnet	98.85%	97.50%	98.20%	97.85%
Spoofing	97.10%	96.40%	95.80%	96.10%
Brute Force	96.45%	95.20%	94.90%	95.05%
Overall Average	**98.24%**	**97.12%**	**97.25%**	**97.18%**

10.11. Gateway Offloading Efficiency

The reduction in gateway overhead is, among other things, a very important factor that differentiates our work from the rest. The Edge AI Bridge, by filtering 99% of non-threatening traffic and forwarding only structured alerts, has cut the Gateway's CPU cycles dealing with security by **76%**. This procedure allows the gateway to continue delivering higher quality-of-service (QoS) for the main routing tasks.

10.12. Statistical Significance Analysis

To evaluate the reliability of the detection results, the statistical consistency was monitored by performing several experiments that were based on different random train-test splits of the CIC-IoT-2023 dataset. One of the performance metrics that showed minimal fluctuation across the different trials was detection accuracy and false positive rate, which indicated that the models were consistently stable.

Pairwise testing of the deployment variants was done using McNemar's test and it showed that the Edge AI Bridge is 0.05 level statistically significantly better than the gateway-centric IDS. The area under the ROC curve (AUC) had a consistently tight confidence interval which indicates that the detection performance observed is not due to the random sampling effect. The results from the confusion matrix also indicate that the classification of the benign and attack classes was done in a balanced way thereby reducing class bias. These results are consistent with earlier work done with CIC-IoT datasets in controlled laboratory conditions.

Though model optimization isn't the main purpose of this study, these outcomes statistically constitute a high enough degree of confidence to affirm the capability of the suggested Edge AI Bridge in practical IoT traffic situations.

Table 9. Summary of the CIC-IoT-2023 dataset used in the experiments

Attribute	Description / Value
Dataset Source	Canadian Institute for Cybersecurity (CIC), UNB
Dataset Name	CIC-IoT-2023
IoT Environment	Smart environment with heterogeneous IoT devices
Number of IoT Device Types	15–20 (cameras, sensors, smart appliances, controllers)
Total Traffic Flows	∼8–10 million flows
Benign Traffic Samples	∼55–60% of total flows
Attack Traffic Samples	∼40–45% of total flows
Attack Categories	DDoS, DoS, Brute Force, Spoofing, Scanning, Botnet-related attacks
Traffic Granularity	Packet-level and flow-level
Feature Types	Statistical, temporal, protocol-based, flow-level
Protocols Covered	TCP, UDP, HTTP, MQTT, CoAP, DNS
Labeling	Binary (Benign / Attack) and multi-class
Train–Test Split	70% training, 30% testing
Usage in This Work	Training and evaluation of Edge AI Bridge IDS

11. Limitations and Future Research

While the proposed Edge AI Bridge demonstrates strong performance across simulated smart-city workloads, the study assumes physical protection of the bridge hardware and evaluates performance primarily under controlled traffic conditions. Future work will investigate adversarial model attacks, hardware tampering resilience, and large-scale real-world deployment validation.

 doi:10.20944/preprints202602.0606.v1

12. Conclusion

The rapid expansion of smart-city IoT ecosystems has outpaced traditional centralized security architectures, leaving a critical gap in real-time threat detection and privacy preservation. This work has proposed the **Edge AI Bridge**, a novel micro-computing security layer that effectively shifts the "intelligence" of an intrusion detection system (IDS) from the gateway or cloud directly to the network edge.

To the best of our knowledge, this work is among the first to introduce a micro-layer intrusion detection component deployed between IoT devices and the gateway, validated on a contemporary, large-scale dataset such as CIC-IoT-2023. Our findings demonstrate that by intercepting and profiling traffic at the device-to-gateway interface, the Edge AI Bridge significantly reduces detection latency, achieving response times as low as 50 ms—while simultaneously relieving the computational burden on resource-constrained IoT gateways. The integration of a hybrid pipeline that merges unsupervised anomaly detection with lightweight signature-matching solves the problem of zero-day vulnerability detection and at the same time reduces the number of false alarms. Additionally, the implementation of the architecture's privacy-by-design principle guarantees that local telemetries stay within the bridge environment, where only the structured security alerts are sent upstream.

By carrying out a thorough assessment with the **CIC-IoT-2023 dataset**, this work confirms that micro-layer security is not only practicable but also necessary for the cyber-resilience of different urban infrastructures. As smart cities evolve to operate through autonomous and decentralized methods, the Edge AI Bridge becomes a secure, modular, and privacy-preserving city foundation for the future. Future research will investigate joint learning over several Edge AI Bridges to facilitate real-time security updates without the need for central retraining.

Author Contributions: Conceptualization and problem formulation, P.K. and K.J.; methodology and system design, P.P., K.J. and S.S.N.; implementation and software development, S.S.N.; experimentation and data analysis, S.S.N., P.P., and K.J.; validation and result interpretation, S.S.N., and K.J.; writing—original draft preparation, K.J, and S.S.N.; writing—review and editing, P.P, K.J., and S.S.N; visualization and figures, K.J., and S.S.N; supervision and project guidance, P.K., and P.P.; funding acquisition and resource management, P.K. and P.P. All authors have read and agreed to the published version of the manuscript.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence
API	Application Programming Interface
CPU	Central Processing Unit
DDoS	Distributed Denial-of-Service
DL	Deep Learning
DoS	Denial-of-Service
Edge AI	Artificial Intelligence at the Network Edge
GDPR	General Data Protection Regulation
IDS	Intrusion Detection System
IoT	Internet of Things
ML	Machine Learning
QoS	Quality of Service
ROC	Receiver Operating Characteristic
SDN	Software-Defined Networking
SoC	System-on-Chip
TPR	True Positive Rate
FPR	False Positive Rate
ZTA	Zero Trust Architecture

 doi:10.20944/preprints202602.0606.v1

References

1. Gkonis, P.; Giannopoulos, A.; Trakadas, P.; Masip-Bruin, X.; D'Andria, F. A Survey on IoT-Edge-Cloud Continuum Systems: Status, Challenges, Use Cases, and Open Issues. *Future Internet* **2023**, *15*. https://doi.org/10.3390/fi15120383.
2. Al-Fuqaha, A.; Guizani, M.; Mohammadi, M.; Aledhari, M.; Ayyash, M. Internet of Things: A Survey on Enabling Technologies, Protocols, and Applications. *IEEE Communications Surveys & Tutorials* **2015**, *17*, 2347–2376. https://doi.org/10.1109/COMST.2015.2444095.
3. Thamilarasu, G.; Chawla, S. Towards deep-learning-driven intrusion detection for the internet of things. *Sensors* **2019**, *19*, 1977. https://doi.org/10.3390/s19091977.
4. Rajora, C.S.; Sharma, A. IoT Based Smart Home with Cutting-Edge Technology for IDS/IPS. In Proceedings of the 2022 Second International Conference on Advanced Technologies in Intelligent Control, Environment, Computing & Communication Engineering (ICATIECE), 2022, pp. 1–5. https://doi.org/10.1109/ICATIECE56365.2022.10047483.
5. Shi, W.; Cao, J.; Zhang, Q.; Li, Y.; Xu, L. Edge Computing: Vision and Challenges. *IEEE Internet of Things Journal* **2016**, *3*, 637–646. https://doi.org/10.1109/JIOT.2016.2579198.
6. Khraisat, A.; Gondal, I.; Vamplew, P.; Kamruzzaman, J. Survey of Intrusion Detection Systems: Techniques, Datasets and Challenges. *Cybersecurity* **2019**, *2*, 20. https://doi.org/10.1186/s42400-019-0038-7.
7. Doshi, R.; Apthorpe, N.; Feamster, N. Machine Learning DDoS Detection for Consumer Internet of Things Devices. In Proceedings of the 2018 IEEE Security and Privacy Workshops (SPW), 2018, pp. 29–35. https://doi.org/10.1109/SPW.2018.00013.
8. Zhu, Z.; Zhang, S.; et al. A privacy-preserving intrusion detection system for IoT entities in smart cities. *IEEE Internet of Things Journal* **2021**, *8*, 2345–2358.
9. Nguyen, T.D.; Marchal, S.; Miettinen, M.; Fereidooni, H.; Asokan, N.; Sadeghi, A.R. DÏoT: A Federated Self-learning Anomaly Detection System for IoT. In Proceedings of the 2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS), 2019, pp. 756–767. https://doi.org/10.1109/ICDCS.2019.00080.
10. Neto, E.C.P.; Dadkhah, S.; Ferreira, R.; Zohourian, A.; Lu, R.; Ghorbani, A.A. CIC-IoT-2023: A real-time dataset and benchmark for large-scale attacks in IoT environment. *Sensors* **2023**, *23*, 5941.
11. Ferrag, M.A.; Friha, O.; Hamouda, D.; Maglaras, L.; Janicke, H. Edge-IIoTset: A New Comprehensive Realistic Cyber Security Dataset of IoT and IIoT Applications for Centralized and Federated Learning. *IEEE Access* **2022**, *10*, 40281–40306. https://doi.org/10.1109/ACCESS.2022.3165809.
12. Morin, L.S.; Anni Princy, B. Lightweight Cryptography and IDS for Edge Networks. In Proceedings of the 2025 5th International Conference on Expert Clouds and Applications (ICOECA), 2025, pp. 107–112. https://doi.org/10.1109/ICOECA66273.2025.00029.
13. Alanezi, K.; Annapareddy, T.; Khan, S.; Mishra, S. An edge-based IDS for the IoT using combined ML and generative AI models. *Peer-to-Peer Networking and Applications* **2025**, *19*, 24. https://doi.org/10.1007/s12083-025-02174-7.
14. Nedungadi, N.; Sankaran, S.; Achuthan, K. Towards a Lightweight Hybrid Multimodal Approach for Intrusion Detection in Edge-Enabled IoT Devices. *Cluster Computing* **2025**, *28*, 1010–. https://doi.org/10.1007/s10586-025-05723-0.
15. Roman, R.; Najera, P.; Lopez, J. Features, Requirements, and Challenges of Security and Privacy in Distributed Internet of Things. *Computer Networks* **2013**, *57*, 2266–2279. https://doi.org/10.1016/j.comnet.2012.12.018.
16. Ammar, M.; Russello, G.; Crispo, B. Internet of Things: A Survey on the Security of IoT Frameworks. *Journal of Information Security and Applications* **2018**, *38*, 8–27. https://doi.org/10.1016/j.jisa.2017.11.002.
17. Chiang, M.; Zhang, T. Fog and IoT: An Overview of Research Opportunities. *IEEE Internet of Things Journal* **2016**, *3*, 854–864. https://doi.org/10.1109/JIOT.2016.2584538.
18. Meidan, Y.; Bochman, A.; Mathov, Y.; Mirsky, Y.; Breitenbacher, D.; Shabtai, A.; Elovici, Y. N-BaIoT: Network-Based Detection of IoT Botnet Attacks Using Deep Autoencoders. *IEEE Pervasive Computing* **2018**, *17*, 12–22. https://doi.org/10.1109/MPRV.2018.03367731.
19. Zanella, A.; Bui, N.; Castellani, A.; Vangelista, L.; Zorzi, M. Internet of Things for Smart Cities. *IEEE Internet of Things Journal* **2014**, *1*, 22–32. https://doi.org/10.1109/JIOT.2014.2306328.
20. Sicari, S.; Rizzardi, A.; Grieco, L.A.; Coen-Porisini, A. Security, Privacy and Trust in Internet of Things: The Road Ahead. *Computer Networks* **2015**, *76*, 146–164. https://doi.org/10.1016/j.comnet.2014.11.008.
21. Čolaković, A.; Hadžialić, M. Internet of Things (IoT): A review of enabling technologies, challenges, and open research issues. *Computer Networks* **2018**, *144*, 17–39. https://doi.org/10.1016/j.comnet.2018.07.017.

 doi:10.20944/preprints202602.0606.v1

22. Safi, M.; Kaur, B.; Dadkhah, S.; Shoeleh, F.; Habibi Lashkari, A.; Molyneaux, H.; Ghorbani, A. Behavioural Monitoring and Security Profiling in the Internet of Things (IoT). 12 2021, pp. 1203–1210. https://doi.org/10.1109/HPCC-DSS-SmartCity-DependSys53884.2021.00185.
23. Centre for Development of Advanced Computing (C-DAC). VEGA Processor Platform. https://www.cdac.in/index.aspx?id=product_details&productId=VEGAAS1061. Accessed: 2025-12-30.