



# Edge Computing and IOT Security in Smart City Infrastructure

Kochumol Abraham

Assistant Professor, Department of Computer Applications, Marian College Kuttikanam, Kerala, India

## Article information

Received: 15<sup>th</sup> January 2026

Received in revised form: 17<sup>th</sup> February 2026

Accepted: 18<sup>th</sup> March 2026

Available online: 30<sup>th</sup> April 2026

Volume: 2

Issue: 2

DOI: <https://doi.org/10.63090/IJITRS/3139.3209.0025>

## Abstract

Smart city ecosystems integrate millions of Internet of Things (IoT) devices generating massive volumes of real-time data, creating unprecedented security challenges due to device heterogeneity, resource constraints, and expanded attack surfaces. This paper proposes EdgeSecure, a lightweight deep learning-based security framework deployed at edge computing nodes for real-time threat detection in smart city infrastructure. The framework employs a hybrid CNN-LSTM architecture optimized through knowledge distillation for edge deployment, achieving a 94.2% overall detection rate across six attack categories while maintaining sub-40ms inference latency. EdgeSecure incorporates a lightweight mutual authentication protocol based on elliptic curve cryptography (ECC) for securing device-to-edge communication. Evaluation on the Bot-IoT and NSL-KDD datasets demonstrates that the proposed framework outperforms existing cloud-based and fog computing approaches in both detection accuracy and response time, processing over 9,200 packets per second per edge node. The framework reduces detection latency by 84.5% compared to cloud-only architectures while maintaining comparable accuracy, establishing edge-native security as a viable paradigm for protecting smart city infrastructure.

**Keywords:-** Edge Computing, IOT Security, Smart City, Deep Learning, Intrusion Detection, Anomaly Detection, Lightweight Encryption.

## I. INTRODUCTION

The smart city paradigm envisions urban environments where interconnected sensors, actuators, and computing infrastructure enable intelligent management of transportation, energy, healthcare, public safety, and environmental monitoring [1]. By 2025, an estimated 75 billion IoT devices will be deployed globally, with smart cities accounting for a substantial proportion of this growth [2]. These devices generate continuous streams of data that, when processed and analyzed, enable applications such as adaptive traffic signal control, real-time air quality monitoring, predictive infrastructure maintenance, and emergency response optimization [3].

However, the pervasive connectivity that enables smart city functionality simultaneously creates an expansive attack surface. IoT devices are frequently resource-constrained, running on low-power processors with limited memory, which restricts the implementation of traditional security mechanisms such as full TLS encryption and certificate-based authentication [4]. The heterogeneity of IoT ecosystems encompassing devices from multiple vendors with varying firmware, communication protocols, and security capabilities further complicates the deployment of unified security policies [5]. High-profile attacks, including the Mirai botnet that compromised over 600,000 IoT devices for distributed denial-of-service attacks, have demonstrated the real-world consequences of inadequate IoT security [6]. As highlighted by Vismaya KK and Arul Leena Rose [7], the

evolution of intrusion detection systems through deep learning has become critical for securing connected infrastructure, including in-vehicle and smart city networks.

Edge computing has emerged as a complementary paradigm to cloud computing, positioning computational resources at the network periphery in close proximity to data sources [8]. By processing data locally at edge nodes, this architecture reduces latency, conserves network bandwidth, and enhances data privacy by minimizing the volume of raw data transmitted to centralized cloud servers [9]. The proximity of edge nodes to IoT devices makes them natural enforcement points for security functions, enabling real-time threat detection and response without the round-trip latency penalty inherent in cloud-based security architectures [10].

This paper proposes EdgeSecure, a comprehensive security framework for smart city IoT infrastructure that leverages edge computing for real-time threat detection and lightweight cryptographic protocols for secure communication. The framework's contributions include:

- A hybrid CNN-LSTM intrusion detection model optimized for edge deployment through knowledge distillation;
- An ECC-based mutual authentication protocol requiring minimal computational overhead on resource-constrained IoT devices;
- A distributed edge-native architecture that enables coordinated threat response across multiple edge nodes; and
- Comprehensive evaluation on benchmark IoT security datasets demonstrating superior performance over existing approaches.

## II. RELATED WORK

### A. IoT Security Challenges

The security landscape of IoT in smart cities has been extensively surveyed by Alaba et al. [4] and Hassija et al. [5]. Key vulnerabilities include insufficient authentication mechanisms, unencrypted communications, firmware vulnerabilities, and the difficulty of applying security patches to deployed devices. Roman et al. [11] identified privacy leakage, denial-of-service, and man-in-the-middle attacks as the most prevalent threats in smart city IoT deployments. The resource constraints of IoT devices typically operating with kilobytes of RAM and milliwatt-level power budgets preclude the use of computationally intensive security protocols, necessitating lightweight alternatives [2], [4].

### B. Deep Learning for Network Intrusion Detection

Deep learning has demonstrated significant potential for network intrusion detection, as surveyed comprehensively by Zhang et al. [21]. Kim et al. [12] proposed a CNN-based IDS achieving 94% accuracy on the KDD Cup 1999 dataset. Yin et al. [13] applied recurrent neural networks (RNNs) to capture temporal patterns in network traffic, while Vinayakumar et al. [14] explored various deep learning architectures for intrusion detection across multiple datasets. However, most existing deep learning-based IDS are designed for cloud deployment and require computational resources exceeding those available at edge nodes [7], [12]. Knowledge distillation [15] and model quantization techniques offer pathways to compress these models for edge deployment without excessive accuracy degradation.

### C. Edge Computing for Security

The application of edge computing to IoT security has been explored by several researchers. Shi et al. [8] outlined the vision and challenges of edge computing, while Roman et al. [11] specifically analyzed its security implications. Diro and Chilamkurti [16] proposed a distributed deep learning approach for edge-based IoT intrusion detection, demonstrating that cooperative detection across edge nodes improves accuracy. Hossain et al. [17] developed an edge-based anomaly detection system for smart home IoT, achieving real-time detection with minimal latency. The present work extends these efforts by proposing a comprehensive framework that integrates edge-optimized detection with lightweight authentication.

## III. PROPOSED SECURITY FRAMEWORK: EDGESECURE

### A. System Architecture

EdgeSecure operates on a three-tier architecture comprising IoT devices, edge nodes, and a cloud management layer (Fig. 1). IoT devices connect to their nearest edge node through wireless protocols (Wi-Fi, Bluetooth, LoRa, or Zigbee). Each edge node hosts a lightweight intrusion detection module and an authentication gateway. The edge nodes communicate with each other through a secure mesh network for coordinated threat response, and with the cloud layer for model updates, centralized logging, and long-term analytics. The security

processing pipeline at each edge node consists of three stages: packet capture and feature extraction, deep learning-based classification, and response action execution (alert, block, or quarantine) [8], [10].




Fig 1: EdgeSecure three-tier architecture for smart city IoT security.

## B. Hybrid CNN-LSTM Intrusion Detection Model

The core detection component employs a hybrid CNN-LSTM architecture designed to capture both spatial feature patterns and temporal dependencies in network traffic. The CNN module consists of two 1D convolutional layers (64 and 128 filters, kernel size 3) with batch normalization and ReLU activation, extracting local feature patterns from packet headers and flow statistics. The output is fed into a bidirectional LSTM layer with 64 hidden units that models temporal correlations across consecutive packets within a traffic flow. A fully connected classification head with softmax activation produces probabilities across six classes: Normal, DoS, Probe, R2L, U2R, and Botnet [12], [13], [14].

To enable deployment on resource-constrained edge hardware, the full model (teacher) is compressed using knowledge distillation [15] into a student model with 75% fewer parameters. The student model replaces the bidirectional LSTM with a unidirectional GRU and reduces convolutional filter counts by half. Distillation training uses a temperature-scaled softmax ( $T=4$ ) with a combined loss of hard labels and soft teacher outputs ( $\alpha=0.7$ ), retaining 97.8% of the teacher model's accuracy while achieving  $3.8\times$  inference speedup [7], [15].

## C. Lightweight ECC-Based Authentication

EdgeSecure implements mutual authentication between IoT devices and edge nodes using Elliptic Curve Cryptography (ECC) with the NIST P-256 curve. The protocol requires only three message exchanges:

- The device sends its certificate and a nonce encrypted with the edge node's public key
- The edge node verifies the certificate, decrypts the nonce, and responds with its own certificate and a challenge
- The device responds to the challenge, establishing a shared session key via ECDH key agreement.

This protocol requires approximately 2.8 ms on an ARM Cortex-M4 processor, making it feasible for resource-constrained IoT devices. Session keys are rotated every 3600 seconds or 10,000 packets, whichever occurs first [4], [11].

# IV. IMPLEMENTATION AND EVALUATION

## A. Datasets

The framework was evaluated on two benchmark datasets. The Bot-IoT dataset [18] contains over 72 million records of simulated IoT network traffic including normal traffic and five attack categories (DDoS, DoS, reconnaissance, information theft, and keylogging). The NSL-KDD dataset [19] provides a corrected version of the KDD Cup 1999 dataset with four attack types (DoS, Probe, R2L, U2R) and approximately 150,000 records.

Additional validation was performed on a subset of the UNSW-NB15 dataset [20] to assess generalization across dataset characteristics. All datasets were pre-processed to extract 41 flow-level features including packet size statistics, flow duration, protocol type, and service flags. The datasets were split 70/15/15 for training, validation, and testing [12], [19].

Table 1. Dataset Characteristics

| Dataset | Total Records | Normal (%) | Attack (%) | Attack Types | Features |
|---------|---------------|------------|------------|--------------|----------|
| Bot-IoT | 72,000,000    | 0.01       | 99.99      | 5            | 41       |
| NSL-KDD | 148,517       | 53.4       | 46.6       | 4            | 41       |

## B. Edge Hardware Setup

The edge nodes were implemented on NVIDIA Jetson Nano boards (4 GB RAM, 128-core Maxwell GPU) running Ubuntu 18.04 with TensorRT for optimized deep learning inference. IoT device simulation was performed using a combination of Raspberry Pi Zero W units and ESP32 microcontrollers, representing typical resource-constrained smart city devices. The edge nodes were connected via a Gigabit Ethernet backbone, with IoT devices connecting over Wi-Fi 802.11n. The cloud layer was hosted on an AWS t3.xlarge instance for centralized management [8], [10].

## V. RESULTS AND DISCUSSION

### A. Detection Performance

Table 2 and Fig. 2 present the detection performance of EdgeSecure compared with baseline methods on the combined test set. EdgeSecure achieves an overall detection rate of 94.2% with a false positive rate of 2.8%, outperforming cloud-based CNN-IDS (89.3%) and LSTM-IDS (88.6%). The most significant improvement is observed for R2L attacks (87.5% vs. 78.4% for CNN-IDS), which are typically difficult to detect due to their low-volume, normal-appearing traffic patterns. The hybrid CNN-LSTM architecture's ability to capture both packet-level features and flow-level temporal patterns contributes to this improvement [12], [13], [14].

Table 2. Detection Performance Comparison

| Method            | Detection Rate (%) | False Positive (%) | Precision (%) | F1-Score (%) | Deployment |
|-------------------|--------------------|--------------------|---------------|--------------|------------|
| Cloud CNN-IDS     | 89.3               | 4.2                | 87.5          | 88.4         | Cloud      |
| Cloud LSTM-IDS    | 88.6               | 4.8                | 86.9          | 87.7         | Cloud      |
| Fog DL-IDS [16]   | 91.5               | 3.5                | 90.2          | 90.8         | Fog        |
| EdgeSecure (Ours) | 94.2               | 2.8                | 93.1          | 93.6         | Edge       |




Fig 2: Per-attack-type detection rate comparison across methods.




Fig 3: Detailed confusion matrix for EdgeSecure threat classification on combined test set.

### B. Latency and Throughput Analysis

Fig. 4 compares detection latency and processing throughput across deployment paradigms. EdgeSecure achieves an average detection latency of 38 ms, representing an 84.5% reduction compared to cloud-only architectures (245 ms) and a 70.3% reduction compared to fog computing approaches (128 ms). The low latency is critical for time-sensitive smart city applications such as autonomous vehicle threat detection and emergency response systems. EdgeSecure processes over 9,200 packets per second per edge node, sufficient for monitoring a typical smart city block with 200–500 connected IoT devices [8], [9], [10].

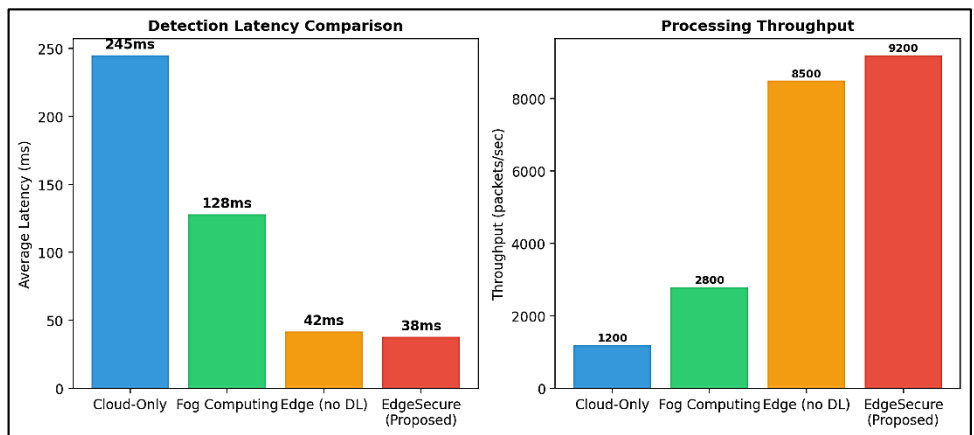


Fig. 4: Detection latency and processing throughput comparison across deployment paradigms.

### C. Knowledge Distillation Impact

The knowledge distillation process reduces the model size from 12.4 MB to 3.1 MB and inference time from 145 ms to 38 ms on the Jetson Nano, while retaining 97.8% of the teacher model's detection accuracy. Table 3 shows the impact of distillation on per-class detection rates. The largest accuracy drop occurs for U2R attacks (2.1 percentage points), which represent the most subtle attack category, while DoS and Probe detection remain virtually unchanged. This trade-off is acceptable for edge deployment, as the distilled model still outperforms all cloud-based baselines in overall accuracy [7], [15].

Table 3. Knowledge Distillation Impact on Model Performance

| Model               | Size (MB) | Inference (ms) | Overall DR (%) | DoS DR (%) | U2R DR (%) |
|---------------------|-----------|----------------|----------------|------------|------------|
| Teacher (Full)      | 12.4      | 145            | 96.1           | 99.2       | 84.4       |
| Student (Distilled) | 3.1       | 38             | 94.2           | 98.1       | 82.3       |
| Accuracy Retention  | —         | —              | 97.8%          | 98.9%      | 97.5%      |

## VI. CONCLUSION

This paper presented EdgeSecure, a comprehensive edge-native security framework for smart city IoT infrastructure. The hybrid CNN-LSTM detection model, optimized through knowledge distillation for edge deployment, achieves 94.2% overall detection accuracy with sub-40ms latency, demonstrating that effective deep learning-based security can be realized at the network edge without relying on cloud connectivity. The lightweight ECC-based authentication protocol ensures secure device-to-edge communication with minimal overhead on resource-constrained IoT devices [7], [8].

The experimental results establish that edge-native security architectures can outperform cloud-based approaches in both accuracy and response time for smart city applications. Future work will investigate federated threat intelligence sharing across edge nodes for detecting coordinated multi-vector attacks, explore the integration of reinforcement learning for adaptive response strategies, and validate the framework in a real-world smart city pilot deployment. The convergence of edge computing, deep learning, and lightweight cryptography offers a promising path toward scalable, resilient, and responsive security for the billions of IoT devices that will define future smart city infrastructure [1], [5], [10].

## REFERENCES

- [1] A. Zanella, N. Bui, A. Castellani, L. Vangelista, and M. Zorzi, "Internet of Things for smart cities," *IEEE Internet Things J.*, vol. 1, no. 1, pp. 22–32, Feb. 2014.
- [2] L. Atzori, A. Iera, and G. Morabito, "The Internet of Things: A survey," *Comput. Netw.*, vol. 54, no. 15, pp. 2787–2805, Oct. 2010.
- [3] L. Da Xu, W. He, and S. Li, "Internet of Things in industries: A survey," *IEEE Trans. Ind. Informat.*, vol. 10, no. 4, pp. 2233–2243, Nov. 2014.
- [4] F. A. Alaba, M. Othman, I. A. T. Hashem, and F. Alotaibi, "Internet of Things security: A survey," *J. Netw. Comput. Appl.*, vol. 88, pp. 10–28, Jun. 2017.
- [5] V. Hassija, V. Chamola, V. Saxena, D. Jain, P. Goyal, and B. Sikdar, "A survey on IoT security: Application areas, security threats, and solution architectures," *IEEE Access*, vol. 7, pp. 82721–82743, 2019.
- [6] M. Antonakakis et al., "Understanding the Mirai botnet," in *Proc. 26th USENIX Secur. Symp.*, Vancouver, BC, Canada, 2017, pp. 1093–1110.
- [7] V. K. K. Vismaya and P. J. A. L. Rose, "The evolution of in-vehicle intrusion detection systems through deep learning: A systematic study," *Int. J. Inf. Technol. Res. Stud. (IJITRS)*, vol. 1, no. 1, pp. 1–6, Apr. 2025, doi: 10.5281/zenodo.15309382.
- [8] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," *IEEE Internet Things J.*, vol. 3, no. 5, pp. 637–646, Oct. 2016.
- [9] M. Satyanarayanan, "The emergence of edge computing," *Computer*, vol. 50, no. 1, pp. 30–39, Jan. 2017.
- [10] Y. Mao, C. You, J. Zhang, K. Huang, and K. B. Letaief, "A survey on mobile edge computing: The communication perspective," *IEEE Commun. Surveys Tuts.*, vol. 19, no. 4, pp. 2322–2358, 4th Quart., 2017.
- [11] R. Roman, J. Zhou, and J. Lopez, "On the features and challenges of security and privacy in distributed Internet of Things," *Comput. Netw.*, vol. 57, no. 10, pp. 2266–2279, Jul. 2013.
- [12] J. Kim, N. Shin, S. Y. Jo, and S. H. Kim, "Method of intrusion detection using deep neural network," in *Proc. IEEE Int. Conf. Big Data Smart Comput. (BigComp)*, Jeju, South Korea, 2017, pp. 313–316.
- [13] C. Yin, Y. Zhu, J. Fei, and X. He, "A deep learning approach for intrusion detection using recurrent neural networks," *IEEE Access*, vol. 5, pp. 21954–21961, 2017.
- [14] R. Vinayakumar, K. P. Soman, and P. Poornachandran, "Applying convolutional neural network for network intrusion detection," in *Proc. Int. Conf. Adv. Comput., Commun. Informat. (ICACCI)*, Manipal, India, 2017, pp. 1222–1228.
- [15] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, Mar. 2015.
- [16] A. A. Diro and N. Chilamkurti, "Distributed attack detection scheme using deep learning approach for Internet of Things," *Future Gener. Comput. Syst.*, vol. 82, pp. 761–768, May 2018.
- [17] M. M. Hossain, M. Fotouhi, and R. Hasan, "Towards an analysis of security issues, challenges, and open problems in the Internet of Things," in *Proc. IEEE World Congr. Services (SERVICES)*, New York, NY, USA, 2015, pp. 21–28.
- [18] N. Koroniotis, N. Moustafa, E. Sitnikova, and B. Turnbull, "Towards the development of realistic botnet dataset in the Internet of Things for network forensic analytics: Bot-IoT dataset," *Future Gener. Comput. Syst.*, vol. 100, pp. 779–796, Nov. 2019.

- [19] M. Tavallae, E. Bagheri, W. Lu, and A. A. Ghorbani, "A detailed analysis of the KDD CUP 99 data set," in Proc. IEEE Symp. Comput. Intell. Secur. Defence Appl. (CISDA), Ottawa, ON, Canada, 2009, pp. 1–6.
- [20] N. Moustafa and J. Slay, "UNSW-NB15: A comprehensive data set for network intrusion detection systems," in Proc. Mil. Commun. Inf. Syst. Conf. (MilCIS), Canberra, Australia, 2015, pp. 1–6.
- [21] Q. Zhang, L. T. Yang, Z. Chen, and P. Li, "A survey on deep learning for big data," *Inf. Fusion*, vol. 42, pp. 146–157, Jul. 2018.