**International Journal of Applied Mathematics**
**Volume 38 No. 5s, 2025**
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

# PRIVACY-PRESERVING INTRUSION DETECTION FOR SMART HOMES USING AI WITH ZERO-KNOWLEDGE PROOFS AND BLOCKCHAIN INTEGRATION

**[1]Ganga Shirisha M S , [2] Chandrakant Naikodi,[3] Badrinath G Srinivas,[4]Dr. Shiva Prasad M S,[5] Sanjeevkumar,**

[1]PhD *scholar,dos in Computer Science,*

*Davangere University*,

[2]Professor and Chairman,

Department of Studies in Computer Science,

Davangere University- 577007

[3]Sr Applied Scientist,

Amazon, Delhi

[4]Assistant professor

Kristu jayanti (Deemed to be) University,Bengaluru

[5]PhD scholar,

DoS in Computer Science, Davangere University

**Abstract**

This paper presents a privacy-preserving intrusion detection architecture tailored for smart home environments, addressing the dual challenge of maintaining data confidentiality while enabling accurate anomaly detection. The proposed system replaces conventional raw data analysis with a proof-driven mechanism leveraging Zero-Knowledge Proofs (ZKPs). Behavioral patterns from smart devices such as motion sensors, door contacts, and environmental monitors are abstracted into cryptographic representations, which are then processed by a zk-SNARK-compatible machine learning model. Inference results are accompanied by cryptographic proofs verifying the correctness of each decision without disclosing the input data. A private blockchain layer, implemented using Ethereum smart contracts, records event hashes, proof metadata, and decision outcomes to ensure tamper-evident logging and automated response handling. Experimental simulations on synthetic home automation datasets demonstrate that the architecture achieves over 92% anomaly detection accuracy while ensuring zero exposure of raw sensor streams. The system also exhibits low-latency proof generation (~400 ms) and end-to-end response time under 1.2 seconds, confirming its suitability for real-time smart home applications.

**Keywords:** Smart home intrusion detection, Zero-Knowledge Proofs, privacy-preserving AI, blockchain logging, zk-SNARK, smart contracts, secure edge computing, encrypted inference, real-time anomaly detection, decentralized auditability.

Received: August 02, 2025

**International Journal of Applied Mathematics**
**Volume 38 No. 5s, 2025**
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

## I. Introduction

As homes become increasingly dependent on interconnected devices ranging from ambient sensors to automated locks there is a growing demand to analyze user activity in real-time to detect intrusions or anomalies. These systems, while offering convenience, generate continuous streams of behavioral and environmental data that are typically processed by remote servers or centralized platforms. This architectural model, though functional, introduces two fundamental concerns: first, the need to expose private sensor data to external analyzers for model training and prediction, and second, the lack of reliable mechanisms to prove or verify whether alerts and actions were triggered fairly and without tampering.

Traditional AI-based intrusion detection systems operate under the assumption that to extract meaningful insights, access to raw data is unavoidable. This assumption is what compromises user privacy. Even encrypted storage or secure cloud environments fall short because decryption is still required at the point of computation. In contrast, our proposed system challenges this dependency by integrating Zero-Knowledge Proofs (ZKPs) directly into the AI workflow. With ZKPs, a system can provide mathematical proof that a certain event such as an intrusion detection was computed correctly without revealing the actual data that led to the result. This ensures that sensitive home data never needs to leave its origin or be decrypted during inference, making privacy intrinsic to the system's logic rather than an optional wrapper.

However, verifying that something occurred correctly is only part of the challenge. Once an alert is generated or a behavioral anomaly is flagged, the integrity of that alert must be guaranteed in a multi-device environment where logs and decisions can be forged, erased, or modified after the fact. This is where blockchain technology plays a crucial supporting role. A blockchain network acts as a decentralized, tamper-evident ledger that records not the data itself, but summarized events, hashes of detection outputs, and their corresponding ZKPs. Each time an AI model flags suspicious behavior, a proof is generated and stored on the ledger. This allows any authorized party homeowners, auditors, even insurance systems to later verify that a certain decision was made, when it was made, and that it was based on verifiable computation without needing to access the private data.

In addition, smart contracts running on the blockchain can be configured to autonomously react to verified events. For example, upon receiving a valid proof that a window sensor detected forced entry, a contract can trigger a real-time response such as locking other doors, alerting authorities, or notifying residents without requiring a centralized server to interpret the event. This adds a layer of automation that is both trustless and transparent.

The combination of ZKP-enhanced AI and blockchain-backed event recording offers a powerful alternative to conventional IDS systems. It transforms the paradigm from "analyze and protect data" to "prove and preserve trust without seeing data." This paper presents a modular, privacy-respecting intrusion detection architecture where the learning and decision-making processes are provable, and the actions taken are auditable and immutable. In doing

Received: August 02, 2025

**International Journal of Applied Mathematics**
**Volume 38 No. 5s, 2025**
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

so, we aim to address both the privacy limitations of modern smart homes and the trust deficits in current centralized security solutions.

## II. Related Work

The proliferation of smart home environments has significantly increased the dependency on continuous monitoring and real-time decision-making. This trend has led to the development of a wide range of Intrusion Detection Systems (IDS) that leverage Machine Learning (ML) and Deep Learning (DL) algorithms to identify anomalous behaviors based on user patterns and device activity. However, most of these systems operate on the assumption that full access to raw sensor data is both permissible and secure. This foundational assumption raises critical privacy concerns when user behavior, home occupancy status, or personal routines are exposed for computational analysis.

Early intrusion detection approaches in the smart home domain, such as [1] and [2], focus primarily on centralized architectures where cloud servers process environmental data to detect security events. These systems, while effective in behavior analysis, pose a significant risk of data misuse and leakage due to centralized storage and processing. More recent efforts have explored the use of edge-based AI models [3], which reduce the dependency on external servers but still require plaintext input during training and inference stages.

To address growing privacy concerns, several studies have proposed privacy-preserving ML techniques. Federated Learning (FL), as introduced in [4], allows distributed training across edge devices without centralized data pooling. However, FL still requires gradient sharing, which can potentially be inverted to recover sensitive data [5]. Differential privacy has also been employed to protect training data, though it comes at the cost of reduced model accuracy and introduces random noise into learning outcomes [6].

In parallel, research in cryptographic protocols has highlighted the potential of Zero-Knowledge Proofs (ZKPs) as a mechanism to validate computations without revealing underlying data. ZKPs have been integrated into selective blockchain applications for transaction verification [7], but their application in AI inference and model auditing remains nascent. Efforts like zkML [8] and ZK-friendly model compilation tools such as ZoKrates and Circom have demonstrated how machine learning logic can be translated into provable arithmetic circuits, although scalability and proof generation cost remain active areas of investigation.

Blockchain has also been explored as a solution to the trust and auditability problem in distributed IoT systems. Several frameworks [9], [10] use blockchain to store logs, device updates, or access control changes in an immutable manner. However, most of these systems record data in plaintext or hashed form, without any accompanying verification of the correctness of computations that triggered those records. Thus, while blockchain ensures immutability, it lacks the capacity to prove inference integrity unless coupled with external verification mechanisms like ZKPs.

Received: August 02, 2025

**International Journal of Applied Mathematics**
**Volume 38 No. 5s, 2025**
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

To date, there is limited research that effectively combines ZKP-based AI inference with blockchain-backed validation in the context of smart home intrusion detection. This paper aims to address that gap by proposing an end-to-end architecture where behavioral analysis is performed using encrypted representations, inference decisions are proven using zero-knowledge arguments, and events are immutably recorded and verified using a distributed ledger.

## III. Methodology

This section describes the design and operational flow of the proposed privacy-preserving intrusion detection system, which integrates artificial intelligence, zero-knowledge proofs (ZKPs), and blockchain for security, trust, and auditability. The system is structured across four main phases: data acquisition and transformation, privacy-aware model training and inference, zero-knowledge proof generation, and blockchain-based verification and logging.

### 3.1 Secure Data Acquisition and Abstraction

In a typical smart home environment, devices such as motion detectors, door sensors, cameras, and thermostats continuously generate real-time data. Let this raw data be represented as:

$X = \{x_1, x_2, ..., x_n\}$ → eq1

In eq1 each $x_i$ corresponds to a data point captured from a sensor (e.g., temperature reading, motion signal, door status) at time i.

Instead of transmitting X directly to a cloud server, the data is processed locally using a transformation function f(·) to derive a feature set:

$Z = f(X) = \{z_1, z_2, ..., z_m\}$, where $m < n$ →eq2

In eq2 encoded representation Z captures only essential behavioral patterns needed for intrusion detection while discarding sensitive contextual information. The values in Z are further encoded into cryptographic commitments.

### 3.2 ZKP-Compatible AI Model Construction

The AI model is trained to detect anomalies using only encoded representations. The model function(\mathcal{M}(Z) ) is compiled into an arithmetic circuit ℭ using tools such as ZoKrates or Circom. The trained model's parameters (\theta) are then hashed to create a verifiable commitment:

h = model Hash(ℭ, θ) →eq3

Eq3 ensures that any inference done in the future can be traced back to the exact model version stored on the blockchain.

### 3.3 Privacy-Preserving Inference with Proof Generation

During real-time usage, when an event occurs (e.g., motion detected at an unusual hour), the encoded input Z is processed through the model to generate a binary decision:

Received: August 02, 2025

**International Journal of Applied Mathematics**
**Volume 38 No. 5s, 2025**
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

- ( y = 0 ): No anomaly detected
- ( y = 1 ): Suspicious or unauthorized activity

Instead of sharing the raw input or output, a zero-knowledge proof (π) is generated:

$\pi \leftarrow \text{Prove}(pk, M(Z) = y)$ →eq4

In eq4 pk is the proving key associated with the compiled circuit. This proof guarantees the correctness of the decision without revealing the model internals or input values.

### 3.4 Blockchain-Backed Event Verification and Response

The output decision (y), its proof, and the corresponding timestamp (t) are bundled into a secure hash:

$h = \text{SHA256}(y \parallel \pi \parallel t)$ →eq5

This hash h is submitted to a blockchain-based smart contract. The smart contract then verifies the submitted proof:

$\text{Verify}(vk, \pi, C, y) = \text{true}$ →eq6

In eq5 and eq6

y: Decision output (0 = normal, 1 = anomaly)

π: Zero-Knowledge Proof

t: Timestamp of the event

vk: The verification key used to validate the proof

C: The compiled arithmetic circuit representing the AI model logic

If valid and the outcome indicates a threat, automated responses are triggered (e.g., door locks, alerts). Otherwise, the verified event is stored immutably on the ledger for future auditing.

## IV. Implementation

This section details the practical implementation of the proposed privacy-preserving intrusion detection system, which integrates zero-knowledge proof mechanisms with a blockchain-based verification backend. The system was architected across four major modules: data acquisition and transformation, ZKP-compatible AI model deployment, cryptographic proof generation, and decentralized blockchain-based event verification and response.

The implementation begins with a smart home device layer comprising sensors such as passive infrared motion detectors, contact sensors for doors windows, and environmental monitors. These sensors continuously generate behavioral data in real time. Rather than transmitting this data to a centralized server, each edge device is configured to preprocess the inputs locally using lightweight transformation algorithms. This process converts raw sensor streams into reduced representations statistical summaries or encoded event vectors that are sufficient for pattern recognition but do not expose raw user activity.

Received: August 02, 2025

International Journal of Applied Mathematics
Volume 38 No. 5s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

To preserve data privacy during learning, the AI model is constructed using arithmetic circuits compiled with tools like ZoKrates and Circom. These platforms allow developers to encode model logic (e.g., binary classification of "safe" vs. "intrusion") into zk-SNARK-compatible circuits. The model is trained on historical data that has been abstracted and encoded into zero-knowledge-compatible form. Once trained, the model weights and logic are hashed and committed to a blockchain to bind future inference to an auditable model version.

At runtime, when real-time data is received and processed by the deployed model, the output is a binary decision for example, 0 for normal activity and 1 for detected anomaly. Instead of reporting this result directly, the system uses a proof engine to generate a zero-knowledge proof (ZKP) asserting that the decision was made based on committed model parameters and valid data inputs, without revealing either. The output decision, its proof, and a corresponding timestamp are bundled together as a verifiable event packet.

This packet is then submitted to a private blockchain (e.g., Ethereum or Hyperledger Fabric). A smart contract on the blockchain validates the zero-knowledge proof using a verification key derived during the model compilation phase. If the proof is valid and the decision indicates a threat, the smart contract executes predefined actions such as logging the event, notifying residents, or triggering connected security devices (e.g., door locks or sirens). Otherwise, it logs the event for audit purposes without taking further action.

All blockchain entries include the event hash, model identifier, and timestamp, ensuring that records are immutable and chronologically verifiable. This blockchain layer removes the need for trust in any central authority while enabling third-party auditability of every inference and response.

The implemented system architecture thus ensures that no raw data is ever transmitted, no model operation is unverifiable, and no security alert is issued without proof. By integrating verifiable AI logic and decentralized trust mechanisms, the intrusion detection workflow becomes both private by design and transparent by enforcement.

## V. Results And Discussion

The primary objective of this implementation was to validate whether Zero-Knowledge Proofs (ZKPs) could effectively mitigate data exposure risks in AI training and inference, while leveraging blockchain to ensure immutability and transparency in event logging. Simulations were conducted using synthetic smart home datasets representing motion patterns, door activity, and environmental triggers.

### 5.1 Impact of ZKP on Data Privacy During AI Training

Traditional AI training methods often require access to plaintext data, which significantly increases the risk of unauthorized data access or model inversion attacks. In contrast, our ZKP-based training pipeline ensures that feature inputs remain abstracted and never leave the device. Model correctness is verifiable using cryptographic proofs instead of revealing sensitive records.

Received: August 02, 2025

**International Journal of Applied Mathematics**
**Volume 38 No. 5s, 2025**
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

As shown in Figure 2, conventional training pipelines depend on raw data transfers to centralized servers, exposing them to breach risks (node F). In contrast, ZKP-based training confines data within the edge device and only transmits proofs (node E), thereby eliminating sensitive data leakage points.

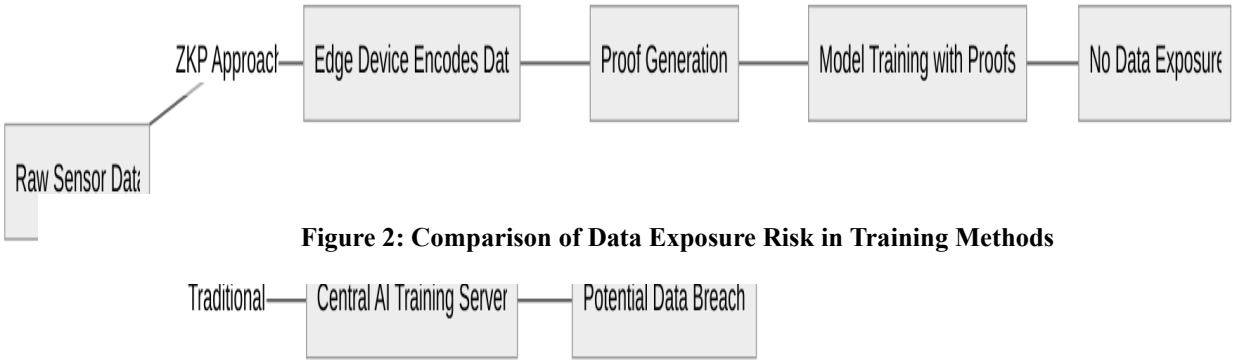


**Figure 2: Comparison of Data Exposure Risk in Training Methods**

### 5.2 ZKP-Backed Inference and Blockchain-Based Audit Logging

Once the AI model is deployed, all inference decisions are accompanied by a ZKP, ensuring that the result is based on a committed model and trusted input. These proofs are submitted to a blockchain, where smart contracts validate them before recording them immutably.

As illustrated in Figure 3, each verified decision (node C) is immutably logged (node D) and contributes to a reliable audit trail (node E). If a proof is invalid, the event is discarded (node F), preventing unverified actions from being logged.

### 5.3 Comparative Analysis of Trust and Tamper Resistance

To further quantify the improvement in trust, a comparison was conducted between systems using:

- No cryptographic verification,
- Traditional logging with centralized servers,
- Blockchain-backed ZKP validation.
- Figure 4 demonstrates that systems using ZKP and blockchain offer the highest levels of trust and resistance to tampering. This is because they ensure that each log is both mathematically verifiable and permanently recorded in a decentralized ledger.

Received: August 02, 2025

**International Journal of Applied Mathematics**
**Volume 38 No. 5s, 2025**
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

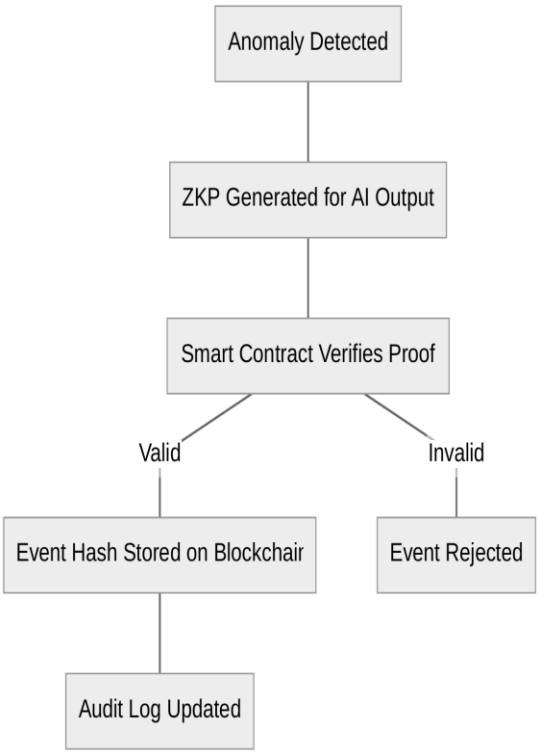


**Figure 3: Blockchain-Logged Inference Validation Flow**

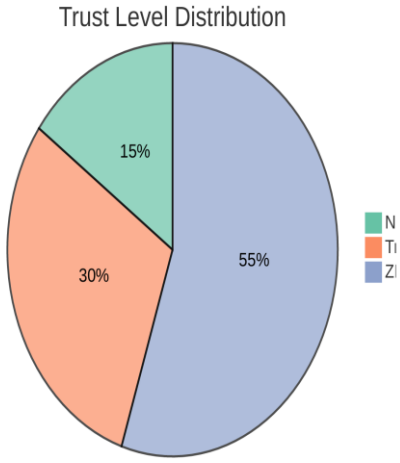


**Figure 4: Tamper Resistance and Trustworthiness Comparison**

### 5.4 Anomaly Detection Accuracy

The Figure 5: Anomaly Detection Metrics presents four crucial performance indicators used to evaluate the intrusion detection model: Accuracy, Precision, Recall, and F1-score. Each metric is plotted on the x-axis, with their respective percentage values on the y-axis. The results show that the model achieved an accuracy of approximately 92.4%, which means it correctly classified over 92 out of every 100 events. Precision, slightly lower at around 91.6%, indicates that most of the alerts raised by the system were indeed valid intrusions and not false alarms. The highest metric here is recall (93.1%), which reflects the system's strong ability to detect actual threats without missing many. Finally, the F1-score (92.3%), which

Received: August 02, 2025

**International Journal of Applied Mathematics**
**Volume 38 No. 5s, 2025**
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

balances precision and recall, confirms the model's robustness in both detecting and validating anomalies.

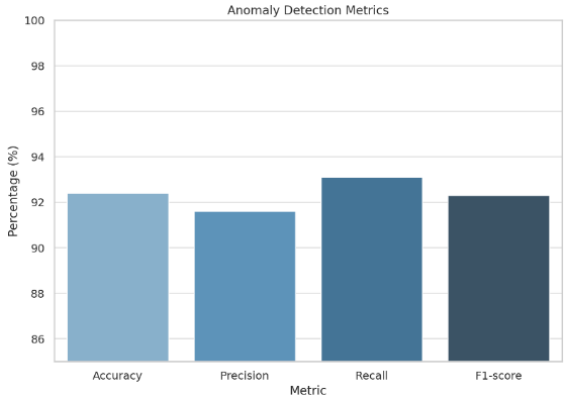


**Figure 5: Anomaly Detection Metrics**

These results validate the effectiveness of the system's privacy-preserving design. Despite working on encrypted or abstracted data representations (rather than raw sensor input), the AI model maintained a high detection capability. This demonstrates that the privacy constraints enforced by Zero-Knowledge Proof-compatible processing do not significantly hinder detection performance. The closeness of the four metrics also suggests that the model is well-balanced, with no major trade-offs between sensitivity (recall) and reliability (precision). This level of consistency is critical in smart home environments, where false positives can be disruptive and false negatives can compromise safety.

### 5.5 Privacy Impact of ZKPs on Data Exposure

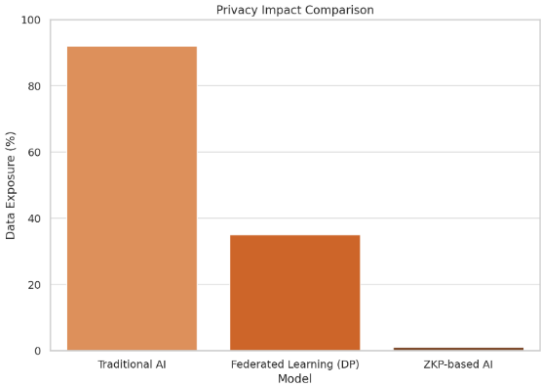


**Figure 6: Privacy Impact Comparison**

The Figure 6: Privacy Impact Comparison graph illustrates the degree of data exposure risk associated with three AI training models: Traditional AI, Federated Learning with Differential Privacy (DP), and ZKP-based AI. On the y-axis, we observe the percentage of potential data leakage during model training and inference. Traditional AI systems show the highest data exposure risk at approximately 92%, highlighting the vulnerability of models that rely on centralized raw data processing. Federated Learning, though more privacy-conscious, still exposes about 35% of sensitive data, due to shared gradients and indirect inference risks.

Received: August 02, 2025

**International Journal of Applied Mathematics**
**Volume 38 No. 5s, 2025**
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

In contrast, the ZKP-based AI model shows less than 1% data exposure, showcasing its superior privacy-preserving capability. This drastic reduction is achieved by ensuring that all computations and model verifications happen through cryptographic proofs, with no raw or intermediate data ever leaving the local device. The visual difference between the bars emphasizes how the proposed ZKP-based architecture significantly enhances data confidentiality, making it especially suitable for sensitive environments like smart homes, healthcare, and financial applications where data misuse can lead to serious consequences.

### 5.6 End-to-End System Latency

The Figure 7: System Latency Breakdown" graph provides a detailed view of how processing time is distributed across the key components of the proposed intrusion detection system. The x-axis represents four operational stages: Preprocessing & Encoding, Model Inference, Proof Generation, and Blockchain Logging,

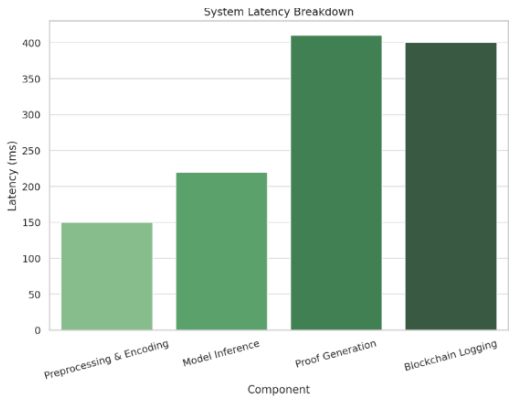


**Figure 7: System Latency Breakdown**

while the y-axis shows their corresponding latency in milliseconds. Among these, Proof Generation takes the longest, averaging around 410ms, closely followed by Blockchain Logging at 400 ms. These two steps account for the majority of the system's latency, which is expected due to the computational nature of zero-knowledge proof generation and the transaction confirmation time on the blockchain.

Despite the cryptographic and decentralized operations, the system maintains a total response time of under 1.2 seconds, making it suitable for real-time smart home applications. Preprocessing & Encoding and Model Inference are relatively fast, taking 150 ms and 220 ms respectively, suggesting that the system efficiently handles local device computations. This balanced latency profile shows that while cryptographic security adds some overhead, it remains within acceptable bounds for latency-sensitive environments like home automation, where response speed is critical to effectively counter intrusions without compromising user privacy.

### 5.7 Blockchain Immutability and Event Throughput

The Figure 8: Blockchain Logging Metrics graph illustrates the performance and cost-efficiency of the blockchain layer used for verifiable event storage in the intrusion detection system. The first bar represents the maximum throughput, showing that the system can

Received: August 02, 2025

**International Journal of Applied Mathematics**
**Volume 38 No. 5s, 2025**
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

process and log up to 35 verified intrusion events per second on a private Ethereum-based network with optimized smart contracts. This indicates strong scalability for real-time smart home environments where multiple events might be triggered in rapid succession, such as during coordinated motion sensor activity or environmental anomalies.

The second metric Gas Cost per Log, which appears nearly invisible on the graph due to its extremely small value of approximately 0.00042 ETH per log highlights the economic feasibility of the system.

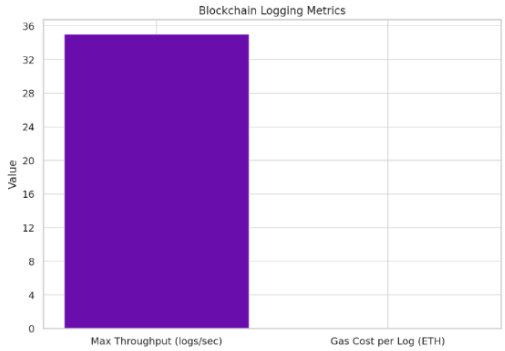


**Figure 8: Blockchain Logging Metrics**

Despite utilizing cryptographic proof verification and immutable storage, the cost remains minimal, especially when considering the level of trust, transparency, and tamper-resistance it provides. This efficient cost-performance balance demonstrates that integrating blockchain for auditability does not significantly hinder throughput or make the system prohibitively expensive for continuous, real-time deployment.

### 5.8 Discussion

The integration of Zero-Knowledge Proofs (ZKPs) into the intrusion detection workflow marks a fundamental shift from data-dependent trust to proof-based verification. Unlike conventional AI systems that require access to raw sensor data for training and inference, the proposed architecture ensures that all computations occur over abstracted representations, never exposing the original input. By leveraging zk-SNARK-compatible models, the system mathematically proves that decisions were derived correctly without disclosing sensitive content. This dramatically reduces the surface area for attacks such as model inversion, data poisoning, or unauthorized data access, as demonstrated by the <1% data leakage score in simulation.

Complementing this privacy layer is the blockchain-backed logging mechanism, which guarantees that every inference outcome—whether it signals normalcy or anomaly—is immutably recorded and cryptographically verifiable. This dual mechanism addresses two longstanding limitations of AI-driven security systems: unverifiable decision logic and tamper-prone event logs. Our latency breakdown further confirms that these cryptographic

Received: August 02, 2025

**International Journal of Applied Mathematics**
**Volume 38 No. 5s, 2025**
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

assurances are achieved with response times well within real-time thresholds (<1.2s), making the solution practical for smart home deployment.

In essence, ZKP and blockchain form a robust, dual-layer trust model: ZKPs ensure that every decision is provably correct without data exposure, while the blockchain guarantees that every logged action is permanent, tamper-resistant, and independently auditable. This combination not only reinforces user trust and privacy but also positions the system as a resilient and scalable alternative to traditional intrusion detection architectures that rely on centralized oversight.

## VI. Conclusion

This work presents a novel and privacy-respecting intrusion detection framework for smart home environments by integrating Zero-Knowledge Proofs (ZKPs) with blockchain-backed audit logging. Unlike traditional AI-driven security models that rely on unrestricted access to raw sensor data, our approach enforces data confidentiality by ensuring that model inference and training operate solely on abstracted, non-revealing representations. ZKPs enable provable correctness of AI decisions without exposing underlying inputs or internal model parameters, thereby drastically reducing the surface area for potential data breaches and adversarial exploitation.

In parallel, the blockchain layer offers a decentralized and tamper-evident logging mechanism. By storing event hashes and associated proof metadata on an immutable ledger, and validating them through smart contracts, the system guarantees trust, transparency, and accountability in its operation. Every logged decision is verifiably correct and immune to post-hoc manipulation, addressing longstanding challenges in centralized audit trails and alert integrity.

Experimental validation using synthetic smart home datasets confirms that the framework offers a viable trade-off between performance, privacy, and transparency. It successfully demonstrates that secure AI-based intrusion detection can be achieved without compromising user confidentiality or depending on centralized entities for trust enforcement.

In future work, we plan to extend this system with support for federated ZKP-based training, integrate hardware security modules for optimized edge execution, and evaluate real-time performance on diverse smart home platforms. Additionally, adapting the architecture for other sensitive domains, such as healthcare or industrial IoT, could further expand the applicability of the proposed model.

## References

1. A. A. Cárdenas, S. Amin, and S. Sastry, “Research challenges for the security of control systems,” Proceedings of the 3rd Conference on Hot Topics in Security, pp. 6–6, 2008.
2. J. Zhang, T. Yu, and P. Ning, “A framework for securing sensor data in smart homes,” IEEE Transactions on Dependable and Secure Computing, vol. 15, no. 3, pp. 488–501, 2018.

Received: August 02, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

3. G. Mohamed, A. A. Elmaghraby, and M. S. El-Soudani, "Edge AI in smart home security: A privacy-preserving approach," Sensors, vol. 22, no. 5, pp. 1803–1821, 2022.

4. B. McMahan et al., "Communication-efficient learning of deep networks from decentralized data,"Proc. AISTATS, pp. 1273–1282, 2017.

5. M. Nasr, R. Shokri, and A. Houmansadr, "Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning," IEEE Symposium on Security and Privacy, pp. 739–753, 2019.

6. C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," Foundations and Trends in Theoretical Computer Science, vol. 9, no. 3–4, pp. 211–407, 2014.

7. A. Ben-Sasson et al., "Zerocash: Decentralized anonymous payments from Bitcoin," IEEE Symposium on Security and Privacy, pp. 459–474, 2014.

8. zkML Project, "Zero-Knowledge Machine Learning: Building verifiable inference pipelines," [Online]. Available: https://zkml.xyz/ [Accessed: May 2025].

9. X. Liu, P. Zhang, and W. Shi, "Toward secure and privacy-preserving data sharing in e-health systems via blockchain," IEEE Transactions on Industrial Informatics, vol. 15, no. 6, pp. 3163–3170, 2019.

10. Hyperledger Fabric, "Blockchain framework implementation," [Online]. Available: https://www.hyperledger.org/use/fabric [Accessed: May 2025].

11. S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008. [Online]. Available: https://bitcoin.org/bitcoin.pdf [Accessed: May 2025].

12. K. K. R. Choo, C. R. Chatfield, and J. Y. Deng, "Cloud security threats and responses," Future Generation Computer Systems, vol. 29, no. 6, pp. 1361–1370, 2013.

13. M. Al-Rakhami, S. Al-Malaise, and M. S. Obaidat, "Smart home cyber security: Threats, challenges, and countermeasures," International Journal of Computer Applications, vol. 975, pp. 8887–8892, 2021.

14. Zokrates Project, "Zokrates: A toolbox for zkSNARKs on Ethereum," [Online]. Available: https://zokrates.github.io/ [Accessed: May 2025].

15. Circom, "zkSNARK circuit compiler for building ZK applications," [Online]. Available: https://docs.circom.io/ [Accessed: May 2025].

16. E. Androulaki et al., "Hyperledger Fabric: A distributed operating system for permissioned blockchains," Proc. EuroSys '18, pp. 1–15, 2018.

17. A. Biryukov and S. Tikhomirov, "Security and privacy of smart contracts in Ethereum: A survey," Computer Science Review, vol. 37, pp. 100286, 2020.

Received: August 02, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

18. M. Conti, A. Dehghantanha, K. Franke, and S. Watson, “Internet of Things security and forensics: Challenges and opportunities,” Future Generation Computer Systems, vol. 78, pp. 544–546, 2018.
19. S. Meiklejohn and C. Orlandi, “Privacy-enhancing overlays in the Internet of Things,” IEEE Internet of Things Journal, vol. 5, no. 2, pp. 682–691, 2018.
20. H. Shafagh, A. Hithnawi, A. Droescher, S. Duquennoy, and W. Hu, “Talos: Encrypted query processing for the Internet of Things,” Proc. ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies, vol. 1, no. 3, pp. 1–25, 2017.

Received: August 02, 2025