Stallion Journal for Multidisciplinary Associated Research Studies
ISSN (Online): 2583-3340
2nd International Conference on Recent Trends in Multidisciplinary Research and Innovation (ICMRI)

A Blockchain-Driven, Machine Learning-Enabled Adaptive Security Framework for IoT Ecosystems

Chandra Bhushan[1] and Dr. Seema Tripathi[2]
[1]Integral & Innovative Sustainable Education (IISE) College, Lucknow, Uttar Pradesh, INDIA.
[2]International Institute for Special Education, Lucknow, Uttar Pradesh, INDIA.

www.sjmars.com || PP. 36-40 || https://doi.org/10.55544/sjmars.icmri.6

Date of Submission: 05-07-2025 **Date of Publication:** 27-07-2025

ABSTRACT

The exponential growth of the Internet of Things (IoT) has revolutionized multiple domains—ranging from smart homes and healthcare monitoring to industrial automation and smart cities. Yet this proliferation of connected, resource-constrained devices has also dramatically expanded the attack surface, exposing networks to data tampering, device impersonation, denial-of-service (DoS) attacks, and unauthorized access. Traditional, centralized security measures struggle to keep pace with the dynamic and heterogeneous nature of IoT environments. In this paper, we propose a hybrid security framework that synergizes blockchain technology and machine learning (ML) to deliver a decentralized, tamper-resistant, and adaptive protection mechanism for IoT ecosystems. Blockchain provides immutable audit trails, decentralized trust, and programmable enforcement via smart contracts, while ML offers real-time anomaly detection and predictive threat analytics. We describe the architecture and workflows of our framework, outline our implementation using a permissioned Hyperledger Fabric network and edge-deployed ML models (including LSTM for sequential anomaly detection), and present simulation results showing over 97% detection accuracy, a false-positive rate below 3%, and acceptable transaction latencies (<1 s) on resource-constrained devices. We conclude that the integration of blockchain and ML yields a resilient security posture that can adapt autonomously to emerging threats, scale to millions of devices, and maintain low overhead on edge hardware.

Keywords- Internet of Things (IoT), Blockchain Security, Machine Learning, Anomaly Detection, Hyperledger Fabric, Smart Contracts, Edge Computing, Adaptive Security.

I. INTRODUCTION

By 2030, IoT devices are projected to exceed 50 billion worldwide, enabling unprecedented data-driven automation in healthcare, transportation, manufacturing, energy, and domestic environments. However, the very characteristics that make IoT so powerful—ubiquity, heterogeneity, and lightweight design—also exacerbate security challenges. Many devices lack hardware-based root of trust, run outdated firmware, and cannot support heavyweight cryptographic protocols. Dependence on centralized cloud servers for authentication and logging introduces single points of failure, and attackers are increasingly leveraging compromised IoT endpoints as vectors for large-scale attacks, such as the Mirai botnet (2016), which leveraged default credentials to conscript hundreds of thousands of cameras and routers into a massive DDoS campaign.

1.1 Overview of IoT Growth and Associated Security Risks

The Internet of Things (IoT) has seen explosive growth in recent years, with estimates projecting over 50 billion connected devices globally by 2030. These range from consumer gadgets (smart speakers, wearables) to critical infrastructure (industrial sensors, medical implants). While this connectivity enables unprecedented levels of automation, data-driven insights, and user convenience, it also dramatically expands the attack surface. Many IoT devices are built with minimal processing power and limited onboard memory, making them unable to support robust security protocols. As a result, they can become entry points for attackers seeking to intercept data, inject malicious code, or commandeer devices into botnets for distributed denial-of-service (DDoS) attacks. Further, heterogeneity in hardware, firmware, and communication protocols means that vulnerabilities in one device type can cascade through an entire network.

1.2 Need for Adaptive, Scalable, and Intelligent Security Frameworks

Traditional security architectures—centered on perimeter firewalls, signature-based intrusion detection, or periodic patching—are ill-suited to the dynamic, large-scale nature of IoT ecosystems. An effective IoT security framework must be:

- It should automatically adjust defenses when new threats emerge or network conditions change, without requiring manual intervention for each device.
- It needs to protect thousands or even millions of endpoints, distributing security functions so as not to overwhelm any single node or central server.
- It should leverage data analytics and machine learning to distinguish between benign anomalies (e.g., firmware updates) and true security incidents, minimizing false alarms and ensuring rapid response.

1.3 Blockchain technology offers three key properties that directly address core IoT security challenges:

- **Decentralization:** Instead of relying on a single, potentially vulnerable server for authentication or data storage, blockchain spreads data across multiple distributed nodes. This removes single points of failure and makes it extremely difficult for an attacker to take down the entire system.
- **Immutability:** Once a transaction (e.g., a device's firmware-update record or a sensor reading) is committed to the chain and confirmed by consensus, it cannot be altered or deleted. This tamper-proof ledger ensures data integrity, enabling trustworthy audit trails and forensics when investigating security incidents.
- **Transparency:** Every participating node holds a copy of the ledger and can verify transactions independently. While sensitive payloads can be encrypted or stored off-chain, metadata and integrity proofs remain visible, fostering trust without sacrificing privacy.

1.4 Machine learning (ML) brings intelligence and adaptability to IoT security

- **Predictive Analytics:** ML models can mine historical device and network data to forecast potential failures or attack trends before they occur. For example, a model trained on past traffic spikes and latency patterns might predict an impending DDoS campaign, triggering preemptive countermeasures.
- **Anomaly Detection:** Unsupervised or semi-supervised ML algorithms (e.g., clustering, autoencoders) can learn the "normal" behavior profiles of devices—typical packet sizes, communication frequencies, or energy consumption patterns. When a device deviates significantly from its learned baseline, the system flags it for further inspection, catching zero-day exploits and novel malware strains that signature-based systems would miss.

The central aim of this research is to architect and validate a hybrid security framework that leverages blockchain for decentralized trust and machine learning for intelligent, data-driven threat detection. In this model:

1. **Device Registration & Authentication:** Smart contracts on a permissioned blockchain verify each device's identity before granting network access.
2. **Immutable Logging:** Every critical event—firmware updates, configuration changes, alert statuses—is hashed and recorded on the blockchain for auditability.
3. **Real-Time Monitoring:** A distributed ML engine (deployed at edge gateways or in the cloud) continuously analyzes streaming telemetry for anomalies, feeding its findings back into the blockchain via new transactions or "alert" flags.
4. **Automated Response:** Upon detection of suspicious behavior, pre-defined smart contracts can revoke a device's credentials, isolate it from the network, or trigger remediation workflows (e.g., over-the-air patch distribution).

II. PROBLEM STATEMENT

IoT ecosystems must satisfy several critical security requirements to remain robust and reliable. First, data integrity is essential to ensure that sensor readings and control commands arrive unaltered from their source to their destination. Equally important is device authentication, which guarantees that only legitimate devices can join and interact on the network. Real-time anomaly detection is needed to identify novel or evolving attack patterns as they emerge, and the solution must also scale to support millions of devices while imposing minimal CPU, memory, and power overhead. Traditional defenses—such as firewalls, VPNs, and signature-based intrusion detection systems—are often inadequate against distributed, zero-day attacks and struggle to adapt to the dynamic topologies of modern IoT deployments. Consequently, there is an urgent need for a security architecture that is adaptive, decentralized, and intelligently proactive.

III. LITERATURE REVIEW

IoT networks face a wide range of complex security challenges that stem from both the diversity of attack vectors and the inherent resource constraints of connected devices. Service availability can be compromised by impulse attacks such as distributed denial-of-service (DDoS) or jamming, while integrity attacks—like replay or man-in-the-middle—can silently tamper with critical sensor readings and control messages. Identity attacks, including spoofing and Sybil strategies, enable adversaries to masquerade as legitimate devices and gain unauthorized access. Traditional security measures—relying on

heavyweight asymmetric cryptography or centralized servers—often introduce prohibitive latency or create single points of failure, making them poorly suited to the dynamic, low-power environments characteristic of IoT deployments.

Blockchain technology offers a promising remedy by providing a decentralized, append-only ledger in which every device transaction—sensor data, configuration updates, and security alerts—is immutably recorded. In permissioned blockchains such as Hyperledger Fabric, Byzantine Fault Tolerant (BFT) consensus algorithms efficiently validate transactions among trusted participants, while smart contracts codify and automatically enforce access-control policies. Prior research has demonstrated blockchain's utility for secure firmware distribution, decentralized device identity management, and tamper-proof audit logging. Nevertheless, purely blockchain-based solutions often struggle with the privacy of on-chain data and the computational overhead required by consensus protocols, particularly when deployed on constrained IoT hardware.

Machine learning (ML) techniques have also been widely explored to bolster IoT security through intelligent threat detection. Supervised, semi-supervised, and unsupervised methods—ranging from Support Vector Machines (SVM) and Random Forests to deep neural networks like LSTM models and autoencoders—can profile normal device behavior and flag deviations indicative of zero-day exploits or emerging attack patterns. While ML systems excel at identifying novel threats, they depend heavily on the availability of high-quality, representative training datasets and can suffer from concept drift as adversaries adapt their tactics over time.

To address the limitations of standalone blockchain or ML approaches, recent work increasingly advocates for hybrid architectures that integrate both technologies. In these models, blockchain ensures the trustworthiness and provenance of telemetry data before it enters ML pipelines, and the ML-generated anomaly alerts are in turn committed back to the ledger to trigger smart-contract-driven remediation workflows. Although this synergy promises end-to-end security and automation, many existing prototypes fall short in practical IoT settings—they often rely on public blockchains with high transaction costs, omit edge or gateway-level deployment of ML for real-time response, or require manual intervention for critical policy updates. Continued research is needed to refine these integrated frameworks for scalable, privacy-preserving, and fully automated IoT security.

IV. PROPOSED FRAMEWORK

Our proposed framework is organized into four cohesive layers, each fulfilling specific roles to ensure secure, intelligent, and resilient IoT communication.

Our framework comprises four seamlessly integrated layers that collectively provide adaptive, decentralized security for IoT ecosystems. At its foundation, the **IoT Device Layer** employs MQTT/CoAP over TLS for secure sensor-to-gateway communication, while edge gateways (e.g., Raspberry Pi) aggregate telemetry, perform initial feature extraction (packet size, inter-arrival times, CPU/memory usage, firmware version), and run lightweight ML inferences locally. Sitting above this, the **Blockchain Security Layer** uses a Hyperledger Fabric permissioned network managed by device manufacturers, the network operator, and an auditor; its smart contracts—DeviceRegistry for registering device identities and keys, AccessControl for verifying permissions and rate-limiting, and AlertContract for logging anomaly alerts—ensure tamper-proof device authentication and audit trails. Concurrently, the **ML Monitoring Layer** employs an LSTM model to detect subtle, time-series anomalies (like slow-drip attacks) along with a Random Forest classifier to identify overt threats (such as volume-based DoS), continuously analyzing the extracted features. The top-level **Decision Layer** fuses blockchain-logged events with real-time ML alerts to trigger automated remediation—smart contracts can revoke suspect devices' credentials or isolate them at the gateway—while providing operators with a unified dashboard showing device health, anomaly history, and immutable blockchain records. The end-to-end workflow starts with device onboarding via a DeviceRegistry transaction, proceeds through encrypted data transmission and on-chain logging by AccessControl, followed by live ML scoring and AlertContract recording when thresholds are exceeded, and culminates in contract-driven containment actions that uphold the integrity and resilience of the IoT network.

V. METHODOLOGY

We evaluated our hybrid security framework using both real-world and synthetic datasets within a smart-home–style testbed. For network anomaly detection, we utilized the labeled UNSW-NB15 dataset and complemented it with Node-RED–generated IoT traffic that included replay and spoofing attacks. Our physical testbed consisted of 20 Raspberry Pi 4 devices acting as virtual sensor nodes, all communicating with a private Hyperledger Fabric network deployed across four Ubuntu VMs. Prior to training our machine learning models, we normalized all feature vectors and applied PCA for dimensionality reduction. We then split the combined dataset into 70% training and 30% testing sets, and used 5-fold cross-validation to fine-tune hyperparameters such as the number of LSTM layers and the depth of the Random Forest. Model performance was measured using accuracy, precision, recall, F1-score, and ROC-AUC to ensure reliable, low-false-positive anomaly detection.

On the blockchain side, we configured a permissioned Fabric network with Practical Byzantine Fault Tolerance (PBFT) running on four ordering nodes to balance throughput and resilience. Block parameters were set to a 1 MB maximum size and a 1 s block interval, providing low-latency logging suitable for IoT workflows. Smart contracts—written in Go and tested via Hyperledger Composer—handled device registration, access control, and alert logging, ensuring each critical event (from onboarding transactions to anomaly alerts) was immutably recorded and automatically enforced without manual intervention.

VI. EXPERIMENTAL OUTCOME

The hybrid framework outperformed both standalone ML and blockchain implementations across multiple metrics. It achieved a detection accuracy of 97.4%—higher than the 94.2% accuracy of the ML-only approach—by leveraging the blockchain's verified data inputs. Its false-positive rate dropped to 2.3%, compared to 6.8% for ML alone, demonstrating that on-chain context helps reduce spurious alerts. Although the hybrid solution introduced a slight increase in transaction latency (0.95 s) versus the blockchain-only layer (0.8 s), it still delivered sub-second performance, with a throughput of roughly 80 transactions per second—comparable to the blockchain-only rate of 85 tx/s and well within typical IoT requirements. Finally, edge gateways in the hybrid setup registered a CPU usage of 65%, only 5% higher than the 60% observed for ML-only operations, an acceptable overhead given the significant security gains.

Challenges and Limitations

Our framework's reliance on the PBFT consensus mechanism introduces notable messaging overhead, which can impede performance as the network scales to hundreds or thousands of nodes. In real-world deployments, this extra communication may increase latency and reduce throughput, particularly in geographically dispersed environments. Additionally, many ultra-low-power IoT devices cannot support local ML inference due to their limited CPU and memory resources; this constraint necessitates a hierarchical architecture in which more capable edge gateways handle the bulk of feature extraction and model evaluation.

On the machine learning side, the risk of model drift is ever-present: as attackers develop new techniques, previously trained models may fail to recognize novel threats. Continuous monitoring, periodic retraining, and validation against fresh data are therefore essential to maintain detection accuracy. At the same time, care must be taken to preserve data privacy**:** only metadata (such as timestamps and hashes) should be recorded on-chain, while sensitive payloads must remain encrypted or stored off-chain, adding complexity to the system's data management strategy.

Finally, our evaluation faced several practical limitations. Testing was confined to a 20-node Raspberry Pi network, leaving industrial-scale performance and resilience to intermittent connectivity unmeasured. Likewise, the synthetic attack scenarios generated via Node-RED may not fully capture the subtle nuances of real-world adversarial behavior. Addressing these gaps will require larger, more varied testbeds, robust offline/online synchronization mechanisms, and richer, field-derived datasets to ensure our framework's robustness under diverse operational conditions.

VII. FUTURE DIRECTION

Future work will focus on enhancing privacy and efficiency by adopting **federated learning**, distributing model training across edge gateways to keep raw data local and reduce central computation overhead. We also plan to investigate **lightweight consensus mechanisms**—such as DAG-based or leaderless protocols like the IOTA Tangle—specifically tailored for resource-constrained IoT environments. To accelerate on-device inference, we will explore **edge AI accelerators** (for example, Google Coral TPUs or NVIDIA Jetson modules) that deliver high-performance ML workloads with minimal power consumption. Finally, we aim to achieve **cross-chain interoperability** so that multiple blockchain networks—both public and private—can seamlessly exchange and verify IoT telemetry, enabling richer data sharing and more robust security across heterogeneous infrastructures.

VIII. CONCLUSION

We have presented a hybrid security framework that marries the tamper-proof, decentralized nature of blockchain with the adaptive intelligence of machine learning to secure IoT ecosystems. Through a permissioned Fabric deployment, edge-hosted LSTM and Random Forest models, and smart-contract-driven access control, our system achieves >97% detection accuracy**,** <3% false **positives**, and sub-second ledger operations on resource-constrained hardware. This research demonstrates that blockchain and ML are complementary: blockchain ensures the trustworthiness of data feeding ML, and ML provides the adaptive, predictive capabilities that static blockchains lack. Our prototype paves the way toward resilient, self-healing IoT networks capable of defending themselves against next-generation cyber threats.

REFERENCES

[1] Christidis, K., & Devetsikiotis, M. (2016). *Blockchains and Smart Contracts for the Internet of Things*. IEEE Access, 4, 2292–2303.
[2] Dorri, A., Kanhere, S. S., & Jurdak, R. (2017). *Blockchain in Internet of Things: Challenges and Solutions*. arXiv preprint arXiv:1608.05187.
[3] Mohammadi, M., Al-Fuqaha, A., Sorour, S., & Guizani, M. (2018). *Deep Learning for IoT Big Data and Streaming Analytics: A Survey*. IEEE Communications Surveys & Tutorials**,** 20(4), 2923–2960.
[4] Sharma, P. K., Chen, J. H., & Park, J. H. (2018). *A Software Defined Fog Node Based Distributed Blockchain Cloud Architecture for IoT***.** IEEE Access, 6, 115–124.
[5] Nguyen, D. C., Pathirana, P. N., Ding, M., & Seneviratne, A. (2021). *Federated Learning for Smart Healthcare: A Survey*. ACM Computing Surveys, 54(7), Article 141.
[6] Reyna, A., Martin, C., Chen, J., Soler, E., & Díaz, M. (2018). *On Blockchain and Its Integration with IoT: Challenges and Opportunities*. Future Generation Computer Systems, 88, 173–190.
[7] Pop, C., Seceleanu, C., & Crăciunescu, R. (2020). *Lightweight Consensus Algorithms for IoT—Applications and Challenges*. IEEE Access, 8, 137164–137190.
[8] Li, X., Jiang, P., Chen, T., Luo, X., & Wen, Q. (2020). *A Survey on the Security of Blockchain Systems*. Future Generation Computer Systems, 107, 841–853.
[9] Al-Kuwaiti, H. (2020). *A Comprehensive Survey on Blockchain for IoT*. IEEE Communications Surveys & Tutorials**,** 22(4), 2529–2556.
[10] Ullah, I., Khan, M. K., Aalsalem, M. Y., & Bangash, Y. A. (2019). *Machine Learning for Cybersecurity in IoT: A Survey*. IEEE Internet of Things Journal, 6(4), 6285–6304.
[11] Han, K., Mao, H., & Dally, W. J. (2016). *Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding*. International Conference on Learning Representations (ICLR)**.**
[12] Zyskind, G., Nathan, O., & Pentland, A. (2015). *Decentralizing Privacy: Using Blockchain to Protect Personal Data*. IEEE Security & Privacy, 16(4), 28–36.
[13] Kang, J., Yu, R., Huang, X., Maharjan, S., Zhang, Y., & Hossain, E. (2019). *Blockchain for Secure and Efficient Data Sharing in Vehicular Edge Computing and Networks*. IEEE Internet of Things Journal, 6(3), 4660–4670.
[14] Zhang, Y., & Chen, M. (2020). *Blockchain and Deep Reinforcement Learning for Secure Resource Allocation in Edge Computing*. IEEE Transactions on Emerging Topics in Computational Intelligence, 4(5), 525–535.
[15] Kshetri, N. (2017). *1 Blockchain's Roles in Meeting Key Supply Chain Management Objectives*. International Journal of Information Management, 39, 80–89.
[16] Dorri, A., Kanhere, S. S., Jurdak, R., & Gauravaram, P. (2019). *Blockchain for IoT Security and Privacy: The Case Study of a Smart Home*. Proceedings of the IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops), 618–623.
[17] Nguyen, G., Kim, K., & Kim, S. (2020). *A Method for Detecting Anomalies in IoT Data Using Convolutional Neural Networks*. Sensors, 20(1), 257.
[18] Li, S., Da Xu, L., & Zhao, S. (2018). *The Internet of Things: A Survey*. Information Systems Frontiers, 17(2), 243–259.
[19] Xu, X., Weber, I., & Staples, M. (2019). *Architecture for Blockchain Applications*. Springer.
[20] Wu, H., Xu, J., & Zheng, Z. (2021). *Secure Machine-Learning-as-a-Service for IoT Devices: A Lightweight Framework*. IEEE Internet of Things Journal, 8(7), 5737–5750.