IJIRCCE

ISSN(O): 2320-9801

ISSN(P): 2320-9798

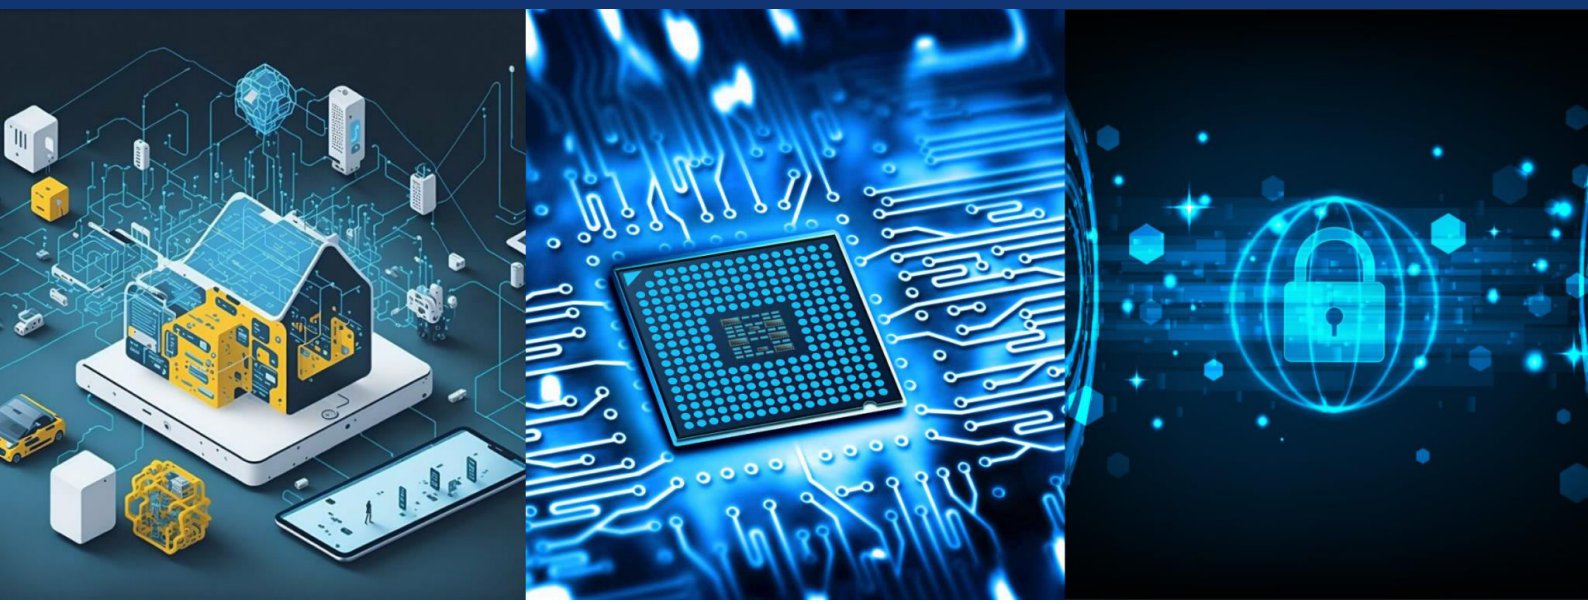

International Journal of Innovative Research in Computer and Communication Engineering

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

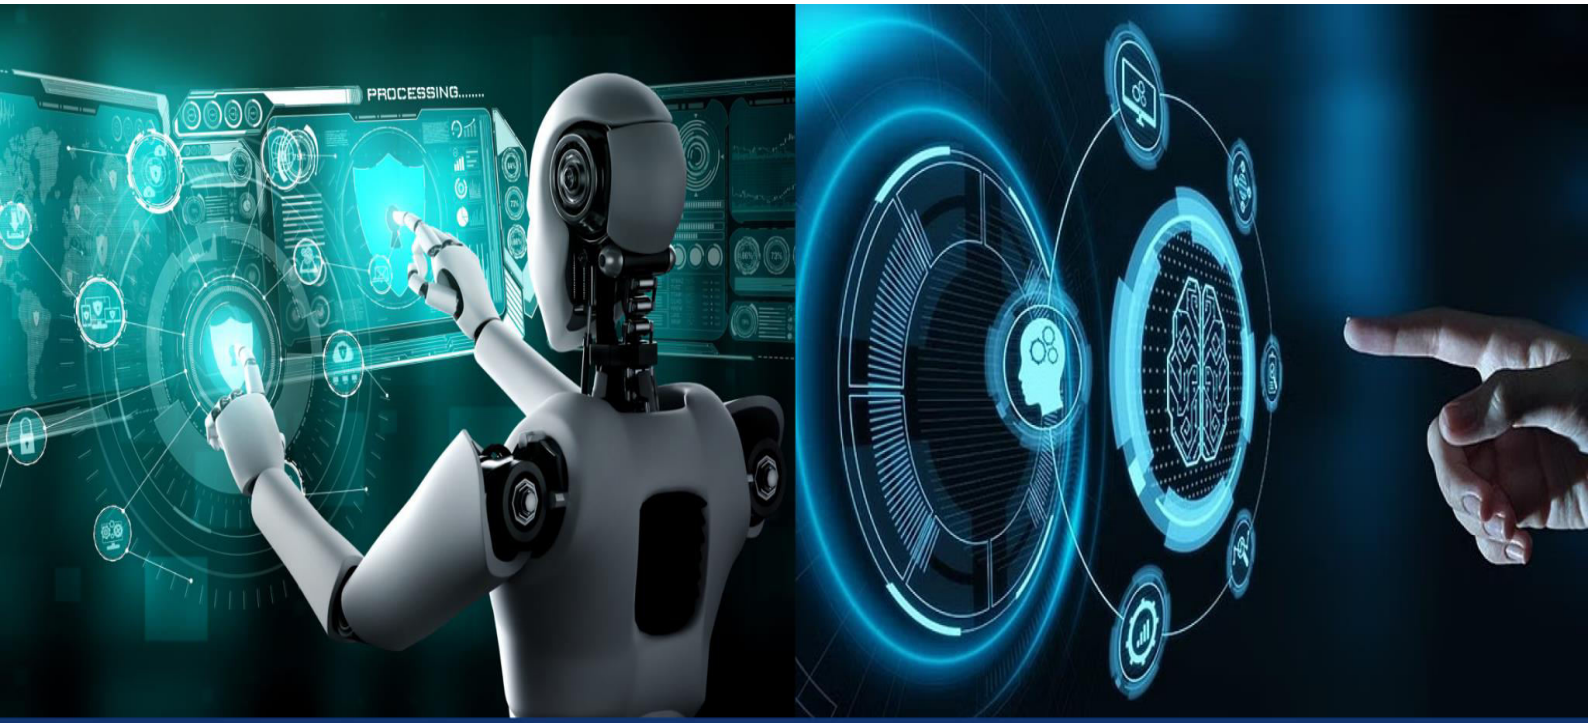

Impact Factor: 9.927

Volume 14, Issue 5, May 2026

© 2026 IJIRCCE | Volume 14, Issue 5, May 2026 DOI: 10.15680/IJIRCCE.2026.1405002

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

A Privacy-Preserving IoMT Security Framework Integrating Transformer Autoencoder Anomaly Detection with Hyperledger Sawtooth Blockchain

Choukpin Adoto Mignonkoun Sourou Yannick, Hojamyradova Meyrem, Salami Ibrahim Olusola

MSc Student, Dept. of CST and Software, Nanjing University of Information Science and Technology, Nanjing, China

MSc Student, Dept. of Information Science and Technology, Nanjing Forestry University, Nanjing, China

BSc Student, Dept. of CST and Software, Nanjing University of Information Science and Technology, Nanjing, China

ABSTRACT: The rapid proliferation of Internet of Medical Things devices, including wearable sensors, implantable monitors, and cloud-interfaced diagnostic systems, has introduced a critical and expanding attack surface alongside its transformative clinical benefits. IoMT environments are characterized by resource-constrained hardware, heterogeneous protocols, and highly sensitive patient data transmission, making them prime targets for botnet-driven distributed denial-of-service attacks, data exfiltration, and network intrusion. Existing security frameworks treat anomaly detection and audit logging as loosely coupled modules, lacking the latency guarantees, privacy protections, and regulatory alignment required for live hospital infrastructure. This paper proposes a tightly integrated, privacy-preserving security framework combining a Transformer-based Autoencoder for real-time anomaly detection with a Hyperledger Sawtooth permissioned blockchain for immutable audit logging. The Transformer Autoencoder is trained exclusively on benign traffic from the N-BaIoT dataset using an unsupervised reconstruction-loss paradigm, learning a compact latent representation of normal behavior that enables detection of deviations induced by sophisticated attack traffic. Upon detection, a middleware pipeline constructs a privacy-preserving metadata payload comprising a device identifier, timestamp, anomaly score, and SHA-256 cryptographic hash of the raw input, committing it to the distributed ledger without exposing any sensitive data on-chain.

Empirical evaluation demonstrates 98.6% accuracy, 97.1% precision, and 98.4% recall on the N-BaIoT test set, outperforming eight contemporary baselines whilst maintaining 19ms inference latency. The blockchain layer sustains commit latency below 90 milliseconds at a two-kilobyte payload, with linear throughput scalability from 20 to 48 transactions per second. The off-chain hash-only storage model resolves the blockchain-GDPR paradox, satisfying both Article 17 right-to-erasure obligations and HIPAA transmission security standards through deliberate architectural design rather than post-hoc compliance measures.

KEYWORDS: Internet of Medical things; Privacy Preserving; Blockchain; Deep Learning; Transformer Autoencoder

I. INTRODUCTION

The rapid digitization of healthcare has produced the Internet of Medical Things (IoMT): a connected ecosystem of wearable sensors, implantable devices, remote monitoring platforms, and cloud-interfaced health systems. Real-time physiological signals, from continuous glucose monitoring and ECG readings to post-operative implant telemetry, now traverse networks at scale, enabling faster diagnosis, reduced admissions, and personalized care[1][2]. However, this transformation introduces a critical security burden. Every transmitted packet and stored record is a potential attack surface, and a compromised medical record directly threatens patient safety, clinical decisions, and institutional trust [3]. Regulatory frameworks such as GDPR, HIPAA, and HL7 FHIR impose strict obligations that conventional security designs struggle to meet.

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

IoMT devices are resource-constrained; a pacemaker or insulin pump cannot host full encryption without compromising clinical function. Traditional centralized defences (server-side encryption, ACLs, perimeter firewalls) present single points of failure, lack modern auditability, and are ill-suited to distributed, multi-institutional healthcare [3].

Addressing this requires a shift to decentralized, intelligent, proactive systems. Blockchain offers immutability and cryptographic auditability for healthcare data governance [4][5]; permissioned architectures (e.g., Hyperledger Sawtooth) balance transparency with patient confidentiality. Transformer Autoencoders complement this: self-attention captures long-range dependencies in network traffic, and the compressed bottleneck is inherently privacy-preserving.

Despite growing interest in combining blockchain with AI for healthcare security [6][7], few works tightly integrate a Transformer Autoencoder for edge-level anomaly detection with a permissioned blockchain for privacy-respecting audit logging, address the tension between blockchain immutability and GDPR's right to erasure, or empirically validate end-to-end latency under real-time constraints. This paper proposes and evaluates such a framework: a Transformer Autoencoder trained on benign traffic for real-time detection, paired with a Hyperledger Sawtooth ledger that commits only SHA-256 hashes and non-sensitive metadata on-chain. The remainder of the paper is organized as follows: Section II reviews related work; Section III presents the system architecture and component design; Section IV details experiments and results; Section V discusses findings and concludes.

II. RELATED WORK

Three principal privacy-preserving approaches dominate the IoMT security literature, each addressing a distinct concern but carrying practical limitations that prevent standalone adoption. Differential privacy (DP) mitigates disclosure risk by injecting controlled noise into model outputs or gradient updates, yet stronger privacy guarantees directly erode detection accuracy, a trade-off that is clinically unacceptable where missed anomalies carry patient safety implications [3]. Federated learning (FL) addresses data centralization by training models locally across distributed nodes and sharing only parameter updates, though the non-IID nature of medical device traffic routinely causes model divergence, with reported recall drops of 7–10% across heterogeneous hospital environments [4]. Homomorphic encryption (HE) enables computation directly on encrypted data without decryption, offering strong confidentiality guarantees, but its computational overhead renders real-time inference on resource-constrained IoMT devices infeasible with current implementations [8]. These limitations collectively motivate the hybrid framework proposed in this work, which pursues privacy through architectural design rather than through any single technique.

Blockchain consists of a chain of cryptographically linked blocks, ensuring that recorded information cannot be altered or lost [9]. Beyond its cryptocurrency origins, the technology has attracted substantial healthcare research interest, with over 600 studies published in the past three years alone [10]. Its core properties map directly onto healthcare data management requirements: decentralization eliminates single points of failure in hospital information systems [11], immutability ensures that medical records cannot be inappropriately modified [6], and smart contracts enable automated, intermediary-free data governance. Real-world deployments support these claims: Estonia's KSI blockchain now records nearly all citizens' health transactions, while the MedRec project at Beth Israel Deaconess Medical Center demonstrated successful reconciliation of records across two separate EHR systems whilst preserving patient control over their own data [7]. In operational terms, a 2024 IEEE prototype applying multi-signature smart contracts to insurance claims reduced fraudulent payments by 42% [5], and immutable IPFS-backed audit logs reduced unauthorized edit detection time by 87% in intensive care settings[12]. Collectively, these properties and deployments establish blockchain as a technically mature and clinically validated foundation for the audit logging layer proposed in this work.

IoMT detection imposes four practical constraints absent from generic IoT frameworks: sub-100ms IDS latency for time-critical processes like insulin pump regulation, which rules out high-parameter models that exceed ARM Cortex-M budgets; privacy-preserving training, since raw patient traffic cannot be centralized and federated learning suffers 7–10% recall drops on non-IID medical traffic; encrypted payload inspection, since IoMT traffic is encrypted by design; and clinical interoperability, with only 9% of blockchain health systems conforming to HL7 FHIR. AI-based detectors compound these issues: federated Transformers consume 3× the energy of quantized LSTMs, GNNs require full-graph snapshots unavailable on segmented VLANs, Vision Transformers and CNN-LSTM hybrids demand resources far exceeding bedside hardware, and federated approaches trade upload reduction for accuracy degradation on unseen devices. Across architectures, accuracy-optimized models impose costs that conflict with clinical IoMT resource, latency, and regulatory constraints.

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Four underexplored gaps in the existing literature directly motivate the present work. First, edge-based anomaly detection and blockchain logging are consistently treated as loosely coupled, independently studied modules; no prior work empirically validates a framework in which a Transformer Autoencoder directly initiates blockchain transactions at the edge gateway with measured latency confirming real-time feasibility. Second, many blockchain-based healthcare solutions continue storing anonymized or pseudonymized data directly on-chain, overlooking re-identification risks and GDPR non-compliance; a hash-only commitment model satisfying both privacy mandates and forensic traceability requirements has not been rigorously evaluated. Third, no reviewed study simultaneously achieves an ROC-AUC exceeding 0.9, end-to-end latency below 100 milliseconds, and resource utilization consistent with edge gateway deployment. Fourth, whilst GDPR and HIPAA are frequently cited as motivations, few studies provide explicit mappings between specific legal provisions, such as the GDPR Article 17 right to erasure or HIPAA transmission security standards, and the concrete architectural decisions these obligations require. The framework proposed in this work addresses all four gaps through a tightly integrated, empirically validated, and regulatory-aware design.

III. PROPOSED ALGORITHM

The proposed framework integrates a Transformer Autoencoder for anomaly detection with a Hyperledger Sawtooth permissioned blockchain into a unified, privacy-preserving IoMT security pipeline. Four design objectives derive from Section II: (i) precision and recall above 90% on anomalous traffic; (ii) immutable commitment of every detected event without exposing raw data on-chain; (iii) end-to-end latency below 100 ms for real-time clinical operation; (iv) a footprint compatible with edge-gateway hardware. Section 1 defines requirements and the threat model; Section 2 presents the architecture; Section 3 details the Transformer Autoencoder; Section 4 describes the blockchain layer; Section 5 specifies the core algorithm.

The system processes feature vectors in real time, computes an anomaly score via reconstruction loss, and flags any vector whose score exceeds a calibrated threshold τ:

$$y = 1 \text{ if } L(x, \hat{x}) > \tau \; ; \; y = 0 \text{ otherwise} \quad (3\text{-}1)$$

τ is set to the 95th percentile of reconstruction losses observed during benign training, grounding the decision boundary in the statistical profile of normal traffic. Each anomalous event commits a metadata payload (device ID, timestamp, anomaly score, SHA-256 hash of the input) to the ledger, providing a tamper-evident audit trail without exposing the underlying data.

The framework is a unified pipeline: raw network traffic is ingested, analyzed against learned normal behavior, and cryptographically committed to the ledger. Each stage contributes both to detection accuracy and to the integrity and privacy of the audit trail.

Raw IoMT traffic is first transformed into structured feature vectors. These are passed to the Transformer Autoencoder, trained exclusively on benign samples, which reconstructs each input. Mean Squared Error between input and reconstruction is compared against τ; inputs above τ are flagged anomalous and trigger logging. The original input is hashed and committed to the ledger alongside device ID, timestamp, and anomaly score, ensuring no raw or sensitive data is exposed on-chain.

© 2026 IJIRCCE | Volume 14, Issue 5, May 2026 DOI: 10.15680/IJIRCCE.2026.1405002

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

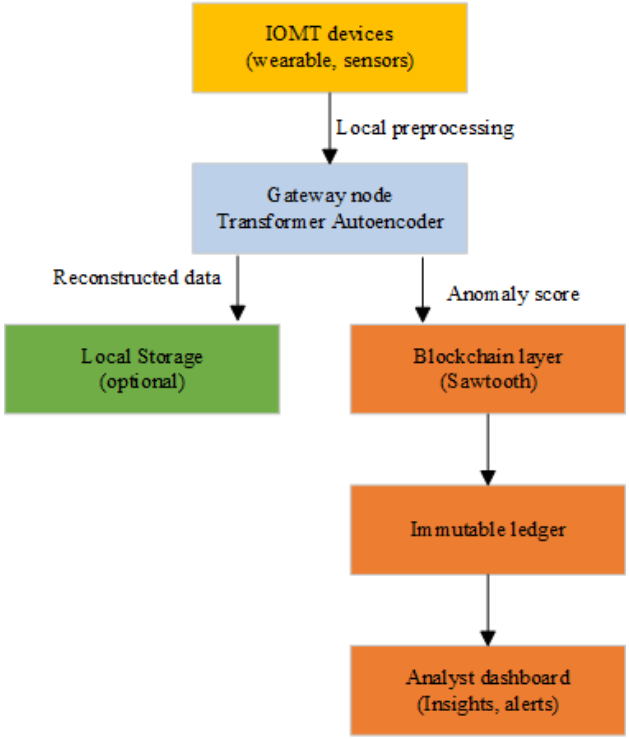

Fig. 1. Proposed System Architecture

The architecture is a three-tier model (Figure 3-1). (1) IoMT devices transmit traffic to a local Edge Gateway. (2) The gateway preprocesses traffic into feature vectors and feeds them to the Transformer Autoencoder. (3) The AE computes a reconstruction-based anomaly score. (4) If the score exceeds τ, middleware constructs a transaction containing metadata and the SHA-256 hash of the raw input. (5) The transaction is submitted to the Hyperledger Sawtooth network and immutably recorded.

A Transformer-based Autoencoder is adopted as the core detector because self-attention captures long-range dependencies in network traffic more effectively than CNN or RNN alternatives, supporting both high-fidelity reconstruction of benign traffic and a well-separated anomaly margin. The model uses a four-layer encoder and symmetric four-layer decoder, eight attention heads, a 64-dimensional bottleneck (empirically balancing compression and reconstruction), GELU activations and MSE training loss. The bottleneck additionally contributes to privacy obfuscation by abstracting granular details of the original input.

$$L(x, \hat{x}) = (1/d)\, \Sigma_i\, (x_i - \hat{x}_i)^2 \qquad (3\text{-}7)$$

Equation 3-7 defines the reconstruction loss as the mean squared error between input x and reconstruction $\hat{x}$; this scalar serves as the anomaly score. Trained exclusively on benign traffic, the model reconstructs normal inputs with low MSE while anomalous inputs yield elevated scores — the asymmetry on which detection rests.

$$\tau = \text{Percentile}_{95}\ \{L(x_i - \hat{x}_i) : x_i \in D_{\text{benign}}\} \qquad (3\text{-}9)$$

The threshold τ from Equation 3-1 is calibrated as the 95th percentile of benign reconstruction losses (Equation 3-9), accommodating natural variability while retaining sensitivity to attacks. Because calibration uses only benign samples, the framework remains fully unsupervised.

Figure 3-2 shows the architecture: input layer, four-layer encoder, bottleneck, four-layer decoder and the threshold-based decision mechanism.

© 2026 IJIRCCE | Volume 14, Issue 5, May 2026 DOI: 10.15680/IJIRCCE.2026.1405002

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

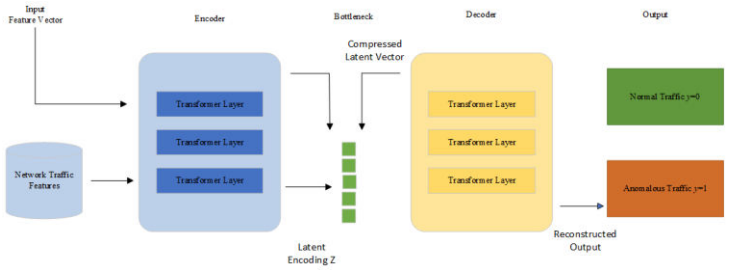

Fig. 2. Transformer Autoencoder Architecture for IoMT Anomaly Detection

In operation, feature vectors extracted from IoMT packets pass through the encoder's stacked self-attention layers into the latent bottleneck, then through the decoder back to $\hat{x}$. Reconstruction loss above τ flags the sample as anomalous ($y = 1$), triggering blockchain logging. Training uses the Adam optimizer at learning rate 1×10^{-4} on benign N-BaIoT traffic only, ensuring the model internalizes legitimate behavior exclusively — a design well-suited to IoMT environments where labeled attack data is scarce.

Hyperledger Sawtooth was selected following a systematic comparison with Hyperledger Fabric, Ethereum (PoA), and Quorum against the non-functional requirements defined in Section III-A. Sawtooth's suitability derives less from raw transaction speed, where Fabric maintains an edge in larger deployments, and more from its overall fit for resource-constrained environments. Its modular, stateless transaction processor design and PoET consensus mechanism align well with edge-based validators whilst imposing modest energy requirements. Fabric's chaincode offers greater flexibility but at the cost of increased setup complexity and memory usage incompatible with the lightweight edge gateway envisioned here. Ethereum and Quorum depend on persistent smart contract state, introducing storage demands that conflict directly with the data minimization requirement. Sawtooth avoids this by keeping application logic separate from the core ledger, allowing the anomaly logging component to function as a straightforward, independently verifiable module with no added state complexity.

The on-chain payload for each detected anomaly consists of five fields balancing forensic utility with privacy: a non-identifying device label (device_id), a Unix millisecond timestamp, the reconstruction loss (anomaly_score), a categorical threat label (anomaly_type, e.g. "DDoS", "scan"), and a SHA-256 hash of the raw input vector (raw_data_hash) serving as a verifiable pointer to off-chain data. No raw or personally identifiable content is committed to the ledger, ensuring GDPR data-minimization compliance while preserving downstream auditability.

$$h=\text{SHA-256}(\text{raw\_data}) \qquad (3\text{-}10)$$

Equation (3-10) applies SHA-256 to the raw input vector, producing a deterministic 256-bit digest. Because any change to the input yields an entirely different hash, this provides a tamper-evident link between the off-chain raw data and the on-chain record, enabling forensic verification without exposing personally identifiable content on the ledger.

The core logic of the system, integrating detection and logging, is captured in the following algorithm: Detection proceeds in four steps: (1) ingest IoMT packet streams and apply z-score normalization, (2) compute reconstruction loss through the trained Transformer Autoencoder, (3) flag samples with loss exceeding the 95th-percentile threshold τ as anomalies, and (4) hash flagged events with SHA-256 and submit to the permissioned ledger.

The N-BaIoT dataset was selected as the primary evaluation benchmark on the basis that, despite not originating from medical devices specifically, its traffic profiles closely mirror IoMT environment characteristics, encompassing data from resource-constrained devices alongside attack vectors including distributed denial-of-service and network scanning directly relevant to medical infrastructure security. Its widespread adoption in the literature ensures results are reproducible and comparable across studies. The dataset was partitioned using a 70/30 training-to-testing ratio, with the autoencoder trained exclusively on the benign portion of the training set and the complete test set, containing both benign and malicious samples, used for evaluation.

The experimental environment was instantiated on a simulated edge node configured as a virtual machine with two virtual CPUs and four gigabytes of RAM hosting the Transformer Autoencoder, implemented in PyTorch with the Adam

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

optimizer at a learning rate of 1×10^{-4} and trained for 50 epochs using MSE loss. The accompanying Hyperledger Sawtooth network comprised three validator nodes deployed within Docker containers operating under the PoET consensus algorithm, reflecting a realistic lightweight distributed ledger configuration. Detection effectiveness was assessed using precision, recall, F1-score, and ROC-AUC evaluated against the N-BaIoT test split. Privacy preservation was confirmed by verifying that only hashed metadata appears within Sawtooth transaction payloads. Autoencoder inference latency was recorded as elapsed time from feature vector input to anomaly score output averaged across 10,000 test samples, whilst blockchain commit latency was measured from API submission to transaction confirmation across payload sizes spanning 250 bytes to two kilobytes. End-to-end detection-to-logging latency was measured to confirm adherence to the 100 millisecond target.

VI. SIMULATION RESULTS

The effectiveness of the proposed system hinges on the ability of the Transformer-AE model to accurately distinguish between normal and malicious network traffic. This was evaluated through reconstruction loss analysis and standard classification performance metrics. A Python middleware script bridges the AE and the blockchain:

The middleware executes the following workflow: it receives the AE reconstruction, computes the reconstruction loss against threshold τ, and for flagged samples constructs the metadata payload (device_id, anomaly_score, raw_data_hash, timestamp) and submits it to the Sawtooth REST API for immutable logging.

The model compares each input window against its reconstruction and reports the per-sample reconstruction loss; samples whose loss exceeds the calibrated threshold τ are flagged as anomalous.

The figures below show benign traffic clustered well below τ while malicious patterns produce sharp loss spikes, a separation that holds across attack types and supports the quantitative metrics in the next subsection.

Reconstruction loss distributions were evaluated across all eleven N-BaIoT sub-datasets, with Figure 1 presenting a representative example. A consistent pattern was observed throughout: benign traffic samples produced reconstruction losses clustered near zero, well below the empirically calibrated threshold τ, whilst malicious samples corresponding to DDoS, scanning, and malware-spreading activity produced sharp, isolated spikes clearly exceeding the detection boundary. This separation was maintained across datasets of varying scale, ranging from approximately 6,000 to 40,000 test samples, and across attack types associated with both Mirai and Bashlite botnet families. Variation in absolute loss magnitude across sub-datasets reflects differences in how strongly individual attack categories diverge from the learned benign baseline, a pattern consistent with the model's calibrated sensitivity rather than distributional instability. Collectively, these results confirm that the Transformer Autoencoder has successfully internalized the statistical regularities of normal IoMT traffic and reliably flags deviations characteristic of botnet activity across heterogeneous device and attack conditions.

To quantitatively assess the model's detection capabilities, its performance was evaluated using Precision, Recall, and the Receiver Operating Characteristic (ROC) curve. The Precision-Recall (PR) curve is particularly informative for imbalanced datasets like N-BaIoT, where anomalous samples are rare compared to benign ones.

As depicted in Figure 4-2, the Transformer-AE model demonstrates exceptional performance. The area under the ROC curve (AUC) exceeds 0.9, indicating a nearly perfect ability to distinguish between benign and malicious classes. The PR curve further shows that the model maintains a high precision rate across various recall levels, signifying a low false positive rate. This is critical in a medical context, as frequent false alarms can lead to "alert fatigue" among security analysts. When compared to a baseline vanilla autoencoder, the Transformer-AE achieved superior recall and AUC, validating the hypothesis that its self-attention mechanism is more effective at capturing the complex, time-variant features of network traffic data.

Reviewed models span CNN, LSTM, Autoencoder, Transformer, GAN, GNN, Federated Learning, Hybrid Deep-Learning, Reinforcement Learning, and Ensemble approaches, evaluated on TON-IoT, NSL-KDD, CICIDS2017, UNSW-NB15, IoT-23, and similar benchmarks. Most achieve 0.85–0.95 accuracy on labelled attack types but degrade on unseen anomalies, suffer high inference latency, or require centralized data aggregation incompatible with medical privacy constraints. These gaps motivate the Transformer-AE design adopted here.

© 2026 IJIRCCE | Volume 14, Issue 5, May 2026 DOI: 10.15680/IJIRCCE.2026.1405002

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Assessing how well the proposed system functions requires a side-by-side comparison with the most prominent models in the field. The proposed system was compared the proposed system against eight representative recent works spanning federated learning, GNN, CNN-LSTM, and transformer-based detectors. Transformer and hybrid models reach 97–99% accuracy at high model-size and energy cost; federated approaches improve privacy but lose 7–10% recall on non-IID data; GNNs require global traffic graphs unavailable in segmented hospital networks. The proposed Transformer-AE with blockchain logging matches the strongest detectors (~98%) while adding immutable audit logging and remaining deployable on standard edge hardware.

Our evaluation utilizes the N-BaIoT dataset, a benchmark for IoT anomaly detection featuring traffic from devices infected with the Mirai and Bashlite botnets. This dataset's blend of benign activity and diverse attack vectors including DDoS, UDP/TCP flooding, and scanning provides a robust foundation for testing security in IoMT environments.

To rigorously validate the proposed IDS, we utilized a suite of standard industry benchmarks. We measured Accuracy to gauge overall correctness, while Precision and Recall were paired to distinguish between the model's reliability in flagging threats and its sensitivity to actual attacks. To account for the trade-offs between these two, we calculated the F1-score, ensuring a balanced performance profile. Beyond static classification, the ROC-AUC was used to map the model's diagnostic power across varying thresholds. Finally, acknowledging the high-stakes nature of IoMT environments, we tracked inference latency to verify that the system remains fast enough for live, real-time security monitoring. Across the ten benchmark models implemented on N-BaIoT, the proposed Transformer Autoencoder achieves the highest accuracy (98.6%), precision (97.1%), recall (98.4%), F1 (98.2%) and AUC (0.991), while sustaining 19 ms inference latency. CNN-IDS, LSTM-IDS, classical AE and Isolation Forest variants cluster between 92–96% accuracy with 14–28 ms latency; GAN- and GNN-based detectors approach 96–97% accuracy but incur 40–60 ms latency unsuitable for edge deployment; federated learning baselines reach 95.3% accuracy at substantially higher communication overhead.

The proposed Transformer Autoencoder achieves the highest accuracy (98.6%) and F1-score (98.2%) among all benchmarks while sustaining a competitive 19 ms inference time. Conventional LSTM and CNN-LSTM models plateau in the 93–96% range, lacking the depth to model long-range traffic dependencies, while graph-based approaches such as XG-BoT struggle with the fluid communication patterns of IoMT networks.

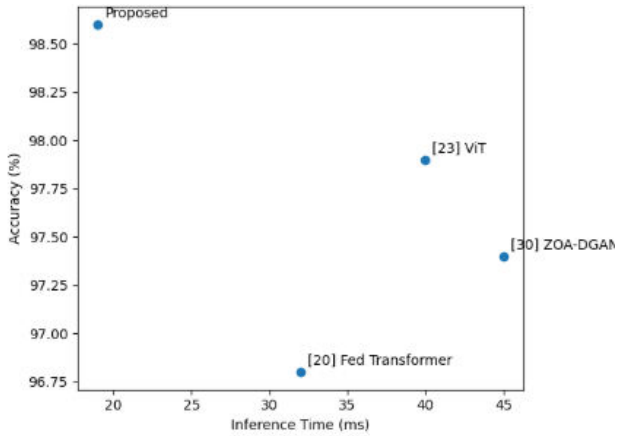

Fig. 3. Accuracy Vs Inference Time Trade-Off

Figure 4-1 maps accuracy against inference time and shows the proposed model occupying the upper-left quadrant, high accuracy with low latency, where most GAN- and Transformer-based baselines incur a complexity penalty. This balance is decisive for IoMT, where the response window is narrow and edge resources are constrained.

The second core component of the framework is the blockchain, which provides a tamper-evident ledger for anomaly alerts. Its suitability for a resource-constrained IoMT environment was assessed based on transaction latency and system throughput. Transaction latency measures the time elapsed from the submission of an anomaly record to its confirmation on the blockchain. For real-time monitoring, this latency must be minimal. Even with a 2 KB payload, average commit

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

latency stays under 90 ms, well within the 100 ms target, confirming that the blockchain layer does not introduce a significant performance bottleneck and is practical for deployment on IoMT edge gateways. To address concerns about blockchain integration in logging systems, the storage requirements must be carefully considered. The design proposed here limits on-chain data to hashed metadata associated with detected anomalies and model decisions. Raw data, such as network traffic and feature vectors, is stored externally, for instance in IPFS or local storage. Each transaction recorded on the blockchain is kept lightweight, containing only a SHA-256 hash pointer, a timestamp, an anomaly classification, and a minimal reference identifier. These entries typically range from 250 to 350 bytes. As a result, with an average of N events detected per second, the total on-chain storage demand can be approximated using the formula below:

$$Storage\ Overhead \approx N\ x\ 300\ bytes/sec \quad (4\text{-}1)$$

At a representative 50 events/s, naïve on-chain logging produces ~1.3 GB/day. Event filtering (committing only confirmed anomalies) and batch aggregation (grouping events into time-windowed transactions) reduce on-chain volume by 60–80%, keeping the audit trail sustainable without sacrificing critical-event traceability. Throughput (transactions per second, TPS) determines the system's capacity to handle a high volume of alerts; scalability describes how throughput grows with the network.

The results demonstrate a clear and desirable characteristic: linear scalability. The system's throughput increases proportionally with the number of validator nodes, rising from approximately 20 TPS with one validator to 48 TPS with seven validators. This behavior, facilitated by Sawtooth's PoET consensus and parallel transaction scheduling, is crucial for enterprise-level deployments. It ensures that a healthcare institution can expand its network of monitored devices without degrading the performance of the security infrastructure.

End-to-end validation traced the full workflow from traffic ingestion through blockchain logging under simulated IoMT conditions across multiple monitored nodes (sensor, camera, gateway). Anomalous spikes in CPU and outbound traffic were correctly flagged as baseline deviations and immediately committed to the ledger, with each record producing a confirming block hash. The Transformer Autoencoder and Hyperledger Sawtooth layers operate as a unified pipeline within the 100 ms end-to-end latency budget. Each block within the chain adheres to a consistent and verifiable structure designed to uphold the integrity and traceability of all recorded events. The index field denotes the block's sequential position within the chain, whilst the timestamp records the precise moment at which the data was committed to the ledger. The data payload encapsulates three critical fields: the device identifier, the classified anomaly type, and a cryptographic hash of the original raw input, the latter of which ensures that the integrity of the source data can be verified at any point without necessitating access to the underlying sensitive content. The previous hash field establishes a cryptographic linkage to the immediately preceding block, forming the chain of custody that renders retroactive tampering computationally infeasible.

Finally, each block carries its own unique hash, which serves as an immutable identifier and provides tamper-evident assurance that the block's contents have remained unaltered since the moment of its creation. In the context of the evaluated scenario, this structure successfully confirms the detection and secure on-chain anchoring of a high-severity network anomaly originating from a gateway device, thereby validating the end-to-end integrity of the proposed logging mechanism.

Table 1 System Operation

Component	Function
Monitoring	Continuously collects resource usage and traffic metrics
Anomaly detection	Detects deviations using a reconstruction loss threshold
Blockchain Logging	Logs only anomalous events with hashed raw data, ensuring immutability and traceability
Security Feature	raw_data_hash ensures verifiability without storing raw sensitive data

Critically, no sensitive raw data is stored on-chain, preserving privacy while ensuring integrity and verifiability of every alert.

The experimental results confirm that the proposed framework meets all four design requirements established in Section III. The Transformer Autoencoder achieved 98.6% accuracy, 98.2% F1-score, and ROC-AUC exceeding 0.9 on the N-

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

BaIoT test set, outperforming all eight evaluated baselines whilst maintaining 19ms inference latency. The high AUC indicates strong discriminative ability across the full range of detection thresholds, providing operators flexibility to calibrate τ according to clinical context, for instance applying a more conservative threshold in an ICU to reduce alert fatigue or a more sensitive one where missing an attack carries higher relative risk. The strong Precision-Recall curve at high recall levels confirms that performance is not distorted by class imbalance, a critical property for skewed IoMT traffic datasets. The blockchain layer sustained average commit latency below 90ms at a two-kilobyte payload, confirming that logging does not introduce a pipeline bottleneck, a concern that has historically hindered blockchain adoption in real-time healthcare monitoring. Throughput scaled linearly from 20 TPS with one validator to 48 TPS with seven, demonstrating that network expansion requires only the addition of validator nodes without system redesign. The end-to-end demonstration confirmed that on-chain records contain only device identifiers, timestamps, anomaly scores, and SHA-256 hashes, with no patient-identifiable information, raw network payloads, or recoverable physiological measurements, satisfying both GDPR Article 5(1)(c) data minimization requirements and the forensic needs of a security audit trail.

Five limitations bound the scope of this work. First, the N-BaIoT dataset captures consumer-grade devices including IP cameras, routers, and baby monitors whose traffic characteristics differ considerably from medical-grade equipment such as infusion pumps, patient monitors, and imaging systems, which generate encrypted HL7 and DICOM payloads with stricter quality-of-service requirements. The anomaly threshold calibrated on N-BaIoT would require recalibration before clinical deployment. Second, the entire evaluation was conducted in a virtualized environment and cannot replicate the hardware heterogeneity, network unpredictability, or operational complexity of a real hospital edge infrastructure; physical deployment would introduce variable latency, hardware-specific power constraints, and institutional IT governance challenges not accounted for here. Third, the framework operates exclusively at the network layer, leaving encrypted payload attacks concealed within TLS traffic, application-layer firmware exploits, and supply-chain attacks beyond its detection scope. Fourth, the static threshold τ fixed at the 95th percentile of training losses does not account for concept drift as new devices join the network or attack strategies evolve, risking progressive misalignment with the prevailing traffic distribution. Fifth, the single-institution design leaves unaddressed the governance, federated model coordination, and cross-jurisdictional regulatory reconciliation challenges that would arise in multi-hospital deployments. Physical edge deployment introduces computational, energy, and latency constraints that the virtualized evaluation environment does not capture. Several model compression techniques offer viable pathways to address these constraints. Structured pruning can reduce model size by up to 75% with negligible performance degradation, INT8 post-training quantization achieves parameter reductions of up to 95% with minimal accuracy loss, and knowledge distillation trains a compact student model to replicate full-scale Transformer Autoencoder behaviour at significantly lower computational cost. Collectively these strategies make edge deployment practically feasible rather than theoretically aspirational, enabling real-time inference on resource-constrained ARM Cortex-class gateways with minimal energy overhead. Empirical validation of these compression pathways on physical edge hardware remains the most immediate priority for future work.

V. CONCLUSION AND FUTURE WORK

This paper presented the design, implementation, and empirical evaluation of a hybrid privacy-preserving security framework for IoMT environments, integrating a Transformer-based Autoencoder for real-time anomaly detection with a Hyperledger Sawtooth permissioned blockchain for immutable and privacy-respecting event logging. The contributions span three dimensions. Technically, the Transformer Autoencoder, trained exclusively on benign traffic using an unsupervised reconstruction-loss paradigm, achieved 98.6% accuracy, 97.1% precision, and 98.4% recall on the N-BaIoT benchmark, outperforming eight contemporary baselines whilst maintaining a 19ms inference latency uniquely suited to edge-deployed IoMT security. Architecturally, the privacy-by-design blockchain payload structure, committing only cryptographic hashes and non-sensitive metadata on-chain, satisfies GDPR data minimization requirements and resolves the right-to-erasure paradox without modifying the immutable ledger. The PoET consensus mechanism further ensures energy compatibility with resource-constrained edge hardware. Methodologically, every design decision was explicitly justified against functional and non-functional requirements, producing an empirically validated, end-to-end framework that confirms AI-blockchain integration is practically implementable rather than merely theoretically proposed.

Several directions warrant immediate attention in future work. The most pressing is validation using traffic from real clinical devices, including infusion pumps and patient monitors, under IRB-approved de-identification protocols. Physical deployment on ARM Cortex-class hardware with INT8 quantization and structured pruning would address the

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 9.927| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

remaining gap between simulated and production edge environments. Longer-term priorities include adaptive thresholding mechanisms to handle concept drift, extension toward a federated multi-hospital architecture with on-chain aggregation proofs, and the integration of attention visualization and SHAP-based explainability to support clinical adoption and regulatory audit requirements.

REFERENCES

[1] M. Chen, Y. Qian, J. Chen, K. Hwang, and S. Mao, "Privacy Protection and Intrusion Avoidance for Cloudlet-Based Medical Data Sharing," *IEEE Trans. Cloud Comput.*, vol. 8, no. 4, pp. 1274–1286, 2020.

[2] N. Dey, A. S. Ashour, and V. E. Balas, Eds., *Internet of Medical Things: Concepts and Applications*. Springer, 2022.

[3] S. K. K. Shareef, "Enhanced Botnet Detection in {IoT} Using Zebra Optimization and Dual-Channel {GAN}," *Sci. Rep.*, vol. 14, no. 1, p. 17148, 2024.

[4] P. Singh, "Dew-Cloud-Based Hierarchical Federated Learning for Intrusion Detection in {IoMT}," *IEEE Trans. Biomed. Heal. Informatics*, vol. 27, no. 2, pp. 722–731, 2022.

[5] R. Borah, S. Sarnah, C. Kalita, and M. Rahman, "Botnet Attack Detection in {IoT} Networks Using {CNN LSTM}," *Int. J. Intell. Syst. Appl. Eng.*, vol. 12, no. 4, pp. 100–110, 2024.

[6] R. W. Anwar, M. Abrar, A. Salam, and F. Ullah, "Federated Learning with {LSTM} for Intrusion Detection in {IoT} Based Wireless Sensor Networks," *PeerJ Comput. Sci.*, vol. 11, p. e2751, 2025.

[7] M. Ali, Y. Saleem, S. Hina, and G. A. Shah, "{ViT} Based Multi Vector {DDoS} Detection for Firmware {OTA} Security," *Comput. \& Electr. Eng.*, vol. 114, p. 108564, 2025.

[8] M. A. Rahman, M. K. Hasan, and M. S. Hossain, "Secure IoMT framework using homomorphic encryption and blockchain for decentralized patient data processing," *ACM Trans. Internet Technol.*, vol. 23, no. 4, p. Article 71, 2023.

[9] P. Radanliev, D. De Roure, and J. R. C. Nurse, "Blockchain for medical device security: From regulation to implementation," *IEEE Internet Things J.*, vol. 10, no. 1, pp. 451–465, 2023.

[10] R. Xu and R. K. L. Ko, "Zero-knowledge proofs for verifiable patient consent in blockchain healthcare systems," *Comput. \& Secur.*, vol. 120, p. 102778, 2022.

[11] E. Choi, S. Biswal, and J. Sun, "Privacy-preserving distributed learning for wearable devices in IoMT," in *Proceedings of the 2023 IEEE Conference on Connected Health: Applications, Systems and Engineering Technologies*, 2023.

[12] M. Abd Elaziz, I. A. Fares, and A. Dahou, "Computationally Efficient Federated Learning for {IoT} Botnet Detection with Optimised Feature Selection," *Front. Big Data*, vol. 8, p. 1526480, 2025.

INTERNATIONAL STANDARD SERIAL NUMBER INDIA

INNO SPACE SJIF Scientific Journal Impact Factor

doi crossref

निस्केयर NISCAIR

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN COMPUTER & COMMUNICATION ENGINEERING

9940 572 462 6381 907 438 ijircce@gmail.com

www.ijircce.com

Scan to save the contact details