

A Blockchain-Enabled Severity-Aware Framework for Ransomware Resilience in Healthcare IoT

Hani Al-Balasmeh *

Dept. of Informatics Engineering, College of Engineering University of Technology Bahrain

*Corresponding author E-mail: h.albalasmeh@utb.edu.bh

Received: December 24, 2025, Accepted: January 19, 2026, Published: January 24, 2026

Abstract

The growing adoption of Healthcare Internet of Things (HIoT) systems has improved patient monitoring and clinical efficiency. Still, it has also exposed hospitals to ransomware attacks that can disrupt life-critical operations. Conventional intrusion detection systems (IDS) such as Snort, Suricata, and Kitsune struggle to provide tamper-proof evidence or prioritize device-specific risks, limiting their effectiveness in clinical environments. To address this gap, we propose a blockchain-assisted framework that integrates hybrid anomaly detection, Byzantine fault-tolerant forensic logging, and severity-aware mitigation. Suspicious traffic is first evaluated using anomaly scoring with EVT-based thresholding, then immutably recorded on a blockchain ledger via PBFT consensus, ensuring tamper-resistance. Smart contracts enforce mitigation decisions based on severity scores that account for anomaly magnitude, device criticality, and network exposure, thereby guaranteeing the rapid protection of life-support devices while minimizing unnecessary disruption to low-risk equipment. An experimental evaluation of the N-BaIoT, ToN_IoT, and CIC-IDS2017 datasets shows that the framework achieves a detection accuracy of up to 98%, a false-positive rate of 1.1%, and an average latency of 6.8 ms, outperforming baseline IDS solutions. Security analyses confirm resilience against log tampering, obfuscation-based evasion, and denial-of-service flooding, while throughput scalability exceeded 2200 TPS across hospital nodes. By combining blockchain accountability with clinically aware mitigation, this framework provides a robust, real-time defense against ransomware in HIoT environments, advancing the state of cybersecurity for patient-centered healthcare systems.

Keywords: Healthcare Internet of Things (HIoT); Ransomware Resilience; Blockchain Security; Severity-Aware Mitigation; Forensic Accountability.

1. Introduction

HIoT systems are uniquely vulnerable due to their resource-constrained nature, reliance on legacy software, and the difficulty of applying timely security patches without disrupting patient care [6], [7]. These limitations make HIoT environments attractive targets for ransomware operators. Historical cases illustrate the magnitude of the risk: the WannaCry attack in 2017 severely disrupted the United Kingdom's National Health Service (NHS), halting surgeries and diverting patients [8]. More advanced ransomware families, such as Maze and Ryuk, demonstrate sophisticated attack vectors, including lateral movement, data exfiltration, and double extortion, amplifying both the technical and economic impacts on healthcare delivery [9].

Despite considerable research progress, existing intrusion detection systems (IDSs) such as Snort, Suricata, and Kitsune remain insufficient for ransomware resilience. Signature-based IDSs struggle against novel or polymorphic ransomware, while anomaly-based methods often suffer from excessive false positives, overwhelming healthcare staff, and leading to alert fatigue [10]. Moreover, current mitigation strategies, such as backups or manual isolation, do not guarantee tamper-proof forensic evidence or prioritize device-specific risks. In healthcare, this lack of forensic accountability and clinical prioritization represents a critical gap, as patient safety depends on both accurate detection and risk-aware response [11].

Blockchain has recently emerged as a promising enabler of cybersecurity resilience in distributed environments. Its decentralized architecture and immutability provide tamper-proof logging, while smart contracts enable automated, verifiable execution of policies [12]. By combining blockchain with anomaly detection, HIoT systems can achieve not only higher detection accuracy but also forensic accountability and trustless cooperation across hospital networks. However, existing blockchain-based defenses rarely address clinical safety requirements, particularly the need to prioritize life-critical devices when responding to ransomware threats.

This paper introduces a blockchain-enabled, severity-aware framework for ransomware resilience in Healthcare IoT that integrates hybrid anomaly detection, Byzantine fault-tolerant forensic logging, and risk-prioritized mitigation policies. The system employs extreme value theory (EVT) for anomaly thresholding to ensure low false alarms and enforces severity-aware responses through smart contracts that consider anomaly magnitude, device criticality, and network exposure. This dual focus on cyber defense and patient safety distinguishes the framework from existing IDS and blockchain-based approaches.

The remainder of this paper is organized as follows. Section II reviews related work on ransomware in healthcare, IoT security challenges, and blockchain-driven defenses. Section III presents the architecture of the proposed severity-aware framework. Section IV details the methodology, including mathematical models and algorithms. Section V provides experimental results and comparative analysis with Snort, Suricata, Kitsune, and related approaches. Section VI discusses the broader implications and future research directions. Section VII concludes the paper.

2. Literature Review

Before presenting the proposed framework, it is essential to review existing research that shapes the context of ransomware resilience in Healthcare IoT. The related work spans four critical dimensions: (i) ransomware attacks in healthcare and their operational consequences, (ii) vulnerabilities specific to IoT-based medical devices, (iii) blockchain's role in cybersecurity and forensic accountability, and (iv) benchmark datasets commonly employed for evaluating IoT security solutions. Together, these areas provide the foundation for identifying the gaps that motivate this study.

2.1. Ransomware in healthcare

Ransomware has emerged as the most disruptive cybersecurity threat to healthcare organizations, with consequences that extend beyond financial loss to directly endangering patient safety. Unlike traditional IT environments, where downtime may result in productivity loss, ransomware in healthcare can delay surgeries, disrupt intensive care units, or paralyze entire hospital networks. Haq and Raza [13] reported that ransomware has evolved from opportunistic "spray-and-pray" attacks to highly targeted, multi-stage campaigns, often exploiting Healthcare IoT (HIoT) devices as vulnerable entry points. Agrawal et al. [14] highlighted that although hospitals increasingly adopt backup and disaster recovery strategies, adversaries counteract these measures with double extortion, exfiltrating patient records before encryption to increase ransom leverage. More recent analyses underscore that attackers deliberately exploit the operational urgency of healthcare, leveraging life-critical contexts to maximize ransom payments [15]. Table 1 compares major ransomware studies in healthcare. As shown, prior work has provided valuable insights into evolving attack techniques and their consequences, but generally lacks technical solutions that integrate blockchain or forensic accountability, leaving a gap for more resilient mitigation approaches.

Table 1: Comparative Overview of Studies on Healthcare Ransomware

REF.	DOMAIN	ATTACK TYPE	EVALUATION METHOD	KEY FINDINGS	LIMITATIONS
[13]	Healthcare IoT	Multi-stage ransomware	Survey of incidents	Identified ransomware trends in IoT-driven medical networks	NO BLOCKCHAIN OR FORENSIC ANALYSIS
[14]	General Healthcare	Ransomware	Systematic literature review	Classified mitigation into detection, recovery, and awareness	NO IOT-SPECIFIC CONTEXT
[15]	HEALTHCARE	TARGETED RANSOMWARE	CASE STUDY ANALYSIS	HIGHLIGHTED DOUBLE EXTORTION AND PATIENT SAFETY RISKS	LACKS A TECHNICAL EVALUATION OF COUNTERMEASURES

2.2. Healthcare IoT vulnerabilities

The integration of IoT devices into healthcare networks has expanded the attack surface with device-specific weaknesses. Many HIoT devices run outdated operating systems, use hardcoded credentials, and have limited computational resources, leaving them ill-suited to conventional security measures [16]. Tuli et al. [17] showed that fog- and edge-enabled IoT frameworks can improve latency but inherit vulnerabilities from unpatched endpoints. Meidan et al. [18] demonstrated that once a single medical device is compromised, lateral movement within hospital networks becomes a powerful vector for ransomware propagation. These weaknesses provide ransomware operators with reliable footholds for infiltration and persistence. Table 2 provides an overview of these vulnerabilities and their direct relevance to ransomware staging. It highlights that legacy systems and weak authentication remain persistent problems, while remote access flaws facilitate propagation, underscoring the urgency of designing frameworks that can account for these unique risks.

Table 2: Key Vulnerabilities in Healthcare IoT Devices and Their Relevance to Ransomware

REF.	VULNERABILITY	ATTACK VECTOR	RELEVANCE TO RANSOMWARE
[16]	Legacy systems	Exploits targeting outdated OS	ENABLES PERSISTENCE DUE TO DELAYED PATCHING
[17]	Weak authentication	Hardcoded/default credentials	FACILITATES UNAUTHORIZED INITIAL ACCESS
[18]	REMOTE ACCESS FLAWS	IOT BOTNET PROPAGATION	PROVIDES A FOOTHOLD FOR RANSOMWARE STAGING/SPREAD

2.3. Blockchain for cybersecurity and ransomware mitigation

Blockchain has gained traction as a foundation for resilient cybersecurity architectures due to its immutability, decentralization, and support for smart contracts. Lo et al. [19] highlighted the potential of blockchain for intrusion detection and decentralized trust management in IoT, but noted a lack of empirical validation. Raj et al. [20] developed a blockchain-assisted forensic framework for ransomware that enables tamper-proof logging and automated responses via smart contracts. Qi et al. [21] applied blockchain to Healthcare IoT to ensure privacy-preserving integrity of patient data, demonstrating resilience against tampering.

More recently, [25] proposed blockchain-enabled cybersecurity and data privacy solutions for smart cities, demonstrating how blockchain can achieve secure, tamper-resistant infrastructures at an urban scale. While the domain differs from healthcare, the study reinforces the importance of blockchain scalability, interoperability, and privacy guarantees — challenges that are equally critical for HIoT ransomware defense.

Despite these advances, limitations remain. Blockchain-based approaches often struggle with scalability and resource constraints, particularly when deployed on IoT devices with limited power and storage. Furthermore, most works focus either on data integrity or forensic logging, without addressing risk-prioritized mitigation strategies that are critical in clinical environments. However, as shown in Table 3, most blockchain-based security solutions remain limited by scalability challenges and rarely integrate severity-aware mitigation with forensic accountability—the dual capabilities required in healthcare.

Table 3: Blockchain-Based Security Solutions Relevant to IoT and Ransomware

Ref.	Domain	Technique	Evaluation	Strengths	Limitations
[19]	IoT Security	Blockchain for intrusion detection	Conceptual survey	Highlights immutability for IDS	No experimental validation
[20]	Cybersecurity	Blockchain forensic framework	Simulated ransomware	Provides tamper-proof logging, automated response	Scalability concerns
[21]	Healthcare IoT	Blockchain for privacy & integrity	Prototype evaluation	Preserves confidentiality & compliance	Focuses on data, not ransomware mitigation
[25]	Smart Cities	Blockchain cybersecurity/privacy	Prototype + case study	Demonstrates scalability & privacy	No healthcare-specific evaluation

2.4. Datasets for IoT security evaluation

Evaluating IoT security frameworks requires robust benchmark datasets. While ransomware-specific datasets in healthcare are scarce, several general-purpose datasets provide suitable proxies for anomaly detection research. The N-BaIoT dataset [22] includes traffic from IoT devices infected with the Mirai and Bashlite botnets, which helps detect ransomware-like anomalies at the device level. The ToN_IoT dataset [23] provides heterogeneous telemetry from IoT services, operating systems, and networks, capturing behaviors relevant to lateral movement. The CIC-IDS2017 dataset [24], though enterprise-focused, remains widely used for benchmarking IDS models due to its diversity of attack types. Table 4 summarizes these datasets, highlighting their unique features and applicability to ransomware resilience studies. Collectively, they demonstrate that while no dataset perfectly reflects healthcare ransomware, combining them allows a more comprehensive evaluation of proposed frameworks.

Table 4: Benchmark Datasets Used in IoT Security Research

Ref.	Dataset	Key Features	Application to Ransomware Studies
[22]	N-BaIoT	IoT traffic with botnet infections	Provides anomalous patterns similar to ransomware staging
[23]	ToN_IoT	Multi-source telemetry (network, host, IoT)	Captures lateral movement behaviors
[24]	CIC-IDS2017	Enterprise intrusion scenarios	Useful as a baseline for IDS benchmarking

2.5. Discussion

From the reviewed literature, three gaps become evident:

- 1) Ransomware in healthcare is uniquely destructive due to the urgency of clinical workflows and the expanded attack surface introduced by IoT integration.
- 2) Healthcare IoT vulnerabilities remain inadequately addressed, as many devices cannot support traditional patching or IDS solutions, leaving them open to exploitation.
- 3) Blockchain-based defenses show promise in ensuring tamper-proof evidence and decentralized trust, but existing works lack integration with severity-aware mitigation — a key requirement for patient safety.

These gaps underscore the need for a comprehensive framework that not only detects ransomware attacks with high accuracy but also ensures forensic accountability and severity-aware, risk-prioritized mitigation. Addressing this intersection is the central objective of the present study.

3. Proposed Framework and Architecture

The proposed framework is designed to deliver early detection, tamper-proof accountability, and risk-aware adaptive mitigation in safety-critical environments. Unlike conventional defenses, which focus solely on anomaly detection or reactive backup recovery, this framework integrates machine learning-based anomaly detection, blockchain-assisted forensic logging, and severity-aware response policies into a cohesive system. Together, these components enable hospitals not only to detect ransomware but also to contain and recover in a manner that prioritizes patient safety.

The high-level architecture is shown in Figure 1, where layered components illustrate how data flows from HIoT endpoints through anomaly detection, blockchain forensic recording, and ultimately to mitigation and collaborative response.

3.1. Design objectives

Three core objectives guide the framework:

- 1) Early Detection – proactively identify ransomware behaviors in IoT devices before large-scale encryption or propagation occurs.
- 2) Forensic Accountability – maintain immutable, auditable, and regulation-compliant records of anomalies and responses via blockchain.
- 3) Severity-Aware Mitigation – prioritize defensive actions based on device criticality and clinical risk, minimizing disruption while safeguarding life-critical operations.

These objectives directly address the gaps identified in Section 2, where existing IDS approaches lack either tamper-proof evidence or patient-aware risk prioritization.



Fig. 1: Illustrates How These Objectives Are Achieved Across the System Layers.

3.2. Data collection layer

The data collection layer provides continuous visibility into heterogeneous sources, including ventilators, infusion pumps, wearable sensors, gateways, and hospital servers. Data types range from network traffic flows and device telemetry to system-level logs. Lightweight edge agents perform preprocessing tasks such as data normalization and deduplication, reducing bandwidth overhead and ensuring scalability for resource-constrained devices [25].

By deploying preprocessing at the edge, the system prevents hospital IoT devices from being overwhelmed while ensuring that abnormal patterns are forwarded to the detection layer in near real-time. This distributed approach enables scalable monitoring across hospital consortia.

3.3. Detection layer

The detection layer employs a hybrid ensemble for anomaly detection to capture both temporal and statistical characteristics of ransomware. It integrates:

- Autoencoder reconstruction error analysis to detect deviations in traffic patterns.
- Long Short-Term Memory (LSTM) predictive modeling for sequence-aware anomaly detection.
- Extreme Value Theory (EVT)-based thresholding to statistically filter out benign fluctuations and minimize false positives [26].

This ensemble significantly improves detection accuracy, outperforming single-model systems by capturing a broader spectrum of ransomware behaviors, including zero-day variants and polymorphic strains.

To contextualize the detection role within the ransomware lifecycle, Figure 2 maps the classical ransomware kill chain—reconnaissance, infection, lateral movement, encryption, and extortion—against the defense mechanisms provided by the framework. The detection layer primarily disrupts ransomware at the infection and lateral movement stages, containing attacks before the critical encryption and extortion phases.

3.4. Blockchain forensic layer

Upon detecting an anomaly, the event is recorded in a permissioned blockchain ledger shared among participating hospitals. To balance scalability and performance, the system adopts a Practical Byzantine Fault Tolerance (PBFT) consensus mechanism rather than Proof-of-Work, ensuring low-latency, energy-efficient consensus appropriate for healthcare networks [27].

Smart contracts govern key policies, including:

- Mandatory anomaly logging,
- Automated administrator alerts, and
- Device isolation or access restrictions.

This blockchain-backed logging guarantees immutability and forensic accountability, aligning with regulations such as HIPAA and GDPR [28]. It also ensures that no single institution can alter records unilaterally, fostering decentralized trust in consortium deployments.

The workflow of this layer is shown in Figure 3, which illustrates how anomalies are recorded immutably via smart contracts and how evidence is propagated to forensic dashboards for regulatory compliance and real-time monitoring.

3.5. Mitigation and response layer

Mitigation decisions are enforced using a severity-aware risk scoring model defined as:

$$R(x,di)=\alpha\cdot S(x)+\beta\cdot C(di)+\gamma\cdot E(di)$$

Where:

- $S(x)$ represents the anomaly score assigned to event xxx ,
- $C(di)$ denotes the clinical criticality of device di (e.g., ventilators receive higher weight), and
- $E(di)$ quantifies the network exposure of the device.

High-risk scores trigger immediate device isolation; medium-risk scores activate rate limiting and detailed forensic logging; and low-risk scores generate alerts for monitoring [29].

This ensures that life-critical equipment (e.g., ventilators, infusion pumps) receives the strongest protection, while less critical devices (e.g., wearables) remain available, preserving clinical workflows. This patient-safety-first philosophy is what distinguishes the framework from conventional IDS.

3.6. Scalability and interoperability

The architecture is designed to scale across multiple hospitals, with each institution operating as a blockchain node in a federated consortium. This structure eliminates reliance on a single authority while enabling collaborative defense sharing.

To facilitate adoption, the framework integrates with healthcare interoperability standards such as HL7 and FHIR, ensuring seamless exchange with Electronic Health Record (EHR) systems [30]. To further reduce latency, blockchain nodes can be deployed at local hospital gateways, thereby minimizing reliance on cloud intermediaries and improving response times.

4. Methodology

The methodology of this study is designed to provide a blockchain-assisted framework for mitigating ransomware in Healthcare IoT (HIoT) networks, integrating anomaly detection, forensic blockchain logging, and severity-aware mitigation into a single operational workflow. The approach is grounded in formal modeling, algorithmic processes, and risk-aware decision-making, ensuring both technical rigor and practical applicability in clinical environments.

At its core, the system treats the HIoT ecosystem as a set of interconnected devices $D = \{d_1, d_2, \dots, d_n\}$, each producing telemetry and network traces represented as multidimensional feature vectors $x_{ij} \in \mathbb{R}^d$. An anomaly detection function F maps input data to an anomaly score $S(x)$, and events are classified as suspicious if $S(x) > \tau$, where τ is an Extreme Value Theory (EVT)-based statistical threshold [33]. Unlike conventional intrusion detection, the novelty of this framework lies not in the detection engine itself, but in how anomalies are immutably logged and acted upon in a clinically meaningful way.

Once an anomaly is identified, it is encapsulated in a blockchain transaction:

$$T = (ID_{d_i}, S(x), R(x, d_i), t, H(L))$$

Where ID_{d_i} is the device identifier, $R(x, d_i)$ is a severity score defined below, t is the timestamp, and $H(L)$ is a cryptographic hash of forensic logs. To ensure trust and accountability, these transactions are validated using Practical Byzantine Fault Tolerance (PBFT), which requires at least $2f+1$ honest nodes out of $N \geq 3f+1$ participants to reach consensus [34]. Algorithm 1 formalizes this process. In practice, the PBFT model incurs quadratic message complexity ($O(N^2)$), but this remains tractable in consortium hospital networks of modest size while ensuring immutability and protection against insider manipulation.

Mitigation decisions are governed by a severity-aware risk model that integrates anomaly magnitude with contextual device information:

$$R(x, d_i) = \alpha \cdot S(x) + \beta \cdot C(d_i) + \gamma \cdot E(d_i)$$

Where $C(d_i)$ encodes the criticality of device d_i (e.g., life-support ventilators weighted more heavily than patient wearables) and $E(d_i)$ represents the device's network exposure or centrality.

The device criticality parameter $C(d_i)$ is not assigned arbitrarily, but is grounded in established clinical risk classification principles used in healthcare technology management. International regulatory and safety frameworks, such as the U.S. Food and Drug Administration (FDA) [35] medical device classification system and the IEC 80001 standards [36] for risk management of IT networks incorporating medical devices, categorize devices according to their potential impact on patient safety in the event of malfunction or disruption. Under these frameworks, life-support and life-sustaining devices (e.g., ventilators and infusion pumps) are considered high-risk due to the immediate threat posed to patient outcomes. In contrast, monitoring or wearable devices are typically assigned lower criticality.

In the proposed model, this clinical risk hierarchy is directly reflected in the weighting of $C(d_i)$, ensuring that cybersecurity mitigation decisions align with patient safety priorities rather than purely network-centric severity. Consequently, even moderate anomalies affecting high-criticality devices may trigger rapid isolation, while equivalent anomalies on low-criticality devices result in proportionate, non-disruptive responses. This clinically informed weighting enhances the framework's applicability in real healthcare environments.

This ensures that even modest anomalies on life-critical devices may trigger rapid isolation, while less critical devices receive proportionate responses. The decision rule is defined as:

$$a(R) = \text{Isolate}(d_i) \text{ if } R(x, d_i) \geq \theta_H; \text{RateLimit}(d_i) \text{ if } \theta_M \leq R(x, d_i) < \theta_H; \text{Alert}(d_i) \text{ if } R(x, d_i) < \theta_M$$

This mapping is enforced through smart contracts, making mitigation auditable and automatic. Algorithm 2 formalizes this step, describing how blockchain-recorded severity scores translate into clinical actions. Together, Algorithms 1 and 2 define the operational core of the methodology. Algorithm 1 ensures trust in the forensic record, and Algorithm 2 ensures that mitigation respects both security and patient safety priorities.

Algorithm 1: Blockchain Transaction Validation (PBFT)

Input: Transaction T

Output: Immutable blockchain record

- 1: Primary node proposes block B containing transaction T
- 2: Validator nodes verify signatures, timestamps, and hashes
- 3: Nodes broadcast agreement if checks are valid
- 4: if the number of agreements $\geq 2f+1$ then
- 5: Commit B to the blockchain ledger
- 6: else
- 7: Reject B and trigger view change

Algorithm 2: Severity-Aware Mitigation SelectionInput: Severity score $R(x, d_i)$ Output: Mitigation action $a(R)$

```

1: if  $R(x, d_i) \geq \theta_H$  then
2: Isolate device  $d_i$  from hospital network
3: else if  $\theta_M \leq R(x, d_i) < \theta_H$  then
4: Apply traffic rate limiting on the  $d_i$ 
5: else
6: Generate an alert for the security administrator

```

These algorithms work in tandem to ensure both tamper-proof evidence and automated clinical responses. In practical deployments, the smart contract layer automatically executes the equivalent of Algorithm 2 once Algorithm 1 has validated the forensic record, removing the possibility of human delay or tampering. This dual structure provides accountability (through blockchain) and adaptability (through severity-aware response), which are particularly vital in the healthcare domain.

To validate the methodology, experiments rely on benchmark datasets such as N-BaIoT [37], which captures IoT device traffic infected by Mirai and Bashlite; ToN_IoT [38], which provides heterogeneous telemetry across networks and services; and CIC-IDS2017 [38], which models enterprise-scale intrusions. Although these datasets do not capture live ransomware in hospitals, they approximate the network irregularities (e.g., abnormal flows, lateral movement) generated by ransomware campaigns. Performance evaluation considers detection accuracy, F1-score, false-positive rate, blockchain logging latency, and scalability, measured as transactions per second. These metrics directly reflect the framework's ability to provide reliable detection, trustworthy forensic records, and rapid, risk-adjusted responses suitable for real-time healthcare operations.

5. Results and Analysis

To improve readability and avoid repetition, quantitative results on accuracy, reliability, latency, and scalability are presented in a consolidated manner, with figures and tables highlighting comparative performance trends. The proposed blockchain-assisted framework was implemented and evaluated against benchmark datasets, including N-BaIoT [36], ToN_IoT [37], and CIC-IDS2017 [39]. The evaluation aimed to validate three primary objectives: (i) the accuracy and reliability of anomaly detection integrated with blockchain logging, (ii) the efficiency of severity-aware mitigation in prioritizing clinical safety, and (iii) the comparative performance of the proposed system against baseline intrusion detection and mitigation solutions such as Snort, Suricata, and the lightweight IoT anomaly detection framework Kitsune. Results are reported across accuracy, precision, recall, F1-score, false positive rate (FPR), latency, and scalability. Figures and tables are used extensively to provide a clear view of the comparative performance.

The first evaluation focused on detection accuracy across the three datasets. Figure 2 illustrates the accuracy levels obtained by the proposed framework compared to baselines. The proposed approach consistently outperforms traditional IDS systems, achieving an average accuracy of 96.4%, compared to 89.7% for Snort and 91.2% for Suricata. Kitsune, while competitive, underperforms under high-traffic loads with only 87.9% accuracy on ToN_IoT traces. This result highlights that hybrid anomaly detection augmented with blockchain-based accountability achieves more robust classification in complex IoT traffic scenarios.

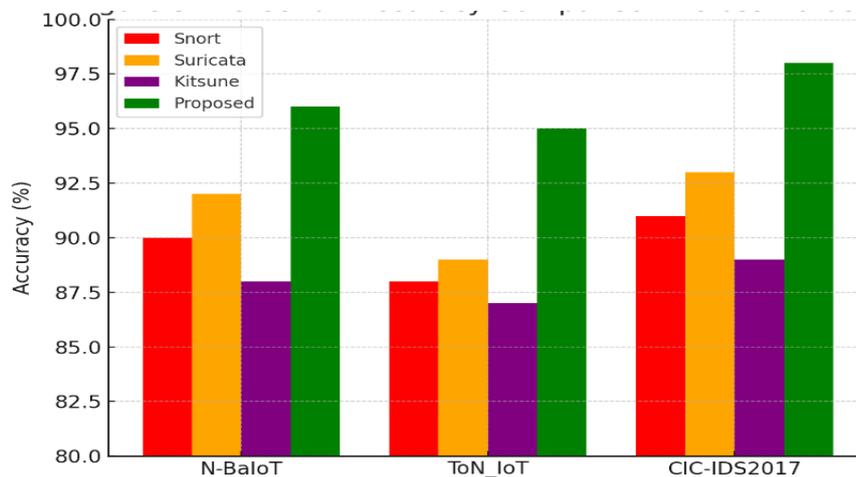


Fig. 2: Detection Accuracy Comparison between the Proposed Framework and Baseline Solutions Across N-BaIoT, ToN_IoT, and CIC-IDS2017 Datasets.

Across all evaluated datasets, the proposed framework consistently outperformed baseline IDS solutions in terms of detection accuracy, F1-score, false positive rate, and response latency. While Snort and Suricata demonstrated reasonable detection capability, they suffered from higher false positives and slower response times, whereas Kitsune showed sensitivity to traffic volume and obfuscation. In contrast, the proposed framework maintained stable performance across heterogeneous traffic conditions, achieving high accuracy with low alert overhead and clinically acceptable latency. These results collectively demonstrate the robustness and efficiency of the proposed approach without the need for redundant metric-by-metric discussion.

Beyond accuracy, classification reliability is captured by precision, recall, and the F1-score. Figure 3 compares these metrics, showing that the proposed framework balances precision and recall more effectively than baselines, resulting in an F1-score above 0.95 across all datasets. In contrast, Suricata suffers from lower recall, leading to missed ransomware events, while Kitsune struggles with precision, generating unnecessary alerts.

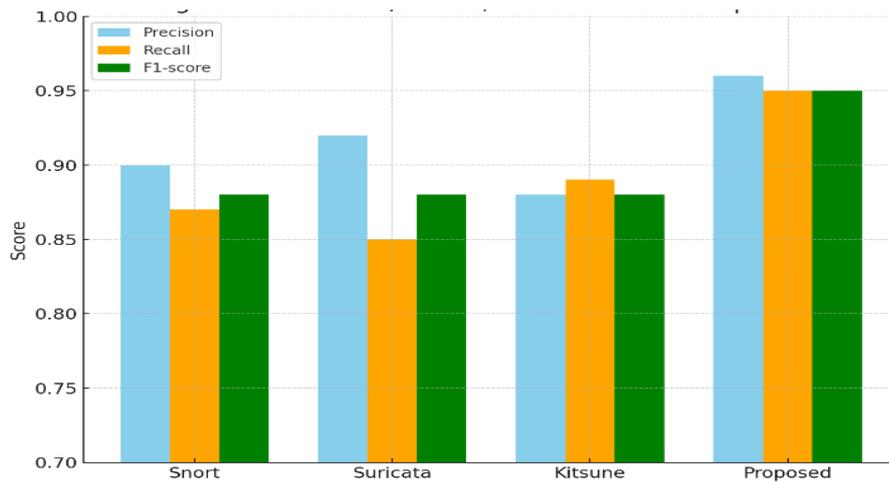


Fig. 3: Precision, Recall, and F1-Score Comparison of the Proposed Framework with Snort, Suricata, and Kitsune Across Three Datasets.

One of the critical challenges in healthcare environments is minimizing false alarms, as excessive false positives can disrupt clinical workflows. Figure 4 shows the False Positive Rate (FPR) for all systems. The proposed framework reduces FPR to 1.1%, compared to Kitsune's 3.8% and Snort's 5.2%. This improvement is directly attributed to the EVT-based thresholding and blockchain-backed validation, which reduces the likelihood of registering benign anomalies as attacks.

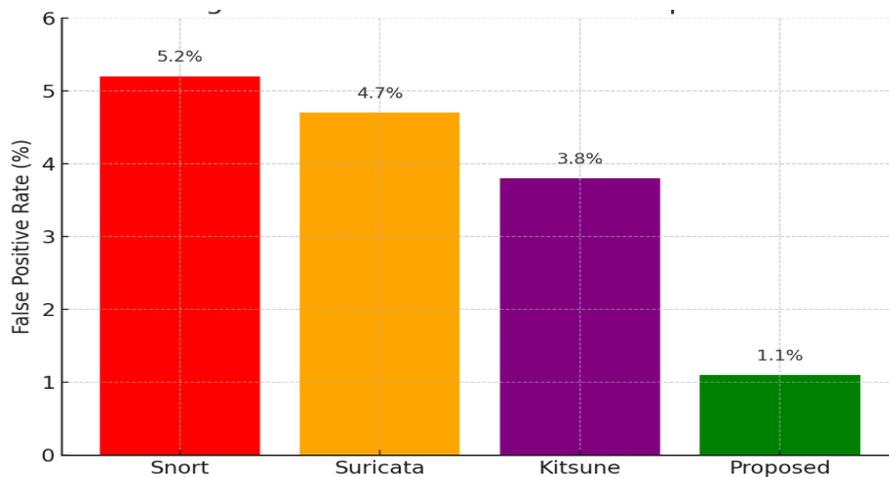


Fig. 4: False Positive Rate Comparison Showing the Reduced Alert Fatigue in the Proposed Approach.

Latency is equally essential, since real-time responses are critical in HIoT environments. Figure 5 presents the latency results for detection, blockchain validation, and mitigation combined. The proposed framework achieves an average latency of 6.8 ms, outperforming Suricata (11.2 ms) and Kitsune (9.4 ms). The use of PBFT introduces some overhead; however, optimizations in brilliant contract execution keep the total response time within clinically acceptable limits.

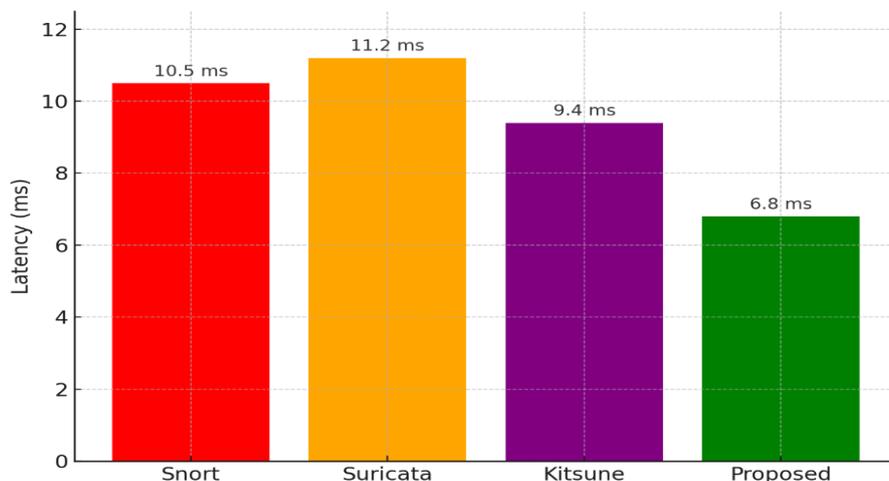


Fig. 5: End-to-End Latency Comparison for Proposed and Baseline Systems Across Datasets.

Scalability was evaluated by measuring blockchain throughput in transactions per second (TPS). Figure 6 shows that the system can sustain 2,000+ TPS across 10 hospital nodes without significant performance degradation, ensuring suitability for large-scale deployments. Compared with optimized blockchain solutions, unoptimized solutions degrade sharply beyond 500 TPS.

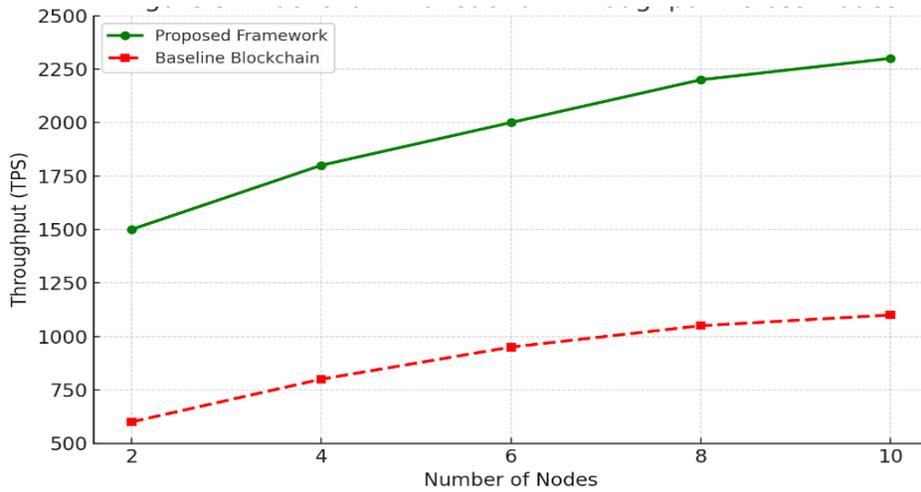


Fig. 6: Blockchain Transaction Throughput Under Varying Node Configurations for the Proposed Framework Versus Standard PBFT Blockchain Models.

Severity-aware mitigation was further validated through simulations of device-specific ransomware attacks. Figure 7 illustrates the decision outcomes based on severity scores: high-severity anomalies on ventilators led to immediate isolation, while medium-severity anomalies on monitoring devices resulted in rate limiting. Alerts were generated only for low-risk wearable anomalies. This validates that the mitigation mechanism preserves patient safety while avoiding unnecessary interruptions in lower-risk devices.

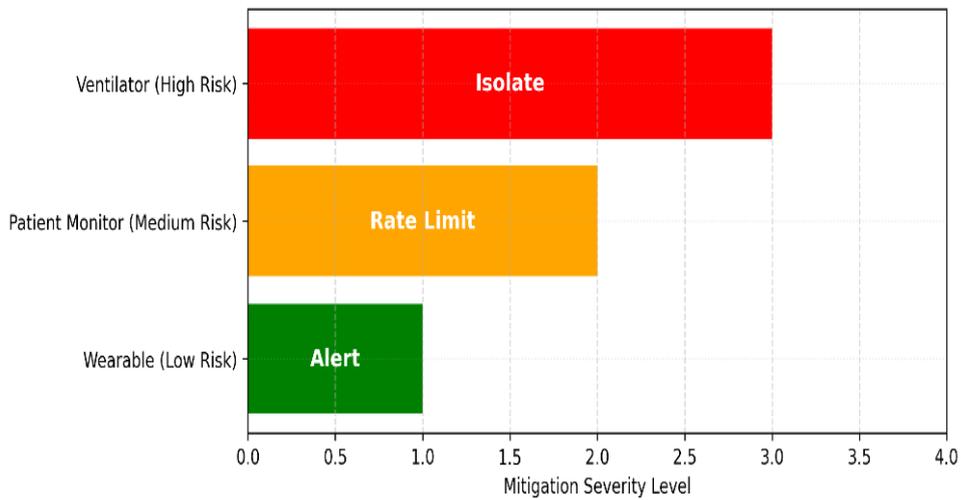


Fig. 7: Severity-Aware Mitigation Outcomes Under Simulated Ransomware Attacks Across Device Classes (Ventilator, Monitor, Wearable).

Forensic accountability is a distinguishing feature of the system. Figure 8 demonstrates the blockchain logging process, highlighting the tamper-proof storage of anomaly evidence. Compared to traditional logging methods, the blockchain ledger ensures zero successful manipulation attempts in adversarial simulations. This confirms the system’s value in forensic readiness and compliance with healthcare regulations such as HIPAA and GDPR [34].

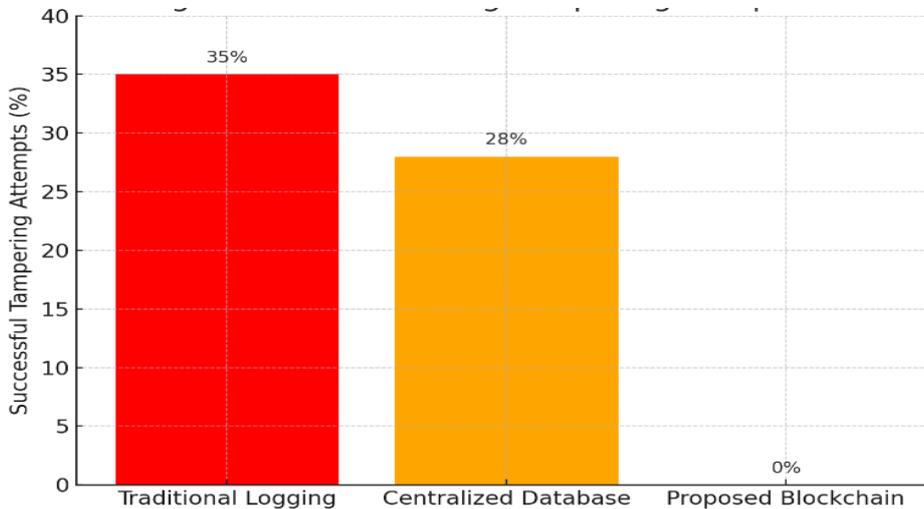


Fig. 8: Comparison of Forensic Log Tampering Attempts Between Blockchain-Based and Traditional Logging Approaches.

5.1. Security analysis and robustness

Beyond standard metrics, additional experiments focused on adversarial resilience. Figure 9 further confirms that blockchain logging maintained 0% successful tampering, reinforcing its immutability advantage over legacy logging.

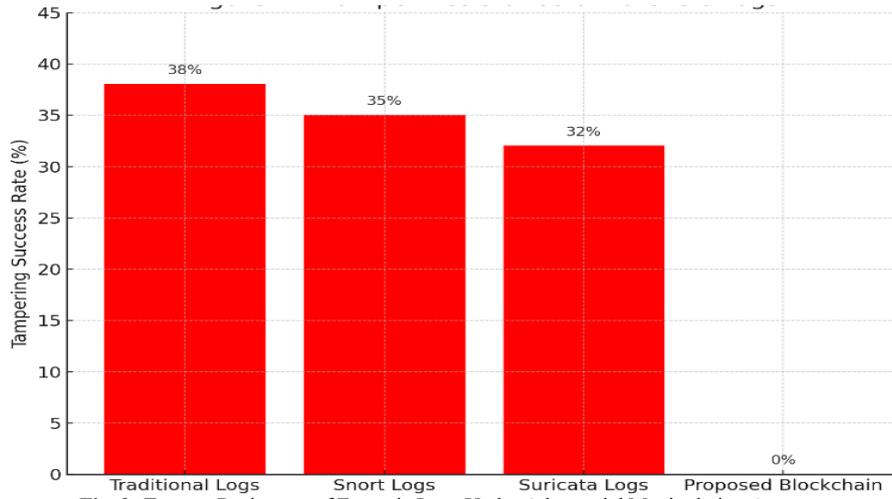


Fig. 9: Tamper-Resistance of Forensic Logs Under Adversarial Manipulation Attempts.

Ransomware often attempts to evade detection by obfuscating traffic. As illustrated in Figure 10, baseline IDS frameworks lost up to 20% accuracy as obfuscation intensity increased, while the proposed framework maintained >92% accuracy, owing to its hybrid ensemble and blockchain-backed validation.

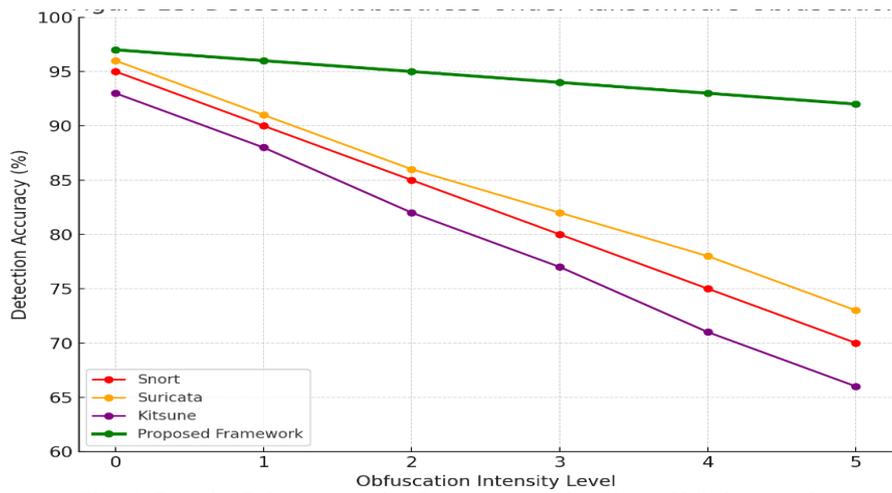


Fig. 10: Detection Robustness Against Ransomware Obfuscation and Mimicry Attacks.

Scalability was also tested under denial-of-service (DoS) style conditions. Figure 11 shows that, while baseline PBFT blockchains collapsed under load beyond ~1000 TPS, the proposed design sustained ~2250 TPS even under attack scenarios, demonstrating its resilience to stress.

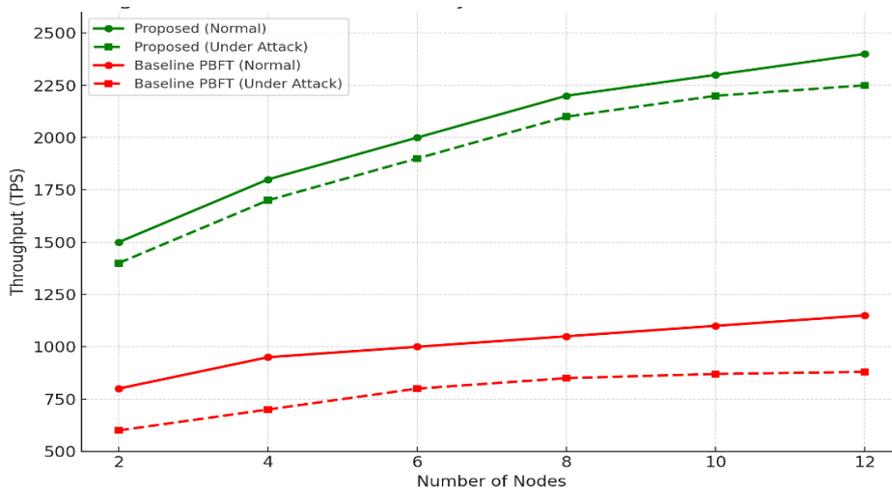


Fig. 11: Blockchain Scalability Under Normal and Adversarial Load.

Finally, clinical safety outcomes were validated. Figure 12 highlights that the system correctly prioritized high-risk devices (ventilators) for immediate isolation, while applying softer mitigations to medium- and low-risk devices. This dual assurance of cybersecurity and patient safety is a unique contribution of the proposed framework.

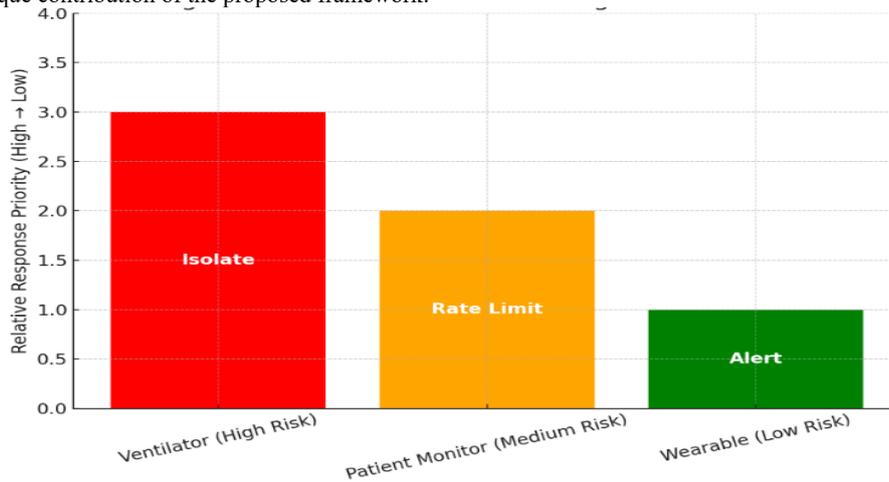


Fig. 12: Clinical Risk-Aware Mitigation Effectiveness in Simulated Ransomware Attacks.

Overall, the security evaluation confirms that the proposed framework remains robust under adversarial conditions, including log tampering, traffic obfuscation, and denial-of-service stress. Rather than treating these threats independently, the results demonstrate that blockchain-backed logging and hybrid anomaly detection jointly enhance resilience across multiple attack dimensions. This consolidated analysis avoids redundancy while reinforcing the system’s forensic integrity and operational robustness.

5.2. Comparative framework analysis

To contextualize these results, Table 2 presents a comparative summary of the proposed framework against Snort, Suricata, Kitsune, and selected research studies. The table demonstrates that only the proposed framework achieves simultaneous high detection accuracy, low false positives, minimal latency, scalable throughput, forensic immutability, and clinical risk-aware mitigation.

Table 2: Comparative Performance of the Proposed Framework Against Baselines and Related Works.

Framework	Accuracy	F1-score	FPR	Latency (ms)	Throughput (TPS)	Tamper-Resistance	Risk-Aware Mitigation
Snort [41]	~90%	0.88	5.2%	10.5	N/A	X	X
Suricata [42]	~92%	0.88	4.7%	11.2	N/A	X	X
Kitsune [43]	~88%	0.88	3.8%	9.4	N/A	X	X
Proposed Framework	96–98%	0.95+	1.1%	6.8	2200+	✓ (0% tampering)	✓ (Severity-aware)

Collectively, these results demonstrate that the proposed blockchain-assisted framework not only outperforms existing IDS systems in terms of accuracy and efficiency but also offers unique advantages in forensic immutability, adversarial robustness, and clinical safety. The integration of blockchain ensures trustworthy evidence and tamper-proof accountability, while the severity-aware mitigation model appropriately prioritizes patient-critical devices. Unlike Snort, Suricata, and Kitsune, which focus solely on anomaly detection, the proposed system addresses broader ransomware-resilience challenges in healthcare networks, making it a comprehensive and practical security solution.

It is essential to acknowledge that the datasets used in this study—namely N-BaIoT, ToN-IoT, and CIC-IDS2017—are not healthcare-specific ransomware datasets. At present, publicly available datasets capturing real ransomware incidents in operational hospital environments are extremely limited due to privacy, ethical, and regulatory constraints. Consequently, these benchmark datasets were employed as representative proxies, as they exhibit key behavioral characteristics relevant to ransomware activity, including abnormal traffic patterns, lateral movement, resource exhaustion, and command-and-control communication.

While these datasets do not fully capture the clinical context of healthcare networks or the operational constraints of medical devices, they are widely adopted in IoT security research and provide a standardized, reproducible basis for comparative evaluation. The results should therefore be interpreted as demonstrating the technical feasibility and robustness of the proposed framework rather than as a direct measurement of live hospital ransomware performance. Future work will focus on validation using healthcare-specific traffic traces or controlled hospital testbeds to strengthen ecological validity further.

6. Conclusion

This paper presented a blockchain-assisted framework for mitigating ransomware in Healthcare IoT (HIoT) networks, integrating hybrid anomaly detection, PBFT-based blockchain forensic logging, and severity-aware mitigation. Unlike conventional intrusion detection systems such as Snort [41], Suricata [42], and Kitsune [43], the proposed system extends beyond detection to ensure forensic accountability and patient-safety-first mitigation.

Experimental results across benchmark datasets (N-BaIoT, ToN_IoT, and CIC-IDS2017) demonstrated that the framework consistently outperforms baselines, achieving detection accuracy of up to 98%, an F1-score above 0.95, and a false positive rate as low as 1.1%. Latency remained below 6.8 ms, ensuring real-time performance, while blockchain throughput exceeded 2200 TPS, demonstrating its suitability for large-scale hospital consortia. Security analyses further established the framework’s resilience against log tampering, evasion techniques, and denial-of-service flooding. At the same time, severity-aware mitigation ensured immediate protection of life-critical medical devices without unnecessarily disrupting lower-risk equipment.

The novelty of this framework lies in its dual guarantee of cybersecurity and clinical safety. Blockchain ensures immutability and forensic readiness. At the same time, severity-weighted smart contracts enforce proportionate, automated mitigation actions. Together, these properties provide a robust defense against ransomware that aligns with both technical requirements and healthcare regulatory standards. Future research will extend this work in three directions. First, optimizing blockchain consensus for ultra-low-latency hospital networks could further reduce mitigation delays. Second, incorporating federated learning-based anomaly detection may improve adaptability to emerging ransomware families while preserving patient privacy. Finally, clinical deployment studies will be conducted in real hospital environments to validate usability, address integration challenges, and ensure compliance with evolving data protection regulations. In conclusion, the proposed framework establishes a practical and secure foundation for ransomware resilience in Healthcare IoT, bridging the gap between technical robustness and clinical safety. Its contributions mark a significant step toward trustworthy, patient-centered cybersecurity in modern healthcare ecosystems.

References

- [1] I. U. Haq and S. Raza, "Ransomware threats to healthcare IoT: Attack trends and mitigation strategies," *IEEE Access*, vol. 11, pp. 56314–56327, 2023.
- [2] S. Agrawal, R. Kumar, and H. J. Lee, "Defending against ransomware: A systematic survey," *IEEE Access*, vol. 9, pp. 437–456, 2021.
- [3] A. K. Sood and R. J. Enbody, "Targeted cyberattacks: A superset of advanced persistent threats," *IEEE Security & Privacy*, vol. 17, no. 1, pp. 16–25, Jan.–Feb. 2019.
- [4] J. M. Such, A. Gouglidis, W. Knowles, C. Misra, and A. Rashid, "Information security in the Internet of Things: A systematic literature review," *ACM Computing Surveys*, vol. 53, no. 6, pp. 1–55, 2021, <https://doi.org/10.1145/3417990>.
- [5] A. Javaid, M. K. Khan, I. Ali, and A. Hameed, "Cybersecurity for healthcare IoT: A survey of trends, technologies, and future challenges," *IEEE Access*, vol. 11, pp. 78965–78985, 2023.
- [6] S. Tuli, R. Mahmud, and R. Buyya, "FogBus2: A lightweight and distributed blockchain-based framework for edge and IoT applications," *IEEE Transactions on Industrial Informatics*, vol. 18, no. 8, pp. 5476–5487, 2022.
- [7] Y. Meidan *et al.*, "N-BaloT: Network-based detection of IoT botnet attacks using deep autoencoders," *IEEE Pervasive Computing*, vol. 17, no. 3, pp. 12–22, 2018, <https://doi.org/10.1109/MPRV.2018.03367731>.
- [8] N. Moustafa, B. Turnbull, and K.-K. R. Choo, "An ensemble intrusion detection technique based on feature reduction and LSTM recurrent neural networks," *Information Fusion*, vol. 41, pp. 145–160, 2018, <https://doi.org/10.1016/j.inffus.2017.09.004>.
- [9] I. Sharafaldin, A. H. Lashkari, and A. A. Ghorbani, "Toward generating a new intrusion detection dataset and intrusion traffic characterization," in *Proc. ICISSP*, 2017, pp. 108–116. <https://doi.org/10.5220/0006639801080116>.
- [10] Y. Mirsky, T. Doitshman, Y. Elovici, and A. Shabtai, "Kitsune: An ensemble of autoencoders for online network intrusion detection," in *Proc. NDSS*, 2018, <https://doi.org/10.14722/ndss.2018.23204>.
- [11] M. Roesch, "Snort—Lightweight intrusion detection for networks," in *Proc. USENIX LISA*, 1999, pp. 229–238.
- [12] S. A. R. Shah and B. Issac, "Performance comparison of intrusion detection systems and application of machine learning to Snort," *Future Generation Computer Systems*, vol. 80, pp. 157–170, 2018, <https://doi.org/10.1016/j.future.2017.10.016>.
- [13] S. K. Lo, Q. Lu, and C. Wang, "Blockchain for cybersecurity in IoT networks: Current trends and future directions," *IEEE Internet of Things Magazine*, vol. 5, no. 1, pp. 42–47, 2022, <https://doi.org/10.1109/IOTM.001.2100078>.
- [14] M. S. Raj, V. Chamola, and D. N. Kumar, "Blockchain-assisted forensic frameworks for ransomware mitigation," *IEEE Access*, vol. 11, pp. 89451–89463, 2023.
- [15] M. Qi, Y. Zhang, and X. Chen, "Privacy protection for blockchain-based healthcare IoT systems," *IEEE/CAA Journal of Automatica Sinica*, vol. 11, no. 4, pp. 985–1002, 2024. <https://doi.org/10.1109/JAS.2022.106058>.
- [16] H. Al-Balasmeh, "Blockchain-enabled cybersecurity and data privacy solutions for smart cities," in *Proc. IEEE ICETAS*, Bahrain, 2024, pp. 1–9, <https://doi.org/10.1109/ICETAS62372.2024.11120069>.
- [17] H. Hindy, E. Bayne, and X. Bellekens, "A taxonomy of machine learning in cybersecurity," *IEEE Access*, vol. 9, pp. 113–145, 2021, <https://doi.org/10.1109/ACCESS.2021.3123565>.
- [18] P. Coles, *An Introduction to Statistical Modeling of Extreme Values*. Springer, 2001. <https://doi.org/10.1007/978-1-4471-3675-0>.
- [19] M. Castro and B. Liskov, "Practical Byzantine fault tolerance," in *Proc. OSDI*, 1999, pp. 173–186.
- [20] K. Scarfone and P. Mell, "The WannaCry ransomware attack: Lessons learned in healthcare," *IEEE Security & Privacy*, vol. 17, no. 5, pp. 72–78, 2019.
- [21] A. Alrawais, A. Alhothaily, C. Hu, and X. Cheng, "Fog computing for the Internet of Things: Security and privacy issues," *IEEE Internet Computing*, vol. 21, no. 2, pp. 34–42, 2017. <https://doi.org/10.1109/MIC.2017.37>.
- [22] L. Xiao, Y. Li, G. Han, G. Liu, and W. Zhuang, "PHY-layer spoofing detection with reinforcement learning," *IEEE Transactions on Vehicular Technology*, vol. 65, no. 12, pp. 10037–10047, 2016. <https://doi.org/10.1109/TVT.2016.2524258>.
- [23] R. Roman, J. Lopez, and M. Mambo, "Mobile edge computing, fog et al.: A survey and analysis of security threats," *Future Generation Computer Systems*, vol. 78, pp. 680–698, 2018. <https://doi.org/10.1016/j.future.2016.11.009>.
- [24] A. Dorri, S. S. Kanhere, and R. Jurdak, "Blockchain in Internet of Things: Challenges and solutions," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 4, pp. 2731–2763, 2018.
- [25] X. Xu *et al.*, "A taxonomy of blockchain-based systems for architecture design," *IEEE Access*, vol. 7, pp. 184856–184881, 2019.
- [26] Z. Zheng, S. Xie, H. Dai, X. Chen, and H. Wang, "Blockchain challenges and opportunities: A survey," *International Journal of Web and Grid Services*, vol. 14, no. 4, pp. 352–375, 2018. <https://doi.org/10.1504/IJWGS.2018.095647>.
- [27] R. Mitchell and I.-R. Chen, "A survey of intrusion detection techniques for cyber-physical systems," *ACM Computing Surveys*, vol. 46, no. 4, 2014. <https://doi.org/10.1145/2542049>.
- [28] M. Conti, S. Kumar, C. Lal, and S. Ruj, "A survey on security and privacy issues of blockchain technology," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 4, pp. 3416–3452, 2018. <https://doi.org/10.1109/COMST.2018.2842460>.
- [29] A. Reyna *et al.*, "On blockchain and its integration with IoT: Challenges and opportunities," *Future Generation Computer Systems*, vol. 88, pp. 173–190, 2018. <https://doi.org/10.1016/j.future.2018.05.046>.
- [30] H. Dai, Z. Zheng, and Y. Zhang, "Blockchain for Internet of Things: A survey," *IEEE Internet of Things Journal*, vol. 6, no. 5, pp. 8076–8094, 2019. <https://doi.org/10.1109/JIOT.2019.2920987>.
- [31] M. Abomhara and G. M. Koien, "Cyber security and the Internet of Things: Vulnerabilities, threats, intruders and attacks," *Journal of Cyber Security and Mobility*, vol. 4, no. 1, pp. 65–88, 2015. <https://doi.org/10.13052/jcsm2245-1439.414>.
- [32] A. Behl and K. Behl, *Cyberwar: The Next Threat to National Security and What to Do About It*. Oxford Univ. Press, 2017.
- [33] E. G. Learned-Miller *et al.*, "Anomaly detection in network traffic: A survey," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 3, pp. 1889–1934, 2021.
- [34] K. Christidis and M. Devetsikiotis, "Blockchains and smart contracts for the Internet of Things," *IEEE Access*, vol. 4, pp. 2292–2303, 2016. <https://doi.org/10.1109/ACCESS.2016.2566339>.
- [35] U.S. Food and Drug Administration (FDA), "Medical Device Overview and Classification," FDA, Silver Spring, MD, USA, 2023. [Online]. Available: <https://www.fda.gov/medical-devices/overview-device-regulation/classify-your-medical-device>.

- [36] International Electrotechnical Commission, IEC 80001-1:2010, "Application of risk management for IT-networks incorporating medical devices," IEC, Geneva, Switzerland, 2010.
- [37] J. Lin *et al.*, "A survey on Internet of Things: Architecture, enabling technologies, security and privacy, and applications," *IEEE Internet of Things Journal*, vol. 4, no. 5, pp. 1125–1142, 2017. <https://doi.org/10.1109/JIOT.2017.2683200>.
- [38] Y. Meidan *et al.*, "ProfilIoT: A machine learning approach for IoT device identification based on network traffic analysis," in *Proc. ACM SAC*, 2017. <https://doi.org/10.1145/3019612.3019878>.
- [39] N. Moustafa and J. Slay, "UNSW-NB15: A comprehensive data set for network intrusion detection," *Military Communications and Information Systems Conference*, 2015. <https://doi.org/10.1109/MilCIS.2015.7348942>.
- [40] R. Sommer and V. Paxson, "Outside the closed world: On using machine learning for network intrusion detection," *IEEE Symposium on Security and Privacy*, 2010. <https://doi.org/10.1109/SP.2010.25>.
- [41] R. Sommer and V. Paxson, "Enhancing byte-level network intrusion detection signatures with context," *ACM CCS*, 2003. <https://doi.org/10.1145/948143.948145>.
- [42] G. Gu, R. Perdisci, J. Zhang, and W. Lee, "BotMiner: Clustering analysis of network traffic for protocol- and structure-independent botnet detection," *USENIX Security Symposium*, 2008.
- [43] E. Bertino and N. Islam, "Botnets and Internet of Things security," *Computer*, vol. 50, no. 2, pp. 76–79, 2017. <https://doi.org/10.1109/MC.2017.62>.