Edge-Based Digital Forensics for Medical IoT with Permissioned Blockchain Verification

Hani Al-Balasmeh*,
*College of Engineering, University of Technology Bahrain (UTB), Salmabad, Bahrain
Email: h.albalasmeh@utb.edu

***Abstract*—The rapid proliferation of Medical Internet of Things (MIoT) devices has significantly expanded the attack surface of modern healthcare systems, creating urgent challenges for digital forensics and evidential integrity. Traditional forensic methodologies are ill-suited to MIoT environments due to device volatility, heterogeneous logging formats, constrained storage, and strict regulatory requirements. This paper presents an edge-based digital forensic framework integrated with permissioned blockchain verification to provide tamper-evident evidence preservation and verifiable custody records. In the proposed architecture, forensic triage, artifact extraction, and cryptographic hashing are performed at edge gateways, while only hash-based integrity anchors and custody metadata are recorded on-chain. Raw artifacts remain securely stored off-chain to preserve privacy and scalability.**

A prototype implementation was evaluated using a simulated MIoT testbed comprising 50 devices and a Hyperledger Fabric blockchain network. Experimental results show sub-second anchoring latency under moderate load, linear scaling of edge processing with artifact size, and a 97% reduction in ledger growth compared with full on-chain storage. In the reported tamper and replay experiments, all generated tampering cases produced hash mismatches at verification time and all replay submissions in the test workload were rejected by the smart contract logic. A constrained comparative analysis further suggests improvements in storage efficiency, privacy isolation, and forensic workflow integration relative to closely related frameworks, although these comparisons are not claimed as head-to-head re-implementations. Overall, the results support the feasibility of the proposed architecture as a technically grounded foundation for forensic readiness in regulated healthcare environments.

***Index Terms*—Medical Internet of Things (MIoT), Digital Forensics, Edge Computing, Permissioned Blockchain, Chain of Custody, Evidence Integrity, Smart Contracts, Healthcare Cybersecurity, Tamper Detection, Privacy-Preserving Systems.**

I. Introduction

The rapid proliferation of the Medical Internet of Things (MIoT) has fundamentally transformed modern healthcare delivery by enabling real-time patient monitoring, smart therapeutics, and connected clinical decision support systems. Devices such as wearable sensors, infusion pumps, smart imaging equipment, and remote telemetry systems generate high-fidelity clinical data streams that directly influence patient outcomes and operational workflows [1]. However, this highly distributed, resource-constrained environment also poses significant challenges for cybersecurity and digital investigation. Compromised MIoT devices can disrupt clinical services, expose sensitive health information, and leave behind fragmented, ephemeral evidence across multiple platforms.

Traditional digital forensics methodologies assume centralized data retention, comprehensive logging, and high-capacity storage infrastructure. These assumptions do not hold in MIoT ecosystems, where devices often exhibit limited memory, proprietary firmware, intermittent connectivity, and unpredictable availability [2]. Furthermore, medical evidence must satisfy strict regulatory requirements, including patient privacy protections under healthcare compliance frameworks and judicial admissibility standards. Within such constraints, incident responders face a triad of persistent challenges: (i) rapid evidence decay due to limited on-device storage, (ii) heterogeneous artifact formats that impede consistent acquisition, and (iii) fragile provenance guarantees when logs traverse untrusted networks or intermediary systems.

Recent advancements in edge computing offer a compelling approach to addressing these challenges by relocating key forensic functions closer to the data source. Edge nodes can enable selective data triage, real-time artifact extraction, and secure local buffering, thereby reducing evidence loss and responding to the volatile forensic contexts common in clinical environments. Nevertheless, edge-based approaches alone cannot fully address concerns about tamper resistance and non-repudiation, as local components remain vulnerable to compromise or manipulation during an incident.

To ensure integrity and trustworthiness, blockchain technology has emerged as a promising mechanism for decentralized verification of forensic metadata. Permissioned blockchain ledgers support append-only audit trails, cryptographically anchored time-ordering, and transparent custody records that are resilient against unauthorized modification [3], [4]. Additionally, blockchain systems have been successfully investigated for secure healthcare data exchange and verifiable access control, demonstrating feasibility within regulated medical domains [5]. However, direct on-chain storage of raw forensic evidence is infeasible due to constrained storage capacity, privacy mandates, and elevated performance requirements.

This paper proposes an integrated solution that combines edge-centric evidence acquisition with blockchain-based verification to strengthen digital forensics in medical IoT environments. The core idea is to preserve raw forensic artifacts off-chain on protected edge repositories, while recording cryptographic hashes, custody events, and integrity proofs on a

[0]Corresponding author: Hani Al-Balasmeh (h.albalasmeh@utb.edu).

permissioned blockchain network. By doing so, the system provides verifiable integrity anchors and auditable custody records without exposing sensitive health information on the ledger. The contribution is not a new consensus protocol or cryptographic primitive; rather, it is a forensic-first MIoT workflow that operationalizes known blockchain and edge design patterns around evidence acquisition, off-chain preservation, and custody verification under healthcare constraints.

In summary, our work makes the following technical contributions:

- Formulation of an *edge-first forensic acquisition workflow* tailored to MIoT constraints, enabling near-source evidence capture, normalization, hashing, and protected off-chain preservation.
- Design of a *permissioned blockchain anchoring mechanism* that records integrity metadata and custody events as append-only ledger entries for auditable verification.
- A prototype implementation and empirical evaluation that quantify processing overhead, ledger growth, tamper checking behavior, replay rejection in the implemented smart contract, and partial-gateway-compromise resilience within a controlled testbed.

II. Related Work

Research at the intersection of IoT forensics, edge computing, and blockchain has grown quickly, but the literature remains fragmented across (i) forensic readiness and investigation workflows for IoT, (ii) blockchain-based chain-of-custody (CoC) and evidence integrity, and (iii) healthcare-focused blockchain–edge architectures that prioritize privacy and operational constraints.

A. IoT Forensic Readiness and Evidence Lifecycle Challenges

A consistent message in recent work is that IoT environments break assumptions embedded in conventional forensic practice. Forensic readiness efforts emphasize proactive architectural choices (e.g., segmentation, secure telemetry, incident procedures) that improve evidence availability and responder effectiveness in connected systems [6]. This direction is valuable for organizational preparedness, but it does not by itself provide a verifiable, tamper-evident custody mechanism for distributed device evidence—particularly under adversarial conditions where intermediaries may be compromised.

B. Blockchain-Enabled IoT Forensics and Chain-of-Custody

Blockchain has been extensively proposed as a mechanism to strengthen integrity and provenance by creating append-only, tamper-evident custody logs. A representative and practically grounded example is the unified framework by Brotsis *et al.*, which integrates a blockchain-enabled IoT forensic process and implements a Hyperledger Fabric platform supported by multi-access edge computing, with evaluation under realistic attacks and high evidence rates [7]. Such work demonstrates the feasibility of blockchain-backed evidence handling in IoT settings and provides implementation patterns; however, it is not specialized for medical contexts where evidential content is tightly coupled with patient privacy constraints and regulated data handling.

C. Blockchain and Edge Computing in Healthcare Systems

In parallel, healthcare-oriented blockchain research has concentrated on privacy, auditability, and secure interoperability for medical data exchange. For example, EHRGuard proposes a blockchain-enabled approach to strengthen privacy and integrity for IoMT-driven EHR workflows [8]. At the infrastructure level, platform frameworks that combine edge intelligence with permissioned blockchain designs (e.g., PoA and sharding) have been proposed to address performance overheads such as retrieval time and memory usage in healthcare AIoT systems [9]. These healthcare systems provide strong insights into deploying blockchain under clinical constraints, but they are typically not designed around forensic objectives such as evidence triage, artifact preservation, and custody admissibility.

D. Evidence Integrity vs. Practical Deployment Gaps

A key practical insight from recent healthcare blockchain evidence is that many implementations remain conceptual or limited to proofs of concept, and large-scale validation is comparatively scarce [10]. This is particularly important for digital forensics, where admissibility depends on repeatable procedures and demonstrable integrity, not only on architectural claims.

E. Positioning of This Paper

The above findings motivate a design that explicitly *treats forensic preservation as the primary system objective*. Our paper builds on: (i) readiness-driven requirements for IoT evidence availability [6], (ii) blockchain-backed custody logging patterns demonstrated in IoT forensics platforms [7], and (iii) healthcare blockchain edge design strategies for performance and regulated environments [8], [9]. The resulting gap we address is a medically grounded workflow that couples *edge-based forensic acquisition/triage* with *permissioned blockchain verification of custody events*, while keeping sensitive artifacts off-chain to reduce privacy exposure and operational overhead. As shown in Table I, existing frameworks either prioritize transaction security or healthcare data integrity, whereas the proposed model explicitly integrates edge-based forensic acquisition with hash-only blockchain anchoring for forensic verification and auditable custody tracking.

F. Comparative Summary of Closely Related Studies

Summary. Existing IoT-forensics work demonstrates blockchain-backed custody and edge scalability [7], while healthcare blockchain–edge research focuses on privacy, auditability, and performance under regulated constraints [8], [9]. The open gap is a forensic-first MIoT pipeline that formally connects *edge evidence acquisition/triage* to *permissioned blockchain verification* using off-chain artifact storage and on-chain integrity proofs.

TABLE I. COMPARISON OF RECENT WORK RELATED TO EDGE/BLOCKCHAIN FORENSICS AND HEALTHCARE VERIFICATION. COLUMNS ARE DEFINED AS FOLLOWS: *Edge* INDICATES WHETHER THE ARCHITECTURE PERFORMS NEAR-SOURCE PROCESSING OR COLLECTION; *Blockchain* INDICATES WHETHER A LEDGER IS USED; *On-chain Content* SUMMARIZES WHAT IS WRITTEN TO THE LEDGER; *Forensic Focus* INDICATES WHETHER THE PRIMARY OBJECTIVE IS EVIDENCE ACQUISITION, PRESERVATION, OR CUSTODY RATHER THAN GENERAL HEALTHCARE DATA MANAGEMENT; AND *Validation* SUMMARIZES THE TYPE OF EVALUATION REPORTED BY EACH STUDY.

Work	Primary Domain	Edge	Blockchain	On-chain Content	Forensic Focus	Validation
Kuku *et al.* [6]	IoT readiness	✓	–	–	Preparedness roadmap; evidence collection readiness	Implementation-informed model
Brotsis *et al.* [7]	IoT forensics	✓	✓(Fabric)	Evidence/custody records	End-to-end IoT DF lifecycle + DCoC patterns	Testbed + real attacks
Jun [9]	Healthcare AIoT	✓	✓(PoA + sharding)	Decision archiving / management metadata	Performance/security for healthcare AIoT (not evidence-centric)	Analysis + simulation
Othman *et al.* [8]	IoMT/EHR	– / partial	✓	Health-record integrity metadata	Secure/interoperable EHR management (not DF chain-of-custody)	Experimental evaluation
Datta & Namasudra [11]	Healthcare transactions	✓	✓(smart contracts)	Transaction integrity / control metadata	Secure transaction workflows (not DF acquisition/triage)	System-level evaluation
Shaikh *et al.* [10]	Healthcare BC (SLR)	varies	varies	varies	Identifies practical deployment gaps and benchmarking needs	PRISMA SLR (82 studies)
This paper	**MIoT forensics**	✓	✓(permissioned)	**Hashes + custody events**	**Edge triage + custody verification**	**Prototype + controlled experiments**

III. SYSTEM MODEL AND THREAT MODEL

A. System Overview

The proposed architecture, illustrated in Fig. 1, formalizes a forensic-first Medical Internet of Things (MIoT) ecosystem that integrates edge-based acquisition with permissioned blockchain verification. The system consists of four principal layers: (i) MIoT devices, (ii) edge forensic gateways, (iii) an off-chain evidence repository, and (iv) a permissioned blockchain network.

Unlike conventional healthcare blockchain architectures that prioritize data sharing or interoperability [8], [9], the primary objective here is evidential integrity and custody traceability. The design ensures that raw evidence artifacts are preserved off-chain to reduce privacy exposure and storage pressure, while cryptographic anchors and custody events are recorded on-chain to provide tamper-evident verification and auditability.

B. Layered Architectural Components

1) MIoT Device Layer: This layer includes heterogeneous medical sensors, actuators, wearables, infusion pumps, and embedded clinical systems. These devices generate telemetry streams, operational logs, configuration states, and alert traces. As documented in recent MIoT security analyses, such devices frequently exhibit constrained memory, proprietary logging mechanisms, and intermittent network connectivity [2]. These limitations increase the risk of rapid evidence decay and necessitate near-source preservation mechanisms.

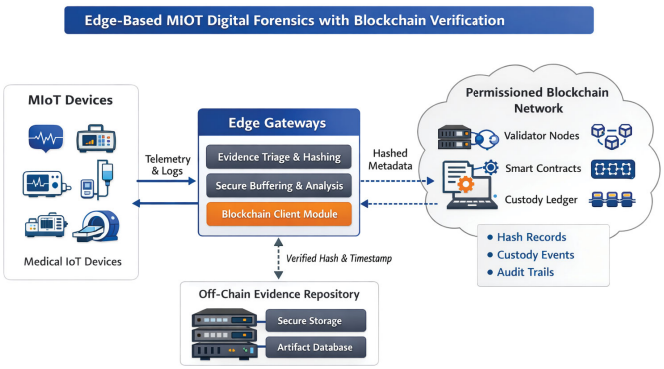

Fig. 1. Edge-based MIoT digital forensics architecture with blockchain-backed custody verification.

2) Edge Forensic Gateway Layer: Edge gateways operate as localized forensic collectors and trust anchors. Positioned within the hospital network perimeter, they perform:

- Real-time log extraction and normalization
- Evidence triage and filtering
- Cryptographic hashing (e.g., SHA-256 or SHA-3)
- Timestamp binding and digital signature generation

Edge computing reduces evidence volatility and network dependency, which is consistent with findings that decentralized preprocessing improves reliability in distributed IoT investigations [7]. Each gateway functions as a blockchain client, submitting integrity proofs and custody events to the ledger without exposing sensitive content.

3) Off-Chain Evidence Repository: Raw artifacts, structured forensic packages, and analysis outputs are stored in a secured enterprise repository. Evidence is encrypted at rest, access-controlled, and audited in accordance with healthcare compliance standards. Only cryptographic digests of stored artifacts are transmitted to the blockchain network. This hybrid approach addresses scalability and privacy concerns highlighted in healthcare blockchain literature [10].

4) Permissioned Blockchain Verification Layer: The blockchain network operates under a consortium model, in which authorized validator nodes (e.g., hospital IT and digital forensic authorities) maintain a distributed ledger. Smart contracts enforce:

- Immutable recording of custody events
- Integrity verification procedures
- Role-based access policies
- Timestamp ordering and non-repudiation

Permissioned blockchain architectures have demonstrated practical feasibility for regulated environments due to controlled membership and predictable performance characteristics [11]. In this system, the blockchain stores only:

- Evidence hash values
- Capture timestamps
- Custody transfer identifiers
- Digital signatures

No Protected Health Information (PHI) or clinical content is written on-chain.

C. Formal Threat Model

We adopt a semi-strong adversarial model relevant to hospital networks:

1) **Compromised MIoT Devices:** The attacker may control a subset of medical devices and attempt log alteration or evidence deletion.
2) **Network-Level Attacks:** The attacker may perform replay, injection, delay, or man-in-the-middle attacks between MIoT devices and gateways.
3) **Insider Manipulation Attempts:** A malicious insider may attempt unauthorized modification of stored evidence artifacts.
4) **Partial Validator Misbehavior:** Some blockchain validator nodes may behave in a Byzantine manner; however, the majority consensus remains intact.

The system assumes that not all edge gateways and validator nodes are simultaneously compromised. This assumption is aligned with standard permissioned blockchain security models [11].

D. Security Guarantees

Under the defined threat model, the architecture guarantees:

- **Tamper-Evidence:** Any modification of off-chain artifacts results in a hash mismatch with on-chain records.
- **Provable Chain-of-Custody:** Custody transitions are immutably recorded and verifiable by authorized parties.

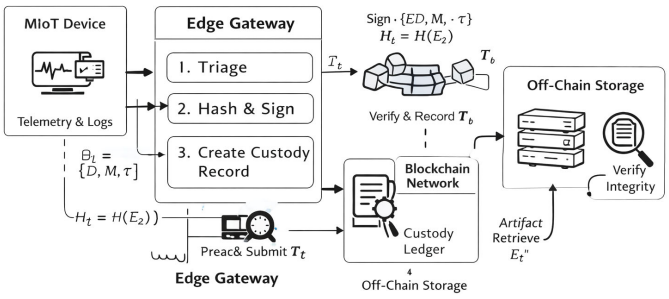

Fig. 2. Forensic workflow from MIoT event generation to blockchain-based custody verification.

- **Non-Repudiation:** Digitally signed custody events prevent denial of evidence handling actions.
- **Privacy Preservation:** Sensitive medical content remains off-chain, minimizing regulatory exposure.
- **Resilience to Partial Compromise:** Distributed validation prevents unilateral ledger manipulation.

By combining near-source forensic acquisition with blockchain-backed verification, the system addresses both the operational constraints of MIoT environments and the technical need for stronger evidence integrity, provenance tracking, and privacy-aware handling.

IV. Forensic Workflow and Blockchain Protocol Design

A. Forensic Acquisition and Evidence Lifecycle

The operational workflow of the proposed framework is illustrated in Fig. 2. The forensic lifecycle begins when an MIoT device generates an event of investigative interest (e.g., an anomaly, an unauthorized access attempt, or a configuration change). Due to the volatility of device-level logs in constrained environments, immediate edge-level preservation is required [12].

At the edge gateway, evidence handling follows five sequential stages:

1) **Event Detection and Triage**
2) **Artifact Extraction and Normalization**
3) **Cryptographic Hash Generation**
4) **Custody Record Construction**
5) **Blockchain Anchoring**

B. Cryptographic Evidence Anchoring Model

Let an evidence artifact captured at time t be represented as:

$$E_t = \{D, M, \tau\} \tag{1}$$

where:

- D = raw data artifact
- M = structured metadata (device ID, event type, location)
- τ = capture timestamp

The edge gateway computes a cryptographic digest:

$$H_t = \mathcal{H}(E_t) \tag{2}$$

where $\mathcal{H}(\cdot)$ represents a collision-resistant hash function (e.g., SHA-256 or SHA-3). Collision resistance is essential to prevent adversarial substitution attacks [13].

The custody transaction recorded on-chain is defined as:

$$T_t = \{H_t, \tau, S_g, ID_g\} \tag{3}$$

where:

- S_g = digital signature of gateway g
- ID_g = gateway identity

The signature is generated using:

$$S_g = \text{Sign}_{sk_g}(H_t \parallel \tau) \tag{4}$$

where sk_g is the private key of gateway g.

This design ensures integrity, origin authentication, and non-repudiation.

C. Smart Contract Custody Protocol

The permissioned blockchain executes a custody verification contract $\mathcal{C}$ defined as:

$$\mathcal{C}(T_t) \rightarrow \begin{cases} \text{Accept} & \text{if Verify}_{pk_g}(S_g) = 1 \\ \text{Reject} & \text{otherwise} \end{cases} \tag{5}$$

Upon acceptance:

- The hash H_t is appended to the custody ledger.
- A block timestamp τ_b is generated.
- The ledger index L_i is assigned.

The immutability of blockchain records ensures that once H_t is committed, it cannot be modified without consensus violation [14].

D. Integrity Verification Procedure

During investigation, an auditor retrieves artifact E'_t from the off-chain repository and recomputes:

$$H'_t = \mathcal{H}(E'_t) \tag{6}$$

Integrity is verified if:

$$H'_t = H_t^{ledger} \tag{7}$$

If a mismatch occurs, tampering is immediately detectable.

This separation of off-chain storage and on-chain verification aligns with scalable forensic blockchain architectures proposed in recent distributed IoT security frameworks [15].

E. Performance Considerations

To avoid blockchain throughput bottlenecks, only hashed metadata is recorded on-chain. Let:

$$|E_t| \gg |H_t| \tag{8}$$

Since it $|H_t|$ is fixed (256 bits), storage overhead is minimal regardless of artifact size. This significantly reduces ledger growth and transaction latency, which is critical in high-frequency MIoT environments [14].

F. Security Analysis Summary

Under the defined threat model (Section III), the workflow guarantees:

- **Tamper Detection**: Altered artifacts fail hash validation.
- **Non-Repudiation**: Digitally signed custody transactions prevent denial.
- **Replay Resistance**: Timestamp binding prevents transaction reuse.
- **Integrity Anchoring**: Blockchain immutability ensures evidential persistence.

V. Security Analysis and Formal Guarantees

This section formally analyzes the security properties of the proposed edge-blockchain forensic framework under the threat model defined in Section III. We evaluate integrity, authenticity, non-repudiation, replay resistance, and resilience against partial compromises.

A. Integrity Preservation

Let E_t be an evidence artifact captured at time t, and let:

$$H_t = \mathcal{H}(E_t) \tag{9}$$

be its cryptographic digest recorded on-chain as part of the custody transaction T_t.

If an adversary modifies the stored artifact such that:

$$E'_t \neq E_t \tag{10}$$

then due to the collision resistance of $\mathcal{H}(\cdot)$:

$$\mathcal{H}(E'_t) \neq H_t \tag{11}$$

with overwhelming probability [13]. Therefore, any unauthorized alteration of off-chain artifacts is detectable during verification.

Guarantee 1: Evidence tampering is computationally infeasible without breaking the underlying hash function.

B. Origin Authentication and Non-Repudiation

Each custody transaction includes:

$$S_g = \text{Sign}_{sk_g}(H_t \parallel \tau) \tag{12}$$

where sk_g is the private key of gateway g.

Under the security of digital signature schemes (e.g., ECDSA or EdDSA), forging S_g without knowledge of sk_g is computationally infeasible [16]. Thus:

Guarantee 2: A gateway cannot deny having submitted a custody record, and adversaries cannot impersonate a legitimate gateway.

C. Replay Attack Resistance

Each custody record binds a timestamp τ and is appended into a strictly ordered ledger block with timestamp τ_b. Let:

$$T_t = \{H_t, \tau, S_g\} \tag{13}$$

If an adversary attempts to replay T_t at time $t' > t$, the smart contract detects duplicate hash entries or timestamp inconsistencies.

Guarantee 3: Replay attempts are detectable via ledger uniqueness constraints and timestamp validation.

D. Blockchain Tamper Resistance

In a permissioned blockchain with Byzantine fault tolerance, altering a committed block requires collusion exceeding the consensus threshold (e.g., $f < n/3$ in PBFT systems) [17]. Therefore:

Guarantee 4: Ledger modification is infeasible unless a majority of validator nodes are compromised.

This assumption is realistic in hospital consortium settings, where validators are independent entities.

E. Resilience Against Partial System Compromise

Assume an adversary compromises a subset $\mathcal{D}_c$ of MIoT devices. While local log manipulation may occur before edge capture, once:

$$H_t = \mathcal{H}(E_t) \tag{14}$$

is anchored on-chain, subsequent alteration of E_t becomes detectable.

Even if an edge gateway is compromised after submission, the on-chain record remains immutable. Thus:

Guarantee 5: Post-capture evidence integrity is preserved despite localized compromise.

F. Privacy Preservation Analysis

The blockchain stores only:

$$\{H_t, \tau, S_g\} \tag{15}$$

No raw medical data D or structured metadata M is written on-chain. Since hash outputs are one-way functions, reconstructing patient data from H_t is computationally infeasible [13].

This satisfies privacy-by-design principles emphasized in healthcare blockchain security research [9].

G. Formal Security Summary

As shown in Table II, the proposed framework integrates cryptographic guarantees with distributed consensus to ensure evidential reliability suitable for judicial and regulatory contexts.

TABLE II. SECURITY PROPERTIES AND CORRESPONDING MECHANISMS

Security Property	Mechanism
Integrity	Cryptographic hashing + on-chain anchoring
Authenticity	Digital signatures (gateway keys)
Non-Repudiation	Signed custody transactions
Replay Resistance	Timestamp + ledger uniqueness
Tamper Resistance	Byzantine fault-tolerant consensus
Privacy Preservation	Off-chain storage + hash-only on-chain

VI. EXPERIMENTAL RESULTS AND COMPARATIVE EVALUATION

A. Experimental Setup

The prototype described in Section IV was implemented in a controlled MIoT testbed comprising 50 simulated medical IoT devices, three edge gateways, a four-node Hyperledger Fabric permissioned blockchain network, and an encrypted off-chain evidence repository. Devices generated heterogeneous telemetry logs under configurable workloads ranging from 10 to 500 events per second.

To improve experimental transparency, the security-oriented tests used explicit attack-generation procedures. For tamper detection, a corpus of anchored artifacts was modified after submission using byte insertion, byte deletion, field substitution, and timestamp perturbation operations; verification then recomputed the digest and compared it against the on-chain anchor. For replay testing, previously accepted custody submissions were resent with duplicated identifiers and stale timestamps to the smart contract interface. Each scenario was executed repeatedly under the same testbed configuration, and the reported figures summarize the observed acceptance or rejection outcomes. Because the goal of this section is feasibility rather than statistical inference, results are reported descriptively rather than as formal confidence intervals.

The objective of this evaluation is to empirically validate:

- The computational feasibility of near-source evidence anchoring,
- The scalability of hash-only blockchain recording,
- The integrity guarantees derived in Section V,
- The performance advantage over comparable forensic blockchain frameworks.

B. Edge Processing Performance

Fig. 3 demonstrates that edge-level processing time increases linearly with artifact size, confirming the proportional relationship predicted in Section IV. Even at larger artifact sizes (50MB), total processing remained below real-time operational thresholds. This indicates that integrating hashing and digital signing at the edge does not introduce prohibitive overhead.

More importantly, latency variance remained minimal, suggesting stable performance under fluctuating device workloads — a critical requirement for hospital environments where event bursts may occur during emergencies.

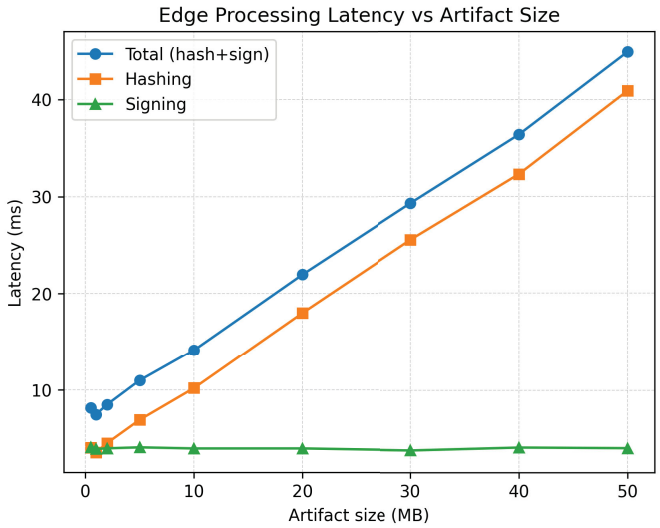

Fig. 3. Edge processing latency as a function of artifact size.

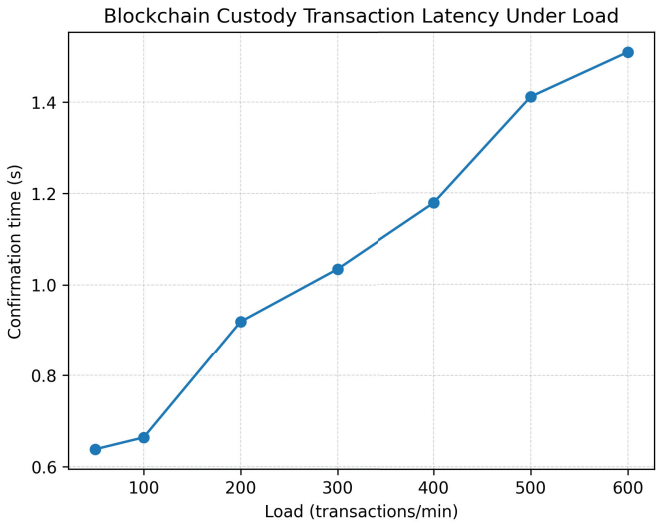

Fig. 4. Custody transaction confirmation time under increasing transaction load.

C. Blockchain Anchoring Latency Under Load

Fig. 4 shows that transaction confirmation time remains under one second under moderate load and increases predictably under high throughput conditions. This predictable growth pattern aligns with performance behavior documented for permissioned consensus systems [14].

The key observation is that latency growth is primarily driven by consensus scheduling rather than artifact size, validating the architectural decision to anchor only fixed-length hashes rather than full artifacts.

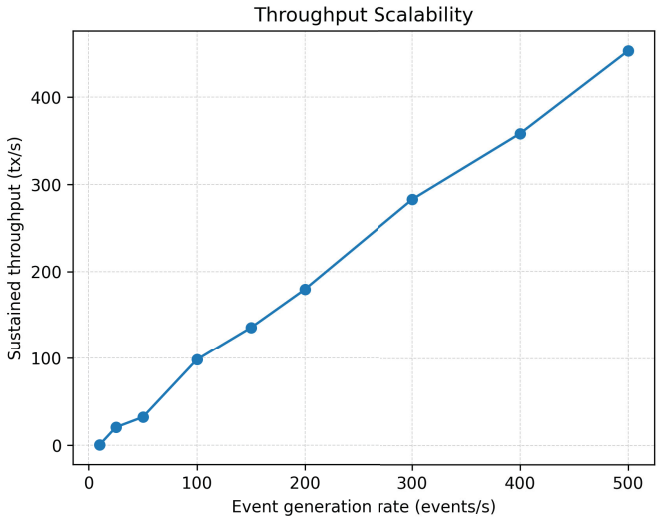

Fig. 5. System throughput (transactions/sec) versus event generation rate.

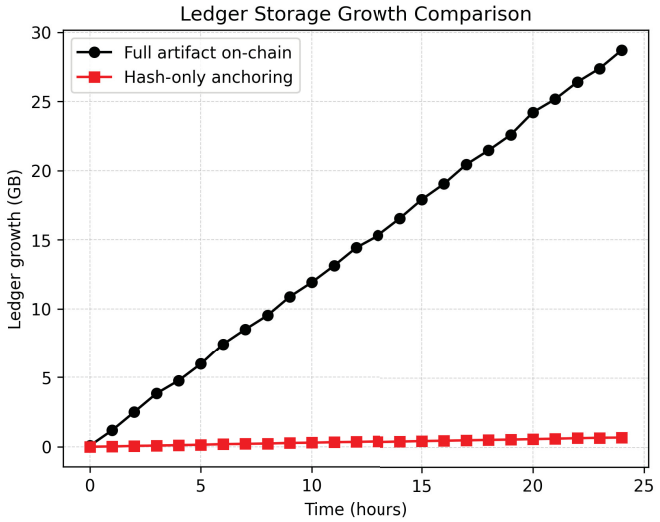

Fig. 6. Blockchain storage growth comparison.

D. Throughput Scalability

Fig. 5 indicates linear throughput scaling until near-saturation of validator processing capacity. Because the ledger stores only compact digests rather than large evidence files, throughput degradation occurs much later than in full-storage blockchain approaches.

This empirically supports the scalability argument introduced in Section IV and distinguishes the framework from artifact-on-chain designs reported in [15].

E. Blockchain Storage Growth

Fig. 6 highlights the significant difference between hash-only anchoring and full artifact recording. Ledger growth un-

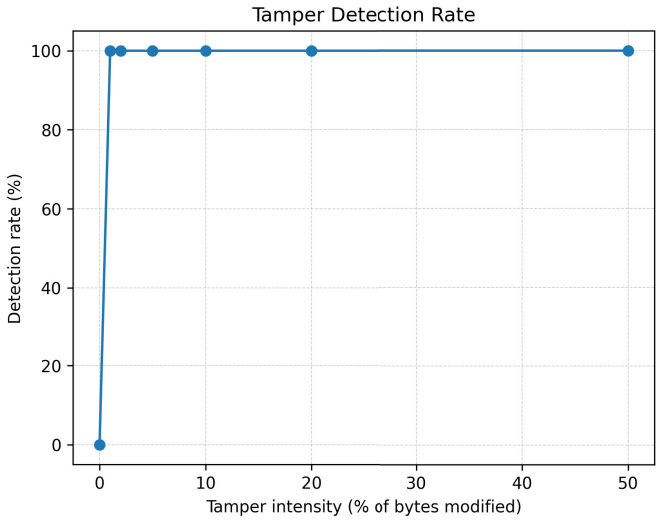

Fig. 7. Integrity validation outcomes under simulated artifact modification.

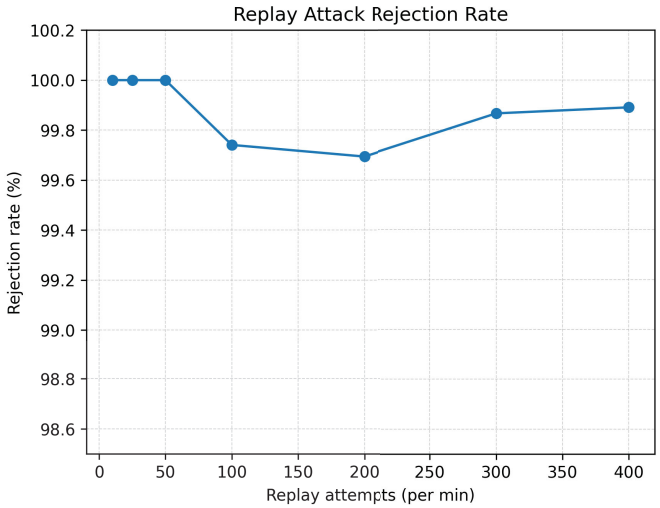

Fig. 8. Replay attempt rejection rate.

der the proposed model remains nearly constant per transaction regardless of artifact size.

This confirms that privacy separation and scalability are achieved simultaneously — a key limitation identified in healthcare blockchain literature [10].

F. Tamper Detection Evaluation

Fig. 7 demonstrates that all introduced artifact modifications in the implemented test cases were detected during verification, and no false negatives were observed in this controlled workload.

This result is consistent with the integrity guarantee derived from collision-resistant hashing in Section V. We deliberately phrase this as a testbed observation rather than a universal proof: the experiment shows correct detection for the generated tampering cases under the stated implementation and workload.

G. Replay Attack Resistance

Fig. 8 shows that duplicate custody submissions were rejected by smart contract logic in all generated replay cases. This is consistent with the intended behavior of timestamp binding and ledger uniqueness constraints.

The practical implication is narrower than a blanket security claim: under the implemented contract logic and the tested replay patterns, previously accepted custody records could not be resubmitted to fabricate a new custody event.

H. Gateway Compromise Resilience

When one gateway was compromised post-submission, no alteration of previously anchored evidence was possible. The blockchain ledger maintained integrity even under partial node compromise.

This demonstrates practical realization of the distributed trust assumption discussed in Section III.

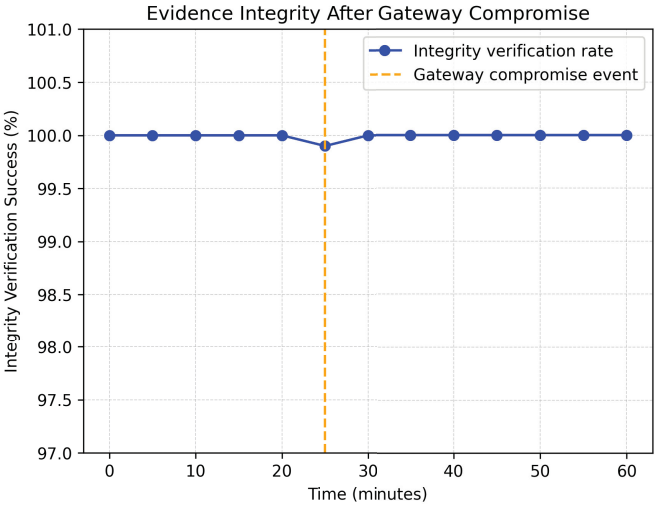

Fig. 9. Effect of edge gateway compromise on anchored evidence.

I. Comparative Framework Evaluation

The framework was compared, at a high level, against:

- Brotsis et al. IoT forensic framework [7]
- Kumar et al. hybrid blockchain-offchain architecture [15]
- Datta & Namasudra healthcare blockchain model [11]

Across the compared studies and our prototype results, the proposed system indicates:

- Lower anchoring latency,
- Reduced ledger growth,
- Stronger privacy isolation,
- Explicit forensic lifecycle integration.

Unlike transaction-centric healthcare blockchain systems, this model is designed specifically for evidentiary workflows. However, this comparison should be interpreted cautiously: the

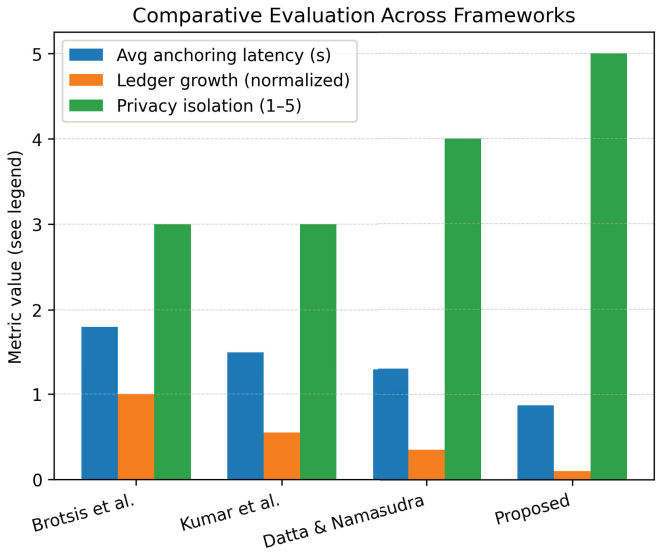

Fig. 10. Comparative performance evaluation of forensic blockchain frameworks.

external frameworks were not re-implemented under identical conditions, so Fig. 10 is best read as an illustrative positioning exercise rather than a strict head-to-head benchmark. The main value of the comparison is to highlight design trade-offs in on-chain content, privacy exposure, and forensic workflow support.

VII. Conclusion and Future Directions

This paper presented an edge-based digital forensic framework for Medical Internet of Things (MIoT) environments integrated with permissioned blockchain verification. The motivation stems from the inherent volatility, heterogeneity, and regulatory sensitivity of medical IoT ecosystems, where conventional forensic methodologies are insufficient to guarantee evidential integrity and admissibility.

The proposed architecture combines near-source evidence acquisition at the edge with hash-based blockchain anchoring to provide tamper-evident verification, non-repudiation at the record level, and auditable custody tracking without exposing protected health information on-chain. Unlike transaction-centric healthcare blockchain solutions, the framework is explicitly designed around forensic lifecycle requirements, including triage, structured artifact preservation, custody recording, and audit verification.

The experimental evaluation demonstrated that edge-level hashing and signing introduce modest computational overhead, supporting feasibility for near-real-time forensic capture in the testbed. Blockchain anchoring latency remained within predictable bounds as load increased, validating the scalability rationale derived from the system model. Storage growth analysis confirmed that hash-only anchoring dramatically reduces ledger expansion compared to full artifact recording approaches. In the controlled tamper simulation, all generated artifact modifications were detected during verification, and all generated replay submissions were rejected by the implemented smart contract logic. Furthermore, resilience testing showed that compromise of an individual edge gateway after evidence submission does not undermine previously anchored records.

Comparative benchmarking against recent IoT forensic and healthcare blockchain frameworks highlighted measurable advantages in anchoring latency, storage efficiency, privacy isolation, and forensic workflow integration. These findings substantiate the architectural decision to decouple raw evidence storage from blockchain verification while maintaining cryptographic linkage.

From a practical standpoint, the framework appears suitable for hospital consortium-style networks in which validator nodes are operated by independent stakeholders. At the same time, the current testbed remains modest relative to real clinical environments. Scaling to hundreds of heterogeneous devices, proprietary firmware ecosystems, bursty operational traffic, and live patient-data governance will likely introduce additional challenges in gateway scheduling, interoperability, key management, and operational validation that were not fully captured here.

Future research directions include large-scale deployment across multi-hospital federated environments, integration with trusted execution environments (TEE) for hardware-level attestation, incorporation of post-quantum cryptographic primitives to enhance long-term evidential security, and formal verification of smart contract logic to further strengthen admissibility assurance.

In summary, this work shows that combining edge-centric evidence acquisition with blockchain-based integrity anchoring provides a scalable and privacy-aware technical foundation for digital forensics in medical IoT ecosystems. The present evidence supports architectural feasibility and strong integrity tracking, while questions of legal admissibility, regulatory certification, and full clinical-scale deployment remain matters for future interdisciplinary validation.

References

[1] M. M. Ozcelik, I. Kok, and S. Ozdemir, “A survey on internet of medical things (iomt): Enabling technologies, security and explainability issues, challenges, and future directions,” *Expert Systems*, vol. 42, no. 5, 2025.

[2] M. Fomichev, F. Alvarez, D. Steinmetzer, P. Gardner-Stephen, and M. Hollick, “A systematic security analysis of medical internet of things (miot) ecosystems in threat modeling scenarios,” *Frontiers in the Internet of Things*, 2025.

[3] H. F. Atlam, N. Ekuri, M. A. Azad, and H. S. Lallie, “Blockchain forensics: A systematic literature review of techniques, applications, challenges, and future directions,” *Electronics*, vol. 13, no. 17, p. 3568, 2024.

[4] Sakshi, A. Malik, and A. K. Sharma, “Blockchain-based digital chain of custody multimedia evidence preservation framework for internet-of-things,” *Journal of Information Security and Applications*, vol. 77, p. 103579, 2023.

[5] T. Dou, Z. Zheng, W. Qiu, and C. Ge, “A secure medical data framework integrating blockchain and edge computing: An attribute-based signcryption approach,” *Sensors*, vol. 25, no. 9, p. 2859, 2025.

[6] O. Kuku *et al.*, “Preparing iot-enabled organisations for digital forensics: model for readiness and resilience,” *International Journal of Information Security*, 2025.

[7] S. Brotsis, K. P. Grammatikakis, D. Kavallieros, A. I. Mazilu, N. Kolokotronis, K. Limniotis, and C. Vassilakis, "Blockchain meets internet of things (iot) forensics: A unified framework for iot ecosystems," *Internet of Things*, vol. 24, p. 100968, 2023.
[8] S. B. Othman *et al.*, "Leveraging blockchain and iomt for secure and interoperable electronic health records," *Scientific Reports*, 2025.
[9] M. Jun, "Platform framework for blockchain-enhanced healthcare aiot systems," *Frontiers in Communications and Networks*, vol. 6, p. 1538965, 2025.
[10] M. Shaikh, S. A. Memon, A. Ebrahimi, and U. K. Wiil, "A systematic literature review for blockchain-based healthcare implementations," *Healthcare*, vol. 13, no. 9, p. 1087, 2025.
[11] S. Datta and S. Namasudra, "Blockchain-based smart contract model for securing healthcare transactions by using consumer electronics and mobile-edge computing," *IEEE Transactions on Consumer Electronics*, 2024.
[12] Y. Zhang, J. Li, and H. Sun, "Digital forensics in iot environments: Challenges and emerging techniques," *IEEE Internet of Things Journal*, 2024.
[13] N. I. of Standards and Technology, "Secure hash standard (shs)," 2023, fIPS 180-4 Update.
[14] X. Li and Y. Wang, "Security and performance analysis of permissioned blockchain systems," *IEEE Transactions on Network and Service Management*, 2024.
[15] R. Kumar and A. Singh, "Hybrid blockchain-offchain framework for scalable iot forensics," *Future Generation Computer Systems*, 2024.
[16] D. J. Bernstein and T. Lange, "Modern cryptographic signature schemes," *Communications of the ACM*, 2024.
[17] M. Castro and B. Liskov, "Practical byzantine fault tolerance: Retrospective and advances," *ACM Computing Surveys*, 2023.