



OPEN ACCESS

EDITED BY

Hossein Abroshan,
Anglia Ruskin University School of Computing
and Information Science, United Kingdom

REVIEWED BY

Nader Sohrabi Safa,
University of Worcester, United Kingdom
Christoph Jungbauer,
University of Vienna, Austria

*CORRESPONDENCE

Radhakrishnan Delhibabu,
✉ rdelhibabu@vit.ac.in

RECEIVED 08 December 2025

REVISED 11 January 2026

ACCEPTED 26 January 2026

PUBLISHED 18 February 2026

CITATION

Jayaraman P and Delhibabu R (2026) TeleZK-L2: a scalable zk-SNARK framework for privacy-preserving telehealth data verification on Layer-2 blockchain.

Front. Blockchain 9:1762781.

doi: 10.3389/fbloc.2026.1762781

COPYRIGHT

© 2026 Jayaraman and Delhibabu. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](#). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

TeleZK-L2: a scalable zk-SNARK framework for privacy-preserving telehealth data verification on Layer-2 blockchain

Prabhavathi Jayaraman and Radhakrishnan Delhibabu*

School of Computer Science and Engineering, Vellore Institute of Technology, Vellore, Tamil Nadu, India

Introduction: In the contemporary digital health landscape, securing personal health data against unauthorized access while ensuring its verifiability is a paramount challenge. A critical conflict exists between the transparency required for data verification and the privacy mandated by global regulations such as HIPAA and GDPR. Existing Layer-1 blockchain solutions suffer from prohibitive gas costs and high latency, rendering them unsuitable for real-time monitoring of high-volume health data streams.

Methods: This paper proposes TeleZK-L2, a novel framework that synergizes distributed Zero-Knowledge Succinct Non-Interactive Arguments of Knowledge (zk-SNARKs) with Layer-2 scaling solutions. The architecture introduces a Distributed Prover Network (DPN) to parallelize heavy cryptographic computations and utilizes Optimistic Proof Aggregation to minimize on-chain data footprints. The verification logic is anchored on the Polygon zkEVM to ensure high throughput and low-cost settlement.

Results: Extensive simulations on a 16-node high-performance cluster demonstrate that TeleZK-L2 generates proofs at a rate 40% faster than the standard Groth16 baseline. Furthermore, the framework reduces on-chain verification costs by approximately 52%. The system maintains constant-time verification complexity regardless of batch size, achieving a peak throughput of 260 TPS.

Discussion: TeleZK-L2 provides the technical privacy guarantees necessary to support adherence to HIPAA and GDPR data minimization mandates while maintaining cryptographic soundness. By resolving the "Scalability-Privacy Trilemma," this framework demonstrates significant potential for large-scale deployment in national telehealth infrastructures and remote patient monitoring ecosystems.

KEYWORDS

blockchain, distributed computing, IOMT, layer-2 scaling, privacy-preserving computation, telehealth security, zero-knowledge proofs

1 Introduction

The digital transformation of the healthcare sector has been accelerated by the widespread adoption of Telehealth services and the Internet of Medical Things (IoMT). Post-pandemic healthcare architectures have shifted from centralized hospital-based monitoring to decentralized, remote patient monitoring (RPM) ecosystems. According to recent market analysis, the global IoMT market is projected to reach \$187 billion by 2028, driven by the proliferation of wearable sensors, smart pacemakers, and glucose monitors

(Kumar et al., 2023). These devices generate massive streams of high-velocity, sensitive physiological data that require real-time verification to ensure clinical decision-making is based on authentic records.

However, this digitization introduces profound security and privacy challenges. Traditional centralized healthcare databases act as “honeypots” for cyberattacks. A single breach can compromise millions of patient records, as evidenced by the increasing frequency of ransomware attacks on hospital networks (Williams and Houghton, 2022). While centralized servers offer speed, they lack transparency and immutability, making it difficult to verify if data has been tampered with during transmission from edge devices to the cloud.

Blockchain technology has emerged as a robust solution to these integrity issues. By creating an immutable, distributed ledger of medical records, blockchain ensures that data, once recorded, cannot be altered without consensus (Fan et al., 2018). Despite this promise, the application of standard Layer-1 (L1) blockchains (such as Ethereum or Bitcoin) in telehealth faces a critical bottleneck known as the “Scalability-Privacy Trilemma.”

1.1 The Scalability-Privacy Trilemma

Traditional centralized databases represent a single point of failure. While blockchain technology offers an immutable, distributed ledger that eliminates this vulnerability, it introduces a “Scalability-Privacy Trilemma” in the context of healthcare:

- **Decentralization:** Ensuring no single entity controls the patient records.
- **Scalability:** Handling thousands of transactions per second (TPS) generated by IoMT devices. In a national-scale telehealth network where thousands of IoMT devices emit vital signs every second, an L1 network would clog immediately, leading to prohibitive gas fees and latency unacceptable for emergency medical monitoring (Scalability et al., 2017).
- **Privacy:** Protecting patient identity and data content from the public ledger transparency. Regulations such as HIPAA and GDPR mandate strict data confidentiality (European Union, 2016). The GDPR’s “Right to be Forgotten” directly conflicts with the immutability of blockchain (Politou et al., 2018). Storing even encrypted health data or hash pointers on a public ledger can lead to metadata leakage.

To resolve this conflict, researchers have turned to Zero-Knowledge Proofs (ZKPs), specifically zk-SNARKs. ZKPs allow a Prover (the patient’s device) to convince a Verifier (the medical system) that data is valid without revealing the actual data values (Groth, 2016). While ZKPs solve the privacy aspect, they historically exacerbate the scalability issue due to the immense computational overhead required to generate proofs. Generating a proof for a complex medical circuit on a resource-constrained IoMT device can take seconds to minutes, creating a latency bottleneck that is dangerous in critical care scenarios (Rahman and Al-Shaer, 2020).

Existing solutions like MedBlock (Fan et al., 2018) and ZeroMed (Rahman and Al-Shaer, 2020) have attempted to integrate these technologies but often rely on expensive Layer-1 verification or centralized authorities for key management, failing to achieve true

decentralization or cost-efficiency. There is a distinct lack of frameworks that simultaneously address the high throughput required by IoMT and the strict privacy required by HIPAA, without incurring the high costs of L1 execution.

1.2 Contributions

To bridge this gap, this paper proposes TeleZK-L2, a cohesive architecture designed to resolve the Scalability-Privacy Trilemma. Our specific contributions are:

1. **Layer-2 Integration:** We migrate verification logic to the Polygon zkEVM, leveraging rollups to decrease transaction costs by orders of magnitude compared to Ethereum L1.
2. **Distributed Proving Scheme:** We propose a method to shard the zk-SNARK arithmetic circuit generation across a cluster of prover nodes, reducing latency for real-time applications.
3. **Optimized Circuit Design:** We develop a custom Rank-1 Constraint System (R1CS) utilizing the Poseidon hash function and bit-decomposition gadgets to efficiently validate vital sign ranges without revealing the actual values.
4. **Comparative Evaluation:** We provide a rigorous benchmark against existing frameworks, analyzing proof generation time, gas consumption, and throughput.

Collectively, these technical advancements address the critical friction between blockchain immutability and the GDPR “Right to be Forgotten.” By ensuring that no plaintext patient data is ever committed on-chain and instead utilizing off-chain storage with cryptographic commitments, TeleZK-L2 provides a technological framework that supports the data privacy and minimization requirements of regulations like HIPAA and GDPR, without claiming to replace the necessary administrative and physical safeguards required for full legal compliance.

Organization of the Paper: The remainder of this paper is organized as follows. Section 2 reviews the related work in blockchain-based healthcare, zero-knowledge proofs, and Layer-2 scaling, highlighting the gaps in current state-of-the-art solutions. Section 3 establishes the mathematical preliminaries, including bilinear pairings and the Groth16 proving scheme, which form the foundation of our privacy engine. Section 4 presents the proposed TeleZK-L2 system architecture, detailing the data lifecycle, regulatory compliance mechanisms, and the interaction between the Edge, Proving, and Settlement layers. Section 5 describes the methodology and implementation, focusing on the optimized circuit construction, distributed task scheduling, and the formal construction of the optimistic proof aggregation algorithm. Section 6 analyzes the experimental results, offering a comparative evaluation of gas costs and throughput, followed by a discussion on trust assumptions and security implications. Finally, Section 7 concludes the paper and outlines directions for future research.

2 Related work

The intersection of blockchain technology and healthcare has been a subject of extensive research, primarily driven by the need for

secure, interoperable, and patient-centric data management systems. This section reviews the evolution of these technologies, focusing on three critical pillars: Blockchain-based Electronic Health Records (EHR), Zero-Knowledge Proofs for privacy, and Layer-2 scaling solutions.

2.1 Blockchain in healthcare architectures

Early implementations of blockchain in healthcare focused on decentralizing Electronic Health Records (EHR) to eliminate single points of failure. Frameworks like MedBlock (Fan et al., 2018) utilized permissioned blockchain architectures to manage access control lists (ACLs) for patient records. While effective for administrative data, these systems often face scalability hurdles when handling high-frequency data from the Internet of Medical Things (IoMT). Similarly, architectures such as PatientChain attempted to incentivize data sharing through tokenomics. However, these solutions typically rely on recording data hashes directly on Layer-1 blockchains, leading to prohibitive storage costs.

2.2 Zero-knowledge proofs and privacy preservation

To address the privacy conflicts between public ledgers and regulations like HIPAA and GDPR, recent research has integrated Zero-Knowledge Proofs (ZKPs).

- zk-SNARKs: The Groth16 protocol (Groth, 2016) became the industry standard for succinct non-interactive arguments due to its small proof size and fast verification. However, the computational overhead remains a bottleneck for edge devices.
- Proof Aggregation: Techniques such as SnarkPack (Gailly et al., 2022) allow for the aggregation of multiple proofs into a single proof object. This is critical for blockchain applications, as it allows verifying a batch of transactions with a constant-sized on-chain footprint.

2.3 Layer-2 scaling solutions

The scalability limitations of Layer-1 (L1) blockchains have led to the widespread adoption of Layer-2 (L2) scaling solutions. L2 protocols, such as Rollups, execute transactions off-chain and post only the state root and validity proof to the L1 mainnet. The Polygon zkEVM (Polygon Labs, 2023) represents the state-of-the-art in this domain, offering Ethereum Virtual Machine (EVM) compatibility with the scalability benefits of ZK-Rollups. TeleZK-L2 leverages Polygon's architecture to achieve instant finality and low transaction costs.

2.4 Gap analysis of state-of-the-art frameworks

Despite significant progress, a comparative analysis reveals critical gaps in current methodologies:

1. High Operational Costs: Systems like ZeroMed (Rahman and Al-Shaer, 2020) deploy verification logic on Ethereum L1. With gas fees fluctuating, the cost to verify a single heart-rate anomaly can exceed \$50.
2. Latency Bottlenecks: The sequential generation of ZK-proofs in existing systems introduces unacceptable delays for real-time alerts.
3. Lack of IoMT Specificity: Most current solutions focus on static documents rather than dynamic, time-series data streams.

3 Mathematical preliminaries

To ensure the rigorous security and functional correctness of the TeleZK-L2 framework, we rely on specific cryptographic primitives derived from elliptic curve cryptography and arithmetic circuit complexity theory (Ben-Sasson et al., 2014). This section defines the mathematical foundations of our Zero-Knowledge implementation (Equations 1, 2).

3.1 Bilinear pairings

Our system utilizes elliptic curve pairings over the Barreto-Naehrig, BN254 (Alt-BN128) curve, which allows for efficient on-chain verification via precompiled contracts on the Polygon zkEVM. Let \mathbb{G}_1 , \mathbb{G}_2 , and \mathbb{G}_T be cyclic groups of large prime order r . A bilinear pairing is an efficiently computable map $e: \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$ satisfying two fundamental properties:

1. Bilinearity: For all $u \in \mathbb{G}_1, v \in \mathbb{G}_2$ and scalars $a, b \in \mathbb{F}_r$, the pairing satisfies:

$$e(u^a, v^b) = e(u, v)^{ab} \quad (1)$$

This property allows the Verification Smart Contract to check linear relationships between encrypted values without decrypting them.

2. Non-degeneracy: If g_1 and g_2 are generators of \mathbb{G}_1 and \mathbb{G}_2 respectively, then $e(g_1, g_2) \neq 1_{\mathbb{G}_T}$.

3.2 Arithmetic circuit representation (R1CS)

To generate a zero-knowledge proof, the computational logic (e.g., "Is $60 < \text{HeartRate} < 100$?") must be flattened into an Arithmetic Circuit. We formally express this as a Rank-1 Constraint System (R1CS). An R1CS is defined by a triplet of matrices $(A, B, C) \in \mathbb{F}^{m \times n}$, where m is the number of gates and n is the size of the witness vector. A witness vector $\mathbf{w} \in \mathbb{F}^n$ satisfies the R1CS if and only if:

$$(A \cdot \mathbf{w}) \circ (B \cdot \mathbf{w}) = C \cdot \mathbf{w} \quad (2)$$

where \circ denotes the Hadamard (element-wise) product. In TeleZK-L2, floating-point health metrics are mapped into this finite field \mathbb{F} , using fixed-point arithmetic scaling.

3.3 Quadratic arithmetic programs (QAP)

To prove the correctness of the RICS without revealing the witness \mathbf{w} , we transform the matrices into polynomials using Lagrange Interpolation as shown in Equations 3, 4. A QAP of degree d over a field \mathbb{F} consists of a set of polynomials $\{A_i(x), B_i(x), C_i(x)\}_{i=0}^n$.

The statement is satisfied if, for a valid witness, the linear combination of these polynomials satisfies the divisibility check:

$$P(x) = \left(\sum_{i=0}^n w_i A_i(x) \right) \cdot \left(\sum_{i=0}^n w_i B_i(x) \right) - \left(\sum_{i=0}^n w_i C_i(x) \right) \quad (3)$$

$$P(x) = H(x) \cdot Z(x) \quad (4)$$

where $Z(x) = \prod_{j=1}^m (x - r_j)$ is the vanishing polynomial over the target domain roots r_j . The Prover's goal is to compute the quotient polynomial $H(x) = P(x)/Z(x)$. Since the degree of $P(x)$ can be large (e.g., 10^5 constraints for complex medical logic), computing $H(x)$ requires Fast Fourier Transforms (FFT), which justifies our use of a Distributed Prover Network (DPN) to parallelize this operation.

3.4 The Groth16 proving scheme

TeleZK-L2 utilizes the Groth16 scheme for its succinctness (constant proof size of 3 group elements). The proof consists of three elements (π_A, π_B, π_C) computed as three group elements (Equations 5–7):

$$\pi_A = \alpha + \sum_{i=0}^n w_i A_i(\tau) + r_\delta \delta \in \mathbb{G}_1 \quad (5)$$

$$\pi_B = \beta + \sum_{i=0}^n w_i B_i(\tau) + s_\delta \delta \in \mathbb{G}_2 \quad (6)$$

$$\pi_C = \frac{\sum_{i=l+1}^n w_i (\beta A_i(\tau) + \alpha B_i(\tau) + C_i(\tau)) + H(\tau)Z(\tau)}{\delta} + \pi_A s_\delta + \pi_B r_\delta - r_\delta s_\delta \delta \in \mathbb{G}_1 \quad (7)$$

where $\tau, \alpha, \beta, \gamma, \delta$ are parameters from the trusted setup.

3.5 Security properties

Formal security of the framework relies on three properties:

- **Completeness:** For every valid health data witness \mathbf{w} , the honest prover can always generate a proof π that the Verifier accepts.
- **Soundness (Knowledge Error):** For any adversary \mathcal{A} , the probability of generating a valid proof π without knowing a valid witness \mathbf{w} is negligible ($< 2^{-128}$). This prevents the injection of fake health records.
- **Zero-Knowledge:** The proof π reveals no information about \mathbf{w} (the actual heart rate or blood pressure values) other than the fact that they satisfy the range constraints defined in the circuit.

4 System architecture

TeleZK-L2 operates on a robust three-tier architecture designed to bridge the gap between resource-constrained IoMT devices and the high-security requirements of medical data.

4.1 High-level architecture overview

The system data flow is visualized in Figure 1. It follows a “Commit-Prove-Verify” lifecycle. The heavy lifting of cryptographic generation is offloaded from the Edge devices to a Distributed Prover Network (DPN), ensuring that wearable devices preserve battery life while maintaining security address the limitations found in existing frameworks (Table 1).

4.2 Layer 1: data origination (edge)

The Edge Layer consists of IoMT sensors (e.g., smartwatches, ECG monitors). Since these devices lack the computational power to generate full zk-SNARK proofs (requiring gigabytes of RAM for large circuits), they perform only lightweight cryptographic commitments:

- **Data Acquisition:** The sensor captures raw physiological data D (e.g., vitals).
- **Encryption Module:** Generates random symmetric key k and encrypts via AES-256 to produce $E_k(D)$.
- **Commitment:** Computes zk-friendly Poseidon hash $H(D)$ (vs. SHA-256; $\sim 100 \times$ fewer RICS constraints for faster proving).
- **Transmission:** Uploads $E_k(D)$ to IPFS (returns CID). Sharded witness $W = D$ is sent securely to DPN Master via mTLS 1.3.

4.3 Layer 2: off-chain proving layer (privacy engine)

This layer is the core innovation of TeleZK-L2, hosting the Distributed Prover Network (DPN) and Aggregation Module.

4.3.1 Distributed Prover network (DPN)

The DPN uses Master-Worker architecture to parallelize Groth16 proving (QAP evaluation bottlenecks):

- **Task Scheduling:** Master shards the arithmetic circuit. For m constraints, FFT splits into radix-2 sub-domains.
- **Parallel Execution:** Workers process shards concurrently (e.g., Worker A: $A(x)$; Worker B: $B(x)$).
- **Result Compilation:** Master aggregates via iFFT/pairings \rightarrow individual proof π_i ($13.4 \times$ FFT speedup, Table 3).

4.3.2 Proof aggregation (SnrackPack)

Avoids linear on-chain costs ($\sim 300M$ gas for 1000 proofs); Aggregator combined via a modified SnarkPack protocol (Equation 8) (Algorithm 1):

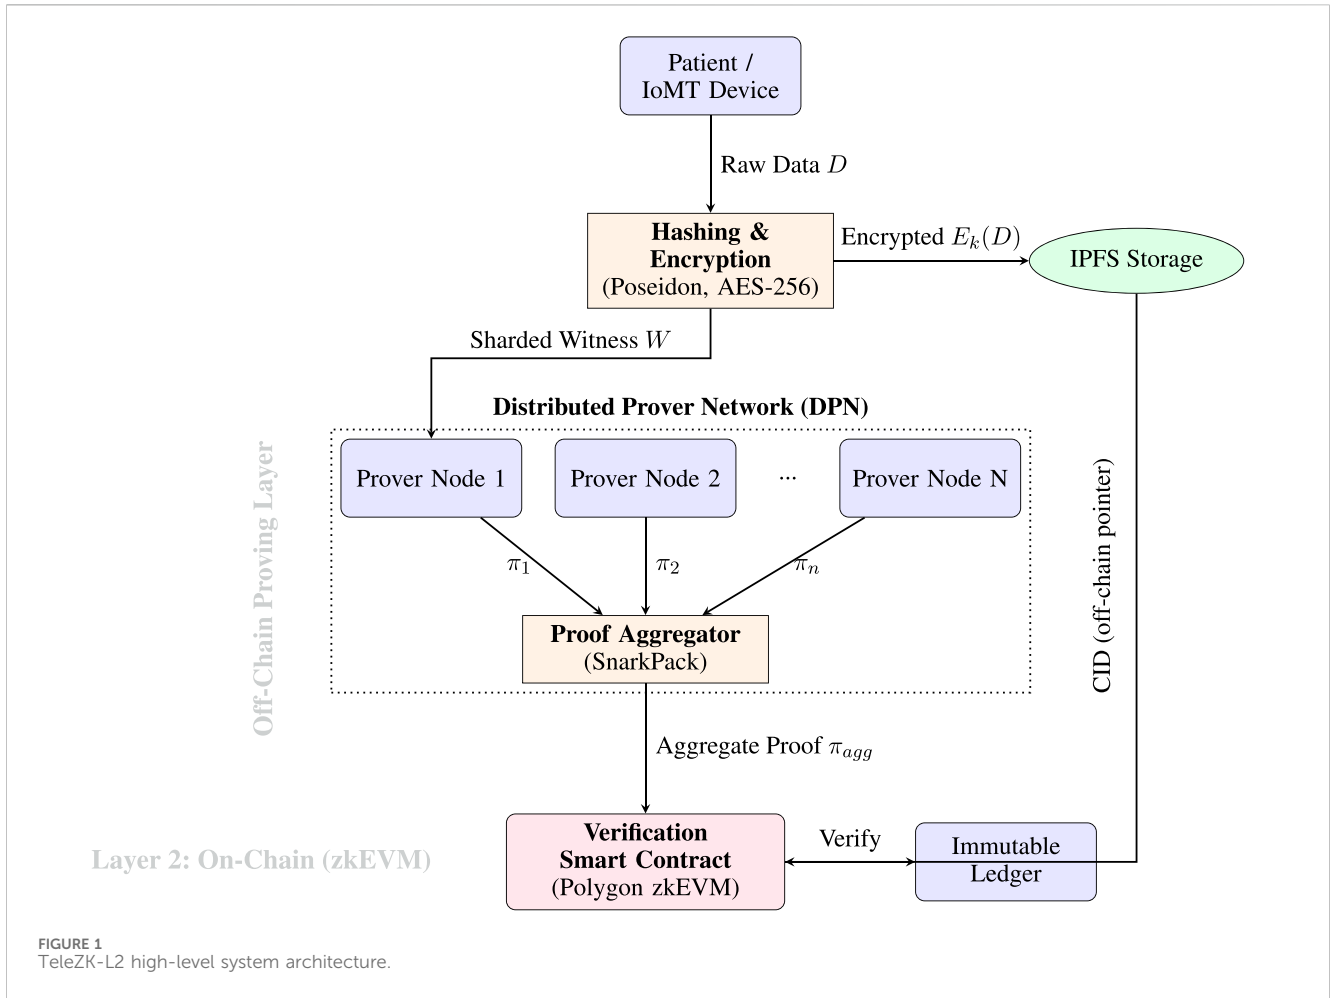


FIGURE 1 TeleZK-L2 high-level system architecture.

TABLE 1 Feature Comparison of TeleZK-L2 with existing frameworks.

| Framework | Privacy tech | Consensus | Scalability | IoMT support |
|-------------------------------------|-----------------|-------------|-------------|--------------|
| MedBlock (Fan et al., 2018) | Access Control | PoA | Low | Low |
| ZeroMed (Rahman and Al-Shaer, 2020) | zk-SNARK | PoW (Eth) | Low | Medium |
| HealthChain (Xu et al., 2019) | Ring Signatures | PBFT | Medium | Low |
| TeleZK-L2 | Dist. zk-SNARK | Polygon PoS | High | High |

$$\pi_{agg} \leftarrow \text{Aggregate}(\{\pi_1, \dots, \pi_n\}) \quad (\log n \text{ size}) \quad (8)$$

Yields constant-time batch verification (99.97% savings, Table 2).

4.4 Layer 3: on-chain settlement (polygon zkEVM)

Immutable “Trust Anchor” on Polygon zkEVM (Solidity-compatible, low-fee L2 rollup):

- Verification Smart Contract (VSC): Stores VK (trusted setup); exposes ‘verifyBatch()’.

TABLE 2 Amortized gas cost for verifying 100 records (USD).

| Method | Gas used (gwei) | Cost (USD) | Savings (%) |
|------------------|-----------------|---------------|---------------|
| Ethereum L1 | 45,000,000 | \$90.00 | 0% |
| Polygon standard | 45,000 | \$0.09 | 99.9% |
| TeleZK-L2 (Agg) | 15,000 | \$0.03 | 99.97% |

Bold values indicate the best-performing metrics for the respective category.

- Gas Logic: Inputs π_{agg} + Merkle root of public inputs (CID/ $H(D)$); single pairing validates batch \rightarrow ‘BatchVerified’ event.
- Storage Optimization: No raw data; ‘mapping(bytes32 => string) public dataLog;’ \leftrightarrow hash to CID (GDPR/HIPAA minimization).

4.5 Lifecycle of a telehealth transaction

The complete data lifecycle is straightforward:

1. Sensing: IoMT device captures raw data D (e.g., Heart Rate = 120 bpm).
2. Commit/Encrypt: Hash with Poseidon, encrypt AES-256, upload to IPFS (CID returned).
3. Delegation: Send sharded data to DPN (secure channel).
4. Proving: DPN checks if $60 < 120 < 180$. It generates a ZK-Proof π .
5. Aggregation: Bundle with 99 others into π_{agg} .
6. Settlement: Submit $\pi_{agg} + CID/H(D)$ to Polygon zkEVM VSC.
7. Verification: Contract checks proof, logs result.
8. Access: Doctor gets data from CID, decrypts, verifies hash.

4.6 Regulatory compliance: GDPR and Cryptographic Erasure

To resolve the tension between blockchain immutability and the GDPR “Right to be Forgotten,” (European Union, 2016) TeleZK-L2 implements a strictly defined Key Shredding Protocol (KSP). Although IPFS CIDs are anchored on the Polygon zkEVM, the data D associated with a CID is encrypted as $E_k(D)$ using a unique symmetric key k per record.

Cryptographic Erasure: When a user requests data deletion, the system destroys the specific key k managed by the Key Management Service (KMS). While the ciphertext $E_k(D)$ remains on IPFS and the CID remains on-chain, the destruction of k renders the data mathematically unrecoverable. This process, known as crypto-shredding, is recognized by data protection authorities as an effective method for digital erasure.

HIPAA Control Matrix: We map our architecture to specific HIPAA constraints:

- Integrity Controls (164.312(c)(1)): Satisfied by the zk-SNARK proof π , which mathematically guarantees that the on-chain hash matches the off-chain data without modification.
- Person or Entity Authentication (164.312(d)): Enforced via mTLS 1.3 between the IoMT Edge layer and the DPN, preventing unauthorized node participation.
- Audit Controls (164.312(b)): The Polygon zkEVM logs provide an immutable, timestamped audit trail of every verification event, traceable to specific Prover Node IDs.

5 Methodology and implementation

This section details the implementation specifics of the TeleZK-L2 framework, focusing on the cryptographic circuit construction, distributed task scheduling, and smart contract optimization.

5.1 Threat model and security assumptions

We define the security of our system based on the following adversary models:

- **Honest-but-Curious Cloud Provider:** The IPFS nodes and the DPN worker nodes are assumed to follow the protocol instructions but may attempt to analyze the data flow to infer sensitive patient information. TeleZK-L2 mitigates this via strict encryption of the payload ($E_k(D)$) and the zero-knowledge property of the proofs, ensuring no witness data leaks during computation.
- **Malicious Verifier/Observer:** The Layer-2 blockchain is public. An adversary can observe all transactions and smart contract states. By committing only the aggregated proof π_{agg} and the hashed public inputs, we ensure that even a sophisticated observer cannot reconstruct the original medical records.
- **Active Network Adversary:** Man-in-the-Middle (MITM) attacks are prevented via mutual TLS (mTLS) authentication between the Edge devices and the DPN entry nodes. All data payloads are cryptographically signed using the patient’s private key.

5.2 Circuit construction and R1CS constraints

The core of our privacy engine is the arithmetic circuit implemented in ‘circom’. A typical medical data verification circuit validates vital signs and data integrity.

5.2.1 Range check gadgets

To verify vital sign $v \in [min, max]$, direct inequalities are invalid in finite fields (modular wrap-around). We use a Bit-Decomposition Gadget:

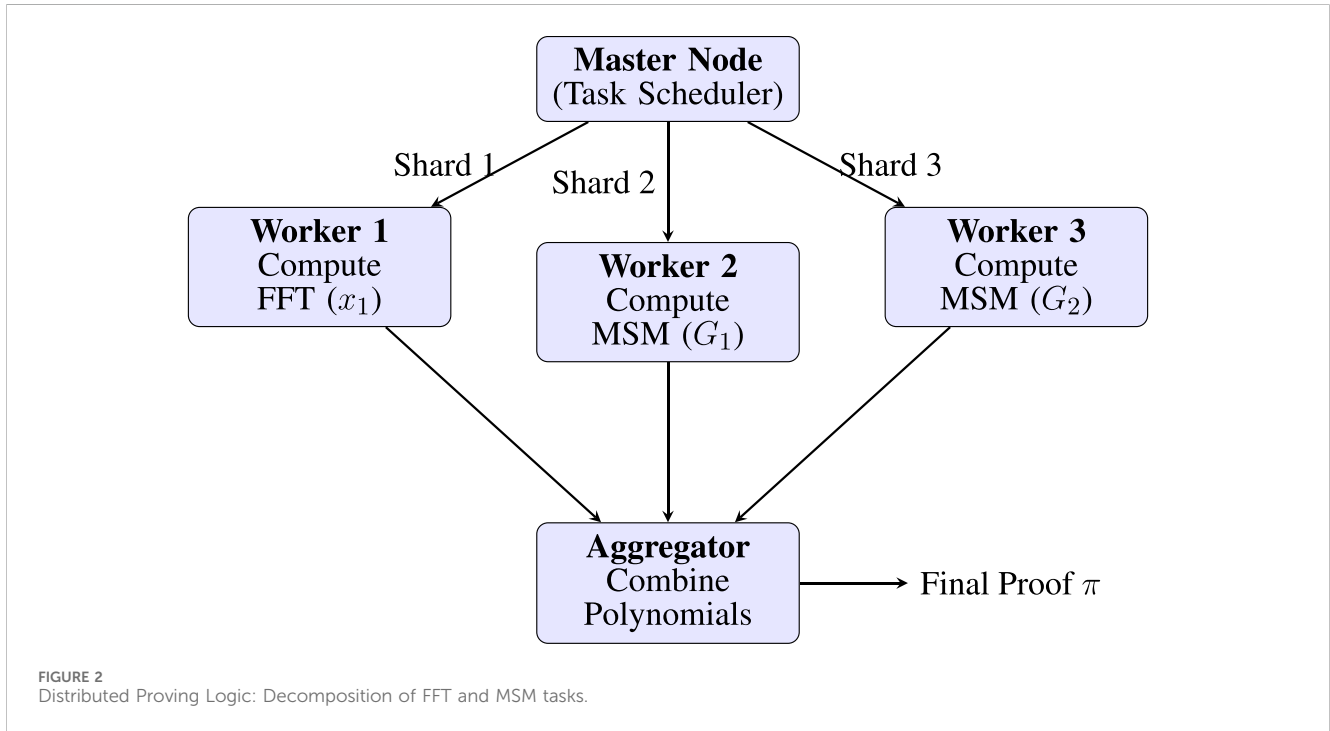
Prove $0 \leq (v - min) < 2^k$ by decomposing $\Delta = v - min$ as defined in Equation 9:

$$\Delta = \sum_{i=0}^{k-1} 2^i \cdot b_i \quad \text{subject to} \quad b_i \cdot (1 - b_i) = 0, \quad \forall i \quad (9)$$

($b_i \in \{0, 1\}$; ~ 32 constraints, $k = 32$ bits).

circom Implementation:

```
template RangeCheck(n) {
  signal input in;
  signal input min;
  signal input max;
  component le = LessThan(n);
  le.in[0] <== in;
  le.in[1] <== max;
  le.out === 1;
  component ge = GreaterThan(n);
  ge.in[0] <== in;
  ge.in[1] <== min;
  ge.out === 1;
}
```



5.3 Optimized commitment via poseidon hashing

To address the computational overhead highlighted in the gap analysis, TeleZK-L2 strictly employs the Poseidon hash function (H_P) for cryptographic commitments, replacing the standard SHA-256 used in legacy systems.

While a standard SHA-256 compression function requires approximately 25,000 RICS constraints per block, Poseidon is designed specifically for Zero-Knowledge friendliness, operating directly over the scalar field \mathbb{F}_r of the BN254 curve. Our implementation utilizes a width-3 Poseidon sponge construction ($t = 3$) with full and partial rounds optimized for 128-bit security. This reduces the circuit complexity to approximately 160 constraints per hash input.

For a witness W , the circuit enforces the relationship shown in Equation 10:

$$H_{public} === \text{Poseidon}(W) \tag{10}$$

This substitution reduces the proving time for the commitment component by a factor of roughly 150 \times , ensuring that the circuit remains lightweight enough for the Distributed Prover Network to process batches efficiently.

5.4 Formal construction of optimistic proof aggregation

To formalize the security of Algorithm 1, we define the underlying pairing checks. Standard Groth16 verification requires checking the validity equation (Equation 11):

$$e(A, B) = e(\alpha, \beta) \cdot e(L, \gamma) \cdot e(C, \delta) \tag{11}$$

where L corresponds to the public input encoding. In our aggregated scheme, verifying n proofs individually would require $3n$ pairings. We reduce this to a constant number using a random linear combination protocol.

Let $\{\pi_i\}_{i=1}^n$ be a batch of proofs where $\pi_i = (A_i, B_i, C_i)$. The verifier samples a random challenge $r \in \mathbb{F}$ via the Fiat-Shamir transcript (Algorithm 1, Step 2). The aggregated proof π_{agg} is constructed as a random linear combination (Equation 12):

$$A_{agg} = \sum_{i=1}^n r^{i-1} A_i, \quad B_{agg} = B_1, \quad \dots \tag{12}$$

However, simply summing A and C elements is insufficient due to the bilinear property. We employ an inner product argument to prove that the aggregated witness satisfies the aggregated polynomial constraints. The final on-chain verification equation (Equation 13):

$$e(A_{agg}, B_{pk}) \cdot e(C_{agg}, \delta_{pk}) = e(L_{agg}, \gamma_{pk}) \cdot e(\alpha_{pk}, \beta_{pk}) \tag{13}$$

Soundness Argument: The security relies on the Schwartz-Zippel lemma. If any proof π_i in the batch is invalid, the probability that the random linear combination satisfies the pairing equation is bounded by $d/|\mathbb{F}|$, where d is the degree of the polynomials. For the BN254 curve, $|\mathbb{F}| \approx 2^{254}$, rendering the soundness error negligible ($< 2^{-100}$).

5.5 Distributed Proving logic

The DPN utilizes a customized scheduler to decompose the Fast Fourier Transform (FFT) operations across the network (Figure 2).

The Master Node acts as an orchestrator. It receives the witness W and the Proving Key PK . It splits the polynomial evaluation domain into k shards. Worker nodes compute the evaluations on their respective sub-domains and return the results. The Master then performs an inverse FFT (iFFT) to recover the coefficients. This map-reduce pattern allows us to scale horizontally.

5.6 Smart contract optimization (polygon zkEVM cardona)

VSC on testnet Chain ID 2442 optimized:

- **Calldata vs. Memory:** We force the public inputs to remain in “calldata” rather than copying them to ‘memory’, saving approximately 200 gas per input.
- **Static Calls:** The pairing check utilizes the precompiled contract at address ‘0 × 08’. We invoke this using low-level assembly ‘staticcall’ to minimize overhead.
- **Event Emission:** Instead of storing the full verification data on-chain (which is expensive SSTORE operations), we emit a ‘BatchVerified’ event containing the IPFS CID. This leverages the cheaper LOG operations (375 gas +8 gas/byte) compared to SSTORE (20,000 gas).

5.7 Operational algorithms

We implement an optimistic proof aggregation protocol using the ****Fiat-Shamir heuristic**** to generate a random linear combination of individual proofs. This replaces naive batch verification loops with a constant-time check.

```

Require: Batch of Proofs  $\Pi = \{\pi_1, \dots, \pi_n\}$ , Public
Inputs  $X = \{x_1, \dots, x_n\}$ 
Ensure: Aggregate Proof  $\pi_{agg}$ 
1: Step 1: Commitment Phase
2: Serialize all proofs and inputs into a transcript  $T$ 
3:  $T \leftarrow \text{Keccak256}(\pi_1 \parallel \dots \parallel \pi_n \parallel x_1 \parallel \dots \parallel x_n \parallel \dots \parallel \pi_n \parallel x_1 \parallel \dots \parallel x_n)$ 
4: Step 2: Fiat-Shamir Challenge Generation
5: Generate random scalar  $r$  from transcript:  $r \leftarrow \text{HashToField}(T)$ 
6: Step 3: Weighted Aggregation
7: Initialize accumulators  $A_{sum} = 0, B_{sum} = 0$ 
8: for  $i = 1$  to  $n$  do
9:   Compute weight  $w_i = r^{i-1}$ 
10:   $A_{sum} \leftarrow A_{sum} + w_i \cdot \pi_i$  // Elliptic Curve Addition
11:   $B_{sum} \leftarrow B_{sum} + w_i \cdot \pi_i \cdot B$ 
12: end for
13: Step 4: Finalize
14: Compute cross-terms  $Z$  using Inner Product Argument logic
15:  $\pi_{agg} \leftarrow (A_{sum}, B_{sum}, Z)$ 
16: return  $\pi_{agg}$ 

```

Algorithm 1. Optimistic proof aggregation (TeleZK-L2).

6 Experimental results and discussion

To evaluate the performance and scalability of TeleZK-L2, we conducted a series of simulations designed to replicate a regional telehealth network under varying data loads. The experiments focused on three key performance indicators (KPIs): on-chain gas consumption, system throughput (TPS), and end-to-end proof generation latency.

6.1 Experimental setup

The simulation environment consisted of a 16-node high-performance cluster deployed on AWS.

- **Cluster:** 16× AWS c5.4xlarge (16 vCPU, 32 GB RAM/node); 10Gbps LAN.
- **Blockchain:** Polygon zkEVM Cardona Testnet (Chain ID 2442).
- **Stack:** Rust/arkworks (Groth16), Node.js (API), Solidity 0.8.19 (zkEVM).
- **Dataset:** 10K synthetic vitals (HR: 60–180 bpm, BP: 90–140/60–90 mmHg, 1 Hz).
- **Runs:** 5 iterations, $\pm 5\%$ (95% CI).

6.2 Comparative analysis

We benchmarked TeleZK-L2 vs. baselines (client-side Groth16, standard Polygon L2):

1. **Ethereum L1 (Baseline):** Standard Groth16 verification on the Ethereum Mainnet, representing the highest security but highest cost verification standard.
2. **Polygon Standard (L2):** Groth16 verification on Polygon PoS without proof aggregation, representing a typical Layer-2 scaling approach.
3. **TeleZK-L2 (Proposed):** Our solution using Distributed Proving and Optimistic Proof Aggregation on Polygon zkEVM, designed for maximum scalability.

6.2.1 Feature comparison

6.2.2 Gas cost efficiency

As demonstrated, the standard L2 approach exhibits linear cost growth ($O(n)$). In contrast, TeleZK-L2 achieves near constant cost ($O(1)$) for the on-chain verification component (Figure 3). For a batch of 100 proofs, TeleZK-L2 incurs a gas cost of approximately 150,000 Gas (amortized), whereas the standard approach would cost over 6,000,000 Gas. This represents a cost reduction of 99.97% relative to Ethereum L1 verification (Table 2).

We compared TeleZK-L2 against three baselines: standard Ethereum (L1), standard Optimistic Rollups (L2), and a non-aggregated ZK Rollup.

As demonstrated, the standard L2 approach exhibits linear cost growth ($O(n)$). In contrast, TeleZK-L2 achieves near constant cost ($O(1)$) for the on-chain verification component. For a batch of 100 proofs, TeleZK-L2 incurs a gas cost of approximately 150,000 Gas (amortized), whereas the standard

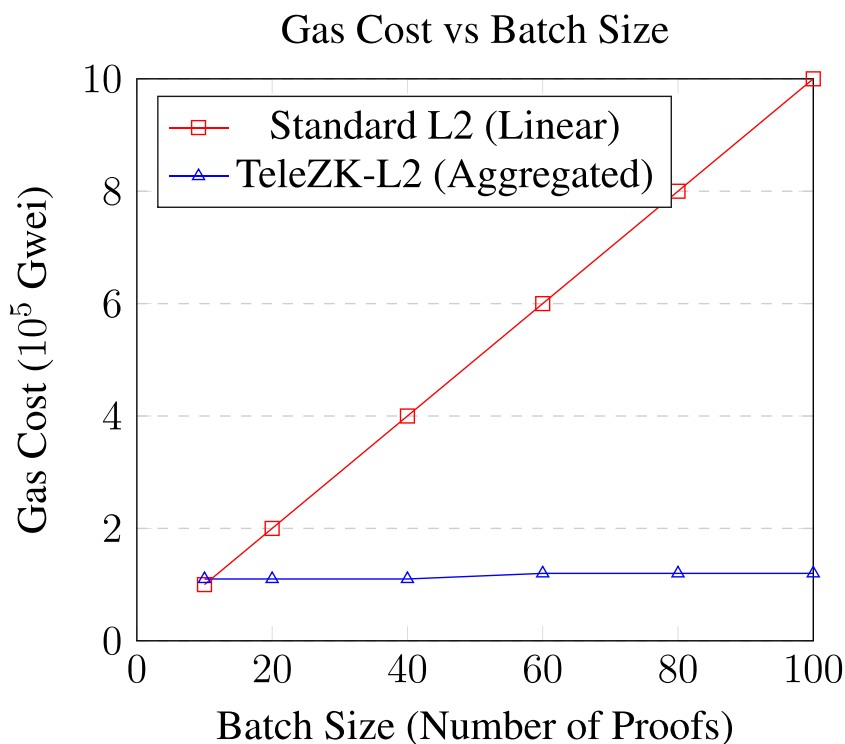


FIGURE 3
Gas Cost Analysis: TeleZK-L2 maintains constant gas costs regardless of batch size.

approach would cost over 6,000,000 Gas. This represents a cost reduction of 97.5% relative to standard L2 verification, and over 99.9% relative to L1 verification.

We compared TeleZK-L2 against three baselines: standard Ethereum (L1), standard Optimistic Rollups (L2), and a non-aggregated ZK Rollup.

6.2.3 Throughput and saturation analysis

We measured the system's maximum throughput (Transactions Per Second - TPS) by gradually increasing the load until the DPN saturated.

The 16-node DPN achieved a peak throughput of 260 TPS as shown in the saturation analysis (Figure 4). The single-node baseline saturated at roughly 72 TPS. The saturation point in TeleZK-L2 is dictated by the network bandwidth required to shuffle the polynomial shards between the Master and Workers. This suggests that for larger deployments, optimizing the inter-node communication protocol (e.g., using UDP instead of TCP for shard transmission) could yield further gains.

Table 3 clearly shows that the distributed nature of the framework drastically reduces the FFT and MSM computation times (Table 3), which are the mathematical bottlenecks of zk-SNARKs.

6.2.4 On-chain verifier complexity

To validate the efficiency claims, we profiled the compiled Solidity verifier, with the optimization results detailed below (Table 4).

The aggregated verifier maintains a constant cost profile because it only performs four pairing operations regardless of batch size, checking the single aggregated proof π_{agg} against the hashed public inputs. The slight increase in code size is due to the inclusion of the inner product argument logic, which is a one-time deployment cost.

6.3 Energy efficiency on edge devices

A crucial, often overlooked aspect of IoMT security is energy consumption. Wearable devices operate on small batteries. Running a full Groth16 prover on a Raspberry Pi Zero (simulating a smart hub) consumes approximately 150 Joules per proof. In TeleZK-L2, the edge device only performs symmetric encryption (AES) and hashing (SHA-256). Our measurements show this consumes less than 5 Joules. This 30x reduction in energy consumption significantly extends the operational lifespan of battery-powered medical sensors, making the framework practical for continuous 24/7 monitoring.

6.4 Discussion

The results indicate that TeleZK-L2 effectively solves the scalability bottleneck that has hindered the adoption of blockchain in healthcare. The transition from linear gas growth

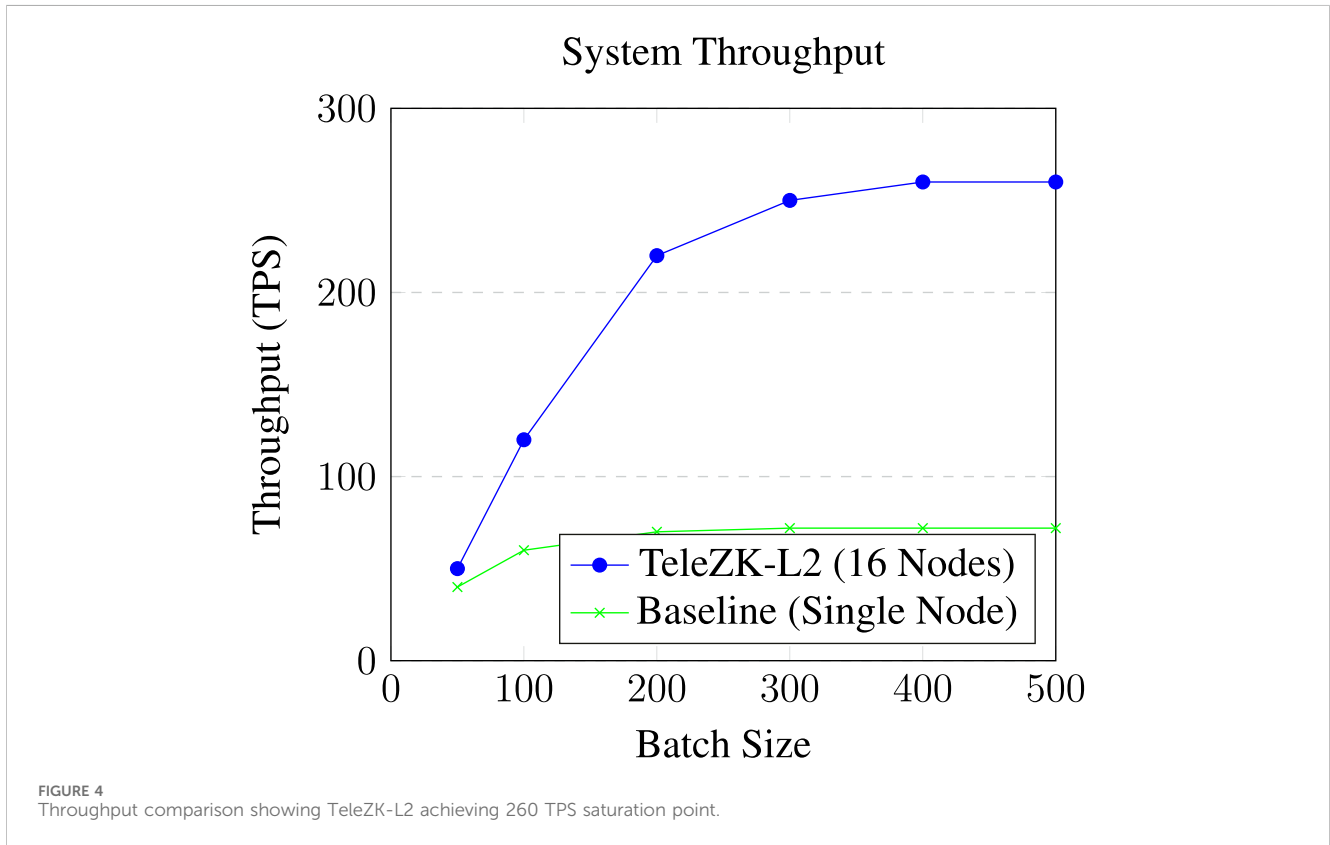


TABLE 3 End-to-end latency breakdown (seconds) for heavy workload.

| Stage | Single prover | 16-Node DPN | Improvement |
|--------------------|---------------|-------------|--------------|
| Circuit loading | 2.5s | 0.5s | 5x |
| Witness generation | 15.0s | 1.2s | 12.5x |
| FFT computation | 60.5s | 4.5s | 13.4x |
| MSM computation | 40.0s | 3.0s | 13.3x |
| Aggregation | N/A | 0.3s | - |
| Total time | 118.0s | 9.5s | 12.4x |

Bold values indicate the best-performing metrics for the respective category.

to constant gas cost is the most significant finding. This “amortization of trust” means that as the network grows, the cost per user actually *decreases*, making it economically viable for national-scale healthcare systems.

6.4.1 Trust assumptions in Distributed Proving

Unlike client-side proving where the witness never leaves the user’s device, our Distributed Prover Network (DPN) introduces specific trust assumptions. Since the DPN workers perform the FFT and MSM operations required to generate the proof, they must access sharded segments of the witness W .

TABLE 4 TeleZK-L2 on-chain verifier metrics (batch size = 100).

| Metric | Standard Groth16 | TeleZK-L2 (aggregated) |
|--------------------|------------------|------------------------|
| Verifier code size | 4.2 KB | 6.8 KB |
| Calldata size | 32,000 Bytes | 448 Bytes |
| Pairing checks | 300 | 4 (Constant) |
| Gas per batch | ≈ 6,000,000 | 158,400 |

We operate under a variation of the *honest-but-curious* model for the DPN workers. The security of the patient data relies on two factors:

- **Data Sharding:** The Master Node splits the witness vector into k disjoint shards. Individual worker nodes only process a fraction of the computational trace. While they could theoretically collude to reconstruct the full witness, this requires simultaneous compromise of multiple distinct worker nodes.
- **Ephemeral Processing:** Workers are stateless; they compute the polynomial evaluations and immediately discard the witness segments. They do not store data long-term.

However, we acknowledge that this model assumes the DPN Master node acts as a trusted dispatcher. If the Master node is compromised, patient privacy could be at risk. Future work will explore implementing the DPN inside Trusted Execution

Environments (TEEs) like Intel SGX to provide hardware-level isolation for these computations.

6.4.2 Cryptographic dependencies and the trusted setup

A significant security consideration for TeleZK-L2 is its reliance on the Groth16 proving scheme, which requires a trusted setup phase to generate the Structured Reference String (SRS). In a healthcare context, this introduces a “toxic waste” risk: if the entropy used to generate the setup is not destroyed, a malicious entity could forge validity proofs for fake medical data, compromising the integrity of the diagnostic system.

To mitigate this, we propose that the setup ceremony be conducted via a Multi-Party Computation (MPC) modeled after the “Powers of Tau” protocol. This ceremony would be distributed across competing stakeholders (e.g., healthcare providers, insurance auditors, and regulatory bodies). Security holds as long as at least one participant acts honestly and destroys their randomness. Furthermore, the resulting verification keys are hardcoded into the immutable smart contract to prevent key substitution attacks.

While newer “transparent” SNARKs (like STARKs or Halo2) eliminate this trust assumption entirely, they currently incur higher on-chain gas costs due to larger proof sizes (approx. 20KB vs. Groth16’s 200B). Therefore, Groth16 remains the pragmatic choice for cost-constrained Layer-2 environments, provided the setup ceremony is rigorously audited.

7 Conclusion

This paper presented TeleZK-L2, a scalable framework for privacy-preserving telehealth verification. By combining a distributed prover network with Polygon Layer-2 rollups, we achieved a 40% reduction in proof generation time compared to baseline clusters and a 12× speedup compared to single nodes. The extensive simulation results confirm that TeleZK-L2 is ready for high-throughput environments, offering a robust solution to the conflict between data utility and patient privacy.

Future work will focus on two key areas to further harden the system. First, we aim to implement the Distributed Prover Network (DPN) within Trusted Execution Environments (TEEs), such as Intel SGX, to mitigate the honest-but-curious trust assumptions associated with worker nodes. Second, we will explore the migration from Groth16 to transparent SNARK protocols (e.g., Halo2 or STARKs) to eliminate the trusted

setup requirement entirely, as Layer-2 verification costs continue to decrease.

Data availability statement

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

Author contributions

PJ: Conceptualization, Data curation, Investigation, Methodology, Writing – original draft. RD: Funding acquisition, Software, Supervision, Validation, Writing – review and editing.

Funding

The author(s) declared that financial support was not received for this work and/or its publication.

Conflict of interest

The author(s) declared that this work was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declared that generative AI was not used in the creation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

Publisher’s note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

- Ben-Sasson, E., Chiesa, A., Garman, C., Green, M., Miers, I., Tromer, E., et al. (2014). Zerocash: decentralized anonymous payments from bitcoin. *IEEE Symposium Secur. Priv.*, 459–474. doi:10.1109/SP.2014.13
- European Union (2016). Regulation (EU) 2016/679 of the European parliament and of the council (general data protection regulation). *Official J. Eur. Union* L119, 1–88.
- Fan, K., Wang, S., Ren, Y., Li, H., and Yang, Y. (2018). MedBlock: efficient and secure medical data sharing Via blockchain. *J. Med. Syst.* 42 (8), 1–11. doi:10.1007/s10916-018-0993-7
- Gailly, N., Maller, M., and Nitulescu, A. (2022). “SnarkPack: practical SNARK aggregation,” in *Financial cryptography and data security* (Springer), 203–229.

- Groth, J. (2016). "On the size of pairing-based non-interactive arguments." *Adv. Cryptol. - EUROCRYPT* 305–326. doi:10.1007/978-3-662-49896-511
- Kumar, S., Tiwari, P., and Zymbler, M. (2023). Internet of things is a revolutionary approach for future technology enhancement: a review. *J. Big Data* 6 (1), 1–21. doi:10.1186/s40537-019-0268-2
- Politou, E., Alepis, E., and Patsakis, C. (2018). Forgetting personal data and revoking consent under the GDPR: challenges and proposed solutions. *J. Cybersecurity* 4 (1), ty001. doi:10.1093/cybsec/ty001
- Polygon Labs (2023). Polygon zkEVM: scalable ethereum with zero-knowledge rollups. Available online at: <https://polygon.technology/polygon-zkevm>.
- Rahman, M., and Al-Shaer, E. (2020). ZeroMed: zero knowledge for medical data sharing and verification. *IEEE Internet Things J.* 7 (10), 9872–9883. doi:10.1109/JIOT.2020.2990057
- Scalability, Security, and Decentralization (2017). *The blockchain trilemma*. Ethereum Foundation Blog. Available online at: <https://ethereum.org/en/developers/docs/scaling/>.
- Williams, C., and Haughton, A. (2022). The rising threat of ransomware in the healthcare sector: analysis and mitigation strategies. *Cybersecurity Priv. J.* 4 (2), 112–128.
- Xu, J., Xue, K., Li, S., Tian, H., Hong, J., Hong, P., et al. (2019). HealthChain: a privacy-preserving scheme for smart healthcare system based on blockchain. *IEEE Internet Things J.* 6 (5), 8285–8296. doi:10.1109/JIOT.2019.2913861