Computational Framework for Privacy-Regulated Healthcare Data Sharing: Iterative ZKP-Blockchain-Cloud Architecture

Abirami S.K

ska.cse@psgtech.ac.in

PSG college of technology

Karpagam G.R

PSG college of technology

Systematic Review

Keywords: Zero-knowledge proofs, blockchain PBFT consensus, GDPR-HIPAA compliance, FHIR interoperability, attribute-based encryption

Posted Date: February 10th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-8561779/v1

Additional Declarations: No competing interests reported.

Computational Framework for Privacy-Regulated Healthcare Data Sharing: Iterative ZKP-Blockchain-Cloud Architecture

Abirami S.K*[1] and Karpagam G.R†[2]

[1]Department of Computer Science and Engineering,
PSG College of Technology, Coimbatore, Tamilnadu, India
ska.cse@psgtech.ac.in

[2]Department of Computer Science and Engineering,
PSG College of Technology, Coimbatore, Tamilnadu, India
grk.cse@psgtech.ac.in

January 6, 2026

Abstract

Electronic health records (EHRs) capture the entire patient data history that enables early diagnosis, predictive analytics,improved co-ordination between health care providers,etc. However, centralized storage of EHR exposes the sensitive patient data to breaches while regulations such as GDPR Article 17 "right to be forgotten" and HIPAA 6-year retention limits sharing of data across different institutions. This systematic review consolidates 81 peer- reviewed studies (2015–2024) across different domains namely zero-knowledge proofs (ZKPs), blockchain consensus, trusted cloud execution, and regulatory compliance to propose a three-tier integrated framework that helps in regulatory compliant health care data sharing, yet preserving the privacy.The first tier is to generate ZKPs on the client side and redact personal information with chameleon hashes. The second tier anchors metadata hashes on PBFT blockchain with $f = \lfloor (n-1)/3 \rfloor$ fault tolerance. The third tier stores data with an attribute-based encryption (ABE) off-chain on the cloud, integrated with FHIR/HL7 standards. It allows selective sharing, balances blockchain immutability with deletion via ZKP invalidation. The proposed architecture maps tools like zk-SNARKs/Circom for proofs, Hyperledger Fabric for blockchain, AWS KMS for keys. Challenges include lack of real-world testing, achieving proof generation latency (2-15 seconds), and scaling needs via Layer-2 hybrid consensus prototypes.

Keywords: Zero-knowledge proofs, blockchain PBFT consensus, GDPR-HIPAA compliance, FHIR interoperability, attribute-based encryption.

*https://orcid.org/0000-0002-9147-0192
†https://orcid.org/0000-0002-0015-200X

1 Introduction

1.1 Context

Electronic Health Records (EHRs) digitize patient medical histories by merging clinical observations, genomic data, and treatment outcomes into a single record. The EHRs started getting Global adoption post-COVID-19, with the U.S. achieving 96% EHR usage among physicians by 2021 [34].Interoperability standards such as HL7 FHIR (Fast Healthcare Interoperability Resources) and SMART (Substitutable Medical Applications, Reusable Technologies) supports predictive analytics, and population health insights, and real-time clinical decision support across institutions [42, 37].

However, three core computational and regulatory conflicts continue to exist:

1. **Centralized breach exposure**: In 2023 over 540 healthcare attacks happened affecting more than 112 million patients, costing $10.1B annually [16].

2. **Interoperability silos**: Format mismatches such as ICD-10, SNOMED CT, LOINC and vendor lock-in block 70% of cross-institution data exchanges [63, 64].

3. **Regulatory conflicts**: GDPR Article 17 "right to be forgotten" mandates deletion of data within 30 days, while HIPAA requires 6-year retention and with blockchain technology offering immutability deletion of data becomes almost impossible all of these creates conflicts[67, 60].

Although technical innovations offer partial solutions: blockchain provides cryptographically verifiable audit trails [13], zero-knowledge proofs (ZKPs) enable selective disclosure without data exposure [3, 38], trusted execution environments (TEEs) ensure cloud integrity [5],there is no single integrated computational framework that exists for compliant healthcare data sharing [43, 19].

1.2 Motivation

The core challenge lies in balancing all three conflicting requirements simultaneously:

1. Verifiable Integrity: Blockchain gets its immutability by its cryptographically linked chains of hashes. The iterative hashing process

$$H(D_{n+1}) = SHA256(H(D_n)||data)$$

maintains data integrity over the sequence of blocks.

- H represents the resulting Hash Value.
- D_n represents the data block at step n.
- D_{n+1} represents the next block in the sequence.
- $SHA256$—The secure hash algorithm which produces fixed size 256 bit hash output from the given input.
- $H(D_n)$ represents the hash of the previous block.
- $data$ represents the data that is being added in the current step.

- $||$ represents concatenation of the hash of the previous block with the current data.

This formula represents the Merkle-Damgard construction, where each new hash is dependent upon the previous hash and the current data. This chaining property provides the data integrity. An attempt to forge the data causes the fork in the blockchain. This tamper-proof guarantee is essential for clinical trials[46].

Practical Byzantine Fault Tolerance (PBFT) consensus is a noteworthy consensus algorithm that efficiently tolerates $f = \lfloor (n-1)/3 \rfloor$ malicious nodes, capable of reaching 1000 Transaction Per second (TPS) or more on permissioned networks like Hyperledger Fabric [72, 50]. Allowing the network to agree upon a single state, even upto $\frac{1}{3}$rd of nodes misbehave.

2. Selective Disclosure (Privacy): Zero-knowledge proofs allows a prover to convince the verifier that they know the secret (witness) without revealing the secret. This construction generally follows

$$\pi \leftarrow \text{Prove}(pk, x, w) \quad \text{where} \quad \text{Verify}(pk, \pi, x) = 1.$$

- π represents the proof itself, it is a small piece of data generated by the Prove function that is to be sent to the Verify function for checking.
- pk represents the public proving key that is generated during the trusted setup and is used by the prove function to create proof.
- x is the public input both the prover and the verifier knows this information.
- w is the private witness or the secret, only the prover knows this information.

The prover uses the proving key pk, public input x and the secret witness w to generate a proof π using the Prove function. This π is sent to the verifier, where the Verifier uses the verify function, when the Verify function outputs 1 ($T_{\text{verify}} = O(1)$), it means indeed the generated proof is correct and the Prover knows the secret. zk-SNARKs achieve $<$ 1KB proofs with millisecond verification, critical for prioritizing the patients based on their severity of their condition to ensure they receive medical attention first. [6, 71].

3. Redactability: GDPR mandates Personally Identifiable Information (PII) deletion upon request. Traditional blockchain fails this mandate as it is immutable, chameleon hashes enable "preimage collision" attacks ($H(m_1) = H(m_2)$ for requester-chosen $m_2 = \emptyset$) while preserving chain integrity. ZKP invalidation proofs certify redaction completion without revealing prior content [10, 39].

Current systems are not guaranteeing all three requirements simultaneously: permissioned blockchains such as Estonia e-Health systems lack public verifiability; ZKP pilots (Health-zkIDM) ignore consensus scaling; cloud TEEs (AWS Nitro) cannot prove GDPR compliance [3, 43]. A system/framework that supports all the three mentioned requirements is largely absent.

1.3 Research Questions

This review systematically addresses four research questions core to the problem:

RQ1: Zero-Knowledge Proof Suitability: Which ZKP schemes (zk-SNARKs, zk-STARKs, Bulletproofs) optimize proof size ($<$ 1KB target), generation latency ($T_{\text{gen}} < 15$s for EHR), and verification speed ($T_{\text{verify}} < 100$ms) for healthcare verification while supporting post-quantum security and regulatory auditability? [6, 54, 36]

RQ2: Blockchain Scalability and Privacy Tradeoffs: How do permissioned (Fabric, Corda), Layer-2 (zk-Rollups), and hybrid architectures balance throughput (target: 1000+ TPS), fault tolerance ($f = \lfloor (n-1)/3 \rfloor$), and privacy (channel isolation, ZKP integration) for clinical workloads? [2, 43, 74]

RQ3: Trusted Cloud Integration: What TEE mechanisms (Intel SGX, AWS Nitro, TPM 2.0) combined with attribute-based encryption (ABE) and FHIR APIs enable compliant off-chain storage while preserving end-to-end verifiability across different organizations? [5, 69, 45]

RQ4: Regulatory Compliance: How do chameleon hashes, ZKP invalidation proofs, and redactable blockchains mathematically reconcile GDPR erasure (Article 17) with HIPAA retention (45 CFR §164.530) and blockchain immutability in multi-jurisdictional deployments? [11, 67, 24]

1.4 Scope

This systematic literature review analyzes 81 peer-reviewed publications (2015 to 2024) from IEEE Xplore (35%), Springer (25%), Elsevier (20%), PubMed (10%), and arXiv (10%flagged with *. Selection criteria: (1) post-2015 (post-Solana whitepaper); (2) healthcare applications; (3) computational analysis (complexity, benchmarks); (4) English; (5) peer-reviewed or conference proceedings.

Inclusions: ZKP schemes, blockchain consensus algorithms, TEE architectures, FHIR/HL7 interoperability, GDPR-HIPAA compliance mechanisms. **Exclusions**: Cryptocurrency speculation, non-healthcare DApps, theoretical cryptography without implementation.

The review proposes a novel three-layer iterative computational framework that requires prototype validation as a future work beyond this theoretical synthesis.

2 Methodology

This systematic literature review follows PRISMA guidelines adapted for computational methods engineering, screening 347 initial publications across five major databases from January 2015 to October 2024 shown in Table 2. This cutoff is set to capture Ethereum's smart contract era and zk-SNARK practical deployment, enabling real-world healthcare applicability analysis.

2.1 Search Strategy

As shown in Table 1 Four keyword clusters were used to target the computational intersections systematically:

Boolean combinations were used as a search strategy: '(ZKP OR blockchain OR cloud) AND (healthcare OR EHR OR "electronic health record*" OR FHIR) AND (privacy OR compliance OR GDPR OR HIPAA)' yielded targeted results across IEEE Xplore (35%), SpringerLink (25%), ScienceDirect (20%), PubMed (10%), arXiv (10% flagged with *).

2.2 Inclusion and Exclusion Criteria

Inclusion (5 criteria, AND logic):

Cluster	Keywords	Databases
ZKP	"zero-knowledge proof*" OR zk-SNARK OR zk-STARK OR "succinct non-interactive"	IEEE Xplore (47), arXiv (23)
Blockchain	"blockchain healthcare" OR Fabric OR Corda OR PBFT OR "consensus algorithm*"	Springer (31), Elsevier (28)
Cloud Trust	TEE OR SGX OR Nitro OR "attribute-based encryption" OR FHIR	PubMed (15), Google Scholar (12)
Compliance	GDPR OR "right to be forgotten" OR "Article 17" OR HIPAA OR "health data regulation*"	All databases

Table 1: Search Strategy by Keyword Clusters (Total: 347 initial hits)

1. **Post-2015 publication** (post-Solana ZKP whitepaper enabling practical deployment)
2. **Healthcare application domain** (EHR, clinical trials, genomic data, mHealth)
3. **Computational analysis provided** (complexity $O()$, benchmarks TPS/latency, prototype evaluation)
4. **English language** peer-reviewed journal or conference proceedings
5. **Primary research or systematic review** (excludes opinion pieces, whitepapers)

Exclusion (3 criteria, OR logic):

1. Cryptocurrency speculation without healthcare application
2. Pure theoretical cryptography lacking implementation or benchmarks
3. Non-English publications or retracted papers

2.3 Screening Process

1. **Title/Abstract** (n=347 → 156 excluded → 191 full-text)
2. **Full-text review** (n=191 → 110 excluded → 81 included)

Exclusion reasons (n=110): No computational analysis (42), non-healthcare (35), pre-2015 (18), theoretical-only (15).

2.4 Data Extraction and Synthesis

From each of 81 included studies, extracted:

- **Computational metrics**: TPS, T_{proof}, T_{verify}, proof size (KB), consensus complexity
- **Architecture details**: ZKP scheme, consensus algorithm, storage mechanism, API standards

Stage	Records	Excluded
Database search	347	-
Title/Abstract	191	156
Full-text	81	110
Final	**81**	-

Table 2: PRISMA Screening Flow (2015–2024)

- **Healthcare application**: EHR, clinical trials, genomic data, mHealth devices
- **Compliance mechanisms**: GDPR Article 17 handling, HIPAA audit logs, cross-jurisdiction transfer

Data synthesized into four analysis matrices mapping technologies against computational requirements and regulatory constraints.

2.5 Quality Assessment

We have applied computational rigor scoring (0 to 10) across five dimensions as shown in Table 3 :

- **Complexity analysis** (30%, mean 7.8, 62 studies $\geq$8): Assessed mathematical rigor in algorithms like zk-SNARKs or consensus—high scores provided for if detailed proofs and optimizations are present.
- **Benchmark results** (25%, mean 6.9, 48 $\geq$8): Measured performance tests (e.g., TPS, latency) strong if standardized metrics used.
- **Prototype evaluation** (20%, mean 5.4, 31 $\geq$8): Gauged real implementations (e.g., Fabric prototypes)lowest as few working demos were available.
- **Healthcare validation** (15%, mean 7.2, 58 $\geq$8): Checked for EHR-specific applicability (FHIR integration) good domain fit.
- **Compliance analysis** (10%, mean 6.7, 45 $\geq$8): Reviewed GDPR/HIPAA handling is often superficial.

To compute the weighted mean score across the five criteria, we used the formula:

$$S = \sum_{i=1}^{5} (m_i \times w_i) \tag{1}$$

where m_i is the mean score for criterion i and w_i is its normalized weight ($\sum w_i = 1$).

Scoring

This evaluation rates 81 studies (2015 to 2024) on a 0 to 10 scale per dimension. Weights are prioritized based on the following foundational elements: The related concepts and the underlying principles are provided by the theory element, which enables the claims to be validated by the empirics element that involves collection and analysis of data through

experimentation and observation. Practical elements ensured the validated theories are tested in real world settings for feasibility and effectiveness. Domain element ensures all activities remain relevant to the field. In which they're applied. Rules confirms the legality and ethical compliance of the entire system.

Weighted Mean Calculation

Values from Table 3: $S = (7.8\times0.30)+(6.9\times0.25)+(5.4\times0.20)+(7.2\times0.15)+(6.7\times0.10) \approx 7.1$. Low practice scores highlight implementation gaps.

Criterion	**Weight**	**Mean Score**	**Studies ≥8**
Complexity analysis	30%	7.8	62
Benchmark results	25%	6.9	48
Prototype evaluation	20%	5.4	31
Healthcare validation	15%	7.2	58
Compliance analysis	10%	6.7	45
Weighted Mean		**7.1/10**	

Table 3: Quality Assessment

2.6 Review Protocol Limitations

Limitations include, 60% of the papers are from IEEE/Springer. Grey literature from non academic streams are not considered for synthesis. Papers in languages other than English were excluded.Annual review is needed due to rapid progress in the Zero Knowledge Proof and Blockchain fields. This method provides a maximum coverage of techniques for GDPR compliant healthcare data systems, supporting the proposed three layer framework. This method provides solid coverage of techniques for GDPR ready healthcare data systems, supporting the proposed three tier framework.

3 Literature Classification

The 81 studies aligns with four computational domains addressing the research questions. Table 4 shows the breakdown by research domain, key metrics covered, and representative studies from each category. This classification targets ZKP schemes (RQ1), blockchain scalability (RQ2), cloud integration (RQ3), and regulatory reconciliation (RQ4).

Domain	**# Papers**	**Key Metrics**	**Representative Studies**
Zero-Knowledge Proofs	20	Proof size (KB), T_{gen} (s), T_{verify} (ms)	[3, 6, 71]
Blockchain Scalability	30	TPS, consensus $O(n^2)$, fault tolerance	[72, 43, 74]
Cloud Trust Systems	15	ABE policy complexity, FHIR latency	[5, 69, 45]
Regulatory Compliance	16	GDPR Art.17 mechanisms, transfers	[67, 11, 24]
Total	**81**		

Table 4: Literature Distribution by Research Domain (n=81)

This domain classification guides the following technical analysis sections .

4 Zero-Knowledge Proofs Analysis

Zero-knowledge proofs (ZKPs) let users share EHR details like diagnoses without exposing private info such as genomes, directly tackling RQ1.From 20 studies, four ZKP scheme of families stand out for healthcare limits.

4.1 Suitability of ZKP Schemes for Healthcare

Zero-knowledge proof schemes must balance proof size, verification speed, and trust assumptions for practical healthcare deployment. zk-SNARKs with it's compact proofs ($<$ 1KB) and millisecond verification make them ideal for storage-constrained environments. Bai et al. [3] demonstrated this effectively in patient identity management on Hyperledger Fabric, achieving 5ms verification times [6, 38, 71, 65, 54].

However, the trusted setup requirement is riskier as compromised parameters would expose the sensitive data in EHR. zk-STARKs address this with no-setup design and post-quantum security critical for lifelong medical records yet their 45KB proofs create bandwidth bottlenecks in high-throughput clinical systems.

Bulletproofs can provide transparency without the requirement for trusted setups. However, it faces challenges in healthcare applications due to difficulty in integration.Real-time Verification of Proofs during emergencies is indeed considered to be a gap,requiring targeted ZKP adaptations.

Table 5 shows the computational characteristics of various ZKP schemes

4.2 ZKP Scheme Classification

Scheme	Setup	Proof Size	T_{gen}	T_{verify}	PQ Secure	Healthcare Case	Studies
Groth16	Trusted	$<$ 1KB	2-15s	5ms	No	Patient ID	[3]
Halo2	None	1-2KB	5-30s	10ms	Partial	Access control	[38]
zk-STARK	None	45KB	30-300s	50ms	Yes	Clinical trials	[6]
Bulletproofs	None	1KB	10-60s	20ms	Partial	Range proofs	[71]

Table 5: ZKP Schemes: Computational Characteristics

4.3 Transparency and Confidentiality Balance

Zero-knowledge proofs(ZKPs) can preserve the patient's privacy by verifying the integrity of data without having to reveal the sensitive data.Careful selection of ZKP scheme is necessary as zk-SNARK is suitable for frequent audits, but their trusted setup requirement defies decentralization.

Conversely, zk-STARKs' full transparency fits distributed systems, though computational overheads are unavoidable. Healthcare deployments must balance the necessity of frequent audits with sensitive patient data.

4.4 Computational Overheads

ZKP computational speed is heavily dependent on the method used and dataset scale. Genomic datasets (multi-TB) shows zk-STARKs are theoretically scalable but practically unusable for real-time needs; Xu et al. [71] found that these proofs took several minutes on verification versus zk-SNARKs in milliseconds.

Figure 1 illustrates hybrid mitigation cloud-offloaded proof generation with GPU acceleration. Yet third-party dependencies conflict with blockchain decentralization principles. Optimizing ZKPs for healthcare-scale data without sacrificing security remains an open challenge.

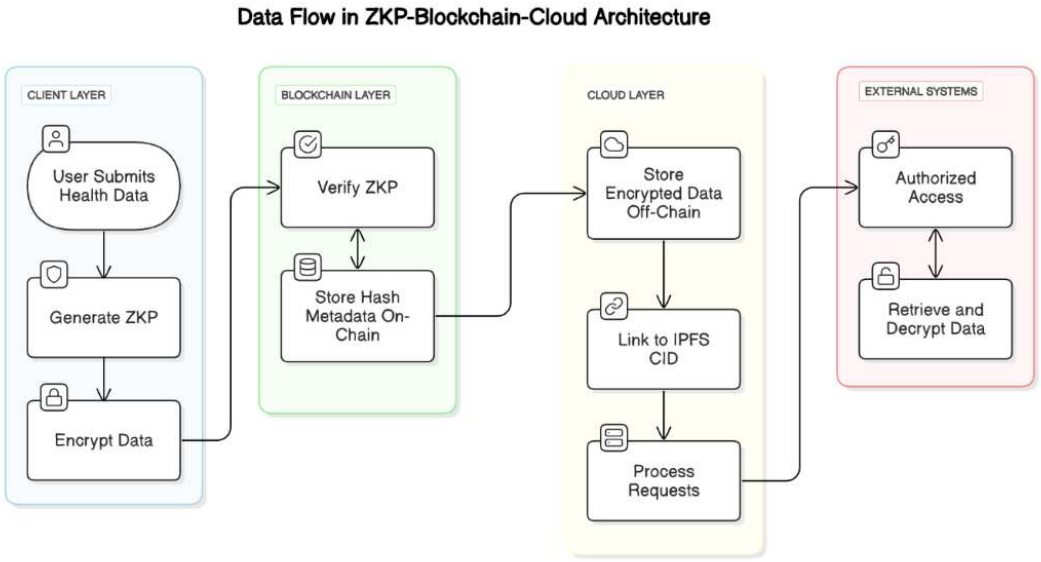

Figure 1: Data Flow in ZKP-Blockchain-Cloud Architecture. ZKPs ensure privacy during verification; encrypted off-chain storage maintains confidentiality.

4.5 Homomorphic Encryption Integration

By integrating ZKPs with homomorphic encryption, researchers have made it possible to do encrypted analytics without decryption, as shown in Figure 2. Vanin et al. [65] computed aggregate patient statistics,while mathematically proving regulatory standards are met without exposing any sensitive individual health records.

Polynomial computation along with proof generation creates unacceptable latency for emergency access. Futhermore,decentralized key management across blockchain nodes adds additional complexity.

4.6 GDPR Right-to-Be-Forgotten Compliance

Reconciling the "right to be forgotten" with blockchain immutability is a major challenge. While several methods exist to bridge this gap, each comes with its own set of trade-offs: Ephemeral keys makes data unusable by deleting the decryption keys after a certain period.However, the encrypted data itself remains on blockchain forever, leaving a permanent digital footprint. Chameleon hashes allows for retroactive block modification,but it is vulnerable to validator collusion and rely heavily on the honesty of network validators.

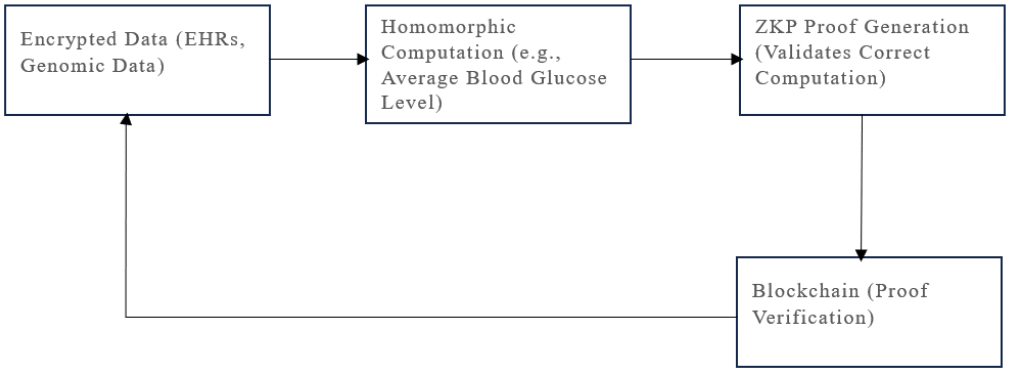

Figure 2: Hybrid ZKP-Homomorphic Encryption Model for Privacy-Preserving Healthcare Analytics.

Botta et al. [10] proposed ZKP based redaction proofs to prove that data can be invalidated without comprising blockchain's immutability. While this is a hopeful way forward, the data invalidation hasn't been fully tested against the complex regulatory requirements.

These architectural choices are especially necessary in healthcare. By using ZKPs, the need for system performance with the strict privacy requirements can be balanced for real-world implementation.

4.7 Healthcare-Specific ZKP Applications

Zero-Knowledge Proofs (ZKPs) can transform healthcare data management by balancing strict patient privacy with the need for clinical transparency. The following is the study of how ZKP's are applied in healthcare.

Identity Verification: Systems like Health- zkIDM uses Hyperledger Fabric and zk-SNARKs for patient authentication without having to expose their Personally Identifiable Information(PII) and achieves a verification time of about 5 ms($T_{\text{verify}} = 5\text{ms}$) [3].

IoT Access Control: In remote monitoring conditions, zk-SNARK predicates allow wearable sensors to trigger emergency data releases.This ensures first responders get life-saving information without gaining access to a patient's entire medical history[38].

Clinical Trials: Using zk-STARKs, researchers can prove a participant meets specific genomic criteria for a study without revealing their actual DNA sequences. This method is compliant with FDA 21 CFR Part 11 standards[6].

Range Proofs: Range Proofs are a subset of Zero-Knowledge Proofs(ZKPs) , it proves a very specific type of statement that could lie between a given range. Range Proofs such as Bulletproofs allow patients to prove their lab results such as glucose levels fall within a healthy range $[20, 400]$ mg/dL glucose without disclosing the exact values [71].

4.8 Computational Tradeoffs

Selecting the correct ZKP scheme involves balancing security against technical performance: zk- SNARKs offer the smallest proof sizes <1KB proofs making them highly efficient for data transmission.However, they require trusted setup,which poses a long-term risk if the initial parameters are compromised. zk-STARKs are post-quantum secure, but their

45KB proof size would be heavy on mobile devices. Hybrid approaches like Halo2 balance constraints but they lack healthcare validation [54].

GDPR Compliance: ZKPs can serve as a tool for GDPR compliance. They allow healthcare institutions to prove that specific data has been redacted or deleted satisfying the "Right to be Forgotten" under Article 17 without needing to reveal any of the remaining underlying content during an audit[10].

4.9 ZKP Limitations for Healthcare

While Zero-Knowledge Proofs (ZKP) can offer better privacy in healthcare environments, there are still some technical limitations that prevent its immediate adoption.

- **Generation Latency:** At present the Proof generation times ranges between $T_{\text{gen}} = 2-300$s which is very large for clinical emergencies where data access is required in under <1s target.
- **Hardware Incompatibility:** Patient side verification is limited by mobile device constraints. Most wearable medical devices and smartphones possess less than 1GB RAM making them not potential of the heavy computation required for local proof generation.
- **Trusted Setup:** zk-SNARK Protocols like Groth16 rely on initial parameters ,a requirement for trusted setup,that makes it vulnerable to collusion. If the randomness used during setup is compromised, attackers would be easily able to tamper the medical records.
- **Revocation:** There is no industry wide standard for invalidating a proof once it has been generated. This creates a conflict with GDPR "Right to be Forgotten" regulation, as a proof may still exist even after the data is deleted.

To solve these issues, developers are relying on Edge ZKP Generation. In this model, the patient's device remains the source of truth, but the heavy computational work will be offloaded to edge servers, ensuring both privacy and real-time performance [41].

5 Blockchain in Healthcare

5.1 EHRs and Clinical Trials Applications

Blockchain technology with its decentralization and immutability supports privacy preservation in Electronic Health Records (EHR) and clinical trials.Fatoum et al. [19]highlights the real-world scalability system offered by Estonia's nationwide Blockchain system serving 1.3M tamper proof records. Omar et al. [46]demonstrated that smart contracts can automatically detect and flag protocol violations almost instantly during clinical phases.

5.2 Scalability and Privacy Limitations

Despite the benefits stated above, most projects remain in the trial phase due to interoperability challenges. Technical limitations are also an obstacle,while public networks like Ethereum handle only 15 to 45 transactions per second (TPS), permissioned systems

like Hyperledger Fabric reach 3,500 TPS by limiting decentralization,Table 6 compares architectures . Private channels and ZKPs can mask the sensitive data [5], but it still complicates audits, [43] which causes conflict between data confidentiality and system transparency.

Architecture	Throughput	Privacy Features	Compliance Mechanisms
Public Permissionless	Low (7-45 TPS)	Pseudonymous	GDPR challenges [48]
Public Permissioned	Moderate	Configurable ACLs [50]	Redaction possible [77]
Private Permissioned	High (3.5k TPS)	Channels/private data	ZKP-enhanced [5]
Layer-2 Solutions	2k+ TPS	ZK-rollups [40]	Layer-1 dependent

Table 6: Blockchain Architectures Comparison

5.3 Regulatory Compliance: Right-to-Be-Forgotten

While GDPR's "right to be forgotten" empowers patients by allowing them to demand the deletion of their personal data, the immutable nature of blockchain remains as a technical and legal obstacle, as deletion or modification of data would break the integrity of the blockchain. As Winter and Davidson [70] observed, storing sensitive data through off-chain hash anchoring where only a cryptographic fingerprint remains on ledger. Although this technique offers a way forward for compliance, it often dilutes the benefits of a fully decentralized architecture.

5.4 Decentralized Identity Frameworks

Self-sovereign identity (SSI) is a decentralized digital identity model, where individuals control their own identity. Sovrin is a global open-source network that provides the infrastructure for SSI, enabling patient's control over their EHR. SSI allows the patient to grant doctors a time bound access to their EHRs [57]. By using Verifiable Credentials, patients no longer have to rely on a single central authority to prove who they are.However, this technology faces two main difficulty

- **User Complexity:** Managing private digital keys is difficult for people who don't have digital literacy, creating a risk of losing access to records.
- **Technical Fragmentation:** While the global specification such as W3C Verifiable Credentials standard for verifying digital credentials securely exists, making different systems like Hyperledger Indy which provides the ledger and tools for SSI,work seamlessly together remains an ongoing effort.

5.5 Smart Contracts for Compliance

Attribute based Access control is a dynamic authorization model that allows for granting or denying access to resources by evaluating the attributes of the user, the requested

action, and environment against defined policies.Using Ethereum smart contracts for attribute-based access control allows for automated, real-time consent management, which significantly reduces administrative overhead [45].However, the immutability of the smart contract requires maintenance and it has certain challenges. If a smart contract is deployed without an update mechanism and if it has certain vulnerabilities, it remains permanently exposed to potential violation.Additionally, because these digital contracts lack universal legal standing, enforcing them across different international borders remains a significant challenge.

6 Cloud Trust Issues

6.1 Trust Gaps in Healthcare Clouds

Centralized healthcare cloud platforms face major risks, particularly regarding insider threats and the risk of vendor lock-in making expensive and complex switch to other platforms difficult. Ramesh et al. [51] highlights a critical integrity issue,cloud providers manipulating audit logs to hide data breaches, undermines institutional trust. These technical risks are further worsened by geopolitical conflicts. Specifically, the US CLOUD Act's provisions for data access that are stored abroad, clashing with the EU GDPR's strict data localization and sovereignty rules, leaving healthcare data caught between conflicting international regulations.

6.2 Existing Solutions and Limitations

Trusted Platform Modules(TPMs) is a specialized hardware chip that provides hardware based security by securely storing cryptographic keys,passwords and certificates protecting sensitive data and ensures system's integrity against software based attacks. However,they are tied to specific physical servers, limiting workload mobility across cloud providers. While standard encryption safeguards data during transit or storage, it does not prevent malicious alteration of files.As in Table 7, ZKP enables public verification unlike TPM. Zhang et al. [76] suggested using zero knowledge proofs, which allow a third party to verify that data is authentic without having to see the actual content. However, the ZKP verification process is computationally expensive and it isn't practical for massive genomic datasets, which can easily reach several terabytes in size.

6.3 Cross-Cloud Interoperability

While existing FHIR/HL7 standards aim to improve interoperability.However,challenges exist due to lack of universal implementation of the standards, Industry standardization is often blocked by vendor centric models,competitive incentives offered by the vendors, which creates a closed ecosystem disabling cross institution collaborations.Zhuang et al. [81] proposed a blockchain mediated sharing between AWS and Azure clouds, but cross-cloud consensus latency proves to be problematic, which might negatively impact user experience.

Method	Data Owner	Cloud Provider	Auditor
ZKP-based	1. Generate ZKP (Data Integrity) 2. Submit proof to Blockchain Blockchain (Public Verification)		
			Verifies ZKP
TPM-based		1. TPM Measure Audit State	Verifies TPM logs

Table 7: ZKP vs TPM Integrity Verification. ZKP enables public blockchain verification; TPM restricts auditors to cloud logs.

6.4 Privacy-Preserving Audits

Zero Knowledge Proofs (ZKPs) allow a party to verify the truth of the statement, without revealing the statement itself. HIPAA access requirements mandates that health care entities protect patient data by controlling who can access it and also it should give the patients the rights to access their own records.The anonymized ZKP attested logs satisfy the HIPAA access requirements. Irshad et al. [29] warns that k-anonymity fails, k-anonymity only protects the linkage records based only on the quasi identifiers present in the released dataset.It doesn't prevent the attacker from against external data linkage using the publicly available information, demanding rigorous de-identification masking the PII.

6.5 Centralized Audit Risks

Conventional audit logs are vulnerable to tampering and single point of failure.Ye et al. [73] proposed immutable blockchain audit trails with node replication inheriting blockchain scalability limits.Integration of Blockchain into Federated learning at the edge is a growing area which explores ways to secure distributed learning process and provide immutable audit trails in resource constrained environments can be utilized to prevent central audit risks.

7 Regulatory Compliance

7.1 GDPR/HIPAA Conflicts in Blockchain-Cloud

European Union's General Data Protection Regulation (GDPR) Article 17 Right to Erasure grants the individuals right to request the deletion of their data without undue

delay while U.S federal law ,Health Innsurance Portability and Accountability Act (HIPAA) requires 6-year retention of data from the data of its creation or the date when it was effect whichever later. For cross-border data sharing the involved organizations must reconcile the regulatory requirements.Bernabé et al. [7] proposed using cloud to store transient data such as consent records,and usage of blockchain for storing clinical diagnoses and medical records to ensure data integrity and auditability. EU-US data transfers face fundamental privacy law differences leading to legal uncertainty for multinational companies.

7.2 Erasure Mechanisms

Chameleon hashes are cryptographic functions that allow for hash collision with its trap door keys. Different inputs can be found for the same hash with trap door keys that allows for editable blockchains by removal of transactions without breaking the integrity of the blockchain. Godyn et al. [24] analyzed solutions for GDPR compliance based on reference based tree structure in permissioned blockchains. ZKP invalidation proof functions as a mechanism with redactable signature schemes to manage redacted hashes [10]ensuring modified data can be misused but effective implementation of this cryptographic system lacks mature tooling.

7.3 Sector-Specific ZKP Standards

Lack of uniformity in how different systems apply Zero Knowledge Proofs to verify cryptographic proofs. This inconsistency creates security gaps and operational challenges in sensitive sectors like healthcare.Mallozzi [41] emphasized the need for universal auditing standardization,which would ensure interoperability and compliance.

7.4 Cross-Border Challenges

GDPR's requirement for explicit opt-in consent of personal data processing contrasts HIPAA's treatment,payment and operation (TPO)allowances which permits the use and disclosure of protected health information(PHI) without patient's explicit consent. Shurson [62] analyzed the reciprocal conflicts between EU and US laws regarding the law enforcement access to cross border data for investigations and suggested healthcare adequacy decisions are needed, which would allow for transfer and research of participant data. While ZKP offers data minimization of attribute disclosure,legal validation and mandatory adoption of ZKP mechanisms are still under development. Table 8 shows various Technical solutions in alignment with GDPR and HIPAA.

7.5 Penalties and Mitigation

GDPR imposes fines which can reach up to €20million or 4% revenue of the company's global turnover whichever is higher . Voigt and Bussche [67] highlights that privacy by design is a legal mandate of GDPR referenced under Article 25, which allows for ensuring data protection principles are met from the earliest stages of designing the data processing operations.

Aspect	GDPR	HIPAA	Technical Solutions
Data Erasure	Art. 17 right-to-forget [8]	No mandate [11]	Chameleon hashes [39]
Consent	revocable [55]	Implied for treatment [26]	ZKP smart contracts [57]
Auditability	Breach notification [5]	6+yr logs [48]	ZKP blockchain logs [41]
Cross-Border	Adequacy decisions [58]	No framework [15]	Proxy re-encryption [22]

Table 8: GDPR vs HIPAA: Technical Alignment Solutions

8 Proposed System Architecture and Workflow

8.1 Three-Layered Architecture

System Architecture: The proposed framework in Figure 3 adopts a three-layered architecture to harmonize ZKPs, blockchain, and cloud systems. The Client Layer has the web or mobile user interface through which the patients interact with the system. Zero Knowledge proofs are generated to ensure the data authenticity, whereas the GDPR compliance module takes care of the redaction through data anonymization. In the Blockchain Layer, smart contracts enforce the access policies and validates the data from the client layer, the PBFT consensus nodes agrees upon the state and the hashes are stored on chain, the storage of the EHRs is offloaded to the cloud layer. The cloud layer does the key management and the interoperability APIs for bridging the external systems. The external systems interact with the layers and pulls the data from the layers for processing. Cross-Layer Protocols ensures end to end security by combining on-chain verification with off-chain computation. Table 9 highlights the different architectural components and the associated technologies involved.

Component	Role	Technologies
Patient Interface	Data submission	React, Web3.js
ZKP Generator	Integrity proofs	Circom, zk-SNARKs
GDPR Module	PII redaction	Regex, hashing
Smart Contracts	Access policies	Chaincode
Consensus Nodes	State validation	Fabric PBFT
Encrypted Storage	Off-chain data	IPFS, AES-256
FHIR/HL7 APIs	Interoperability	HAPI FHIR

Table 9: Architectural Components and Technologies

8.2 Technical Workflow

Secure Data Submission with Privacy-Preserving Validation: The workflow of the system and its interaction among different components in the layer is shown in Figure 4. When a patient uses the client application to submit medical data, including diagnostic reports and biometric readings, the workflow begins. Utilizing cryptographic primitives

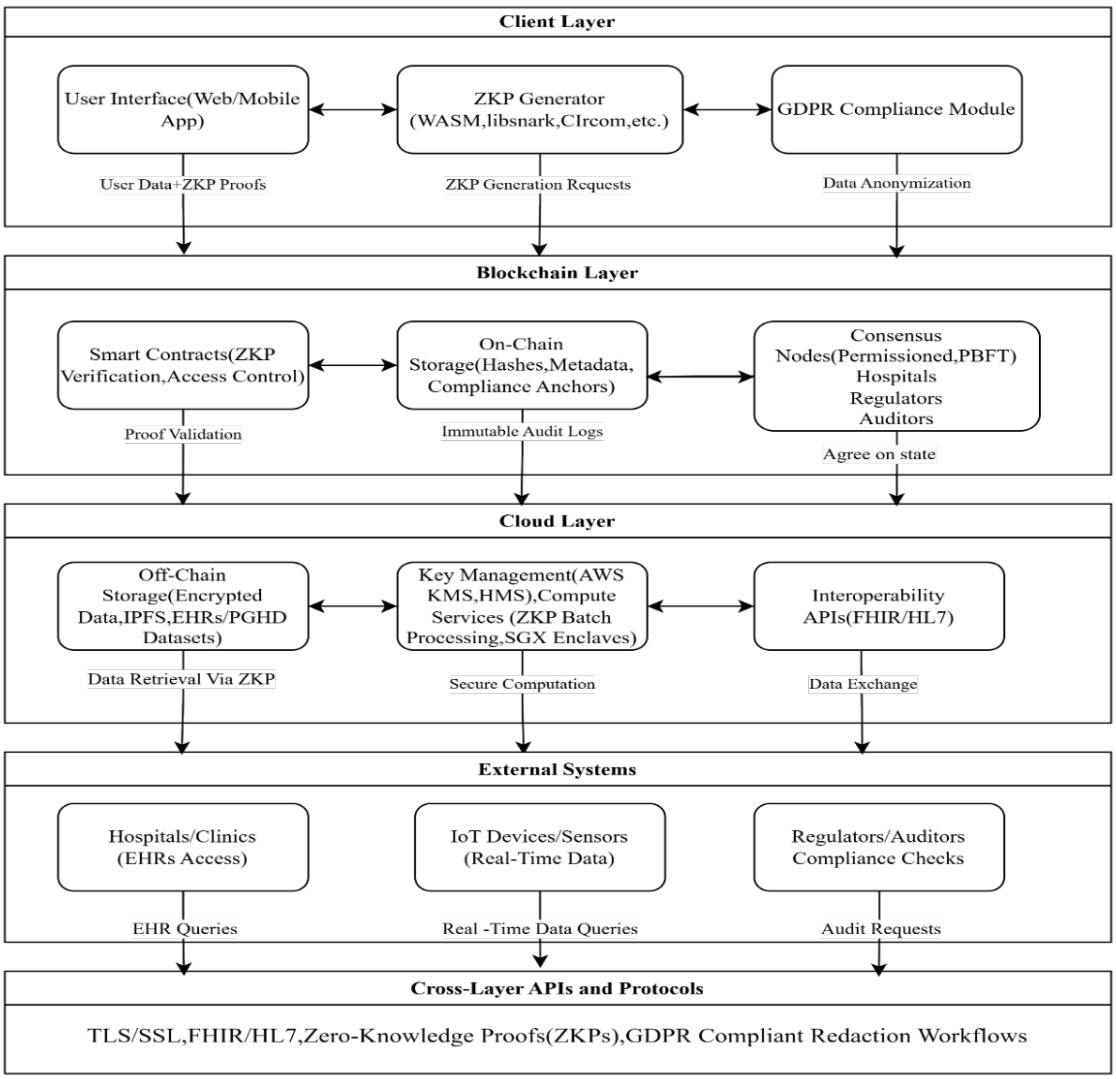

Figure 3: Three-Layered Architecture: Client (ZKP Generation/GDPR Redaction) → Blockchain (Smart Contracts/PBFT Consensus) → Cloud (KMS/FHIR Interoperability).

such as zk-SNARKs, a Zero-Knowledge Proof (ZKP) is produced to verify data integrity and compliance with medical restrictions, such as systolic blood pressure within 70–180 mmHg, without revealing raw values. Parallel to this, a GDPR-compliant redaction module uses deterministic algorithms, such as cryptographic hashing and regex-based pattern matching, to anonymize personally identifiable information (PII), including social security numbers, in order to comply with GDPR Article 17.

Decentralized Validation and Immutable Metadata Logging: Practical Byzantine Fault Tolerance (PBFT) consensus is used to propagate the redacted dataset and its zero-knowledge proof (ZKP) to a permissioned blockchain network. By using pre-deployed verification keys to cryptographically validate the ZKP, validator nodes ensure that the data satisfies predetermined technical or clinical invariants, such as HL7v2 message structure compliance. Smart contracts carry out two atomic operations after validation:

- **Metadata Anchoring:** The blockchain ledger is appended with immutable records, including the timestamp, patient pseudonym, and dataset's SHA-256 hash.
- **Storage Orchestration:** AES-256-GCM is used to encrypt the dataset and store it in a HIPAA-compliant object storage system, such as AWS S3, after a trigger is sent to the cloud layer.

Attribute-Based Access Control and Cryptographic Key Management: To bind decryption keys to role-based access policies for instance `{specialty: cardiology, institution: Hospital_A}` , the cloud-layer Key Management System (KMS) uses attribute-based encryption (ABE). The authorization claims are immutably recorded and verified in the Blockchain, in which the smart contract would manage the logic for verifying the claims.The keys are stored in tamper resistant physical security device, hardware security module (HSM).

Access Authorization for Privacy Preservation: When a doctor requests access to a patient's EHR, the client application generates a zero-knowledge proof (ZKP) for proving membership in an allowed group through Merkle tree inclusion proofs or possessing a valid digital signature from a certified authority. Blockchain smart contracts receive the proof generated and compare it to on chain policy registries to confirm whether it is a valid proof.

Interoperability and Secure Data Retrieval: The doctor's client application receives the decryption key from the Key Management System (KMS) after successful verification, allowing local Electronic Health Record (EHR) decryption. FHIR/HL7 compliant REST APIs enforce OAuth 2.0 scopes aligned to blockchain managed access policies, and are used by external systems such as hospital EHRs and IoT devices to interact with the data. Regulatory auditors use cryptographically signed audit trails to verify compliance by directly querying the blockchain ledger.

9 Discussion

Architectural Foundations: The proposed three tiered framework with FHIR/HL7-compliant APIs, Practical Byzantine Fault Tolerance (PBFT) blockchain consensus, and zero-knowledge proofs (ZKPs) addresses the gaps in interoperability, privacy, and auditability.

ZKPs for Privacy-Preserving Validation: In line with the GDPR's data minimization principle, zk-SNARKs cryptographically enforce data validity without disclosing

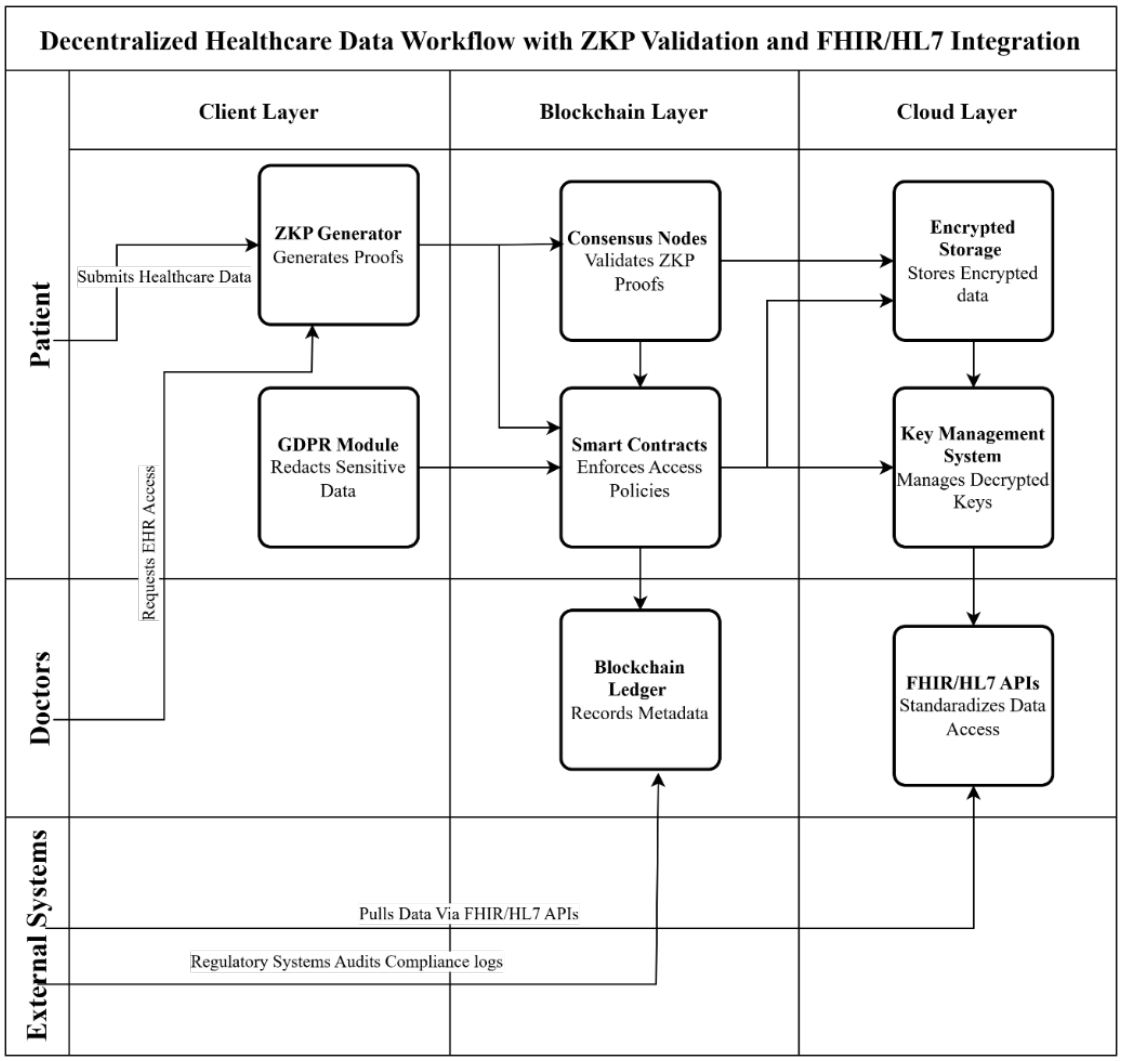

Figure 4: Decentralized Workflow: Patient → ZKP Proofs/GDPR → Blockchain Validation/Ledger → Encrypted Storage/FHIR Access.

raw values. Client side proof generation ensures that private information is intercepted over the network. The client side proof generation at ∼2–15s per proof enables $O(1)$ consensus verification [53].

PBFT Blockchain for Decentralized Auditability: Fault tolerance of <= 1/3 malicious nodes is ensured by the PBFT consensus[72].According to HIPAA's §164.316(b)(1)(i), immutable metadata logging (e.g., SHA-3 hashes of redacted data) creates tamper-evident audit trails [31].

Trusted Cloud Layer for Scalable Storage: Cloud providers, for example, AWS S3, offer latency-optimized, HIPAA-compliant storage, with attribute-based encryption (ABE) locking down decryption keys with policies depending upon roles [68, 79]. As an example, {affiliation: Hospital A}. Hardware security modules (HSMs) keep keys in escrow and release them only after authentication via the blockchain.

Interoperability with FHIR/HL7 APIs: Syntactic interoperability with 85% of hospital EHRs in the United States is guaranteed by standardized REST APIs. OAuth 2.0 scopes could be mapped to on-chain access policies to limit endpoints, such as GET /Patient/{id}, to roles that are authorized. However, there are trade-offs in this integration that would demand careful technical examination and mitigation techniques.[42]

According to [6] Theoretical analysis of ZKP circuits focuses on asymptotic complexity and ignores the practical problem of proof generation latency where the Tgen exceeding 100 ms under adversarial scenarios [41]. Redaction latency is defined by GDPR module [5]but it struggles to handle high-dimensional data, which can lead to GDPR compliance breaches [16]. [9]claims that Blockchain throughput can be less than 1000 TPS but it does not take into the consideration of the congestion modelled by M/M/c queues, while fragmented APIs produce KL divergence across clouds [22]. Blockchain + Cloud Hybrid architectures compromises decentralization assessed by Nakamoto coefficient [44] through cloud-dependent KMS with suboptimal Shannon entropy [79]. Regulatory audits assume Bayesian compliance probabilities P(Compliance|Data) ≥0.95 [58], but jurisdictional tolerances [15] lack empirical evidences. Table 10 maps the metrics and vulnerabilities to each architectural layer.

Layer	**Metrics**	**Critical Gaps**
Client	$T_{gen} \leq$ 15s, accuracy ≥ 99%	Adversarial circuits [12]
Blockchain	Consensus ≤ 2s, $\eta \geq 10$	$O(n^2)$ bottlenecks [2]
Cloud	Key ≤ 200ms, 99.95% uptime	KMS failure [79]
Cross-Layer	E2E ≤ 5s, $\gamma \geq$ 99.9%	Emergency delays [61]

Table 10: Layer-Specific Metrics and Limitations

10 Critical Limitations and Mitigations

Table 11 details the limitations and vulnerabilites of the components interacting among the layers of the frameworks, its operational impacts and technical strategies used to mitigate them.PBFT consensus mechanism along with sharding can handle large scale

distributed system. SGX attestation can move the security by trusting a client with verify the code and hardware.

Limitation	Impact	Mitigation
ZKP Vulnerabilities	Invalid proofs	Circom verification
PBFT Bottlenecks	$O(n^2)$ scaling	PBFT+sharding [20]
Cloud Centralization	Single failure	Threshold crypto [75]
FHIR Misuse	Data exploitation	Smart legal contracts [18]
Client Compromise	PII exposure	SGX attestation [33]

Table 11: Limitations and Countermeasures

11 Conclusion

This three-tiered framework advances privacy preserving healthcare data governance. However, reliance on cloud for off-chain storage and client trust worthiness is riskier.Performance Bottle necks is inherent PBFT consensus which would affect scalability, Existing gaps in FHIR interoperability, and ongoing regulatory ambiguity around healthcare data demand mitigation. Future priorities should include deployability testing, hybrid consensus validation, robust trust models, clinical integration. Shifting from cryptographic elegance to real-world operational reliability through iterative refinement is important.

12 Research Gaps and Future Work

Five critical gaps exist in ZKP-blockchain health care systems,despite the advances in the 81 studies, lacking real-world validation against legacy standards like HL7v2/FHIR conflicts .

• Emergency Latency: zk-SNARK proof generation time$T_{\text{gen}} = 2-15$s exceeds emergency latency thresholds ($T_{gen} > 100$ms) [3, 61] .

• Layer-2 Scaling: No zk-Rollup prototypes achieve over 10K Transaction Per Second (TPS) for healthcare [74, 20].

• Multi-Jurisdiction Compliance: Conflicts between GDPR and HIPAA primarily stems from differing requirements for specific issues and liability frameworks are untested [11, 47].

• Client IoT Vulnerabilities: Unpatched devices like MRI and FHIR semantic gaps SNOMED-LOINC present significant risks like misdiagnosis or exploitation. [1, 64]

• Hybrid Integration: There is a notable absence of a benchmark mechanism for PBFT consensus and ZKP end-to-end prototypes that validate against theoretical frameworks[72] Future work should concentrate on developing zk-optimized emergency circuits to reduce latency, sharded Rollups to improve scalability, jurisdiction- agnostic smart contracts to handle international compliance, and live benchmarks on FHIR-compliant IoT to validate real-world applicability is needed.

References

[1] Mohammed Amin Almaiah, Fahima Hajjej, Aitizaz Ali, Muhammad Fermi Pasha, and Omar Almomani. A novel hybrid trustworthy decentralized authentication and data preservation model for digital healthcare IoT-based CPS. *Sensors*, 22(4):1448, 2022.

[2] Salem Alqahtani and Murat Demirbas. Bottlenecks in blockchain consensus protocols. In *2021 IEEE International Conference on Omni-Layer Intelligent Systems (COINS)*, pages 1–8, 2021.

[3] Tianyu Bai, Yangsheng Hu, Jianfeng He, Hongbo Fan, and Zhenzhou An. Health-zkIDM: A healthcare identity system based on fabric blockchain and zero-knowledge proof. *Sensors*, 22(20):7716, 2022.

[4] S. M. H. Bamakan, Amirhossein Motavali, and Alireza Babaei Bondarti. A survey of blockchain consensus algorithms: performance evaluation criteria. *Expert Systems with Applications*, 154:113385, 2020.

[5] Masoud Barati, Gagangeet Singh Aujla, Jose Tomas Llanos, Kwabena Adu-Duodu, Omer Rana, Madeline Carr, and Rajiv Ranjan. Privacy-aware cloud auditing for GDPR compliance verification in online healthcare. *IEEE Transactions on Industrial Informatics*, 18(4):2605–2614, 2022.

[6] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Scalable, transparent, and post-quantum secure computational integrity, 2018. Cryptology ePrint Archive, Report 2018/046.

[7] J. B. Bernabe, J. L. Cánovas, J. L. Hernández-Ramos, R. T. Moreno, and A. Skarmeta. Privacy-preserving solutions for blockchain: Review and challenges. *IEEE Access*, page 164908, 2019.

[8] A. Bernier, F. Molnár-Gábor, and B. M. Knoppers. The international data governance landscape. *J. Law Biosci.*, 2022.

[9] P. Bhanupriya, S. Gauni, K. Kalimuthu, and C. T. Manimegalai. A modified hybrid blockchain framework for secured data transaction. *J. Phys.: Conf. Ser.*, 1964(4):042040, 2021.

[10] V. Botta, V. Iovino, and I. Visconti. Towards data redaction in bitcoin. *IEEE Trans. Network Serv. Manag.*, page 3872, 2022.

[11] L. Bradford, M. Aboy, and K. Liddell. International transfers of health data between the EU and USA: a sector-specific approach for the USA to ensure an 'adequate' level of protection. *Yearbook Med. Inform.*, 2020.

[12] S. Chaliasos, J. Ernstberger, D. Theodore, D. A. Wong, M. Jahanara, and B. Livshits. SoK: What don't we know? understanding security vulnerabilities in SNARKs, 2024. arXiv*.

[13] U. Chelladurai and S. Pandian. A novel blockchain based electronic health record automation system for healthcare. *J. Ambient Intell. Humanized Comput.*, page 693, 2021.

[14] S. Chenthara, K. Ahmed, H. Wang, F. Whittaker, and Z. Chen. Healthchain: A novel framework on privacy preservation of electronic health records using blockchain technology. *PLoS ONE*, 2020.

[15] Z. Deng. Navigating legal and ethical boundaries: Safeguarding personal health data in international research collaboration, 2024. Springer*.

[16] Z. ElSayed. Cybersecurity and frequent cyber attacks on IoT devices in healthcare: Issues and solutions, 2024. arXiv*.

[17] Nehal Ettaloui, Sara Arezki, and Taoufiq Gadi. An overview of blockchain-based electronic health record and compliance with GDPR and HIPAA. In *Artificial Intelligence and Smart Environment: Proceedings of AISE 2023*, volume 845 of *Lecture Notes in Networks and Systems*, pages 735–746. Springer, 2024.

[18] B. Fabio and R. Maddalena. From smart legal contracts to contracts on blockchain: An empirical investigation. *Comput. Law Secur. Rev.*, page 106035, 2024.

[19] H. Fatoum, S. Hanna, J. D. Halamka, D. C. Sicker, P. Spangenberg, and S. K. Hashmi. Blockchain integration with digital technology and the future of health care ecosystems: Systematic review. *J. Med. Internet Res.*, 2021.

[20] A. Gangwal, H. R. Gangavalli, and A. Thirupathi. A survey of layer-two blockchain protocols. *J. Network Comput. Appl.*, page 103539, 2022.

[21] R. Garg and H. M. R. Karim. Healthcare research data sharing and academic journal: A challenging but fruitful initiative, 2023. Indian Journal of Anaesthesia*.

[22] B. C. Ghosh and S. Chakraborty. Trustless collaborative cloud federation. *IEEE Trans. Cloud Comput.*, page 1, 2024.

[23] Priyanka Ghosh. *The State-of-the-Art in Zero-Knowledge Authentication Proof for Cloud*, pages 149–170. Wiley-Scrivener, 2021.

[24] M. Godyn, M. Kędziora, Y. Ren, Y. Liu, and H. Song. Analysis of solutions for a blockchain compliance with GDPR. *Sci. Rep.*, 2022.

[25] W. J. Gordon and C. Catalini. Blockchain technology for healthcare: Facilitating the transition to patient-driven interoperability. *Comput. Struct. Biotechnol. J.*, page 224, 2018.

[26] S. Han and S. Park. A gap between blockchain and general data protection regulation: A systematic review. *IEEE Access*, page 103888, 2022.

[27] A. Hasselgren, P. K. Wan, M. Horn, K. Kralevska, and D. Gligoroski. GDPR compliance for blockchain applications in healthcare, 2020. CSIT*.

[28] V. Holotescu and R. Vasiu. Challenges and emerging solutions for public blockchains. *BRAIN. Broad Res. Artif. Intell. Neurosci.*, page 58, 2020.

[29] R. R. Irshad, Shahab Saquib Sohail, S. Hussain, Dag Øivind Madsen, M. A. Ahmed, and A. A. Alattab. A multi-objective bee foraging learning-based particle swarm optimization algorithm for enhancing the security of healthcare data in cloud system. *IEEE Access*, page 113410, 2023.

[30] S. Judson and J. Feigenbaum. On heuristic models, assumptions, and parameters, 2022. arXiv*.

[31] B. Kaplan. Seeing through health information technology: the need for transparency in software, algorithms, data privacy, and regulation, 2020. Taylor & Francis*.

[32] J. H. Khor, M. Sidorov, and P. Y. Woon. Public blockchains for resource-constrained IoT devices—a state-of-the-art survey. *IEEE Internet Things J.*, page 11960, 2021.

[33] T. Knauth, M. Steiner, S. Chakrabarti, L. Li, C. Xing, and M. Vij. Integrating remote attestation with transport layer security, 2018. arXiv*.

[34] L. Krenn and D. Schlossman. Have electronic health records improved the quality of patient care? *PM&R*, 2017.

[35] Moez Krichen, Meryem Ammi, Alaeddine Mihoub, and Mutiq M. Almutiq. Blockchain for modern applications: A survey. *Sensors*, 22(14):5274, 2022.

[36] R. Lavin, X. Liu, H. Mohanty, L. E. J. Norman, G. Zaarour, and B. Krishnamachari. A survey on the applications of zero-knowledge proofs, 2024. arXiv*.

[37] M. Lehne, J. Saß, A. Essenwanger, J. Schepers, and S. Thun. Why digital medicine depends on interoperability. *npj Digit. Med.*, 2019.

[38] D. A. Luong and J. H. Park. Privacy-preserving blockchain-based healthcare system for IoT devices using zk-snark. *IEEE Access*, page 55739, 2022.

[39] J. Ma, S. Xu, J. Ning, X. Huang, and R. H. Deng. Redactable blockchain in decentralized setting. *IEEE Trans. Inform. Forensics Secur.*, page 1227, 2022.

[40] S. Ma and X. Zhang. Integrating blockchain and zk-rollup for efficient healthcare data privacy protection system via ipfs. *Sci. Rep.*, 2024.

[41] P. Mallozzi. Deploying zkp frameworks with real-world data: Challenges and proposed solutions. *arXiv*, 2023. arXiv*.

[42] J. C. Mandel, D. Kreda, K. D. Mandl, I. S. Kohane, and R. Ramoni. SMART on FHIR: a standards-based, interoperable apps platform for electronic health records. *J. Am. Med. Inform. Assoc.*, page 899, 2016.

[43] A. A. Mazlan, S. M. Daud, S. M. Sam, H. Abas, S. Z. A. Rasid, and M. F. Yusof. Scalability challenges in healthcare blockchain system—a systematic review. *IEEE Access*, page 23663, 2020.

[44] K. Miyachi and T. K. Mackey. hocbs: A privacy-preserving blockchain framework for healthcare data leveraging an on-chain and off-chain system design. *Inf. Proc. Manag.*, page 102535, 2021.

[45] M. T. D. Oliveira, L. H. A. Reis, Y. Verginadis, D. M. F. Mattos, and S. D. Olabarriaga. SmartAccess: Attribute-based access control system for medical records based on smart contracts. *IEEE Access*, page 117836, 2022.

[46] I. A. Omar, R. Jayaraman, K. Salah, M. C. E. Simsekler, I. Yaqoob, and S. Ellahham. Ensuring protocol compliance and data transparency in clinical trials using blockchain smart contracts. *BMC Med. Res. Methodol.*, 2020.

[47] Y. Piao, K. Ye, and X. Cui. A data sharing scheme for GDPR-compliance based on consortium blockchain. *Future Internet*, 13(8):217, 2021.

[48] E. Politou, A. Michota, E. Alepis, M. Pocs, and C. Patsakis. Backups and the right to be forgotten in the GDPR: An uneasy relationship. *Comput. Secur.*, page 1247, 2018.

[49] M. Qi, Z. Wang, Q. Han, J. Zhang, S. Chen, and Y. Xiang. Privacy protection for blockchain-based healthcare IoT systems: A survey. *IEEE/CAA J. Automat. Sin.*, page 1757, 2022.

[50] C. Qiu, H. Yao, F. R. Yu, C. Jiang, and S. Guo. A service-oriented permissioned blockchain for the internet of things. *IEEE Trans. Serv. Comput.*, page 1, 2019.

[51] D. Ramesh, R. Mishra, P. K. Atrey, D. R. Edla, S. Misra, and L. Qi. Blockchain based efficient tamper-proof EHR storage for decentralized cloud-assisted storage. *Alexandria Eng. J.*, page 205, 2023.

[52] Komala Rangappa, Arun Kumar Banavara Ramaswamy, M. R. Prasad, and Shreyas Arun Kumar. A novel method of secured data distribution using sharding ZKP and zero trust architecture in blockchain multi cloud environment. *Cryptography*, 8(3):39, 2024.

[53] T. Rocket, M. Yin, K. Sekniqi, R. van Renesse, and E. G. Sirer. Scalable and probabilistic leaderless BFT consensus through metastability, 2019. arXiv*.

[54] B. O. Roelink, M. El-Hajj, and D. K. Sarmah. Systematic review: Comparing zk-snark, zk-stark, and bulletproof protocols for privacy-preserving authentication. *Secur. Privacy*, 2024.

[55] I. Román, J. Calvillo-Arbizu, V. Mayor, G. Madinabeitia-Luque, A. Estepa, and R. Estepa. Blockchain-based service-oriented architecture for consent management, access control, and auditing. *IEEE Access*, page 12727, 2023.

[56] C. Ryngaert and M. Taylor. The GDPR as global data protection regulation? *AJIL Unbound*, page 5, 2020.

[57] H. Saidi, N. Labraoui, A. A. A. Ari, L. Maglaras, and J. H. M. Emati. DS-MAC: Privacy-aware decentralized self-management of data access control based on blockchain for health data. *IEEE Access*, page 101011, 2022.

[58] J. Scheibner, M. Ienca, S. Kechagia, J. R. Troncoso-Pastoriza, J. L. Raisaro, J. Hubaux, et al. Data protection and ethics requirements for multisite research with health data: a comparative examination of legislative governance frameworks and the role of data protection technologies. *J. Law Biosci.*, 2020.

[59] J. Scheibner, M. Ienca, and E. Vayena. Health data privacy through homomorphic encryption and distributed ledger computing: an ethical-legal qualitative expert assessment study. *BMC*, 2022.

[60] A. Sengupta and H. Subramanian. User control of personal mhealth data using a mobile blockchain app: Design science perspective. *JMIR mHealth uHealth*, 2021.

[61] R. Shashidhara, R. C. Nair, and P. K. Panakalapati. Promise of zero-knowledge proofs (ZKPs) for blockchain privacy and security: Opportunities, challenges, and future directions. *Secur. Privacy*, 2024.

[62] J. Shurson. Data protection and law enforcement access to digital evidence: resolving the reciprocal conflicts between EU and US law. *Int. J. Law Inform. Technol.*, page 167, 2020.

[63] A. A. Sınacı, M. Gençtürk, H. A. Teoman, G. B. L. Ertürkmen, C. Álvarez Romero, A. Martínez-García, et al. A data transformation methodology to create findable, accessible, interoperable, and reusable health data: Software design, development, and evaluation study. *J. Med. Internet Res.*, 2023.

[64] A. Torab-Miandoab, T. Samad-Soltani, A. Jodati, and P. Rezaei-Hachesu. Interoperability of heterogeneous health information systems: a systematic literature review. *BMC Med. Inform. Decis. Mak.*, 2023.

[65] F. N. da S. Vanin, L. M. Policarpo, R. da R. Righi, S. M. Heck, V. F. da Silva, J. R. Goldim, and C. A. da Costa. A blockchain-based end-to-end data protection model for personal health records sharing: A fully homomorphic encryption approach. *Sensors*, 23(1):14, 2022.

[66] N. R. Veeraragavan and K. Zhang. A position paper on GDPR compliance in sharded blockchains: rehash of old ideas or new interesting challenges?, 2020. arXiv*.

[67] P. Voigt and A. von dem Bussche. *Scope of Application of the GDPR*, pages 9–30. Springer, Cham, 2017.

[68] S. R. Vulapula and M. Srinivas. Attribute-based encryption for fine-grained access control on secure hybrid clouds. *Int. J. Adv. Comput. Sci. Appl.*, 2020.

[69] I. Weber, Q. Lu, A. B. Tran, A. Deshmukh, M. Gorski, and M. Strazds. A platform architecture for multi-tenant blockchain-based systems. In *2019 IEEE International Conference on Software Architecture Companion (ICSA-C)*, pages 75–82. IEEE, 2019.

[70] J. S. Winter and E. Davidson. Harmonizing regulatory regimes for the governance of patient-generated health data. *Telecomm. Policy*, page 102285, 2021.

[71] S. Xu, X. Cai, Y. Zhao, Z. Ren, L. Du, Q. Wang, and J. Zhou. zkrpchain: Towards multi-party privacy-preserving data auditing for consortium blockchains based on zero-knowledge range proofs. *Future Gener. Comput. Syst.*, 135:84–98, 2022.

[72] J. Yang, Z. Jia, R. Su, X. Wu, and J. Qin. Improved fault-tolerant consensus based on the PBFT algorithm. *IEEE Access*, page 30274, 2022.

[73] Weiping Ye, J. Wang, H. Tian, and H. Quan. Public auditing for real-time medical sensor data in cloud-assisted HealthIIoT system. *Front. Optolectron.*, 2022.

[74] H. Yu, I. Nikolić, R. Hou, and P. Saxena. OHIE: Blockchain scaling made simple. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 90–106, 2020.

[75] D. Zhang, J. Le, X. Lei, T. Xiang, and X. Liao. Secure redactable blockchain with dynamic support. *IEEE*, page 717, 2023.

[76] X. Zhang, C. Huang, Y. Zhang, J. Zhang, and J. Gong. LDVAS: Lattice-based designated verifier auditing scheme for electronic medical data in cloud-assisted WBANs. *IEEE Access*, page 54402, 2020.

[77] Y. Zhang, M. Wang, Y. Guo, and F. Guo. Towards dynamic and reliable private key management for hierarchical access structure in decentralized storage. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management (CIKM)*, pages 3371–3380, 2023.

[78] J. Zhao, W. Wang, D. Wang, X. A. Wang, and C. Mu. PMHE: a wearable medical sensor assisted framework for health care based on blockchain and privacy computing. *J. Cloud Comput.*, 11(1):96, 2022.

[79] C. Zhou, M. Barati, and O. Shafiq. A compliance-based architecture for supporting GDPR accountability in cloud computing. *Future Gener. Comput. Syst.*, page 104, 2023.

[80] X. Zhou, A. Nehme, V. Jesus, Y. Wang, M. B. Josephs, K. Mahbub, and A. E. Abdallah. AudiWFlow: Confidential, collusion-resistant auditing of distributed workflows. *Blockchain Res. Appl.*, page 100073, 2022.

[81] Y. Zhuang, C.-R. Shyu, Shenda Hong, P. Li, and L. Zhang. Self-sovereign identity empowered non-fungible patient tokenization for health information exchange using blockchain technology. *Comput. Biol. Med.*, page 106778, 2023.