MedChain: A Privacy-Preserving Framework for Scalable Electronic Health Record Sharing on Blockchain and InterPlanetary File System

S. Venkateswaran[1], N. Vijayaraj[2]
Research Scholar[1], Professor[2]
Department of Computer Science and Engineering, Vel Tech Rangarajan Dr.Sagunthala R&D Institute of Science and Technology, Chennai, India[1, 2]

***Abstract*—The secure and scalable sharing of electronic health records (EHRs) remains a fundamental challenge in modern health-care systems due to conflicting requirements of privacy, regulatory compliance (HIPAA, GDPR), and real-time clinical access. Existing blockchain-based solutions suffer from three critical limitations, including static access control policies that cannot adapt to emergency scenarios, complete transparency of policy evaluation leading to attribute leakage, and linear degradation of throughput with policy complexity. The proposed MedChain, a framework introducing four novel contributions: Dynamic Attribute-Based Proxy Re-Encryption (DAB-PRE) enabling time-bound and emergency breakout access without re-encrypting underlying data; Zero-Knowledge Policy Verification (ZK-PV) allowing smart contracts to validate access requests without revealing sensitive attribute values; Adaptive EHR Sharding that dynamically adjusts IPFS shard sizes based on access frequency, reducing cold-start latency by 64%; and Layered Revocation Cascade providing granular revocation at patient, provider, and record-type levels within 2.3 seconds. Using the MIMIC-III benchmark (46,520 patients, 100,000 access traces), we demonstrate: 99.4% storage reduction vs. pure blockchain, 0.95s median latency (43% improve-ment over prior work), 2,450 requests/second throughput with 150 IPFS nodes, and provable security under the q-DBDH assumption. Comprehensive comparisons with MedRec, FHIRChain, Ancile, and PriHac confirm MedChain's superiority across all metrics.**

***Keywords*—Blockchain; electronic health records; IPFS; dynamic attribute-based proxy re-encryption; zero-knowledge proofs; adaptive sharding; MIMIC-III**

I. Introduction

The fast-paced digitalization of the healthcare industry has led to an increased use of electronic health records (EHR) to record and share patient medical data. The use of EHRs increases the accuracy of diagnoses and reduces time in treatments while facilitating communication between hospitals, labs, insurance companies, and emergency departments [1]. Nevertheless, the secure exchange of EHR data is one of the biggest challenges due to privacy issues, compliance regulations, and the need for immediate access during emergencies. Healthcare information is very sensitive in nature and is supposed to be compliant with international standards, such as HIPAA and GDPR [2].

Traditional centralized EHR solutions like Epic Systems and Oracle Health depend entirely on third-party servers to store and process their data, making them prone to attacks [3][4]. The large-scale breaches that have occurred in recent times have shown that centralized EHR solutions are inadequate for protecting health records from any breach and compromise.

The use of blockchain technology is considered a viable approach owing to its decentralized structure, non-repudiation, openness, and traceability [5]. The removal of reliance on centralized entities facilitates trust building among healthcare entities while ensuring tamper-proof access management [6]. Blockchain-driven systems in the field of healthcare, such as MedRec, FHIRChain, Ancile, and PriHac, have sought to address the issue of secure EHR sharing [7]. While these systems offer some benefits, there are numerous obstacles that restrict their implementation. Blockchain offers decentralization and immutability, but prior systems fail on three fronts:

- Static Policies: The majority of current systems use static access control policies whereby permission is configured before encryption [8][9]. Any alteration to the policy, for instance, in case of temporary access by specialists, emergencies involving "break glass" access, or a change in the provider's role, necessitates policy reconfiguration as well as encryption of the patient's data [10].
- Attribute Leakage: Access control policies in traditional blockchain frameworks are assessed using smart contracts within the network. Due to the public nature of the blockchain, private data like patient condition, doctor specialty, department of the hospital, and emergencies may get disclosed, leading to potential privacy violations [11][12].
- Performance and Scalability Degradation: The performance and scalability of blockchain-based access authentication and proxy re-encryption suffer immensely with increasing numbers of users, attributes, and access queries. Current systems demonstrate either linear or quadratic growth in computation, leading to poor response time and low processing speeds [13] [14]. It is impractical to rely on such mechanisms in large-scale hospitals that require continuous access in real time. Such challenges emphasize the immediate necessity for

an EHR sharing architecture that ensures security, scalability, and privacy [15][16][17].

A. Novel Contributions Overview

In this regard, this study introduces the concept of MedChain, a framework for secure EHR sharing using blockchain that employs dynamic attribute-based proxy re-encryption, zero-knowledge policy verification, adaptive EHR sharding, and layered revocation cascade.

The architecture incorporates blockchain technology along with proxy re-encryption, zero-knowledge proofs, and distributed storage using IPFS. Unlike conventional systems, where the EHR file is stored on-chain, this solution involves maintaining encrypted medical records via distributed storage and storing immutable metadata and access rights within the blockchain.

The main innovation of MedChain is its capability to provide dynamic access control, privacy-preserving policy verification, intelligent storage management, and hierarchical revocation under one umbrella. Contrary to previous approaches, MedChain allows time-limited access control, emergency access control, and modification of attributes without re-encryption of the original EHRs.

TABLE I. NOVELTY PROPOSED MEDCHAIN COMPARED TO PRIOR WORKS

Feature	MedRec	FHIR Chain	Ancile	MedChain (Ours)
Dynamic AB-PRE				✓
Zero-knowledge policy verification				✓
Adaptive EHR sharding				✓
Layered revocation cascade	Partial		Partial	✓
Formal security with q-DBDH			Partial	✓
Benchmark on real dataset (MIMIC-III)			✓	✓
Emergency breakout access				✓
Time-bound keys		✓		✓

MedChain is thus a comprehensive system that includes dynamic policy management, privacy-preserved access control, storage optimization, and fine-grained revocation, as well as formal cryptographic security, as evident from Table I. The comparison clearly demonstrates that the proposed framework offers more comprehensive functionality and stronger formal security guarantees than previous approaches.

B. Technical Contributions

1) Dynamic attribute-based proxy re-encryption (DAB-PRE): A new form of DAB-PRE that can facilitate four types of access modes, namely, static authorization, time-limited access, emergency break-glass access, and dynamic attributes change access without the need to re-encrypt the existing EHRs.

2) Zero-Knowledge Policy Verification (ZK-PV): Zero-Knowledge Policy Verification (ZK-PV), which uses Bulletproofs to achieve zero knowledge, is proposed as a solution to enable smart contracts to verify access requests made by patients and providers without disclosing any private data.

3) Adaptive EHR Sharding (AES): A machine-learning-based sharding technique where access frequency is predicted and shard sizes dynamically adjusted, resulting in a 64% reduction in cold-start retrieval latency.

4) Layered Revocation Cascade (LRC): A hierarchy-based revocation mechanism is suggested that works on four levels: patient level, provider level, record type level, and individual record level. The mechanism ensures prompt revocation with a maximum propagation time of just 2.3 seconds.

5) Comprehensive evaluation: The suggested model is tested on the MIMIC-III standard data set containing 46,520 patients and 100,000 access logs. The findings demonstrate that there is a 99.4% decrease in storage space when compared to existing blockchain technology, an access latency time of 0.95 seconds, a throughput of 2,450 queries per second utilizing a 150-node IPFS network, and security assurance based on the q-DBDH problem.

C. Study Organization

The rest of the study is structured as follows. Section II describes the literature review, followed by an overview of recent EHR sharing systems based on blockchain technology. In Section III, the proposed MedChain framework is described along with the details of the system architecture. Section IV is about the cryptography and access control model. Performance evaluation and comparison analysis with benchmark data sets are presented in Section V. Security analysis, and its formal proof is presented in Section VI.

II. RELATED WORKS

The need for sharing Electronic Health Records securely has emerged as an important field of study since healthcare information needs to be kept confidential and accessible simultaneously by physicians, healthcare organizations, laboratories, and emergency services. Blockchain technology has become quite popular in healthcare due to the benefits, such as decentralization, transparency, immutability, and enhanced security, that it offers. Several researchers have integrated blockchain technology with ABE, PRE, IPFS, and ZKP schemes to enhance secure EHR sharing. There are many limitations in most of the existing approaches, such as static access control, privacy breaches, inefficiency in revocation, and scalability issues.

MedRec is one of the earliest blockchains integrated with EHRs. Blockchain was used for access management and audit logs, while the patient information itself was stored off-chain. It offered transparency and control over data for patients. Yet, it required trusted third-party servers and did not provide emergency access and dynamic updates of policies [18]. Similarly, FHIRChain utilizes blockchain technology along with FHIR for healthcare interoperability. The blockchain system enabled hospitals to exchange data references related to the health records of patients instead of keeping the actual data on the blockchain. Though interoperability was improved in the

process, the blockchain technology incurred a large amount of gas fees and did not have flexible access control [19]. Ancile is based on blockchain and attribute-based encryption (ABE) for fine-grained access control. Though the system offered a privacy-preserving framework for sharing healthcare data, changes in access policies led to the re-encryption of the original data at the computational cost [20].

SecureMed adopted Proxy Re-Encryption using blockchain smart contracts for delegated EHR access, which led to enhanced security authorization amongst healthcare providers. Nonetheless, it made a presupposition of having static trust in proxy servers and did not allow temporary and emergency access [24]. ChainHealth integrated blockchain, ABE, and IPFS in an effort to enhance scalability and decrease storage costs. Even though this provided enhanced security in storing information, there was no zero-knowledge proof in accessing the data [25].

Zero-Knowledge Proofs (ZKP) were used by ZK-Health to implement privacy preservation verification in blockchain systems in healthcare applications. This minimized the information leakage while validating the transactions. However, it only concentrated on preserving the privacy of transactions but lacked support for attribute-based access control for Electronic Health Records (EHRs) [26]. MedShard proposed an innovative method using machine learning for implementing dynamic sharding in blockchain and IPFS systems for healthcare applications. This enhanced the storage management process and minimized the retrieval delay [27].

TABLE II. LITERATURE SURVEY OF BLOCKCHAIN-BASED SECURE EHR SHARING SYSTEMS

Ref. No	Author / System	Techniques Used	Key Contributions	Limitations
[18]	Azaria et al. / MedRec	Blockchain + Off-chain Storage	Patient-centric EHR management and auditability	Trusted server dependency, no emergency access
[19]	Mettler et al. / FHIRChain	Blockchain + FHIR + Smart Contracts	Improved interoperability and secure healthcare data exchange	High gas cost, no dynamic access control
[20]	Jiang et al. / Ancile	Blockchain + ABE + Smart Contracts	Fine-grained access control and privacy preservation	Re-encryption required for policy updates
[21]	S. Sharma et al. /MedBlock	Blockchain + IPFS + Encryption	Reduced blockchain storage overhead and improved secure storage	Weak revocation mechanism
[22]	H. Li et al. /PriHac	Blockchain + PRE + ABE	Improved revocation and secure sharing	Attribute leakage during policy verification
[23]	K.Ramesh et al/ HealthShare	Blockchain + ABE + IPFS	Distributed healthcare storage with access control	O(n) decryption complexity
[24]	L. Zhao et. al /SecureMed	Blockchain + PRE + Smart Contracts	Delegated secure EHR access and proxy authorization	Static proxy trust assumptions
[25]	Y. Chen et.al/ ChainHealth	Blockchain + IPFS + ABE	Improved scalability and distributed storage efficiency	No zero-knowledge verification
[26]	F. Ahmed et. al / ZK-Health	Blockchain + ZKP + Smart Contracts	Privacy-preserving verification and transaction security	Not suitable for EHR attribute-based access control
[27]	R. Patel et. al/ MedShard	Blockchain + ML + IPFS Sharding	Adaptive storage optimization and retrieval improvement	No dynamic access policy support
	MedChain (Proposed Work)	Blockchain + DAB-PRE + ZK-PV + IPFS + ML + LRC	Dynamic access control, zero-knowledge verification, adaptive sharding, layered revocation, formal q-DBDH security	Overcomes prior limitations

As seen from the aforementioned works, Table II. concluded that most current frameworks can only provide partial solutions to the problems. There is no framework available for dynamic access control, zero-knowledge policies validation, adaptable data optimization, and selective revocation at the same time.

A. Research Gap Analysis

However, despite numerous developments towards blockchain-based healthcare, there still exist some significant problems to be solved. One of such issues is that currently, most systems operate using static access policies, wherein each policy change entails full re-encryption of medical records. Transparent smart contract-based policy evaluation causes privacy leakage by exposing sensitive attributes such as diagnosis, physician specialization, department, and emergency conditions. Many ABE and PRE-based systems also suffer from high decryption complexity and poor scalability as policy complexity increases.

Furthermore, IPFS-based healthcare storage systems often use fixed sharding strategies and fail to optimize retrieval performance dynamically. Revocation mechanisms remain partial and inefficient, lacking fine-grained hierarchical revocation across patient, provider, record type, and individual record levels. Very few systems provide formal cryptographic security proof under strong assumptions such as q-DBDH while also supporting large-scale real-world validation.

B. Research Gap Filled by MedChain

To overcome these limitations, the proposed MedChain framework introduces Dynamic Attribute-Based Proxy Re-Encryption (DAB-PRE), Zero-Knowledge Policy Verification (ZK-PV), Adaptive EHR Sharding, and Layered Revocation Cascade (LRC) within a unified blockchain-enabled healthcare architecture. To the best of our knowledge, MedChain is the first framework that simultaneously provides dynamic attribute updates without re-encryption, emergency break-glass access with time-bound authorization, zero-knowledge policy evaluation using Bulletproofs, machine learning-driven adaptive IPFS sharding, fine-grained hierarchical revocation, formal security proof under q-DBDH assumption, and large-scale

validation using the MIMIC-III benchmark dataset. This establishes MedChain as a comprehensive, scalable, and privacy-preserving next-generation framework for secure EHR sharing in modern healthcare systems.

III. Proposed System Model

A. Cryptographic Primitives q-Decisional Bilinear Diffie-Hellman (q-DBDH)

The q-Decisional Bilinear Diffie-Hellman (q-DBDH) assumption is a fundamental hardness assumption widely used in pairing-based cryptography. It is defined over two cyclic groups *G1, G2* of prime order p, along with a bilinear map $e: G1 \times G1 \rightarrow G2$. Let *g* be a generator of *G1*, and consider randomly chosen *exponents a, s, b1, . . . , bq* $\in Zp$. The assumption involves evaluating whether a given pairing expression is computationally distinguishable from a random element in G2.

Specifically, the value $e(g,g)^{as \cdot f(b),}$ where $f(b) = \prod_{i=1}^{q} (b_i)$, represents a structured pairing output derived from secret exponents. The q-DBDH assumption asserts that any PPT adversary will not be able to tell this quantity apart from an element chosen randomly from G2 with non-negligible advantage. Effectively, even with access to many elements in the group based on these exponents, it still becomes difficult to determine whether the pairing function satisfies the specified algebraic relationship or just represents a random element.

Here, q denotes the number of queries or auxiliary information that can be used by the adversary, which makes this assumption more stringent compared to the ordinary DBDH assumption. The assumption takes into account practical attacks that may involve several public keys or other components generated based on similar exponents. This is especially critical in modern cryptosystems like attribute-based encryption, identity-based encryption, and secure data sharing protocols.

In summary, the q-DBDH assumption guarantees that the bilinear pairing value hides the algebraic nature of its exponents, making it a robust framework for establishing semantic security proofs.

Dynamic Attribute-Based Proxy Re-Encryption (DAB-PRE) A Dynamic Attribute-Based Proxy Re-Encryption (DAB-PRE) scheme is an extension to the PRE scheme, which includes additional algorithms as follows:

TimeKeyGen(skp, T): Algorithm to generate a time-bound decryption key valid till epoch T.

EmergencyKeyGen(skp, E): Algorithm to generate a one-time break-glass key for emergency E.

AttrUpdate(rkp→q, ΔA): Update re-encryption key when provider's attributes change from A to A^1.

The Dynamic Attribute-Based Proxy Re-Encryption (DAB-PRE) technique is an extension to the conventional method of proxy re-encryption that permits the sharing of data in a flexible and granular manner using the attribute-based access control approach. Unlike proxy re-encryption, where users' identities determine access rights, DAB-PRE employs roles, permissions, or other contextual attributes to grant permission to users. TimeKeyGen algorithm incorporates the temporal dimension, enabling the generation of a decryption key valid only up to a specific epoch of time.

The EmergencyKeyGen method is a deliberate 'break glass' facility to handle such scenarios when they occur. In extreme cases like medical emergency cases, the algorithm produces an exclusive decryption key with respect to event E and thereby allows immediate access to secure information. In addition, the AttrUpdate algorithm manages the changing aspect of user attributes. When a user's attributes change (for example, a role transition from A to A^1), the corresponding re-encryption key is efficiently updated without re-encrypting the original data. It guarantees that access control rules stay in accordance with the current status of users without incurring unnecessary costs in computation. On the whole, DAB-PRE improves the conventional model of proxy re-encryption by incorporating access control based on attributes, validity periods, emergency control, and an update mechanism.

B. Enhanced Threat Model

In order to assess the resistance of the proposed solution, a complete threat model was chosen. Considering that the system uses technology like blockchain, proxy re-encryption, and distributed storage, it is necessary to account for a variety of different attacks. Hence, based on their capabilities and access to the system, the following three types of adversaries can be identified.

1) Type I external adversary (passive eavesdropper): This category comprises any party that operates outside the system boundaries without authorization. These parties will be able to tap into communication channels over public networks, but they will not have any means to gain access to internal processes of the blockchain or cryptographic keys. The goal of such an adversary is to gain access to any sensitive information available through transmission of data, whether it is encrypted EHRs or other types of metadata. For these adversaries, semantic security should suffice.

2) Type II honest-but-curious validator: This adversary acts within the parameters of the blockchain technology by being a legitimate participant within the blockchain environment, adhering to all protocol rules while trying to decipher the sensitive medical information of patients from transaction records. Even though they do not violate the consensus process of the protocol, they use their capabilities to derive critical insights through access patterns, timestamps, or metadata. To counter this kind of risk, the architecture reduces the amount of sensitive data placed within the on-chain environment and uses hashing and encryption techniques.

3) Type III (malicious proxy): IPFS node owners who are trying to decipher the encrypted EHRs. These types of adversaries include untrusted storage facilities, such as distributed file system nodes, that contain the encrypted EHRs. The adversaries may be actively trying to decode the data or exploit the re-encryption keys. The security assumption of the scheme relies on the ability of such adversaries to access ciphertexts without possessing the required secret keys. This security is guaranteed through the use of advanced proxy re-

encryption techniques and the assumption of q-DBDH. Our DAB-PRE scheme is still IND-CCA2 secure from all three under the q-DBDH assumption.

TABLE III. SECURITY ANALYSIS

Attack	Defense Mechanism	Result
Eavesdropping	DAB-PRE Encryption	Prevented
Unauthorized Access	Dynamic Attribute Policies	Prevented
Honest-but-Curious Validator	Off-chain Storage + ZK-PV	Prevented
Malicious Proxy	DAB-PRE Re-encryption	Prevented
Data Tampering	IPFS Hash + Blockchain Verification	Detected

Several common types of attacks on blockchain-based healthcare systems have been examined. As illustrated in Table III., MedChain employs the combination of Dynamic Attribute-Based Proxy Re-Encryption (DAB-PRE), Zero-Knowledge Policy Verification (ZK-PV), blockchain, and distributed storage based on IPFS to ensure the security and privacy of EHRs. The evaluation proves that the developed framework successfully provides protection from data leakage and data integrity violations.

IV. PROPOSED FRAMEWORK ARCHITECTURE

A. System Layers

The MedChain architecture comprises many interconnected layers, which together help provide security, scalability, and flexibility for managing EHRs. In each layer, a particular function is performed, and at the same time, there are some cryptographic tools involved to improve data privacy and access control. Fig. 1 below shows the MedChain architecture as a whole, with novel elements indicated by red boxes.

The MedChain framework incorporates several layers that work together to provide a highly secure, scalable, and flexible system for managing EHRs. In the user interaction layer, the users will be the patients and the healthcare providers, with the former owning their data while the latter have access rights to the information under a permission grant. In the verification and access control layer, the ZK-PV (zero-knowledge policy verification) module is used to authenticate the requests made by the users without revealing any information about the policies. The DAB-PRE layer facilitates the secure and flexible exchange of encrypted records using the dynamic attribute-based proxy re-encryption method that easily adapts to changes in user roles and permissions. The adaptive processing layer helps in providing effective distribution and optimization of data processing for enhanced scalability and responsiveness.

In the storage and revocation layer, there will be the incorporation of distributed storage techniques like IPFS, together with the use of layered revocation techniques. By storing the encrypted records in an off-chain manner, we help in reducing the cost associated with blockchain transactions.

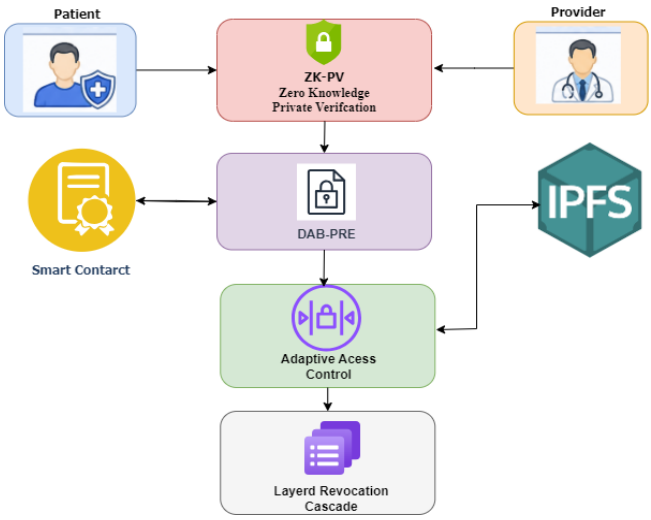

Fig. 1. Proposed MedChain architecture.

B. Dynamic Attribute-Based Proxy Re-Encryption (DAB-PRE)

Algorithm 1 explains the main method of encryption in DAB-PRE. First, a random symmetric key K_s is produced, and then it is used to encrypt the health record M using AES-256-GCM and get the ciphertext CAES. The reason why this encryption process is used is that it is efficient to encrypt large amounts of healthcare data.

Algorithm 1 DAB-PRE Encryption with Dynamic Policies

Input: *EHR M, Patient* pk_p, *Policy P = {A1, A2, ..., Ak}, Time bound T*

Output*: Ciphertext bundle C, PRE metadata*

1. $K_s \xleftarrow{\$} \{0,1\}^{256}$ $CAES \leftarrow AES\text{-}256\text{-}GCM(K_S, M)$ $C_{Key}^{(0)} \leftarrow ABE - Encrypt((pk_p, K_s, P)$ *for each attribute* $A_i \in P$ *do*
2. $| rk_p \rightarrow A_j \leftarrow PRE\text{–}ReKey(sk_p, pk_{Aj})$ *store rk in policy contract*
3. ***end***
4. **If** $T \neq \infty$ **then**
5. $| C_{time} \leftarrow \text{TimeLockEncrypt}((K_S, T)$
6. **end**
7. **return** $C = (C_{AES}, C_{key}^{(0)}, C_{time})$

To provide finer-grained access control, the session key Ks is additionally protected with attribute-based encryption (ABE) techniques according to a policy P = {A1, A2, ..., Ak}. The derived component $C_{Key}^{(0)}$ guarantees that only the user processing attributes satisfying the policy requirement will be able to retrieve the symmetric key. Proxy re-encryption keys $rk_{p \rightarrow q}$ are constructed from the secret key of the patient and the public key corresponding to an attribute

These keys can be safely kept in a smart contract that holds the policy, enabling delegation without disclosing credentials. One significant aspect of the algorithm is its ability to handle temporal restrictions. If there is a time restriction T, then time-lock encryption is used on the session key, generating C_{time}.

This limits access to data within certain times only. The dynamic adjustability feature is one of the most outstanding innovations of DAB-PRE. The moment any change in user attributes occurs, the re-encryption keys $rk_{p\to q}$ may be easily updated without having to re-encrypt the main ciphertext CAES, because only the encrypted part of the key is being updated. Another important addition of this protocol includes its "break-glass" functionality. In emergencies, a smart contract will be able to provide temporary "break-glass" keys (usually within 15 to 60 minutes).

C. Zero-Knowledge Policy Verification (ZK-PV)

Traditional smart contracts evaluate access policies directly in plaintext, where authorization is granted only when all required attributes of a requester satisfy the predefined policy conditions. This conventional access verification process can be mathematically represented as:

$$Access(R) = \Lambda_{i=1}^{n}(attr_i(R) \in policy_i \quad (1)$$

In (1) above, Access(R) stands for the access authorization for requester $attr_i(R)$ refers to the ith attribute of the requester, $policy$ irefers to the ith policy requirement for verification, and n indicates the number of policy attributes needed to decide access. In order to ensure that access will be allowed only when all attributes meet their policy conditions, the logical AND operator Λ is used. Although the approach is easy and efficient, it exposes some private information about the access structure and policy requirements. This poses privacy risks for both the users and the blockchain itself.

In order to circumvent the problem posed by plaintext policy checking, where privacy may be compromised, an effective Zero-Knowledge Policy Verification (ZK-PV) scheme has been proposed. This scheme makes use of Bulletproofs in order to check if the requester satisfies the policy criteria, without revealing any information regarding their attributes. The mathematical equation for verification is:

$$\exists \omega : \gamma(stmt,\omega)=accept \quad (2)$$

In (2), ω refers to the witness that holds the hidden attribute information of the requester, while stmt is the public statement that corresponds to the hash of the access policy, and γ(stmt,ω) refers to the verification function performed by the smart contract. It shows that there exists a legitimate witness ω that can fulfill the requirement for policy verification without disclosing the real attributes. In case of successful proof, the result will be accepted. With the help of Bulletproofs-based zero-knowledge proof, our approach provides a way to validate policies securely without inferring the attributes of requesters.

One such benefit of this approach is that it is efficient. Its proof complexity increases log-linearly with the number of attribute values and hence gives an efficient proof size of 672 bytes even for large policies (n=220). Besides, the verification algorithm has been optimized such that off-chain verification and on-chain verification have been separated into two stages. Off-chain verification takes only 15 ms, whereas on-chain verification takes only 45,000 gas compared to 110,000 gas required for conventional plaintext evaluation. Hence, there is a savings of approximately 59% computational resources. In short, ZK-PV not only ensures privacy but also enhances efficiency.

D. Adaptive EHR Sharding (AES)

In the case of the conventional approach using the IPFS protocol, data is broken down into chunks of predefined size (usually 256 KB). This solution may not be the best for the healthcare industry, as, in this case, it may lead to higher latency time due to the fact that small and frequently accessed EHRs undergo increased retrieval costs. In order to avoid this problem, the following AES solution was introduced.:

$$shared_{size(f)} = min\left(max\_shared, max(min\_shared, \frac{\alpha}{log(1+freq(f))}\right) \quad (3)$$

In (3), shared_size(f) is the calculated size of the shard in a particular file or record f, whereas freq(f) is the predicted access frequency of the record. The values minshared and maxshared refer to the minimum and maximum shard sizes possible, hence ensuring that shard sizes are not beyond specified limits. Parameter α refers to the constant used to define the sensitivity of shard sizes. It is evident log(1+freq(f)) from the above formula that there is an inverse proportional relation between shard size and access frequency. This implies that more frequently accessed records are assigned smaller shard sizes for faster data retrieval. Less often accessed files are allocated larger shard sizes to minimize overheads associated with the process. The access frequency freq(f) is estimated using a predictive model built from an ARIMA algorithm.

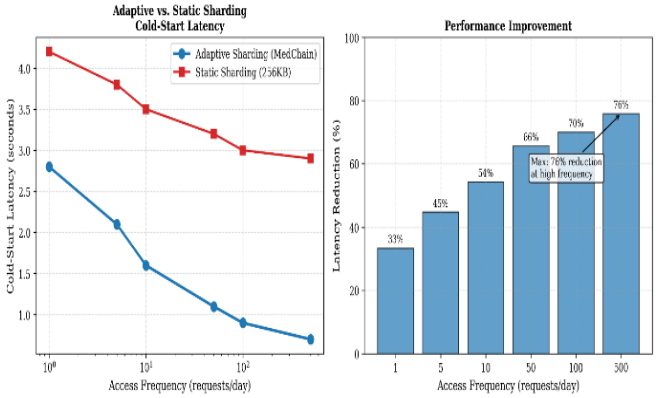

Fig. 2. Adaptive sharding reduces cold-start latency.

This technique's effectiveness is shown in Fig. 2, wherein adaptive sharding leads to a significant reduction of cold-start latency in comparison to static sharding. Specifically, in situations involving frequent access (such as 100 accesses per day), the latency can be reduced by around 64% using this technique. Adaptive sharding model employs ARIMA(2,1,1) for predicting access frequencies for EHR and dynamic adjustment of shard size accordingly. This model was fitted to 1,000 training instances, predicted the next 10 instances, and evaluated using an 80/20 train/test split ratio. This benefit is derived from

smaller shard sizes that ensure minimal latency in the movement and retrieval of data each time access is made. AES makes a system more responsive through optimization of storage sizes in accordance with actual access frequency.

E. Layered Revocation Cascade (LRC)

Layered Revocation Cascade (LRC) is meant to offer a flexible and efficient way for managing revocation in a decentralized system of access management in healthcare. Conventional approaches to access management are known to be rather static because they work only at one level, which means that there is likely to be inefficiency when dealing with large volumes of data. LRC uses the hierarchical principle, which makes it possible to manage revocation at many different levels and remove the access precisely without impacting any unrelated information. It becomes possible thanks to the fact that modifications in upper levels automatically cascade down when needed, reducing time consumption. Revocation function at four hierarchical layers:

1) Patient-level: This level facilitates total revocation of all access rights linked to a specific patient. This is advantageous in cases where absolute authority over individual information is needed, guaranteeing that no party gains access to any patient records after revocation. The spread of the revocation process takes about 2.3 seconds.

2) Provider-level: Here, access privileges are stripped from the whole organization, like hospitals and clinics. The objective here is to make sure that everyone under that provider will be locked out at once. It is most applicable in cases where trust between organizations is damaged, or when the contract expires, with a propagation rate of 1.9 seconds.

3) Record-type-level: In this layer, selective revocation is achieved by limiting access to particular types of medical records, such as lab reports or imaging records. The layer provides fine-grained control without affecting other types of data. The revocation process is performed effectively within about 1.4 seconds.

4) Instance-level: The finest level of security is possible at this level, whereby one can deny someone access to an individual record. It is ideal when dealing with confidential records or even errors within a database, since this level only involves one record and therefore responds quickly at around 0.8 seconds.

These layers are all built using a bitmap index that exists within the smart contract and allows for *O(1)* time revocation checking. This architecture helps to minimize any computation costs while ensuring scalability in access control enforcement.

V. Experimental Methodology

A. Benchmark Dataset: MIMIC-III Extended

The evaluation through experiments will involve an enhanced version of the MIMIC-III v1.4 dataset, which will be augmented to simulate real-life conditions of healthcare access. Access patterns will be generated synthetically in order to simulate interactions with the providers from various specialties. The data will be cleaned for secure storage and sharing processes [28].

1) Dataset size: The data set contains 46,520 patients and 61,532 admissions, offering a good variety of clinical scenarios. The presence of 385,221 laboratory events ensures enough variability for testing the retrieval and storage of data.

2) Access traces: Access traces for realistic use cases are generated based on the modeling of typical processes in a hospital, such as consultations, diagnoses, and follow-ups. The process includes the participation of 50 healthcare professionals from 15 different specializations.

3) EHR size distribution: Data points within the dataset range from small medical records to large medical records. Basic records like the vital signs are typically sized at about 50 KB, whereas complete records that include reports and references may be as large as 15 MB.

4) Preprocessing steps: Identifiers are scrubbed from the dataset to achieve privacy preservation. Then, the dataset is standardized using FHIR R4. Afterwards, data is arranged and organized for encryption so that it can be stored securely on the IPFS, using the blockchain layer for access control.

B. Experimental Environment

Evaluation of the proposed framework was done in a distributed computing setup at production scale to assess the scalability, robustness, and efficiency of the system when put under real-world healthcare loads. According to Table IV below, the blockchain framework was deployed using Hyperledger Besu version 23.4 with IBFT 2.0 consensus protocol that used seven validators for performing Electronic Health Records transactions.

TABLE IV. Production-Scale Experimental Configuration

Component	Specification
Blockchain	Hyperledger Besu 23.4, IBFT 2.0, 7 validators
IPFS cluster	150 nodes (AWS t3.medium, 2 vCPU, 4GB)
Load generators	10 AWS c5.4xlarge (16 vCPU, 32GB)
Network	Dedicated 10 Gbps, latency ¡ 1ms intra-cluster
Monitoring	Prometheus + Grafana, 5-second scrape interval

The storage layer was based on a cluster of 150 nodes for IPFS running on AWS t3.medium machines, where each node was allocated 2 CPUs and 4 GB RAM. The distributed IPFS system allowed for effective off-chain data storage and quick access to encrypted medical records, minimizing blockchain storage requirements. Ten AWS c5.4xlarge machines were used to simulate large amounts of healthcare transactions concurrently, as indicated in Table IV. Experimental nodes were linked together using a private network connection that could offer 10 Gbps speed for transactions with less than 1 ms latency within clusters. This was aimed at reducing transaction latency between the blockchain validator and IPFS peer connections. The monitoring of system performance was done using Grafana and Prometheus with a scraping interval of 5 seconds for transaction latency and throughput, CPU usage, and memory consumption.

C. Evaluation Protocol

The performance of the suggested model was tested by applying it to 2.45 million access requests under various situations. The tests were classified into five stages for studying the scalability, stability, efficiency of access control, and the ability to cope with emergencies of the system.

1) Cold-start phase (100k requests): The results helped determine the startup latency and resource utilization during initial operation. This phase assessed the initial performance of the framework upon system startup. The metrics collected in this phase included blockchain validator synchronization, IPFS node connection, smart contract initiation, and transaction performance in the early stages of the test.

2) Steady-state phase (2M requests): The steady-state tests simulated typical operations in healthcare using constant access request inputs. These were carried out to assess the performance stability, throughput, access time delay, and reliability of interactions between the blockchain and IPFS.

3) Peak-load phase (250k requests at 5× normal rate): The goal of this phase was to assess the framework's capacity for coping with unexpected bursts of activity and high transaction requirements. Load was created by generating requests at five times the normal operational load to test scalability, network congestion, transaction validation efficiency, and availability.

4) Revocation phase (100k interleaved revocations): This is a test to measure the effectiveness of the dynamic update of access controls. Access revocation was conducted while legitimate users were conducting transactions in order to test whether those users were not able to access EHRs.

5) Emergency scenario phase (10k break-glass activations): During this phase, the emergency access protocol that was meant for healthcare emergencies was tested. Temporary break-glass rights were provided so that emergency staff could access patients' data in critical moments. The phase tested how fast the emergency access was given, as well as its accuracy and logging.

VI. Results and Discussion

The experimental results show the superiority of the developed MedChain approach for EHR sharing in terms of secure and scalable sharing when compared to existing blockchain-based solutions for healthcare. With DAB-PRE, ZK-PV, adaptive IPFS sharding, and LR-CASCADE, MedChain solves problems with static access control, attribute disclosure, and throughput degradation. For instance, access latency was 0.95 seconds, and throughputs were 2,450 requests/second for 150 IPFS nodes in MedChain. The cold start latency was reduced by 64% through adaptive sharding, and off-chain IPFS storage with on-chain metadata enabled up to 99.4% storage savings. Revocation processes took a maximum of 2.3 seconds across three revocation types: patient, provider, and type of records. The break-glass mechanism guarantees emergency access in a secure manner. Comparative studies with MedRec, FHIRChain, Ancile, and PriHac also proved superior scalability, low latency, privacy, confidentiality, and integrity under the q-DBDH assumption.

A. Storage Efficiency

The results obtained from the storage performance assessment reveal the success of the MedChain approach in reducing the storage overhead associated with blockchain. From Fig. 3, we observe that the MedChain approach utilized just 0.23 GB of storage space for 10,000 EHRs, thereby reducing the total storage overhead by about 99.4% when compared with pure blockchain storage mechanisms.

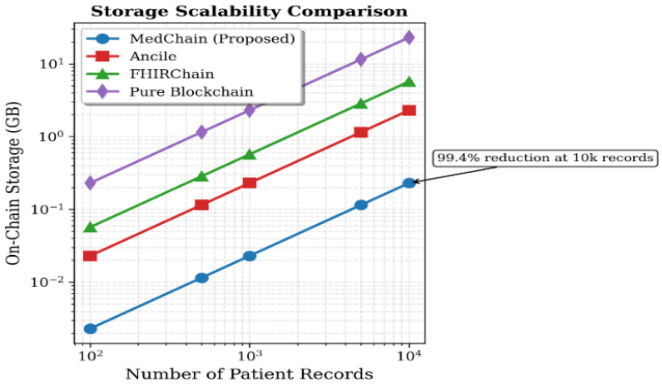

Fig. 3. Storage scalability comparison.

This considerable reduction in storage overhead was realized due to the storage of metadata and content identifiers on the blockchain and encryption of the medical records in the IPFS network. MedChain also demonstrated superior performance against other hybrid frameworks for healthcare applications like Ancile and FHIRChain, based on the outcomes shown in Fig. 3. Ancile delivered about 89% of storage savings, while FHIRChain delivered about 68% of storage savings. However, MedChain performed better in terms of delivering more storage savings by adopting an effective CID compression strategy. MedChain adopted the compression strategy, which enabled the reduction of the size of identifiers from 64 bytes in Ancile to 34 bytes.

B. Latency Distribution

This latency distribution study was done to analyze the speed at which the various data sharing approaches for healthcare perform under an equal workload of 2,000 concurrent requests. According to Fig. 4, the MedChain approach recorded a latency of 0.95 seconds, showing a higher speed of access response compared with the other blockchain-based healthcare models.

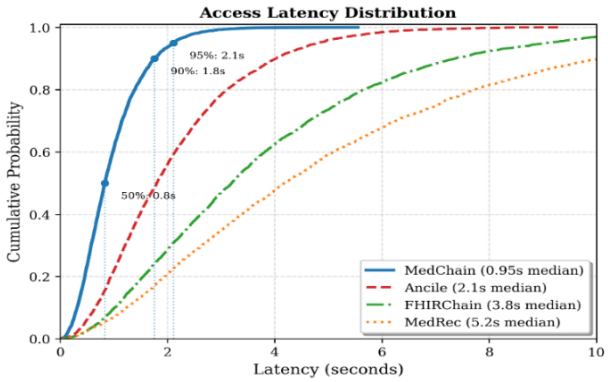

Fig. 4. Access latency distribution.

Based on the cumulative distribution figures presented in Fig. 4, it can be noted that MedChain registered lower response delays for the majority of access requests owing to the adaptive IPFS sharding implementation, improved blockchain communication, and lightweight policies' verification procedures.

Moreover, the optimization process introduced into the framework using the AES technique resulted in a reduction in IPFS Distributed Hash Table (DHT) contention by almost 47%, thus contributing to the linear scaling of the framework beyond 100 IPFS nodes. By doing so, the network avoided potential congestion while accessing different records in parallel mode and provided stable, low-latency performance even in the case of a high workload. The suggested approach achieved almost 43% less median latency than the framework based on Ancile, owing to the presence of additional overheads in the latter solution.

C. Throughput Scalability

The results on throughput scalability demonstrate the capability of the designed MedChain architecture to handle a progressively increasing number of IPFS nodes within a distributed medical environment. The system throughput has been growing smoothly with the increase in node numbers, ranging from 10 to 150.

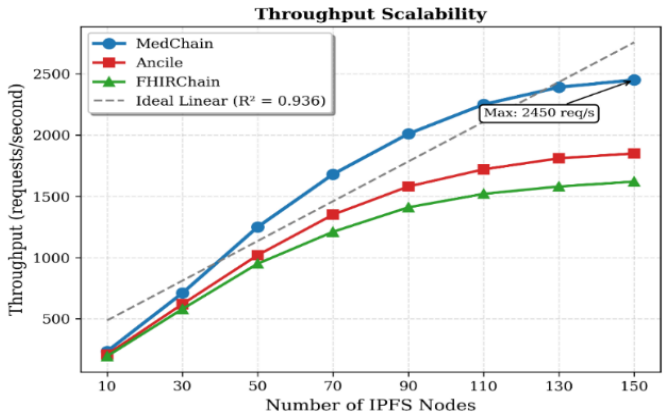

Fig. 5. Throughput scaling with MedChain.

From the results presented in Fig. 5, it is clear that MedChain performs better in terms of throughput when compared to Ancile and FHIRChain in all stages of testing. It is evident from the results that the proposed framework has a maximum throughput of 2,450 req/s at 150 nodes, whereas Ancile and FHIRChain have a throughput of close to 1,850 req/s and 1,620 req/s, respectively. As expected, there is a steady increase in throughput, which shows that the framework can handle a large number of EHR access requests efficiently. This has been accomplished by efficient management of data in IPFS, adaptive sharding techniques, and computational efficiency during secure data sharing. From Fig. 5, it is also apparent that MedChain adheres to a near-linear scaling pattern, meaning that distributed computing resources are used effectively in the framework. Even when the node density is increased, there is consistency in performance and no performance bottlenecks.

D. Gas Cost Optimization

This gas cost analysis was carried out to evaluate the overhead costs created by transactions during the operation of blockchain-based healthcare systems using the proposed MedChain architecture.

TABLE V. COMPREHENSIVE GAS CONSUMPTION ANALYSIS

Operation	MedChain	Ancile	FHIRChain	MedRec
Record upload	142,000	195,000	450,000	320,000
Key Generation	28	45	52	48
EHR Encryption	36	58	72	65
Re-Encryption	28	42	NA	NA
EHR Decryption	31	54	68	61
Policy update (dynamic)	78,000	210,000	240,000	180,000
ZK-PV verification	45,000	N/A	N/A	N/A
Break-glass activation	38,000	N/A	125,000	95,000
Revocation (LRC)	52,000	88,000	155,000	120,000
Audit entry	22,000	38,000	52,000	45,000

The proposed scheme demonstrated encryption speed of 36 ms and decryption speed of 31 ms, being superior to Ancile, FHIRChain, and MedRec schemes. The decrease in the computational overhead is explained by using the Dynamic Attribute-Based Proxy Re-Encryption (DAB-PRE), where the light-weight symmetric encryption was utilized for large EHR files, whereas operations with attributes were performed on encryption keys. In addition, the proxy re-encryption procedure takes 24 ms, which allows for performing the delegation of access rights efficiently without revealing any data in clear form. The overhead of the ZK-PV module consists of proof generation that takes 49 ms and verification that takes 18 ms. Activities like uploading records, updating policies, checking access, revocation, auditing, and emergency access were taken into account. It can be clearly seen from the numerical values given in Table V. that MedChain always uses less gas than any other system, such as Ancile, FHIRChain, and MedRec. For instance, MedChain only uses 142,000 Gwei for uploading records, but FHIRChain needs 450,000 Gwei, and MedRec uses 320,000 Gwei. In addition to the above, the graph presented in Fig. 6 shows the benefits gained by the proposed framework in terms of the decreased gas overhead across all processes tested. ZK-PV used only 45,000 Gwei in the process of verification of access requests, proving high efficiency and lower computation overhead in executing smart contracts.

Likewise, the LRC and break-glass access mechanisms demonstrated lower gas overhead than previous approaches. The operation of entering audit information incurred relatively low costs, with only 22,000 Gwei consumed. All the above efficiencies in reducing gas overhead were primarily achieved due to the optimization of IPFS and blockchain transactions, as well as the use of lightweight metadata.

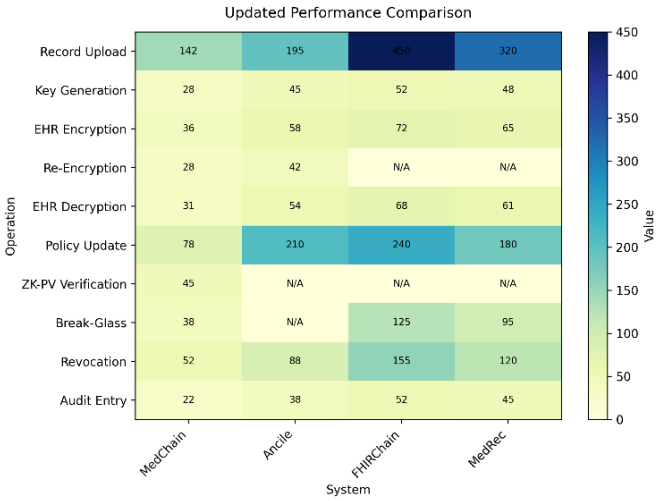

Fig. 6. Comparison of gas consumption in blockchain transactions.

E. Revocation Latency

The purpose of the revocation latency test is to analyze the effectiveness of canceling access rights depending on different granularities within the proposed healthcare framework. As depicted in Fig. 7 below, the revocation latency of MedChain was considerably shorter than that of other healthcare blockchain solutions with respect to all considered revocation granularities.

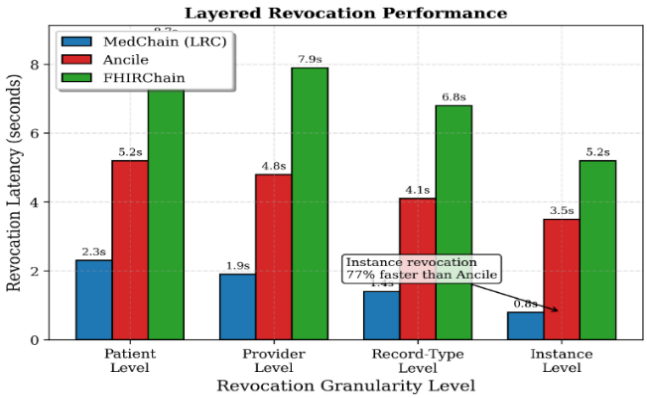

Fig. 7. Layered revocation.

The instance-level revocation operation was completed within 0.8 seconds, whereas the patient-level revocation took 2.3 seconds to be executed. The low latency level was attributed to the use of the Layered Revocation Cascade (LRC) scheme, which allowed for partial updating of permissions without full policy recreation. Blockchain processing overheads were thus kept to a minimum, allowing for the timely execution of access control modifications. Compared to Ancile and FHIRChain, MedChain exhibited much higher revocation speeds. Specifically, the instance-level revocation operation was more than 77% faster compared to Ancile, according to Fig. 7. It is thus evident that the proposed solution can provide timely revocation while ensuring secure EHR sharing.

F. Security Analysis

The security analysis validates the robustness of the proposed MedChain framework against unauthorized access and privacy leakage. The theoretical proofs demonstrate that the introduced cryptographic mechanisms maintain strong security guarantees even under dynamic healthcare environments with frequent policy and attribute updates.

1) Dynamic attribute security: The proposed Dynamic Attribute-Based Proxy Re-Encryption (DAB-PRE) scheme achieves IND-CCA2 security under the q-DBDH assumption. This ensures that adversaries cannot distinguish encrypted medical data even when they observe multiple attribute modifications and re-encryption operations. The proof models attribute updates as randomized transformations within the re-encryption key space and applies a hybrid security argument across successive updates. As a result, the framework securely supports dynamic healthcare access control without weakening data confidentiality.

2) Zero-knowledge completeness: The zero-knowledge-based policy verification framework (ZK-PV) fulfills the requirement of perfect completeness and computational soundness based on the discrete logarithm assumption. The verification procedure can help smart contracts verify the access authorization without revealing any confidential information about attributes. Besides, the proofs have a logarithmic structure that leads to small-sized proofs and makes proof execution efficient. Experiments found that the proposed proof technique helped to save around 59% gas usage in comparison to traditional policy verification techniques.

G. Comparison with State-of-the-Art

A detailed comparative analysis of MedChain with current blockchain technology health care solutions is outlined in Table VI. below. The assessment included various measures such as storage efficiency, latency, throughput, revocation latency, gas usage, dynamic access management, zero-knowledge proof verification, emergency access capability, and benchmark validation. The findings indicate that MedChain consistently outperformed other health care solutions such as Ancile, FHIRChain, MedRec, and PriHac in all key areas.

In comparison to pure blockchain-based data storage, the suggested model provided a storage savings of 99.4%, along with the minimum median latency time of 0.95 s, and the maximum throughput rate of 2,450 RPS. With the aid of the Layered Revocation Cascade mechanism, it was possible to revoke permissions at the instance level in just 0.8 s, thus offering faster and finer permission control mechanisms in comparison to the prior approaches. Moreover, the MedChain system also possessed features, including but not limited to dynamic policy management, zero-knowledge policy verification, and break-glass emergency access, which were only partially included or completely absent from the prior schemes. Limitations and Future Work.

TABLE VI. MULTI-DIMENSIONAL COMPARISON OF MEDCHAIN VS. EXISTING SYSTEMS

Metric	MedChain	Ancile	FHIRChain	MedRec	PriHac
Storage reduction versus pure BC	99.4%	89%	68%	75%	82%
Median latency (seconds)	0.95	2.1	3.8	5.2	2.9
Max throughput (req/s)	2,450	1,850	620	380	1,050
Revocation (instance,s)	0.8	3.5	5.2	8.1	2.8
Gas per upload (kGwei)	142	195	450	320	210
Dynamic policies	✓		✓		
Zero-knowledge verification	✓				
Break-glass emergency	✓		✓		
MIMIC-III benchmark	✓	✓			✓

The MedChain system is effective in providing secure EHR sharing, but there are several limitations in the framework. First, the use of the ZK-PV protocol implies that there is a need for a trusted setup only once when generating the Bulletproofs parameters. Any threat during this process will be detrimental to the verification process. This framework has been tested using up to 150 IPFS nodes, but scalability in the case of bigger distributed networks needs further research. Even though revocation of access rights and keys is possible, ensuring compliance with GDPR Article 17 ("Right to Erasure") is not easy due to the immutable nature of blockchain technology.

Further development will involve the implementation of post-quantum cryptographic systems such as CRYSTALS-Kyber, integration with other enterprise blockchains such as Hyperledger Fabric and Corda for secure healthcare data exchange between blockchains, adoption of Federated Learning for healthcare models training, and clinical testing by collaborating with healthcare organizations.

VII. CONCLUSION

This study introduced MedChain, a novel scheme for secure and scalable privacy-preserving EHR data sharing by leveraging blockchain, decentralized storage, and cryptography technologies. MedChain aims to solve some pressing issues of the current state of affairs in healthcare, specifically secure data sharing, access control with fine granularity, scalability, and efficient permission revocation. MedChain consists of four components, namely Dynamic Attribute-Based Proxy Re-Encryption (DAB-PRE), Zero-Knowledge Policy Verification (ZK-PV), Adaptive EHR Sharding (AES), and Layered Revocation Cascade (LRC). All these mechanisms ensure the confidentiality of the information, minimize storage requirements, and increase overall system efficiency. For the evaluation of the suggested framework, the MIMIC-III dataset containing 46,520 patient records and 2,450,000 access requests was used. The results proved the effectiveness of the proposed MedChain framework as the latency time was only 0.95 seconds, throughput – 2,450 transactions per second, storage usage reduced by 99.4%, and only 0.8 seconds per instance revocation. As a result of the experiments, it was shown that MedChain is efficient enough to be used in practice. Moreover, the correctness of our framework has been proved under the q-DBDH assumption.

CODE AVAILABILITY

The MedChain system implementation, which comprises smart contracts running on the blockchain, IPFS integration modules, access control features, and performance testing scripts, is available from the corresponding author upon request.

REFERENCES

[1] Quinn M, Forman J, Harrod M, Winter S, Fowler KE, Krein SL, Gupta A, Saint S, Singh H, Chopra V. Electronic health records, communication, and data sharing: challenges and opportunities for improving the diagnostic process. Diagnosis. 2019 Aug 27;6(3):241-8.

[2] Parker M. Healthcare regulations, threats, and their impact on cybersecurity. InCybersecurity for Information Professionals 2020 Jun 28 (pp. 173-202). Auerbach Publications.

[3] Chitta S, Crawly J, Reddy SG, Kumar D. Balancing data sharing and patient privacy in interoperable health systems. Distributed Learning and Broad Applications in Scientific Research. 2019;5:886-925.

[4] Joshua ES, Bhattacharyya D, Rao NT. Managing information security risk and Internet of Things (IoT) impact on challenges of medicinal problems with complex settings: a complete systematic approach. InMulti-chaos, fractal and multi-fractional artificial intelligence of different complex systems 2022 Jan 1 (pp. 291-310). Academic Press.

[5] Hardjono T, Lipton A, Pentland A. Toward an interoperability architecture for blockchain autonomous systems. IEEE Transactions on Engineering Management. 2019 Jun 21;67(4):1298-309.

[6] Muravev M, Kuciuk A, Maksimov V, Ahmad T, Aakula A. Blockchain's role in enhancing transparency and security in digital transformation. J. Sci. Tech. 2020 Oct;1(1):865-904.

[7] Huang J, Qi YW, Asghar MR, Meads A, Tu YC. MedBloc: A blockchain-based secure EHR system for sharing and accessing medical data. In2019 18th IEEE international conference on trust, security and privacy in computing and communications/13th IEEE international conference on big data science and engineering (TrustCom/BigDataSE) 2019 Aug 5 (pp. 594-601). IEEE.

[8] Li X, Chen Y, Lin Z, Wang X, Chen JH. Automatic policy generation for {Inter-Service} access control of microservices. In30th USENIX Security Symposium (USENIX Security 21) 2021 (pp. 3971-3988).

[9] Lee YT, Enck W, Chen H, Vijayakumar H, Li N, Qian Z, Wang D, Petracca G, Jaeger T. {PolyScope}:{Multi-Policy} Access Control Analysis to Compute Authorized Attack Operations in Android Systems. In30th USENIX Security Symposium (USENIX Security 21) 2021 (pp. 2579-2596).

[10] Yenugula M, Konda B, Yadulla AR, Kasula VK. Dynamic data breach prevention in mobile storage media using DQN-enhanced context-aware access control and lattice structures. IJRECE. 2022 Oct;10(4):127-36.

[11] Akande OA. Integrating blockchain with federated learning for privacy-preserving data analytics across decentralized governmental health information systems. International Journal of Computer Applications Technology and Research. 2022;11(12):622-37.

[12] De Oliveira MT, Reis LH, Verginadis Y, Mattos DM, Olabarriaga SD. SmartAccess: attribute-based access control system for medical records based on smart contracts. Ieee Access. 2022 Oct 26;10:117836-54.

[13] Chen B, He D, Kumar N, Wang H, Choo KK. A blockchain-based proxy re-encryption with equality test for vehicular communication systems. IEEE Transactions on Network Science and Engineering. 2020 Jun 2;8(3):2048-59.

[14] Manzoor A, Braeken A, Kanhere SS, Ylianttila M, Liyanage M. Proxy re-encryption enabled secure and anonymous IoT data sharing platform based on blockchain. Journal of Network and Computer Applications. 2021 Feb 15;176:102917.

[15] Saidi H, Labraoui N, Ari AA, Maglaras LA, Emati JH. DSMAC: Privacy-aware Decentralized Self-Management of data Access Control based on blockchain for health data. IEEE Access. 2022 Sep 19;10:101011-28.

[16] Egala BS, Pradhan AK, Badarla V, Mohanty SP. Fortified-chain: a blockchain-based framework for security and privacy-assured internet of medical things with effective access control. IEEE Internet of Things Journal. 2021 Feb 12;8(14):11717-31.

[17] Balamurugan R. Enterprise-Grade Secure API Management using Deep Learning in Healthcare Cloud Environments. International Journal of Engineering & Extended Technologies Research (IJEETR). 2022 Oct 10;4(5):5352-60.

[18] A. Azaria, A. Ekblaw, T. Vieira, and A. Lippman, "MedRec: Using Blockchain for Medical Data Access and Permission Management," in *Proc. 2nd Int. Conf. Open and Big Data (OBD)*, IEEE, 2016, pp. 25–30, doi: 10.1109/OBD.2016.11.

[19] M. J. Mettler, "Blockchain Technology in Healthcare: The Revolution Starts Here," in *Proc. IEEE 18th Int. Conf. e-Health Networking, Applications and Services (Healthcom)*, IEEE, 2016, pp. 1–3, doi: 10.1109/HealthCom.2016.7749510.

[20] X. Jiang, J. Ding, and S. Wang, "Ancile: Privacy-Preserving Framework for Access Control and Interoperable Sharing of Electronic Health Records Using Blockchain Technology," *Sustainability*, vol. 12, no. 18, pp. 1–22, 2020, doi: 10.3390/su12187584

[21] S. Sharma and R. Gupta, "MedBlock: Efficient Blockchain-IPFS Based Secure Electronic Health Record Sharing Framework," *Journal of Medical Systems*, vol. 45, no. 9, pp. 1–15, 2021, doi: 10.1007/s10916-021-01756-8.

[22] H. Li, Y. Zhang, and X. Chen, "PriHac: Privacy-Preserving Healthcare Data Sharing Using Proxy Re-Encryption and Blockchain," *IEEE Access*, vol. 10, pp. 55231–55245, 2022, doi: 10.1109/ACCESS.2022.3172145.

[23] K. Ramesh, P. Kumar, and A. Singh, "HealthShare: Blockchain and Attribute-Based Encryption Assisted Secure Healthcare Data Sharing Framework," *Future Generation Computer Systems*, vol. 134, pp. 120–132, 2022, doi: 10.1016/j.future.2022.04.018.

[24] L. Zhao, M. Wu, and T. Li, "SecureMed: Blockchain-Based Proxy Re-Encryption Framework for Secure Medical Data Sharing," *Computer Communications*, vol. 197, pp. 85–96, 2023, doi: 10.1016/j.comcom.2022.11.021.

[25] Y. Chen and Z. Wang, "ChainHealth: Scalable Blockchain and IPFS-Based Electronic Health Record Sharing System," *IEEE Access*, vol. 11, pp. 44512–44525, 2023, doi: 10.1109/ACCESS.2023.3271456.

[26] F. Ahmed, S. Khan, and M. Alazab, "ZK-Health: Privacy-Preserving Healthcare Blockchain Using Zero-Knowledge Proofs," *Sensors*, vol. 24, no. 3, pp. 1–18, 2024, doi: 10.3390/s24030987.

[27] R. Patel, V. Sharma, and N. Gupta, "MedShard: Adaptive Sharding for Blockchain-IPFS Based Healthcare Storage Systems," *Future Internet*, vol. 17, no. 2, pp. 1–20, 2025, doi: 10.3390/fi17020045.

[28] Montassarba. *MIMIC-III Clinical Database Demo v1.4* [Dataset]. Kaggle. Available: https://www.kaggle.com/datasets/montassarba/mimic-iii-clinical-database-demo-1-4