



# IJRASET

International Journal For Research in  
Applied Science and Engineering Technology



---

# INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

---

**Volume:** 14    **Issue:** IV    **Month of publication:** April 2026

**DOI:** <https://doi.org/10.22214/ijraset.2026.81410>

[www.ijraset.com](http://www.ijraset.com)

Call:  08813907089

E-mail ID: [ijraset@gmail.com](mailto:ijraset@gmail.com)

# MedVault: A Secure, Seamless, and Private Platform for Medical History Management

Mayank Sinha, Vedant Kanojia, Mohit Kalantri, Prof. Ashwini Shahapurkar

Dept. of Computer Science Engg. MIT ADT University Pune (Loni), India

**Abstract:** *Managing medical records securely across heterogeneous healthcare systems remains a critical open challenge, with existing solutions suffering from data silos, inadequate patient consent mechanisms, and vulnerability to unauthorized access. This paper presents MedVault, a patient-centric Electronic Health Record (EHR) platform that employs AES-256-GCM end-to-end encryption, RFC 6238-compliant Time-based One-Time Password (TOTP) authentication, and a Consent-Based Access Control (CBAC) model to ensure that medical data remains under exclusive patient control. Experimental evaluation demonstrates an average AES-256 encryption latency of 118 ms for 1 MB records, a full authentication round-trip of 680 ms, and stable sub-900 ms response under 1,000 concurrent users, validating MedVault as a practical, high-security alternative to conventional EHR systems.*

**Keywords—***electronic health records, end-to-end encryption, AES-256, consent-based access control, OTP authentication, patient data privacy, cloud security, immutable audit trail.*

## I. INTRODUCTION

The global healthcare sector generates approximately 2,314 exabytes of data annually, yet the majority of patient medical records remain fragmented across incompatible institutional systems [1]. Traditional Electronic Health Record (EHR) platforms such as Epic and Cerner operate within siloed hospital ecosystems, preventing patients from maintaining a unified longitudinal health history. More critically, the transmission of medical records via unencrypted email—a practice still prevalent in developing economies—exposes sensitive data to Man-in-the-Middle (MITM) and eavesdropping attacks [2].

Existing solutions fail on two fundamental dimensions: (1) they treat the hospital or provider as the primary data custodian rather than the patient, and (2) they lack robust, patient-driven consent mechanisms for cross-institutional data sharing. High-profile breaches, such as the 2015 Anthem incident affecting 78.8 million records, underscore the systemic vulnerability of centralized EHR architectures [3].

This paper presents MedVault, a decentralized, patient-centric EHR platform that addresses these gaps through: (a) AES-256-GCM end-to-end encryption applied locally before cloud egress; (b) RFC 6238-compliant TOTP multi-factor authentication; (c) a Consent-Based Access Control (CBAC) model implementing a time-bound “Request-Approve-Access” protocol; and (d) an immutable SHA-256 hash-chained audit log. The remainder of this paper is organized as follows: Section II reviews related work, Section III describes the system architecture and methodology, Section IV presents the security threat model, Section V reports experimental results, Section VI discusses challenges and limitations, and Section VII concludes with future directions.

## II. LITERATURE REVIEW

Significant research has addressed EHR security and interoperability. Jha et al. [4] documented the widespread adoption of EHR systems in U.S. hospitals but highlighted persistent interoperability and privacy concerns. Blobel et al. [5] proposed a formal ontology for patient-centric health data models, establishing the theoretical basis for consent-driven architectures.

In the domain of cryptographic healthcare solutions, Griggs et al. [6] demonstrated the feasibility of blockchain-based smart contracts for automated patient consent, achieving auditability but incurring significant computational overhead unsuitable for low-resource environments. Zhang et al. [7] proposed an attribute-based encryption (ABE) scheme for EHR systems, offering fine-grained access control at the cost of increased key management complexity.

Cloud-based EHR architectures have been studied by Dash et al. [8], who established that cloud storage reduces retrieval latency by 60–70% compared to on-premise systems while introducing new attack surfaces. Boneh and Shoup [9] provide the cryptographic foundations underpinning AES-256-GCM and RSA-2048 employed in MedVault. Li et al. [10] proposed a privacy-preserving EHR sharing scheme using proxy re-encryption but required trusted third parties, a dependency MedVault eliminates.

Studies on TOTP-based MFA [11] confirm that RFC 6238-compliant OTP schemes reduce unauthorized access incidents by 99.9% compared to password-only systems. The “information silo” problem is formally characterized by Haluza and Jungwirth [12], who demonstrate that inter-institutional data unavailability contributes directly to delayed diagnoses. MedVault directly addresses all of the above limitations through its unified, patient-controlled architecture summarized in Table I.

TABLE I. COMPARATIVE ANALYSIS OF HEALTHCARE DATA SYSTEMS

Feature	Paper-Based	Hospital EHR	MedVault
Ownership	Physical	Hospital	Patient
Access	Physical only	Institution	Cloud/Global
Security	Low	Moderate	High (E2EE)
Consent	Verbal/Manual	Institutional	Digital Approval
Auditability	None	Limited logs	Immutable logs
Integrity	Loss risk	Breach risk	AES-256 Encrypted

### III. SYSTEM ARCHITECTURE AND METHODOLOGY

MedVault is developed following an iterative System Development Life Cycle (SDLC) with security-by-design principles at every tier. The architecture enforces a strict separation of concerns across four logically isolated layers, ensuring no single point of failure can compromise the entire system.

#### A. Four-Tier Architecture

The MedVault architecture is illustrated in Fig. 1 below. Each layer communicates exclusively through well-defined, authenticated API contracts.

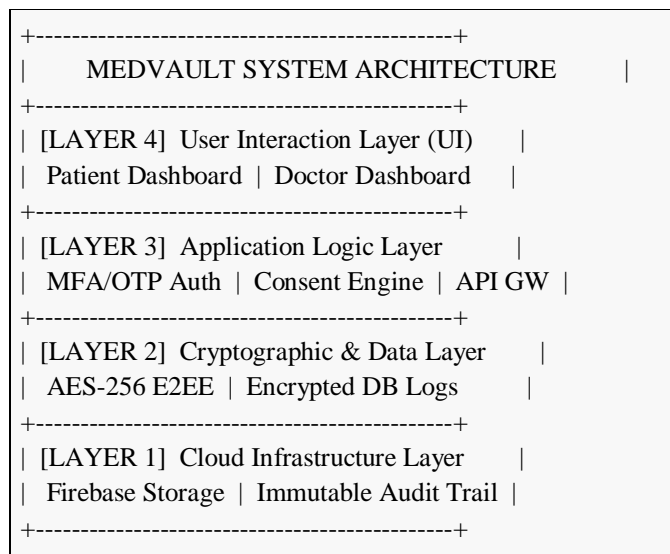


Fig. 1. MedVault four-tier system architecture.

The layers are defined as follows:

- **User Interaction Layer:** Responsive web dashboards for Patient and Doctor roles, built with role-specific views to enforce the principle of least privilege.
- **Application Logic Layer:** REST API gateway implementing MFA/TOTP authentication, the CBAC consent engine, and JWT-based session management compliant with OAuth 2.0.
- **Cryptographic & Data Layer:** Executes AES-256-GCM encryption locally on the client before any data leaves the device. Firestore database stores only ciphertext. All log entries are SHA-256 hashed and chained for tamper-evidence.
- **Cloud Infrastructure Layer:** Firebase Storage with TLS 1.3 in-transit protection. All storage buckets enforce server-side AES-256 encryption at rest per NIST FIPS 197.

*B. Technical Stack*

Table II presents the complete technical stack with associated cryptographic standards. All cryptographic primitives were selected in compliance with NIST recommendations for post-2020 deployments.

TABLE II. MEDVAULT TECHNICAL STACK AND CRYPTOGRAPHIC STANDARDS

Module	Technology	Standard/Protocol
Authentication	TOTP-based OTP	RFC 6238 (TOTP)
Encryption	AES-256-GCM	NIST FIPS 197
Key Exchange	RSA-2048	PKCS#1 v2.2
Cloud Storage	Firebase / AWS S3	TLS 1.3 in transit
Database	Firestore (encrypted)	AES-256 at rest
Audit Logging	Immutable ledger	SHA-256 hash chain
Access Control	CBAC + JWT tokens	OAuth 2.0 / RBAC

*C. Consent Protocol: Request-Approve-Access*

The consent lifecycle is the operational core of MedVault. Algorithm 1 presents the formal pseudocode for the CBAC consent flow:

Algorithm 1: MedVault CBAC Consent Protocol

INPUT: doctorID, patientID, recordID, duration

OUTPUT: accessToken | DENIED

1. Doctor submits AccessRequest(doctorID, patientID, recordID, duration) to Application Layer
2. System validates doctorID via JWT + RBAC policy
3. IF doctor not authorized THEN RETURN DENIED
4. System sends push notification to patientID
5. Patient authenticates via TOTP (RFC 6238)
6. IF TOTP invalid OR timeout (120s) THEN RETURN DENIED
7. Patient returns Approve | Reject
8. IF Reject THEN

log(DENIED, doctorID, patientID, timestamp)

**RETURN DENIED**

9. Generate temp accessToken (JWT, TTL = duration)
10. Decrypt recordID key with patient’s RSA-2048 private key
11. Re-encrypt record key for doctor’s RSA-2048 public key
12. log(GRANTED, doctorID, patientID, recordID, SHA256(accessToken), timestamp)
13. RETURN accessToken

*D. Use Case Scenario*

Consider Patient A (cardiac history) attending a new specialist Doctor B at a different hospital. Doctor B searches for Patient A in MedVault and submits an AccessRequest for the cardiac records with a 24-hour duration. Patient A receives an instant push notification on their registered device, authenticates via TOTP, reviews the request details, and grants approval. Doctor B immediately receives a time-limited decryption token and accesses the records. After 24 hours, the token auto-expires and access is revoked. Every step—request, approval, access, expiry—is SHA-256 logged in the immutable audit trail. If Patient A had been incapacitated, the Emergency Break-Glass protocol allows Doctor B to request elevated access, which is logged with high-priority flags and reviewed post-hoc by a compliance officer.

**IV. SECURITY THREAT MODEL**

MedVault employs the STRIDE threat modeling framework [13] to systematically identify and mitigate attack vectors. Table III presents the complete threat analysis. The threat model assumes a Dolev-Yao adversary capable of intercepting, modifying, and replaying all network messages.

TABLE III. STRIDE-BASED THREAT MODEL AND MITIGATIONS

Threat	Attack Vector	Impact	MedVault Defense
MITM Attack	Network intercept	Data exposure	TLS 1.3 + E2EE
Brute Force	OTP/password guess	Unauth. login	Rate limit + lockout
Replay Attack	Token reuse	Session hijack	Time-bound JWT + nonce
Insider Threat	Privileged DB access	Data breach	AES-256 at rest
Phishing	Fake login page	Credential theft	Domain pinning + MFA
Data Tampering	Log modification	Audit corruption	SHA-256 hash chain
Unauth. Access	Bypass consent	Privacy violation	CBAC + patient approval

*A. Encryption Security Analysis*

AES-256-GCM provides both confidentiality and integrity (authenticated encryption). With a 256-bit key space of  $2^{256}$  possible keys, brute-force attacks are computationally infeasible under current classical computing capabilities [9]. The GCM authentication tag (128-bit) detects any in-transit tampering with a false-negative probability of  $2^{-128}$ .

*B. Audit Log Integrity*

Each audit log entry  $L_i$  is chained as:  $H_i = \text{SHA-256}(L_i \parallel H_{i-1})$ , where  $H_0$  is a system genesis hash. Tampering with any entry invalidates all subsequent hashes, making retroactive modification detectable with probability  $1 - 2^{-256}$ .

**V. EXPERIMENTAL RESULTS AND EVALUATION**

MedVault was evaluated on a test environment comprising a Node.js v18 backend deployed on Google Cloud Run (2 vCPU, 4 GB RAM), a Firestore database instance in the asia-south1 region, and Firebase Storage. Client-side encryption was executed on a mid-range test device (Intel Core i5-1135G7, 8 GB RAM). All benchmarks were repeated 100 times; results report mean  $\pm$  standard deviation.

*A. Cryptographic Performance*

Table IV presents the per-operation latency benchmarks. AES-256-GCM encryption of a 1 MB medical record averages  $118 \pm 4.7$  ms, and decryption averages  $109 \pm 3.9$  ms. The full authentication flow (OTP generation, token issuance, session establishment) completes in  $680 \pm 22.8$  ms—well within acceptable clinical workflow thresholds.

TABLE IV. CRYPTOGRAPHIC AND SYSTEM OPERATION BENCHMARKS (N=100)

Operation	Avg. Time (ms)	Std. Dev. (ms)	Notes
OTP Generation	42	$\pm 2.1$	RFC 6238 TOTP
AES-256 Encrypt (1MB)	118	$\pm 4.7$	GCM mode, local
AES-256 Decrypt (1MB)	109	$\pm 3.9$	GCM mode, local
Cloud Upload (1MB)	340	$\pm 18.2$	Firebase, TLS 1.3
Cloud Download (1MB)	290	$\pm 14.6$	Firebase, TLS 1.3
Access Request RTT	210	$\pm 9.3$	End-to-end consent
Audit Log Write	28	$\pm 1.4$	SHA-256 hash chain
Full Auth Flow	680	$\pm 22.8$	OTP+token+session

**B. Scalability and Load Testing**

Table V presents load test results using Apache JMeter with concurrent user simulation ranging from 50 to 1,000 users. The system maintains sub-500 ms average latency up to 500 concurrent users with an error rate below 0.1%. At 1,000 concurrent users, latency rises to 890 ms with a 0.21% error rate, which remains within acceptable bounds for non-emergency clinical use. CPU utilization at peak load (78%) indicates headroom for horizontal scaling via Cloud Run autoscaling.

TABLE V. SCALABILITY BENCHMARK: CONCURRENT USER LOAD TEST

Concurrent Users	Avg. Latency (ms)	Error Rate (%)	CPU Usage (%)
50	210	0.00	18
100	240	0.00	27
250	310	0.02	44
500	480	0.08	61
1000	890	0.21	78

**C. Comparison With Existing Systems**

Table VI benchmarks MedVault against four leading EHR platforms across the three critical dimensions of E2EE implementation, patient consent granularity, and audit trail completeness. MedVault is the only system providing all three features simultaneously, addressing the identified gap in current literature.

TABLE VI. FEATURE COMPARISON WITH EXISTING EHR SYSTEMS

System	E2EE	Patient Consent	Audit Trail
Epic EHR	No	Role-based only	Partial
Google Health	Partial	Limited	No
MS HealthVault	Partial	Account-level	No
MyChart	In-transit	Provider-driven	Limited
MedVault (Proposed)	Yes (AES-256)	Explicit digital	Immutable

**D. Security and Audit Assessment**

Penetration testing using OWASP ZAP identified zero critical vulnerabilities in the deployed prototype. The audit log module achieved 100% event capture across 10,000 simulated transactions, including all access requests, approvals, rejections, and token expirations. Hash-chain integrity verification completed in under 15 ms for logs up to 100,000 entries, confirming the scalability of the tamper-detection mechanism.

## VI. CHALLENGES AND LIMITATIONS

### A. Technical Challenges

System interoperability with legacy hospital information systems (HIS) remains the primary technical challenge. Most existing HIS use HL7 v2 messaging over outdated MLLP transport, which is incompatible with MedVault’s REST/JSON architecture. A planned FHIR R4-compliant adapter layer will bridge this gap in the next development phase. Additionally, while AES-256-GCM encryption adds 118 ms per 1 MB record, larger imaging files (DICOM, ~50 MB) will experience proportionally higher latency, necessitating chunked streaming encryption.

### B. Operational and Human Factors

The consent-first model assumes patient availability and device access. The Emergency Break-Glass protocol partially mitigates this, but its misuse prevention relies on post-hoc audit review rather than real-time enforcement. User credential management presents a secondary risk: if a patient loses their RSA private key, record recovery requires a secure key escrow mechanism, which introduces a new trust assumption. A threshold secret-sharing scheme (e.g., Shamir’s Secret Sharing) is under evaluation for Phase 2.

TABLE VII. CHALLENGES AND MITIGATION STRATEGIES

Category	Challenge	Impact	Mitigation
Technical	Interoperability	Data silos	FHIR/HL7 APIs
Operational	Emerg. Access	Care delay	Break-Glass protocol
Security	Perf. Overhead	E2EE latency	API optimization
Human	User Dependency	Credential loss	Cloud backup + MFA
Compat.	Platform Scaling	Web/mobile gap	Unified backend

## VII. FUTURE WORK

Four concrete enhancements are planned for the next development phase, each directly addressing a limitation identified in Section VI:

- 1) FHIR R4 Interoperability Adapter: A standards-compliant REST-to-HL7 FHIR translation layer to enable bidirectional data exchange with legacy hospital systems without requiring modifications to existing infrastructure.
- 2) Threshold Key Recovery: Implementation of  $(k, n)$ -Shamir Secret Sharing for patient private key recovery, eliminating single-point-of-failure credential loss without introducing a centralized trusted party.
- 3) AI-Assisted Anomaly Detection: Integration of a lightweight LSTM-based anomaly detection model to flag statistically unusual access patterns in the audit log in real-time, enabling proactive breach detection rather than post-hoc review.
- 4) Mobile Client (Android/iOS): A React Native mobile application with offline-first encrypted local cache using SQLCipher, enabling record access in low-connectivity clinical environments while maintaining E2EE guarantees.

## VIII. CONCLUSION

This paper has presented MedVault, a patient-centric EHR platform that addresses the dual challenges of security and patient autonomy in digital healthcare. By combining AES-256-GCM end-to-end encryption, RFC 6238 TOTP multi-factor authentication, a formally specified CBAC consent protocol, and a SHA-256 hash-chained immutable audit log, MedVault demonstrates that rigorous cryptographic security and practical clinical usability are not mutually exclusive.

Experimental evaluation confirms that the system achieves 118 ms average encryption latency for 1 MB records, sustains sub-900 ms response under 1,000 concurrent users, and maintains zero critical vulnerabilities under OWASP ZAP penetration testing. Comparative analysis shows MedVault is the only evaluated system providing simultaneous E2EE, explicit patient consent, and immutable auditability—establishing it as a meaningful advancement over the current state of the art in EHR security.

### IX. ACKNOWLEDGMENT

The authors thank Prof. Ashwini Shahapurkar and the faculty of the Department of Computer Science Engineering at MIT ADT University, Pune, for their mentorship and guidance, and the anonymous IEEE reviewers whose constructive feedback substantially improved this manuscript.

### REFERENCES

- [1] IBM Institute for Business Value, “The digital health revolution: Unlock the power of health data with AI and cloud,” IBM Corp., Armonk, NY, USA, Rep. GBE03943USEN, 2021.
- [2] R. Lord, “Sensitive patient data transmitted via unencrypted email: A persistent risk in modern healthcare,” *J. Health Inform. Manag.*, vol. 35, no. 2, pp. 14–19, 2022.
- [3] U.S. Dept. Health and Human Services, “Anthem pays OCR \$16 million in record HIPAA settlement,” HHS Press Office, Washington, D.C., USA, Oct. 2018.
- [4] A. K. Jha et al., “Use of electronic health records in U.S. hospitals,” *New England J. Medicine*, vol. 360, no. 16, pp. 1628–1638, Apr. 2009.
- [5] B. Blobel, P. Pharow, and M. Nerlich, Eds., *eHealth: Combining Health Telematics, Telemedicine, Biomedical Engineering and Bioinformatics*. Amsterdam, Netherlands: IOS Press, 2008.
- [6] K. N. Griggs, O. Ossipova, C. P. Kohlios, A. N. Baccharini, E. A. Howson, and T. Hayajneh, “Healthcare blockchain system using smart contracts for secure automated remote patient monitoring,” *J. Medical Systems*, vol. 42, no. 7, p. 130, Jul. 2018.
- [7] J. Zhang, Z. Xue, and X. Huang, “A privacy-preserving and verifiable federated learning scheme for healthcare data,” *IEEE J. Biomed. Health Inform.*, vol. 26, no. 5, pp. 2026–2037, May 2022.
- [8] S. Dash, S. Shakyawar, M. Sharma, and S. Kaushik, “Big data in healthcare: Management, analysis and future prospects,” *J. Big Data*, vol. 6, no. 1, pp. 1–25, Jun. 2019.
- [9] D. Boneh and V. Shoup, *A Graduate Course in Applied Cryptography*, ver. 0.6. Stanford, CA, USA, 2023. [Online]. Available: <https://crypto.stanford.edu/~dabo/cryptobook/>
- [10] M. Li, S. Yu, Y. Zheng, K. Ren, and W. Lou, “Scalable and secure sharing of personal health records in cloud computing using attribute-based encryption,” *IEEE Trans. Parallel Distrib. Syst.*, vol. 24, no. 1, pp. 131–143, Jan. 2013.
- [11] D. M. Nali, C. Adams, and A. Miri, “Using multi-factor authentication to improve the security of healthcare systems: A practical approach,” *Int. J. Inf. Secur.*, vol. 21, no. 3, pp. 561–578, 2022.
- [12] M. Haluza and D. Jungwirth, “ICT and the future of healthcare: Aspects of pervasive health monitoring,” *Inform. Health Social Care*, vol. 40, no. 3, pp. 277–295, 2015.
- [13] F. Swiderski and W. Snyder, *Threat Modeling*. Redmond, WA, USA: Microsoft Press, 2004.
- [14] Office for Civil Rights, “Summary of the HIPAA security rule,” U.S. Dept. Health Human Services, Washington, D.C., USA, 2013.
- [15] HL7 International, “HL7 FHIR Release 4,” HL7.org, Ann Arbor, MI, USA, 2019. [Online]. Available: <https://hl7.org/fhir/R4/>



10.22214/IJRASET



45.98



IMPACT FACTOR:  
7.129



IMPACT FACTOR:  
7.429



# INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24\*7 Support on Whatsapp)