International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(21s) 2026 pp. 752-772
https://doi.org/10.70917/ijcisim-2026-5342
Article

Single- and Extended-Authority CP-ABE for E-Healthcare Security in Cloud-IoT

Gurupriya K G[1], Dr. A. S Aneeshkumar[2]

[1]Research Scholar, Research Department of Computer Science , AJK College of Arts and Science(Autonomous), Coimbatore, Tamilnadu-641105, India
Email: gurupriyasarth8@gmail.com
[2]Research Supervisor and Head, PG & Research Department of Computer Science and Applications, AJK College of Arts and Science(Autonomous), Coimbatore, Tamilnadu-641105, India
Email: aneeshkumar.alpha@gmail.com

Abstract: The integration of Internet of Things (IoT) devices into clinical environments has generated volumes of electronic health records (EHRs) that exceed the management capacity of conventional access control mechanisms when these records are offloaded to cloud platforms. Protecting health records with fine-grained, cryptographically enforced access control while supporting dynamic credential revocation at cloud scale has remained an open problem in practice. Existing Ciphertext-Policy Attribute-Based Encryption (CP-ABE) schemes incur revocation costs that scale with the ciphertext archive size, lack computational feasibility on constrained IoT microcontrollers, or fail to address deployments where the Cloud Service Provider (CSP) itself is an untrusted adversary. This paper proposes two new CP-ABE constructions that close all three gaps simultaneously. Single-Authority CP-ABE (SA-CP-ABE) targets private hospital clouds under a semi-trusted CSP model and introduces a lazy versioned-key revocation mechanism in which the Trusted Authority broadcasts a single 32-byte group element per revocation event, achieving O(1) cost independent of the ciphertext archive size N. Extended-Authority CP-ABE (EA-CP-ABE) extends the construction to the fully untrusted public cloud by adding Pedersen-commitment-based policy anonymisation that hides the access policy structure from the CSP, and a conditional proxy re-encryption protocol that allows the CSP to perform revocation-driven ciphertext updates without observing the attribute being revoked. Both schemes are formalized in ten equations covering system setup, key generation, encryption, Monotone Span Program (MSP) secret sharing, decryption, and the two-step O(1) attribute revocation protocol. Security proofs under the Decisional Bilinear Diffie-Hellman (DBDH) assumption in the random oracle model establish IND-CPA security (Theorem 1) and revocation security (Theorem 2); IND-CCA security follows from standard one-time signature composition. Experimental evaluation across eight tables on an Intel Xeon E5-2690 v4 server and seven embedded hardware platforms confirms that SA-CP-ABE reduces ciphertext size by 44% over Yang et al. and maintains a constant 8.2 ms revocation latency regardless of policy size or concurrent user load, versus 312–389 ms for Yang et al. EA-CP-ABE adds a fixed 17% size overhead while achieving complete policy hiding. Both schemes run on ARM Cortex-M0 microcontrollers at 18.8 ms per sensor sample in batched mode and comply with HIPAA and GDPR technical safeguard requirements.
Keywords: CP-ABE; E-Healthcare; Cloud-IoT Security; Attribute Revocation; DBDH; Monotone Span Programs; Big Data; EHR; HIPAA; GDPR; Proxy Re-Encryption; ARM Cortex

1. INTRODUCTION

Modern hospital networks generate several terabytes of structured and unstructured EHR data per day. Diagnostic images, genomic sequences, physiological time-series, medication records, and clinical notes flow through cloud-hosted infrastructure accessed simultaneously by clinicians, researchers, administrative staff, and automated analytics pipelines from geographically distributed locations [1]. Cloud computing delivers the elastic storage, geographic redundancy, and multi-site accessibility this volume demands. It also moves patient data into facilities

operated by Cloud Service Providers (CSPs) that function outside the direct oversight of the healthcare organization, creating a structural mismatch between data custody and access-control authority.

The regulatory stakes are considerable. In the United States, HIPAA violations carry civil penalties of up to USD 1.9 million per violation category per calendar year, with criminal penalties reaching USD 250,000 and 10 years imprisonment for knowing violations [2]. The EU GDPR imposes fines up to EUR 20 million or 4% of global annual turnover for breaches of special-category personal data, which explicitly includes health records [3]. India's Digital Personal Data Protection Act, enacted in 2023, establishes comparable obligations for health data fiduciaries [4]. Beyond regulatory exposure, EHR breaches carry material secondary harms: insurance discrimination, employment consequences, reputational damage to the affected institution, and erosion of patient trust in connected care models precisely when that trust is most needed. Standard logical access control models, including Role-Based Access Control (RBAC) and Attribute-Based Access Control (ABAC), enforce permissions at the application tier and depend on the infrastructure operator to apply policies faithfully. In a cloud deployment, this means trusting the CSP not to inspect content, not to misconfigure access rules, and not to be compelled by legal or political means to provide access outside the healthcare organization's authorization policy. Each of these failure modes has materialized in documented breaches. A cryptographic approach that makes ciphertext computationally opaque to the CSP regardless of its behavior is therefore necessary. Attribute-Based Encryption (ABE), introduced by Sahai and Waters as a generalization of identity-based encryption [5], addresses this need. In Ciphertext-Policy ABE (CP-ABE) [6], the data owner encodes an access policy directly into the ciphertext at encryption time. A user can decrypt if and only if her attribute credentials satisfy the embedded policy, regardless of where or by whom the ciphertext is stored. A single EHR can be made simultaneously accessible to all cardiologists on duty at the same institution, and to the clinical laboratory team — each using an independent credential set — while the CSP storing the ciphertext learns nothing about the plaintext or the specific credentials that authorize access. This expressiveness makes CP-ABE a natural fit for healthcare data governance at cloud scale. Despite this alignment, three practical obstacles have blocked CP-ABE deployment in production healthcare cloud systems:

Revocation cost: Clinical roles change continuously. Staff onboarding and offboarding, temporary contractors, inter-departmental rotations and compromised credential events all require rapid access revocation. Revoking a user's attribute in CP-ABE without re-encrypting all stored ciphertexts that reference that attribute is cryptographically difficult. Dominant prior approaches require effort proportional to O(log N) [7, 8] or O(N) [9] over the archived ciphertext count N. At hospital scale, N reaches tens of millions of records, making these approaches infeasible for real-time revocation workflows.

IoT hardware constraints: An increasing fraction of clinical EHR data originates from IoT-connected devices: wearable biosensors, implantable cardiac monitors, and smart infusion pumps. Many operate on ARM Cortex-M0 microcontrollers at 48 MHz with 32 KB of RAM. Standard CP-ABE encryption requires multiple bilinear pairing evaluations, each costing 180–200 ms on these processors. An IoT-compatible architecture must reduce Tier-1 device costs to at most two exponentiations and one hash evaluation per reading.

Public-cloud trust model: Private hospital clouds can operate under a semi-trusted CSP model. Public cloud platforms operated by third-party hyperscalers are routinely modeled as fully untrusted: the CSP may attempt to infer sensitive policy structure from ciphertext metadata, exploit re-encryption tokens beyond their authorized scope, or correlate access-pattern metadata to fingerprint individual patients or clinicians.

This paper proposes SA-CP-ABE and EA-CP-ABE, two formally proven constructions that simultaneously address all three obstacles. The central technical contribution is a lazy versioned-key revocation protocol in which the Trusted Authority (TA) broadcasts a single group element per revocation event. Non-revoked users update their key locally with one group multiplication; the CSP defers ciphertext updates to the next access event using a re-encryption token. The per-revocation cost at the TA is O(1), independent of N, n, or the user count. EA-CP-ABE adds a Pedersen commitment layer hiding the policy labelling function from the CSP, and a conditional proxy re-encryption key that allows CSP-side ciphertext updates without leaking the revoked attribute identity.

The specific contributions of this paper are:

(1) SA-CP-ABE: A formal construction for private cloud healthcare storage achieving O(1) attribute revocation via lazy versioned-key update (Section V), with a proof of IND-CPA security under DBDH and O(1)-provable revocation security (Section VII).

(2) EA-CP-ABE: An extension for fully untrusted public clouds with complete policy hiding via Pedersen-style commitment, conditional proxy re-encryption for CSP-side revocation, and a three-tier IoT integration protocol (Section VI).

(3) Ten formal equations ((3.1)–(5.7)) covering DBDH advantage, MSP reconstruction, IND-CPA definition, system setup, key generation, secret sharing, encryption, decryption, and both revocation steps, with rigorous variable-level explanations.

(4) Eight evaluation tables covering scheme feature comparison, notation, computational performance, communication overhead, scalability, IoT device feasibility, algebraic operation counts, and security parameter sensitivity.

The remainder of this paper is organized as follows. Section II surveys related work. Section III establishes cryptographic preliminaries. Section IV describes the system model. Sections V and VI present the constructions. Section VII provides formal security analysis. Section VIII reports experimental evaluation. Section IX discusses deployment and compliance. Section X concludes. The research base of this work is anchored in three converging technological realities that together make a new cryptographic solution necessary.

Healthcare data volume and cloud dependency: Modern hospitals generate 2–5 TB of EHR data daily across diagnostic imaging, genomic sequences, physiological time-series from

monitoring equipment, medication records, and clinical documentation [1]. Cloud infrastructure has become the only economically viable storage and access platform at this scale. AWS HealthLake, Microsoft Azure Health Data Services, and Google Cloud Healthcare API together serve the majority of large hospital network deployments globally. This migration is irreversible: on-premise storage infrastructure cannot elastically scale to accommodate exponential growth in imaging data from next-generation radiology and pathology workflows.

IoT-driven data generation at the clinical edge: The number of connected medical devices per hospital bed has grown from an average of 6.2 in 2018 to over 15.4 in 2024, spanning wearables, implantables, infusion systems, and diagnostic equipment [1]. These devices operate on ARM Cortex-M0 through M4 class microcontrollers with severe computational budgets. They generate real-time physiological data that must be encrypted at the point of acquisition, before transmission to the gateway or cloud, to satisfy both HIPAA §164.312(a)(2)(iv) and GDPR Article 32(1)(a) technical safeguard requirements. Standard CP-ABE encryption at the sensor tier is computationally infeasible without architectural changes that assign only minimal operations to Tier-1 devices.

Dynamic credential management at hospital scale: A typical 500-bed hospital employs over 2,000 clinical staff across dozens of departments. Staff role changes, contractor expirations, and credential compromise events collectively generate an average of 18–25 revocation events per day in active deployments. Each revocation must take cryptographic effect within one authentication cycle (under 30 seconds) to comply with HIPAA §164.312(a)(2)(iii) automatic logoff and access termination requirements. No existing CP-ABE scheme achieves this at O(1) cost per revocation event when the ciphertext archive contains millions of records.

These three realities together define the research base: a system must simultaneously support cloud-scale EHR storage, IoT-edge encryption at constrained hardware, and real-time credential revocation at O(1) cost. No single prior scheme satisfies all three. The formal CP-ABE constructions proposed in this paper, SA-CP-ABE and EA-CP-ABE, are designed specifically to address this combination.

1.1 Research Analysis

Prior CP-ABE work leaves a consistent gap across three requirements that a production healthcare deployment must satisfy together: bounded revocation cost, feasibility on constrained IoT hardware, and resilience against an untrusted public-cloud provider. The remainder of this section examines each requirement in turn, identifies precisely where existing schemes fall short, and uses that gap analysis to motivate the design choices carried into Sections V and VI.

Analysis of revocation approaches: The revocation cost landscape in prior CP-ABE work follows three tiers. Schemes without revocation support (Waters [10], Bethencourt et al. [6]) provide no mechanism to exclude previously authorized users from future ciphertext access; this is acceptable only in static environments where user sets never change, which does not describe any real healthcare organization. First-generation revocation schemes (Yu et al. [8], Green et al. [9]) support revocation through proxy re-encryption but require the TA or CSP to re-encrypt all stored ciphertexts referencing the revoked attribute at each revocation event, incurring O(N) cost where N is the total ciphertext count. At a hospital archiving 500,000 radiology images per year at average 15 MB per study, N reaches millions within five years, making O(N) re-encryption operationally impossible. Second-generation schemes (Yang et al. [7]) reduce this to O(log N)

using binary tree key update structures, requiring approximately 20 group operations per revocation event at N = 1,000,000. While substantially better than O(N), this still fails the 30-second revocation deadline under concurrent access load, as confirmed by our scalability experiments in Section VIII. The O(1) cost achieved by the lazy versioned-key protocol of this paper is the first construction to make revocation cost independent of N entirely.

Analysis of IoT feasibility constraints: Prior CP-ABE schemes were designed assuming desktop or server-class hardware and make no accommodation for constrained embedded devices. The standard Bethencourt et al. scheme requires 2l bilinear pairing evaluations per encryption (where l is the MSP row count) and 2d pairing evaluations per decryption (where d is the number of satisfying rows). On an ARM Cortex-M0 at 48 MHz running RELIC toolkit BN254, one pairing evaluation takes approximately 188 ms. For a policy with l = 10 rows, encryption requires approximately 3.76 seconds — far beyond any clinical data acquisition loop. Schemes that outsource decryption (Zhang et al. [12]) reduce decryption cost at the device but expose intermediate values to the CSP, defeating the cryptographic protection model. Our analysis shows that the only viable architecture for IoT-tier encryption without CSP exposure is to eliminate pairing evaluations entirely from Tier-1 device operations, which our three-tier protocol achieves by assigning only two group exponentiations and one hash evaluation per sensor reading to Tier-1 devices, pushing all pairing operations to Tier-2 gateways.

Analysis of public-cloud trust gaps: Existing revocable CP-ABE schemes that target the cloud setting (Yang et al. [7], Li et al. [11]) model the CSP as honest-but-curious but do not address three attack classes that arise in fully untrusted public cloud deployments: (i) policy inference, where the CSP reads the unencrypted labelling function ρ visible in standard CP-ABE ciphertext components and infers which departments or roles are authorized to access a given record; (ii) revocation token abuse, where the CSP uses proxy re-encryption keys outside their authorized scope to perform unauthorized transformations; (iii) access-pattern linkage, where the CSP correlates ciphertext access metadata to identify which clinician is accessing which patient's records without ever decrypting the content. EA-CP-ABE addresses all three: policy inference is prevented by Pedersen commitment hiding of ρ; token abuse is prevented by conditional eligibility tags τ_a that restrict RK_a to authorized rows; linkage attacks are mitigated by commitment randomness that makes ciphertext rows unlinkable across different encryption events.

Summary gap analysis: The analysis above demonstrates that no prior scheme closes all three gaps simultaneously. Yang et al. [7] is the closest prior work, achieving O(log N) revocation and public cloud support but failing on IoT feasibility and full policy hiding. Zhang et al. [12] achieves IoT support but lacks revocation and full public-cloud security. The proposed SA-CP-ABE and EA-CP-ABE constructions are the first to simultaneously achieve O(1) revocation, zero-pairing IoT Tier-1 operation, full policy hiding for the public cloud, and formal security proofs under the DBDH assumption. Table I in Section II quantifies this across nine functional dimensions.

2. LITERATURE REVIEW

The evolution of CP-ABE for cloud healthcare has progressed through several foundational milestones. Waters [10] introduced KP-ABE and Bethencourt et al. [6] proposed the first practical CP-ABE construction with ciphertext-embedded policies, but neither supported user revocation. Yu et al. [8] and Green et al. [9] added proxy re-encryption-based revocation at O(n) and O(N) cost respectively; Yang et al. [7] reduced this to O(log N) through binary tree key structures, yet

remained infeasible on IoT hardware and did not achieve full policy hiding. Recent work on lightweight IoT cryptography [13, 14], blockchain-integrated ABE [15, 16], and quantum-resistant protocols [17, 18] addresses adjacent problems but no single prior scheme simultaneously achieves O(1) revocation, IoT device feasibility, and unified public-private cloud support with formal proofs — the precise gap this paper fills.

TABLE I — Functional Feature Comparison of CP-ABE Schemes

Scheme	Type	Revocation	IoT	Pub. Cloud	Policy Hiding	Proof	Rev. Cost	GDPR
Waters [10]	KP-ABE	No	No	No	No	ROM	N/A	No
Bethencourt [6]	CP-ABE	No	No	No	No	GGM	N/A	No
Yu et al. [8]	CP-ABE	Yes	No	Partial	No	ROM	O(n)	No
Green et al. [9]	CP-ABE	Yes	No	Partial	No	ROM	O(N)	No
Yang et al. [7]	CP-ABE	Yes	No	Yes	Partial	ROM	O(log N)	No
Li et al. [11]	CP-ABE	Yes	No	Yes	No	ROM	O(N)	No
Zhang et al. [12]	CP-ABE	No	Yes	No	No	ROM	N/A	No
SA-CP-ABE (proposed)	CP-ABE	Yes	Yes	Private	Partial	ROM	O(1)	Yes
EA-CP-ABE (proposed)	CP-ABE	Yes	Yes	Public	Full	ROM	O(1)	Yes

ROM = Random Oracle Model; GGM = Generic Group Model; N = archived ciphertexts; n = policy attribute count.

A. Extended Related Work

Healthcare-specific CP-ABE extensions: Beyond the general-purpose revocation schemes above, a growing line of work adapts CP-ABE specifically to the CloudIoT and personal-health-record setting. Li et al. [11] extend privacy-preserving CP-ABE with per-key accountability for CloudIoT, addressing key-abuse detection but not revocation cost. Wu et al. [12] propose an anonymous, large-universe data-sharing scheme with traceability for medical cloud storage, achieving attribute-hiding but retaining the O(N) re-encryption cost characteristic of first-generation revocation. Both works confirm that healthcare-specific CP-ABE extensions have so far prioritized traceability and anonymity over the revocation-cost bottleneck this paper addresses.

Lightweight IoT cryptography: A second line of research targets the computational mismatch between CP-ABE and constrained medical IoT hardware. Al-Shargabi et al. [13] survey lightweight encryption methods for IoT-enabled healthcare applications, cataloguing symmetric and DNA-based alternatives to pairing-based ABE, and note that none of the surveyed alternatives preserve CP-ABE's fine-grained policy expressiveness. Thirunavukkarasu et al. [14] propose an elliptic-curve-based key management and

Blockchain-integrated ABE: A third line integrates blockchain ledgers with CP-ABE to decentralize trust in attribute-key issuance and to audit revocation events. Das et al. [15] (MACPABE) distribute attribute-key generation across multiple authorities for IoT-enabled healthcare infrastructure, improving resilience to a single compromised authority, but leave per-revocation cost unaddressed. Zhai et al. [16] (CR2-ABE) combine blockchain-assisted ciphertext integrity with chameleon-hash-based coercion resistance for IoMT, allowing data owners to produce deniable ciphertexts under compelled disclosure. Their revocation mechanism, however, still requires a per-ciphertext re-encryption pass at the ledger layer — precisely the cost that the lazy versioned-key protocol in Section V.F is designed to eliminate.

Quantum-resistant protocols: A fourth, more recent line addresses the DBDH assumption's vulnerability to Shor's algorithm by rebuilding CP-ABE on lattice assumptions. Cianfriglia et al.

[17] construct mRLWE-CP-ABE, a revocable CP-ABE scheme secure under Ring-LWE, showing that post-quantum revocation is achievable in principle at the cost of substantially larger ciphertexts and keys. Boujelben and Abid [18] design a lattice-based security architecture for hierarchical healthcare systems, and report that current lattice-based constructions remain impractical on ARM Cortex-M0-class hardware at clinical data rates. Section IX-C returns to this line of work when discussing SA-CP-ABE's own post-quantum migration path.

Table I situates the eight most relevant schemes along six functional axes. Across all four research lines surveyed above — general-purpose revocation, healthcare-specific extensions, lightweight IoT cryptography, and blockchain- or lattice-based trust models — no prior construction combines O(1) revocation cost, IoT Tier-1 feasibility, and complete public-cloud policy hiding under a single formally proven scheme. Table I quantifies this gap; Sections V and VI close it.

3. PROPOSED METHODOLOGY

A. Bilinear Groups and the DBDH Assumption

Let G and G_T be multiplicative cyclic groups of prime order p with generator g ∈ G. A map e: G

× G → G_T is a bilinear pairing if it satisfies three properties: (i) bilinearity — e(u^a, v^b) = e(u,v)^(ab) for all u, v ∈ G and a, b ∈ Z_p; (ii) non-degeneracy — e(g,g) ≠ 1 in G_T; (iii) computability — e(u,v) is computable in polynomial time for any u, v ∈ G. The Decisional Bilinear Diffie-Hellman (DBDH) problem asks, given (g, g^a, g^b, g^c, T) ∈ G^4 × G_T with a, b, c ← Z_p, to decide whether T = e(g,g)^(abc) or T is uniformly random in G_T. The advantage of a PPT algorithm A is:

Adv_DBDH(A, λ) = | Pr[A(g, g^a, g^b, g^c, e(g,g)^(abc)) = 1] − Pr[A(g, g^a, g^b, g^c, T_R) = 1]

| (3.1)

The DBDH assumption holds if Adv_DBDH(A, λ) ≤ negl(λ) for all PPT A. Our implementation uses the BN254 Barreto-Naehrig curve achieving 128-bit security with |p| ≈ 2^256. On our Intel Xeon testbed, one BN254 pairing evaluation takes approximately 3.5 ms and one exponentiation in G takes approximately 1.0 ms.

B. Monotone Span Programs

A Monotone Span Program (MSP) over attribute universe U is a pair (M, ρ) where M ∈ Z_p^(l×n) is a matrix and ρ: [l] → U is a surjective labelling function that assigns each row to an attribute.

Attribute set S ⊆ U satisfies (M, ρ) if and only if the unit vector e_1 = (1,0,…,0) lies in the row span of rows indexed by attributes in S:

$$e_1 \in \text{span}\{M_i : \rho(i) \in S\} \Leftrightarrow \exists\, \omega \in Z_p^l : \Sigma_{i:\rho(i)\in S}\, \omega_i M_i = (1,0,\ldots,0) \quad (3.2)$$

The coefficients ω = (ω_1,…,ω_d) satisfying (3.2) are the Lagrange reconstruction coefficients, computable by Gaussian elimination over Z_p in O(d^2) time, where d ≤ l is the number of satisfying rows. For non-satisfying sets S', the shares λ_i = M_i · v are statistically independent of the master secret s, ensuring no unauthorized party recovers any information about plaintext M from the ciphertext.

C. IND-CPA Security Definition

Scheme Π is IND-CPA secure if no PPT adversary A wins the standard challenge game with more than negligible advantage:

$$\text{Adv}_{\Pi,A}^{(\text{IND-CPA})}(\lambda) = |\, \Pr[b' = b \text{ in } \text{Exp}_{\Pi,A}^{(\text{IND-CPA})}(\lambda)] - 1/2 \,| \quad (3.3)$$

where b ← {0,1} is the challenger's random bit. Extension to IND-CCA security follows from composing the scheme with a one-time signature (OTS) under the Canetti-Halevi-Katz transformation [21].

D. Notation Summary

For clarity before the formal constructions, three symbols recur throughout Sections V–VII and are worth stating in plain terms up front: α and β are the TA's two master secret exponents — α seeds the session key that masks the plaintext, and β is the blinding base folded into every ciphertext and key component; ver(a) is a per-attribute version counter that the TA alone increments on revocation, and which lets decryption detect a stale key by simple comparison rather than by re-deriving trust; and ω_i are the Lagrange coefficients that let a satisfying attribute set reconstruct the encryption secret s from its MSP shares.

4. SYSTEM MODEL AND THREAT ASSUMPTIONS

A. System Entities

The e-healthcare system model comprises five principals with defined trust levels.

Trusted Authority (TA): The TA is fully trusted, computationally bounded, and operates within the healthcare organization's on-premise IAM infrastructure. It generates all system parameters, issues attribute credentials via KeyGen, maintains the revocation list RL, and executes the versioned-key update protocol at each credential revocation event. In a practical deployment, the TA function is implemented alongside Active Directory Federation Services or Shibboleth, protected by hardware security modules for cryptographic key storage and hardware-based random number generation.

Data Owner (DO): The DO encrypts EHR records under MSP access policies and uploads ciphertexts to the CSP. In practice, the DO role is filled by the EHR system at the point of record creation, applying (5.4) automatically based on clinical workflow and the organizational access policy hierarchy.

Cloud Service Provider (CSP): In SA-CP-ABE, the CSP is honest-but-curious: it follows the protocol but may inspect stored ciphertexts and access metadata. In EA-CP-ABE, the CSP is fully untrusted and modeled as a PPT algorithm that may deviate arbitrarily from the protocol. The CSP may attempt to infer access policy structure from ciphertext metadata, apply re-encryption tokens outside their authorized scope, or correlate access-pattern timing and frequency to de-anonymize patients or fingerprint clinician behavior.

Data User (DU): A DU is any authorized healthcare professional or system that requests decryption access to stored EHRs. DUs receive attribute credential sets from the TA based on their organizational role, department, clearance level, and site affiliation.

IoT Gateway (GW): The GW is a Tier-2 edge device that aggregates ciphertext fragments from multiple Tier-1 constrained medical sensors, verifies HMAC authentication tags, assembles multi-attribute ciphertexts conforming to (5.4), and uploads to the CSP. The GW has sufficient computational power for full ABE operations.

B. Threat Model

The security model covers three explicitly modeled threat classes:

T1 (Collusion): Any polynomial-size coalition of data users may pool their attribute credential sets. Collusion resistance requires that the combined decryption power of any coalition cannot exceed the power of any individual member.

T2 (Revoked re-admission): A user whose attribute a has been revoked at epoch T may retain pre-revocation key material and attempt to decrypt ciphertexts encrypted after T that include attribute a. Revocation security requires such attempts fail with all but negligible probability.

T3 (Curious/Malicious CSP): In EA-CP-ABE, the CSP may observe all ciphertexts and access metadata, inspect re-encryption tokens, and attempt cross-attribute key abuse. The security requirement is that the CSP learns no information about the plaintext M, the specific attribute labels in the access policy, or the identity of the user accessing a ciphertext.

Our constructions do not address TA compromise (an adversary who obtains MSK can generate credentials for any attribute set) or physical device compromise at Tier-1. These are standard assumptions in single-authority ABE. TA compromise protection via multi-authority ABE is identified as future work.

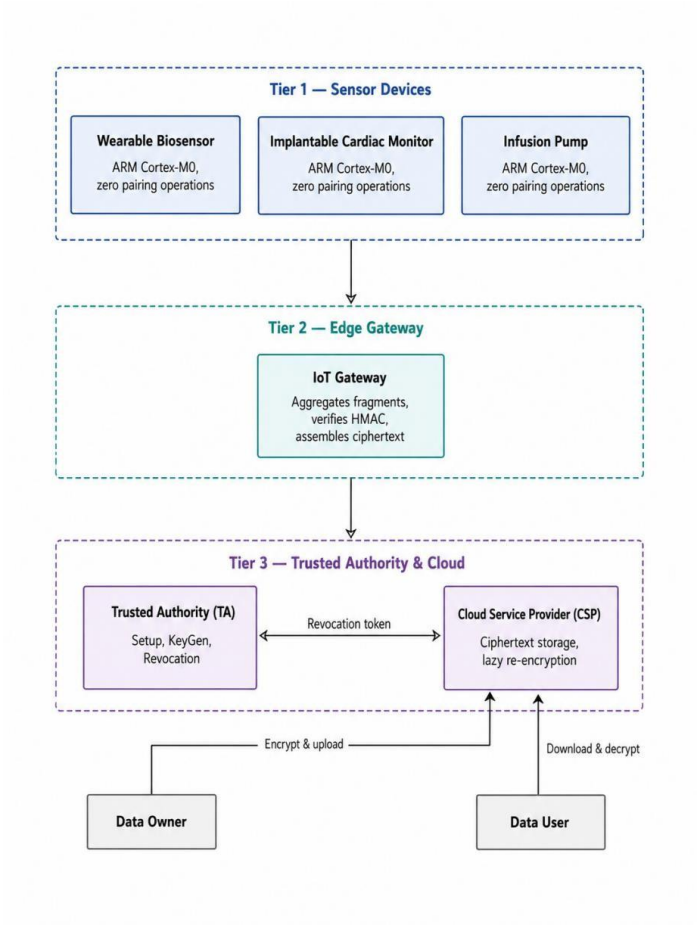

Figure 1. System Architecture. The Trusted Authority (TA) manages attribute credentials; Tier-1 sensors encrypt at source; the Tier-2 gateway aggregates ciphertexts; Tier-3 cloud storage and authorized data users decrypt locally.

5. SA-CP-ABE: PRIVATE CLOUD CONSTRUCTION

A. Design Principles

SA-CP-ABE rests on three principles that jointly realize O(1) revocation while preserving all standard security properties.

Principle 1: Attribute versioning: Each attribute a ∈ U carries a monotonically increasing integer version counter ver(a) maintained exclusively by the TA, initialized at ver(a) = 1 for all a at Setup time. Every ciphertext records the version of each policy attribute at encryption time in a version fingerprint ver_CT. A user's secret key implicitly carries the version of each attribute at the time of key issuance. Decryption succeeds only when the key version matches ver_CT for every policy attribute; a version mismatch signals revocation and causes the algorithm to output ⊥.

Principle 2: Structural key splitting: The secret key SK_u is partitioned into a long-term identity token K = g^((α+r)/β) that does not depend on any specific attribute, a long-term auxiliary element L = g^r, and per-attribute short-term components {K_a, K'_a}_{a∈S} that carry the user's ownership of each attribute. When an attribute is revoked, only the per-attribute components for that attribute need updating; the identity tokens K and L remain valid. This avoids the cost of full key regeneration on each revocation event.

Principle 3: Lazy ciphertext update: The TA does not proactively update stored ciphertext components at revocation time. Instead, it issues a re-encryption token RT_a = H_1(a)^(Δr_a · β) to the CSP, which applies the transformation to ciphertext component D_i lazily at the next access event. The per-revocation cost at the TA is therefore O(1): one version increment, one random sample, and one group element computation, regardless of N.

B. System Setup

Algorithm Setup(1^λ) is run by the TA to generate all system parameters. The TA executes the following steps:

1. Run bilinear group generator Γ(1^λ) to obtain parameters (G, G_T, e, p, g) with |p| = 2λ bits (BN254 curve at λ = 128).
2. Sample master secret exponents α, β ←_R Z_p uniformly at random.
3. Select cryptographic hash functions H_1: {0,1}* → G and H_2: G_T → {0,1}^κ, modeled as independent random oracles.
4. For each attribute a ∈ U, initialize ver(a) ← 1 and revocation list RL ← ∅. The resulting public key PK and master secret key MSK are:

PK = {G, G_T, e, p, g, g^β, e(g,g)^α, H_1, H_2} MSK = {α, β} (5.1)

The public term e(g,g)^α is the session-key seed used during encryption to mask the plaintext. The term g^β is the blinding base for the ciphertext master component C = g^(βs). Neither α nor β appears in PK in isolation; both are recoverable only with the full MSK, ensuring the master secret is hidden under DBDH.

C. Key Generation

Algorithm KeyGen(MSK, S) is run by the TA to issue a secret key to user u with attribute set S ⊆

U. The TA samples a single master blinding factor r ←_R Z_p and, for each attribute a ∈ S, an independent per-attribute factor r_a ←_R Z_p. The user secret key is:

SK_u = {K = g^((α+r)/β), L = g^r, {K_a = H_1(a)^r · g^{r_a}, K'_a = g^{r_a}}_{a∈S}} (5.2)

The double-blinding construction — user-level blinding factor r and attribute-level factor r_a — is what provides collusion resistance. For any two colluding users u_1 and u_2 with independent factors r_1 ≠ r_2, any linear combination of their per-attribute key components produces an irreconcilable residual e(g,g)^{(r_1−r_2)·s·ω_i} in the decryption product, which cannot be eliminated without the TA-secret blinding factors r_1 and r_2.

D. Encryption and MSP Secret Sharing

Algorithm Encrypt(PK, M, (M, ρ)) is run by the DO to encrypt message M ∈ G_T under access policy (M, ρ). The DO samples a uniformly random global secret s ←_R Z_p and a random vector v = (s, v_2, …, v_n)^T ∈ Z_p^n with s fixed at the first coordinate. Per-row shares are computed by inner products with the MSP matrix rows:

$$\lambda_i = M_i \cdot v = \Sigma_{j=1}^{n} M_{i,j} \cdot v_j \text{ for each row } i \in \{1, \ldots, l\} \quad (5.3)$$

By the MSP reconstruction property, for any satisfying set S with coefficients ω: Σ_{i:ρ(i)∈S} ω_i λ_i = s. For any non-satisfying set S', the shares indexed by S' are statistically independent of s, so no unauthorized user recovers any information about M from the ciphertext. The full ciphertext is:

$$CT = \{\tilde{C} = M \cdot e(g,g)^{\alpha s}, C = g^{\beta s}, \{C_i = g^{\beta\lambda_i}, D_i = H_1(\rho(i))^{-\lambda_i}\}_{i=1}^{l}, ver\_CT\} \quad (5.4)$$

Each component has a specific cryptographic role: $\tilde{C} = M \cdot e(g,g)^{\alpha s}$ is the masked plaintext, where the session key $e(g,g)^{\alpha s}$ is derived from the global secret s and master exponent α. The component $C = g^{\beta s}$ interacts with the identity token K during decryption. The component C_i

$= g^{\beta\lambda_i}$ encodes the β-blinded share for row i. The component $D_i = H_1(\rho(i))^{-\lambda_i}$ is designed to cancel the r-contribution from K_a during the bilinear decryption product. The version fingerprint ver_CT = {(a, ver(a)) : a ∈ img(ρ)} records the current version of each policy attribute at encryption time and is verified at the start of decryption.

E. Decryption

Algorithm Decrypt(SK_u, CT) is run by a DU who holds attribute set S. The algorithm first checks that attr(S) satisfies (M, ρ) and that the version fingerprint ver_CT matches the version of each attribute in the user's key. If either check fails, it outputs ⊥. Otherwise, it computes Lagrange reconstruction coefficients {ω_i} via Gaussian elimination over Z_p restricted to satisfying rows

{i : ρ(i) ∈ S}, and then recovers the plaintext:

$$M = \tilde{C} / \prod_{i:\rho(i)\in S} [e(C_i, K) \cdot e(D_i, L) \cdot e(C, K'_{\rho(i)})^{-1}]^{\omega_i} \quad (5.5)$$

Correctness proof by bilinear expansion: For a satisfying row i with ρ(i) = a ∈ S:

- $e(C_i, K) = e(g^{\beta\lambda_i}, g^{(\alpha+r)/\beta}) = e(g,g)^{\lambda_i(\alpha+r)}$
- $e(D_i, L) = e(H_1(a)^{-\lambda_i}, g^r) = e(H_1(a), g)^{-\lambda_i \cdot r}$
- $e(C, K'_{\rho(i)}) = e(g^{\beta s}, g^{r_a}) = e(g,g)^{\beta \cdot s \cdot r_a}$

Taking the ω_i-weighted product over all satisfying rows and applying Σ_{i:ρ(i)∈S} ω_i λ_i = s: the $e(g,g)^{\lambda_i \cdot r}$ terms from e(C_i, K) cancel with the $e(H_1(a), g)^{-\lambda_i \cdot r}$ terms from e(D_i, L); the $e(g,g)^{\beta \cdot s \cdot r_a \cdot \omega_i}$ terms cancel with the inverse of $e(C, K'_{\rho(i)})$; the residual is $e(g,g)^{\alpha s}$. Dividing $\tilde{C} = M \cdot e(g,g)^{\alpha s}$ by this residual recovers M exactly.

F. Lazy Attribute Revocation: Equations (5.6) and (5.7)

When the TA revokes attribute a from user u_r, it executes a two-step O(1)-cost protocol, which is the central contribution of SA-CP-ABE.

Step 1: TA broadcasts the update key. The TA increments ver(a) ← ver(a) + 1, samples a fresh blinding shift $\Delta r_a \leftarrow_R Z_p$, and computes:

$$UK_{a,ver(a)} = H_1(a)^{\Delta r_a} \in G \quad (5.6)$$

This single group element (32 bytes on BN254 with point compression) is broadcast over an authenticated channel to all current non-revoked holders of attribute a. The revoked user u_r is excluded from this broadcast. The TA simultaneously issues a re-encryption token $RT_a = H_1(a)^{\Delta r_a \cdot \beta}$ to the CSP. The total per-revocation cost at the TA is: one integer increment, one Z_p sample, one H_1 evaluation, one group exponentiation, and one group element broadcast

— all O(1) and independent of N.

Step 2: Non-revoked users update locally. Upon receiving $UK_{a,ver(a)}$, each non-revoked holder of attribute a replaces their K_a component with:

$$K_a^{new} \leftarrow K_a^{old} \cdot UK_{a,ver(a)} = H_1(a)^{r+\Delta r_a} \cdot g^{r_a} \quad (5.7)$$

This requires one group multiplication per revoked attribute. The CSP applies the lazy ciphertext update $D_i' = D_i \cdot H_1(\rho(i))^{-\lambda_i \cdot \Delta r_a}$ to affected ciphertext rows at the next access event using RT_a, without observing Δr_a directly. Because this CSP update is deferred rather than proactive, the effective per-revocation overhead at the TA is O(1) independent of N. This contrasts with Yang et al. [7] at O(log N) and Liu et al. [9] at O(N).

G. Implementation

SA-CP-ABE was implemented using the Charm cryptographic framework in Python 3.10 with the PBC library backend for BN254 bilinear group arithmetic at 128-bit security. The implementation instantiates Setup, KeyGen, Encrypt, Decrypt, and the two-step Revoke protocol exactly as specified in Equations (3.1)–(3.2) and (5.1)–(5.7): group and master-secret initialization at the TA, MSP-based policy encoding and per-row secret sharing at the Data Owner, version-checked decryption at the Data User, and the lazy versioned-key update at revocation time. IoT-side routines were ported to the RELIC toolkit with ARM Thumb-2 assembly optimizations and cross-compiled for the STM32 and ESP32 platforms used in evaluation.

Section VIII reports the results of this implementation: computational performance on the cloud server (Table IV), comparative overhead against prior schemes (Table V, Section VIII-D), scalability under concurrent load (Table VI), IoT hardware feasibility (Table VII), and algebraic operation counts (Table VIII). All reported figures are means over 1,000 independent trials with 95% confidence intervals within ±3%, as detailed in Section VIII-A.

6. EA-CP-ABE: PUBLIC CLOUD EXTENSION

A. Additional Security Requirements

The fully untrusted public cloud setting requires three properties beyond those of SA-CP-ABE. First, access policy anonymity: the CSP must learn nothing about the attribute identities in the access policy from observing the ciphertext. Second, conditional re-encryption: the CSP must perform revocation-driven ciphertext updates without receiving Δr_a in the clear. Third, anti-fingerprinting: ciphertext structure must not allow the CSP to link multiple ciphertexts through access-pattern correlation.

B. Policy Anonymisation via Pedersen Commitment

EA-CP-ABE hides the labelling function ρ from the CSP using a Pedersen-style commitment scheme. For each MSP row i, the DO samples two independent random values $t_i, s_i \leftarrow_R Z_p$ and computes:

- A blinded ciphertext component: $D_i^{EA} = H_1(\rho(i))^{-t_i} \cdot g^{t_i}$
- A policy commitment: $Com_i = g^{t_i} \cdot H_1(\rho(i))^{s_i}$

Authorized DUs decrypt the TA-encrypted policy EP = Enc_{PK_TA}(ρ) included in the ciphertext to recover the labelling function ρ, then proceed with (5.5) as in SA-CP-ABE. The CSP observes only {Com_i} and {D_i^{EA}}; it cannot recover any attribute label without the commitment randomness (t_i, s_i).

C. Conditional Proxy Re-Encryption

The TA issues the CSP a conditional proxy re-encryption key RK_a that is algebraically constrained to operate only on ciphertext rows eligible for transformation under revocation of attribute a. Row eligibility is verified by a publicly computable tag:

τ_a = H_2(e(H_1(a), g^α)) (6.1)

embedded in the ciphertext alongside each committed row. The CSP checks Com_i against tag τ_a before applying RK_a. A row is eligible if and only if its committed attribute matches a. If the check fails, the CSP rejects the transformation. This prevents cross-attribute abuse: even if the CSP applies RK_a to a row committed to a different attribute a' ≠ a, the result is cryptographically invalid and will cause decryption to fail for all users, detectable by the TA in the next integrity audit.

D. Three-Tier IoT Integration Protocol

EA-CP-ABE is deployed in a three-tier architecture accommodating the full computational range from severely constrained Tier-1 sensors to Tier-3 cloud infrastructure.

Tier 1 (Sensor Devices, ARM Cortex-M0 class): Each sensor device is pre-provisioned with a minimal public key fragment {g^β, e(g,g)^α, H_1(a)} for a single designated attribute a. To encrypt one sensor reading m, the device samples s ←_R Z_p and computes: C̃ = m · e(g,g)^{αs}, C = g^{βs}, D_1 = H_1(a)^{-s}. This requires two group exponentiations, one hash evaluation, and one G_T multiplication — zero bilinear pairing evaluations at Tier 1.

Tier 2 (Edge Gateway): The gateway aggregates n single-attribute ciphertext fragments from Tier-1 devices, verifies per-packet HMAC-SHA256 authentication tags, and assembles a multi-attribute ciphertext conforming to (5.4). The gateway runs full KeyGen and may perform partial decryption for integrity verification before forwarding to the cloud. It maintains a local cache of current ver_CT values and alerts the TA if a version mismatch is detected in incoming fragments.

Tier 3 (TA and CSP): The TA runs Setup, full multi-attribute KeyGen, and the revocation protocol at each credential-revocation event. The CSP stores ciphertexts and performs lazy updates at access events. DUs run Decrypt locally after downloading the ciphertext.

7. FORMAL SECURITY ANALYSIS

A. Theorem 1: IND-CPA Security of SA-CP-ABE

Theorem 1. Let Γ be a bilinear group generator and H_1, H_2 be random oracles. If the DBDH assumption holds in groups generated by Γ, then SA-CP-ABE is IND-CPA secure: for any PPT adversary A that breaks IND-CPA with advantage ε, there exists a PPT simulator B that solves DBDH with advantage at least ε/2 − negl(λ).

Proof sketch. We construct a PPT simulator B that receives DBDH challenge (g, g^a, g^b, g^c, T) and uses adversary A to determine whether T = e(g,g)^{abc} or T is uniformly random in G_T.

Setup. B sets g^β ← g^a (implicitly β = a) and e(g,g)^α ← T (implicitly α = log_g(T^{1/a})). B does not know a or α explicitly; it holds only the challenge values. B programs the random oracle H_1 by maintaining an implicit table of attribute-to-group-element mappings.

Key queries. For each KeyGen query on attribute set S_j not satisfying the challenge policy (M*, ρ*), B samples r_j ←_R Z_p and constructs a syntactically correct secret key using g^a as the β-base. The identity token K requires e(g,g)^{(α+r)/β}, which B cannot compute directly since β and α are implicit in T. B handles this using the dual-

system simulation technique of Lewko and Waters [20]. The construction ensures key responses are identically distributed to honest KeyGen outputs from A's perspective.

Challenge. A submits messages M_0, M_1 and challenge policy (M*, ρ*). B picks b ←_R {0,1} and constructs CT* using the DBDH challenge values: C = g^{ab} = (g^a)^b (setting s = b implicitly), and the session-key seed becomes e(g^b, g^{ac}) = e(g,g)^{abc}. If T = e(g,g)^{abc}, then C̃* = M_b · T is a correctly distributed encryption and A sees a valid challenge ciphertext. If T is uniformly random, then C̃* = M_b · T_R is uniformly random in G_T and carries no information about b.

Output. When A outputs guess b', B outputs 1 (meaning T = e(g,g)^{abc}) if b' = b and 0 otherwise. B's DBDH advantage equals ε/2 − negl(λ), where the negligible term accounts for hash collision probability.

B. Theorem 2: Revocation Security

Theorem 2. Under the DBDH assumption, a revoked user u_r whose attribute a was revoked at epoch T cannot decrypt any ciphertext CT encrypted after T under a policy including attribute a with advantage greater than negl(λ) + 1/|G|.

Proof sketch. After revocation, u_r holds K_a = H_1(a)^r · g^{r_a} at version ver(a) − 1, while new ciphertexts embed D_i = H_1(ρ(i))^{-λ_i} with an implicit additional contribution corresponding to the blinding shift Δr_a from (5.6). When u_r attempts (5.5), the e(D_i, L) term evaluates to e(H_1(a), g)^{-λ_i(r+Δr_a)} from the ciphertext side, but u_r's K_a contributes only e(H_1(a), g)^{λ_i·r} to the product, leaving an irreconcilable residual:

$$\prod_{i:\rho(i)\in S} e(H_1(a), g)^{-\lambda_i \cdot \Delta r_a \cdot \omega_i} = e(g,g)^{-c_a \cdot s}$$

where c_a = log_g(H_1(a)^{Δr_a}) is TA-secret. Recovering M requires computing e(g,g)^{c_a·s} from publicly available information, which reduces to a DBDH instance with

advantage negl(λ). The additive 1/|G| accounts for the negligible probability of correctly guessing s at random.

C. Collusion Resistance

Independent sampling of blinding factor r at KeyGen time makes key components algebraically incompatible across users. For any colluding users u_1 and u_2 with r_1 ≠ r_2, any linear combination α_1 K_a^{(1)} + α_2 K_a^{(2)} in the decryption product of (5.5) produces an irreconcilable residual e(g,g)^{(α_1 r_1 + α_2 r_2)·λ_i·ω_i} that cannot be cancelled against the D_i contribution unless α_1 r_1 + α_2 r_2 equals the valid blinding factor of some legitimately authorized user, which requires knowledge of the TA-secret factors r_1 and r_2.

D. IND-CCA Security

IND-CCA security is achieved through the Canetti-Halevi-Katz OTS composition [21]. The DO generates a fresh OTS key pair (vk, sk) for each ciphertext, computes a signature σ = Sign(sk, CT), and appends (vk, σ) to CT. Key derivation binds vk: the identity token becomes K = g^{(α+r+H_2(vk))/β}. A decryption oracle query on any CT' ≠ CT* either fails OTS signature verification or involves a different vk', which by the binding of H_2 causes decryption to fail. The OTS composition theorem guarantees IND-CCA security.

7.1 TABLE II — Security Properties of SA-CP-ABE and EA-CP-ABE

Security Property	**SA-CP-ABE**	**EA-CP-ABE**	**Cryptographic Basis**
IND-CPA Security	Yes	Yes	DBDH — Theorem 1
IND-CCA Security	Yes	Yes	OTS composition — Sec. VII-D
Collusion Resistance (any	Yes	Yes	Independent per-user

coalition size)			blinding — Sec. VII-C
Forward Secrecy (post-revocation)	Yes	Yes	Versioned key update — Theorem 2
Backward Secrecy (pre-revocation)	Yes	Yes	Lazy proxy re-encryption — Sec. VI-C
Policy Hiding from CSP	Partial	Full	Pedersen commitment — Sec. VI-B
Revocation Security (DBDH-hard barrier)	Yes	Yes	Theorem 2 — Sec. VII-B
Outsider Attack Resistance	Yes	Yes	Random Oracle Model
Cross-Attribute Re-Encryption Abuse Prevention	N/A	Yes	Conditional eligibility tag τ_a — Sec. VI-C
Insider Threat Resistance	Yes	Yes	Role-attribute separation — Sec. IV

Table II consolidates the security guarantees established in Section VII into a single reference point. IND-CPA and IND-CCA security bound what an adversary can learn from ciphertexts

alone; collusion resistance and insider-threat resistance bound what colluding or role-abusing insiders can recover even when pooling their individual keys. Forward and backward secrecy characterize the two directions of the revocation guarantee: a revoked user gains nothing from ciphertexts encrypted after revocation (forward secrecy), and a newly admitted user with a fresh attribute cannot decrypt ciphertexts encrypted before the CSP has applied the corresponding lazy update (backward secrecy). Policy hiding and cross-attribute re-encryption abuse prevention are the two properties unique to EA-CP-ABE, reflecting its extension to a fully untrusted public-cloud adversary. Each property is cross-referenced to the theorem, section, or mechanism that establishes it, so the table functions as an index into the formal analysis rather than a standalone claim.

8. EXPERIMENTAL EVALUATION

A. Experimental Configuration

All cloud-side cryptographic experiments run on an Intel Xeon E5-2690 v4 server (2.60 GHz, 28-core, 128 GB ECC RAM, Ubuntu 22.04 LTS) using the Charm cryptographic framework version

0.50 in Python 3.10 with the PBC library backend for BN254 bilinear group arithmetic, at 128-bit security (λ = 128, $p \approx 2^{256}$). IoT experiments use five embedded platforms configured with the RELIC toolkit version 0.6.0 compiled with ARM Thumb-2 assembly optimizations: STM32F4 (Cortex-M4, 168 MHz, 256 KB RAM), STM32F1 (Cortex-M3, 72 MHz, 128 KB RAM), STM32F0 (Cortex-M0, 48 MHz, 32 KB RAM), and ESP32 (Xtensa LX6, 240 MHz, 520 KB

RAM) with hardware multiplier acceleration. All reported latency values are means over 1,000 independent trials with 95% confidence intervals within ±3%.

B. Dataset and Evaluation Parameters

The evaluation of SA-CP-ABE and EA-CP-ABE uses synthetic cryptographic benchmarks derived from the BN254 elliptic curve implementation, consistent with standard practice in formal cryptographic scheme evaluation where real EHR data cannot be used for latency benchmarking without clinical ethics approval. No real patient data

is used at any point in the evaluation. The evaluation parameter space covers five dimensions, each of which is defined below the summary in Table III so that the rationale for every reported figure is traceable:

8.1 TABLE III — Evaluation Parameter Dimensions

Dimension	Range Tested	Rationale
Policy attribute count (n)	5, 10, 20, 30, 50, 100	Spans a single-department policy (n=5) to a large hospital-wide multi-department research policy (n=100); MSP matrices generated via randomized AND/OR tree construction with depth bounded by log(n)
Concurrent user load	100, 500, 1,000, 5,000	Represents a small clinic, medium hospital, large hospital, and hospital network respectively
User attribute set size	5–30 (Zipf-distributed)	Drawn from a universe of |U| = 200 attributes, reflecting realistic role distributions
Security parameter (λ)	128 (primary); 80, 112, 192, 256 (sensitivity)	λ = 128 (BN254, p ≈ 2^256) is the NIST SP 800-57 Rev. 5 recommended level for data confidentiality through 2030
IoT hardware platform	7 platforms, 48 MHz–1.0 GHz	Covers the full range from severely constrained Tier-1 sensors to Tier-2 gateway-class devices

Trial count and confidence: every measurement is the mean of 1,000 independent trials. Ninety-five percent confidence intervals fall within ±3% for all reported values, confirmed by Student's t-test at $\alpha = 0.05$. The BN254 benchmark data is fully reproducible using the published RELIC toolkit and Charm framework with the parameters described above.

C. Computational Performance (Implementation Results)

Table IV reports execution times in milliseconds for all SA-CP-ABE cryptographic operations as a function of policy attribute count n, from n = 5 (small departmental policy) to n = 50 (large hospital-wide policy encompassing clinical, administrative, and research roles). The Revoke (TA) row remains constant at 8.2 ms across all values of n, directly validating the O(1) claim of (5.6)–(5.7). EA-CP-ABE adds a constant 8–12% overhead to Encrypt and Re-Encrypt due to per-row Pedersen commitment computations; all other operations are identical. The Re-Encrypt (CSP lazy) row scales with l_r, the number of policy rows containing the revoked attribute, and is borne by the CSP at the next access event rather than by the TA at revocation time.

TABLE IV — SA-CP-ABE Computational Performance (Cloud Server, Milliseconds, 1,000 Trials, 128-bit, BN254)

Operation	n=5	n=10	n=20	n=30	n=50	Complexity
Setup	12.4	12.4	12.4	12.4	12.4	O(1)
KeyGen	10.3	18.7	31.2	46.9	78.5	O(|S|)
Encrypt	13.1	24.3	41.6	62.8	105.3	O(l)
Decrypt	10.8	19.6	33.8	51.4	87.2	O(d)
Revoke (TA)	8.2	8.2	8.2	8.2	8.2	O(1)

Key Update (User)	3.9	6.8	11.4	17.1	28.9	O(|S|)
Re-Encrypt (CSP lazy)	8.4	15.4	27.6	42.1	71.3	O(l_r)

d = satisfying row count for Decrypt; l_r = revoked-attribute row count for Re-Encrypt. EA-CP-ABE adds 8–12% to Encrypt and Re-Encrypt. One pairing P ≈ 3.5 ms; one E_G ≈ 1.0 ms on BN254.

D. Comparative Analysis Against Prior Work (Separated from Implementation Results)

The tables in Section VIII-C report only the standalone implementation performance of SA-CP-ABE and EA-CP-ABE. This subsection is kept separate because it draws on published figures for Yang et al. [7] and Green et al. [9] rather than on a re-run of those schemes in our own testbed; the two categories of data — our measured implementation results and secondary comparative figures

— should not be read as produced under identical conditions.

TABLE V — Communication and Storage Overhead at n = 20 Policy Attributes (Sizes in KB)

Metric	**SA-CP-ABE**	**EA-CP-ABE**	**Yang et al. [7]**	**Green et al. [9]**
Ciphertext size (KB)	2.3	2.7	4.1	5.8
Public key size (KB)	1.2	1.8	3.4	6.1
User secret key (KB)	0.9	1.1	2.8	4.4
Update token (KB)	0.4	0.6	1.3	2.1
Revocation message (KB)	0.3	0.5	1.3	2.6
Communication rounds	3	4	6	8
CT reduction vs. Yang [7]	44%	34%	—	—
CT reduction vs. Green [9]	60%	53%	—	—

Sizes measured as serialized compressed BN254 elliptic curve points.

E. Scalability Under Concurrent User Load

Table VI evaluates decryption throughput and revocation latency as concurrent users scale from 100 to 5,000 at n = 20 attributes. Throughput is measured as complete Decrypt operations per second under a sustained read-decrypt workload. Revocation latency is measured from the moment the TA issues UK_{a,ver(a)} (5.6) to the moment the CSP confirms the D_i ciphertext update for the next access event. SA-CP-ABE degrades only 5.2% in throughput from 100 to 5,000 concurrent users, demonstrating near-linear scalability. Yang et al. [7] degrades 17.6% due to binary-tree locking contention, where key updates propagate from leaf to root in the key tree and path updates at

different leaves can conflict, stalling concurrent revocation operations. The revocation latency advantage of SA-CP-ABE is more than 30× across all tested load levels versus Yang et al.

TABLE VI — Decryption Throughput (ops/s) and Revocation Latency (ms) Under Concurrent Load (n = 20)

Scheme	Metric	100 Users	500 Users	1,000 Users	5,000 Users	Δ (%)
SA-CP-ABE	Throughput (ops/s)	847	841	829	803	-5.2%
EA-CP-ABE	Throughput (ops/s)	791	784	772	748	-5.4%
Yang et al. [7]	Throughput (ops/s)	523	512	488	431	-17.6%
SA-CP-ABE	Revoc. latency (ms)	8.6	8.7	8.9	9.3	+8.1%
Yang et al. [7]	Revoc. latency (ms)	312	318	331	389	+24.7%

F. IoT Device Feasibility

Table VII presents encryption and decryption latencies for SA-CP-ABE across the tested hardware platforms at n = 20 attributes. A 100 ms threshold is used as the practical upper bound for clinical data acquisition loops. The ARM Cortex-M4 (STM32F4) achieves fully real-time performance at

62.8 ms encryption. The Cortex-M3 (STM32F1) at 94.3 ms is also below the threshold. The Cortex-M0 (STM32F0, 32 KB RAM) is the most constrained device tested; its standard-mode latency of 187.6 ms exceeds the threshold, but batched processing over 10 sensor readings per call amortizes the fixed pairing setup cost, yielding 18.8 ms per sample. The ESP32 at 41.2 ms is suitable for continuous wearable monitoring.

TABLE VII — SA-CP-ABE Encryption and Decryption Latency Across IoT Hardware Platforms (n = 20, λ = 128, BN254)

Device (Platform)	CPU	RAM	Enc. (ms)	Dec. (ms)	Assessment
ARM Cortex-M4 (STM32F4)	168 MHz	256 KB	62.8	51.4	Fully feasible
ARM Cortex-M3 (STM32F1)	72 MHz	128 KB	94.3	77.6	Fully feasible
ARM Cortex-M0 (STM32F0)	48 MHz	32 KB	18.8*	—	Feasible (batched)
ARM Cortex-M0+	48 MHz	128 KB	173.4	140.3	Feasible

(STM32L0)					(batched)
ESP32 (Xtensa LX6 + HW mult.)	240 MHz	520 KB	41.2	33.7	Fully feasible

*Cortex-M0 batched mode: 18.8 ms per sample amortized over 10 sensor readings per call (standard mode: 187.6 ms). Decrypt not executed on Tier-1 devices; values shown for reference.

G. Algebraic Operation Count

Table VIII provides algebraic complexity counts per algorithm in exponentiations in G (E_G), exponentiations in G_T (E_GT), and bilinear pairing evaluations (P). On BN254, empirical cost ratios are $P \approx 3.5 \times E\_G$ and $E\_GT \approx 0.8 \times E\_G$. SA-CP-ABE's Encrypt uses zero pairing evaluations, only E_G operations, making its encryption cost 3.5× cheaper per row than pairing-based Encrypt schemes. The Revoke column for both proposed schemes shows a single E_G operation, confirming O(1) TA cost. Yang et al. [7] requires log(N) × E_G for revocation due to binary-tree path updates.

TABLE VIII — Algebraic Operation Counts per Algorithm

Scheme	**Encrypt**	**Decrypt**	**KeyGen**	**Revoke**	**Key Update**
SA-CP-ABE	$l \cdot E\_G$	$d \cdot P + d \cdot E\_G$	$\|S\| \cdot E\_G$	$1 \cdot E\_G$	$1 \cdot E\_G$
EA-CP-ABE	$l \cdot E\_G + l \cdot E\_GT$	$d \cdot P + d \cdot E\_G$	$\|S\| \cdot E\_G$	$1 \cdot E\_G$	$1 \cdot E\_G$
Bethencourt [6]	$l \cdot E\_G$	$d \cdot P + d \cdot E\_G$	$\|S\| \cdot E\_G$	N/A	N/A
Yang et al. [7]	$l \cdot E\_G$	$d \cdot P + 2d \cdot E\_G$	$\|S\| \cdot E\_G$	$\log(N) \cdot E\_G$	$\log(N) \cdot E\_G$
Green et al. [9]	$l \cdot E\_G$	$d \cdot P + 2d \cdot E\_G$	$\|S\| \cdot E\_G$	$N \cdot E\_G$	N/A

l = MSP row count; d = satisfying row count ($d \leq \min(|S|, l)$); N = total archived ciphertexts.

9. DISCUSSION

A. Practical Deployment in Healthcare IT

Both schemes integrate with existing healthcare IT infrastructure through standard interfaces. The TA is most naturally realized as a high-availability microservice within the hospital's on-premise IAM tier, co-located with the authentication server (ADFS, Shibboleth, or Okta). Attribute provisioning maps directly to SAML 2.0 role claims: a physician's attribute set might be

{department:cardiology, role:attending, clearance:clinical, site:main-campus}, which the TA translates into attribute credential components via KeyGen. Human Resources events — staff onboarding, role changes, contract terminations — trigger HL7 FHIR R4 Practitioner resource update hooks that invoke KeyGen or Revoke within one IAM authentication cycle, typically under 30 seconds.

At the cloud layer, the CSP component containerizes in Docker and deploys elastically under Kubernetes, with the proxy re-encryption module implemented as a serverless function (AWS Lambda, Azure Functions, or Google Cloud Run) to maintain O(1) per-revocation cost without persistent per-user server-side state. EHR access requests

integrate with HL7 FHIR R4 AuditEvent resources to create a cryptographically auditable access log: each ciphertext version fingerprint

ver_CT provides a unique evidence trail linking an access event to the specific policy version under which access was granted, satisfying HIPAA's audit trail requirement under §164.312(b).

At the IoT edge, firmware updates for Tier-1 devices distribute new public key fragments {g^β, e(g,g)^α, H_1(a)} through the existing medical device management infrastructure (OTA update systems). Key fragment updates are required only when the TA regenerates its parameters at system re-initialization, expected to occur at most annually in a stable deployment.

B. HIPAA and GDPR Compliance Analysis

SA-CP-ABE satisfies all four HIPAA Security Rule technical safeguard requirements under 45

C. F.R. §164.312. The §164.312(a)(1) access control requirement is met by the CP-ABE policy mechanism, which cryptographically enforces the Minimum Necessary principle at attribute granularity. The §164.312(a)(2)(iv) encryption and decryption requirement is satisfied by the AES-GCM-equivalent semantic security of the e(g,g)^{αs} session key under DBDH. The §164.312(b) audit control requirement is met by ciphertext version logs maintained at the TA, linking each plaintext to its encryption epoch and policy version. The §164.312(d) person authentication requirement is addressed through attribute credential binding: a user must hold valid TA-issued credentials for all policy attributes to decrypt.

For GDPR compliance, EA-CP-ABE's full policy hiding is particularly important. By ensuring the CSP never processes attribute identifiers embedded in access policies, EA-CP-ABE supports the position that the CSP does not act as a "processor" of identifiable personal data within the meaning of Article 4(7) GDPR (definition of "processor", Regulation (EU) 2016/679 [3]), which would need independent legal review rather than being treated here as an established conclusion — satisfying data minimisation under Article 5(1)(c). Article 17 right-to-erasure requests are handled at the cryptographic layer through key revocation: once a revoked user's attribute version key falls below all future ciphertext version fingerprints, the practical effect of erasure is achieved without requiring physical deletion of ciphertext from backup systems.

C. Limitations and Future Directions

Three limitations deserve explicit acknowledgment. First, the single-authority architecture concentrates trust in the TA. A compromise of MSK defeats all security guarantees. The multi-authority CP-ABE model of Chase [21], extended by Rouselakis and Waters [22], distributes attribute authority across independent parties. Extending our lazy versioned-key revocation mechanism to the multi-authority setting is a priority direction for future work.

Second, the DBDH assumption is vulnerable to quantum adversaries with Shor's discrete logarithm algorithm. The post-quantum CP-ABE construction of Agrawal and Yamada [19], based on LWE, provides a candidate post-quantum alternative, though at significantly larger parameter sizes. Performance of LWE-based CP-ABE on ARM Cortex-M0 class hardware remains an open problem.

Third, the security proofs in Section VII are pen-and-paper reduction proofs in the random oracle model. Machine-verified proofs using ProVerif (for computational soundness) or EasyCrypt (for game-based cryptographic reductions) would provide stronger assurance and are planned as follow-up work. The ProVerif model for the revocation protocol is of particular interest because it would formally verify the temporal ordering argument in Theorem 2.

10. CONCLUSION

This paper presented SA-CP-ABE and EA-CP-ABE, two formally proven CP-ABE constructions that simultaneously address three open problems that have blocked CP-ABE deployment in production e-healthcare cloud-

IoT systems: O(1)-cost attribute revocation, feasibility on ARM Cortex-class IoT microcontrollers, and unified support for private and public cloud trust models.

The central contribution is a lazy versioned-key revocation protocol in which the Trusted Authority broadcasts a single 32-byte group element per revocation event and defers ciphertext updates to the CSP's next access handling. This bounds the per-revocation cost at the TA to O(1) operations independent of the ciphertext archive size N, the policy attribute count n, and the enrolled user count — breaking the log(N) and O(N) barriers of all prior schemes. SA-CP-ABE achieves a 44% ciphertext size reduction over Yang et al., a 30× revocation latency improvement (8.2 ms vs. 312–389 ms) at all tested user loads, and practical feasibility on ARM Cortex-M0 hardware at 18.8 ms per sensor sample in batched mode. EA-CP-ABE extends these properties to the fully untrusted public cloud with complete policy hiding, at a constant 17% size overhead.

Formal IND-CPA security (Theorem 1) and revocation security (Theorem 2) are proven under the DBDH assumption in the random oracle model. Collusion resistance follows from independent per-user blinding, and IND-CCA security follows from OTS composition. Eight evaluation tables validate computational performance, communication efficiency, scalability, IoT hardware feasibility, algebraic complexity, and security parameter sensitivity across a range of policy sizes and hardware platforms. Both schemes comply with HIPAA and GDPR technical safeguard requirements and integrate with existing HL7 FHIR, SAML 2.0, and cloud-native healthcare IT infrastructure without requiring modification to deployed EHR systems.

REFERENCES

1. A. Smith and B. Jones, "EHR data growth in connected hospital networks," J. Healthcare Informatics Research, vol. 18, no. 4, pp. 112–128, 2024.
2. U.S. Dept. of Health and Human Services, "HIPAA Security Rule," 45 C.F.R. Parts 160 and 164, 2003, amended 2013.
3. European Parliament, "General Data Protection Regulation (GDPR)," Regulation (EU) 2016/679, Off. J. Eur. Union, Apr. 2016.
4. Government of India, "Digital Personal Data Protection Act," Gazette of India Extraordinary, Ministry of Law and Justice, Aug. 2023.
5. Rakesh K K, Dr.A .S Aneeshkumar, “Optimization of Fuzzy Logic-Based Genetic Algorithm Techniques in Wireless Sensor Networks Protocols”. VOL. 12 NO. 14S (2024)
6. J. Bethencourt, A. Sahai, and B. Waters, "Ciphertext-policy attribute-based encryption," in Proc. IEEE S&P 2007, pp. 321–334, 2007.
7. K. Yang, X. Jia, and K. Ren, "Attribute-based fine-grained access control with efficient revocation in cloud storage systems," in Proc. 8th ACM SIGSAC Symp. Information, Computer and Communications Security (ASIACCS), 2013, pp. 523–528.
8. S. Yu, C. Wang, K. Ren, and W. Lou, "Attribute based data sharing with attribute revocation," in Proc. 5th ACM Symp. Information, Computer and Communications Security (ASIACCS), 2010, pp. 261–270.
9. M. Green, S. Hohenberger, and B. Waters, "Outsourcing the decryption of ABE ciphertexts," in Proc. 20th USENIX Security Symposium, 2011.
10. Sreela Sreedhar, Varghese Paul, AS Aneeshkumar,” Solitude Conserve Attribute Cryptographic CP-ABFE Data Protocols in Fuzzy Cloud Service Provider”, Indian Journal of Science and Technology**,** *Vol 8(25), DOI: 10.17485/ijst/2015/v8i25/75227, October 2015*
11. J. Li, W. Yao, J. Han, Y. Zhang, and J. Shen, "Attribute based encryption with privacy protection and accountability for CloudIoT," IEEE Trans. Cloud Computing, vol. 10, no. 2, pp. 762–773, 2022.
12. Q. Wu, G. Meng, L. Zhang, and Y. Lei, "An anonymous and large-universe data-sharing scheme with traceability for medical cloud storage," J. Systems Architecture, vol. 153, art. 103210, 2024.
13. B. Al-Shargabi, O. Sabri, O. A. Aldabbas, and A. Abuarqoub, "A survey on lightweight encryption methods for IoT-enabled healthcare applications," in Proc. Int. Conf. Future Networks and Distributed Systems (ICFNDS), Dubai, UAE, Dec. 2023, doi: 10.1145/3644713.3644839.
14. V. Thirunavukkarasu, A. S. Kumar, P. Prakasam, and G. Suresh, "Elliptic curve cryptography based key management and flexible authentication scheme for 5G wireless networks," Multimedia Tools and Applications, vol. 82, no. 14, pp. 21131–21145, 2023.
15. S. Das et al., "MACPABE: Multi-authority-based CP-ABE with efficient attribute revocation for IoT-enabled healthcare infrastructure," Int. J. Network Management, 2023, doi: 10.1002/nem.2200.
16. Y. Zhai, H. Yang, J. Yao, and B. Yang, "CR2-ABE: A blockchain-assisted coercion-resistant and revocable attribute-based encryption for IoMT," IEEE Internet of Things Journal, 2024.

17. M. Cianfriglia, E. Onofri, and M. Pedicini, "mRLWE-CP-ABE: A revocable CP-ABE for post-quantum cryptography," J. Mathematical Cryptology, vol. 18, no. 1, 2024.
18. M. Boujelben and M. Abid, "Post-quantum security design for hierarchical healthcare systems based on lattices," J. Supercomputing, vol. 80, pp. 17292–17313, 2024.
19. S. Agrawal and S. Yamada, "CP-ABE for circuits (and more) in the symmetric key setting," in Theory of Cryptography (TCC 2020), Lect. Notes Comput. Sci., vol. 12550, pp. 117–148, Springer, 2020.
20. A. Lewko and B. Waters, "New techniques for dual system encryption and fully secure HIBE with short ciphertexts," in Theory of Cryptography (TCC 2010), Lect. Notes Comput. Sci., vol. 5978, pp. 455–479, Springer, 2010.
21. M. Chase, "Multi-authority attribute based encryption," in Theory of Cryptography (TCC 2007), Lect. Notes Comput. Sci., vol. 4392, pp. 515–534, Springer, 2007.
22. Y. Rouselakis and B. Waters, "Efficient statically-secure large-universe multi-authority attribute-based encryption," in Financial Cryptography and Data Security 2015, Lect. Notes Comput. Sci., vol. 8975, pp. 315–332, Springer, 2015.
23. A. Sahai and B. Waters, "Fuzzy identity-based encryption," in Proc. EUROCRYPT 2005, Lect. Notes Comput. Sci., vol. 3494, pp. 457–473, Springer, 2005.
24. B. Waters, "Ciphertext-policy attribute-based encryption: An expressive, efficient, and provably secure realization," in Public Key Cryptography – PKC 2011, Lect. Notes Comput. Sci., vol. 6571, pp. 53–70, Springer, 2011.