Privacy-Preserving Global Patient Identification Using Cancelable Biometrics and HL7 FHIR

A Framework for Cross-Border Health Information Exchange in Resource-Constrained Settings

The Curely Identity Framework (CIF)

Working Paper — Version 2.0

2026

Abstract

Approximately 850 million people worldwide lack any form of officially recognized identification, a deficit that systematically excludes them from formal healthcare. Closing this gap requires identity infrastructure that can authenticate individuals across jurisdictions without storing recoverable biometric data — a requirement that simple cryptographic hashing cannot satisfy, because biometric captures are inherently noisy and non-deterministic. This paper proposes the Curely Identity Framework (CIF), an architecture that combines ISO/IEC 24745-compliant biometric template protection with HL7 FHIR R4/R5 interoperability to enable patient matching across health systems while preserving privacy.

The framework's core contribution is the explicit treatment of the noise problem in biometric matching. Rather than relying on cryptographic hashes (which collapse under sensor noise), CIF uses cancelable biometrics combined with fuzzy commitment schemes and, optionally, homomorphic encryption for 1:N matching against a protected gallery. We specify a multimodal capture pipeline (fingerprint, iris, face), define indexing strategies for large-scale deduplication on protected templates, and propose realistic accuracy targets (FMR ≤ 10^{-5}, FNMR ≤ 2%, FTE ≤ 1.5% with multimodal fallback). The framework is positioned as a federated, regionally-piloted system aligned with World Bank ID4D principles and WHO SMART Guidelines, rather than as a centralized global registry. We discuss governance constraints, including data sovereignty under GDPR and emerging African Union Convention 108+ regimes, and identify the open research problems — particularly secure 1:N search at population scale — that remain prerequisites to deployment.

Keywords

Biometric template protection, cancelable biometrics, fuzzy commitment, HL7 FHIR, Privacy-by-Design, ISO/IEC 24745, global patient identifier, federated identity, ID4D, digital public infrastructure.

1. Introduction

Reliable patient identification is the precondition for almost everything else in modern healthcare: longitudinal records, immunization schedules, insurance reimbursement, public health surveillance, and continuity of care across providers. Where identification is weak, these functions degrade or fail. The World Bank's Identification for Development (ID4D) initiative estimates that 850 million people — disproportionately women, children, refugees, and the rural poor — lack any form of officially recognized identity (World Bank, 2024). The consequences for healthcare are direct and measurable: unreliable vaccination coverage, untracked chronic disease management, and elevated insurance fraud through duplicate or ghost beneficiaries.

The naïve technical solution — issue every person a unique number tied to their biometrics, store the biometrics centrally, and let any authorized clinician query the database — is unacceptable on three grounds. First, centralized biometric databases are catastrophic single points of failure: a breach compromises a population's biometrics permanently, since people cannot reissue their fingerprints. Second, cross-border data flows are constrained by national data protection regimes that are increasingly assertive about data sovereignty. Third, the populations most in need of identification (refugees, undocumented migrants, marginalized minorities) are also those most at risk from surveillance-capable identity infrastructure.

This paper proposes a framework that attempts to satisfy three constraints simultaneously: (i) provide a stable, unique identifier for a patient across health encounters and jurisdictions; (ii) make the underlying biometric data non-recoverable, even to the system operator; and (iii) interoperate with existing health information systems through open standards. The framework, which we call the Curely Identity Framework (CIF), is not a single algorithm but a layered architecture combining cancelable biometric transforms, fuzzy commitment schemes for template binding, HL7 FHIR for interoperability, and a federated trust model for cross-border recognition.

We are explicit throughout about what is solved, what is approximated, and what remains an open research problem. In particular, efficient privacy-preserving 1:N matching at population scale is not a solved problem, and our proposed approach uses indexing heuristics that trade some privacy for tractability. We argue that this tradeoff is acceptable in the proposed deployment model — federated regional pilots, not a single global database — but acknowledge the limitations.

2. Background and Problem Statement

2.1 The Identification Gap in Healthcare

Health systems have historically issued their own functional identifiers — Medical Record Numbers (MRNs) at the facility level, payer IDs at the insurance level — none of which interoperate across organizational boundaries. Where foundational national IDs exist, they are often not used consistently in clinical settings or are unavailable to undocumented populations. The result is fragmentation: a patient's

record exists as multiple disconnected slices across facilities, with no reliable way to link them. The cost is measurable. Duplicate medical records inflate operational costs by an estimated 8–12% in fragmented systems and contribute to adverse safety events (ECRI Institute, 2017; Pew Charitable Trusts, 2018).

2.2 Why Cryptographic Hashing Alone Does Not Work for Biometrics

It is tempting to treat biometric identification as a password problem and apply standard cryptographic hashing (SHA-256, Argon2) to the captured template. This does not work, for a fundamental reason: biometric captures are noisy. Two presentations of the same finger to the same sensor produce different raw images, and after feature extraction produce different feature vectors. The avalanche property of cryptographic hashes — a single bit flip in the input produces a completely different output — is precisely the wrong property for biometric matching. Hashing two scans of the same finger with SHA-256 yields two unrelated hashes, and the system can never recognize the user again.

The correct framing, formalized in ISO/IEC 24745:2022, is biometric template protection (BTP). A BTP scheme must satisfy three properties: (1) irreversibility — the protected template cannot be inverted to recover the original biometric; (2) unlinkability — protected templates from the same person enrolled in different systems cannot be linked; and (3) renewability — a compromised template can be revoked and reissued from the same biometric. The literature offers several families of schemes that approximate these properties, summarized in Section 4.

2.3 Scope of This Paper

We restrict our scope to the identification and authentication layer of a health information exchange. We do not address the clinical content of the FHIR resources exchanged, the underlying network transport security (assumed to follow standard TLS 1.3 / mTLS practice), or the broader legal harmonization required for cross-border health data flows. Where these adjacent concerns affect feasibility, we flag them explicitly rather than gloss over them.

3. Challenges in Global Patient Identification

The challenges fall into four categories: technical, infrastructural, demographic, and governance. Table 1 summarizes the primary impediments and the mitigations CIF proposes.

Table 1. Challenges and proposed mitigations.

Category	Challenge	CIF Mitigation
Technical	Biometric noise breaks deterministic hashing; cross-sensor variability degrades matching.	Cancelable biometrics with fuzzy commitment; sensor-normalized feature extraction per ISO/IEC 19794.
Technical	Privacy-preserving 1:N matching is computationally expensive at population scale.	Two-stage matching: cancelable index lookup, then secure verification on candidate set (typically <50).

Category	Challenge	CIF Mitigation
Infrastructural	Low connectivity in last-mile clinics; intermittent power; latency-sensitive workflows.	Edge-resident enrollment with offline-capable templates; asynchronous synchronization with conflict resolution.
Demographic	Failure to Enroll (FTE) for manual laborers (degraded fingerprints), elderly, and children under 5.	Multimodal capture (fingerprint + iris + face); alternative credential pathways for FTE cases.
Demographic	Disparate accuracy across skin tone, age, gender — well-documented in face recognition.	Demographic-aware threshold calibration; published per-subgroup FMR/FNMR; preference for fingerprint and iris over face where possible.
Governance	Cross-border data flow restrictions; divergent consent regimes; jurisdictional liability gaps.	Federated architecture with local data residency; FHIR Consent resources; bilateral trust frameworks rather than central registry.
Governance	Risk of function creep into surveillance.	Purpose-bound tokens; query-level audit (FHIR AuditEvent); cryptographic enforcement of "yes/no" responses, not bulk data export.

4. Related Work

4.1 Biometric Template Protection

Biometric template protection emerged as a distinct research area following Ratha, Connell, and Bolle's (2001) seminal work on cancelable biometrics, which introduced the idea of applying a non-invertible, user-specific transform to biometric features before storage. Subsequent work has produced several scheme families, surveyed comprehensively by Jain, Nandakumar, and Nagar (2008) and more recently by Patel, Ratha, and Chellappa (2015). The principal families are:

- **Cancelable biometrics** (Ratha et al., 2001): A non-invertible distortion (e.g., random permutation, BioConvolving) is applied to the feature vector. If compromised, a new transform parameter is issued and re-enrollment is unnecessary if the raw biometric is retained — but in our context we do not retain raw biometrics, which introduces a re-enrollment requirement on revocation.
- **Biohashing** (Teoh, Goh, and Ngo, 2006): Combines biometric features with a user-specific token via random projection. Strong accuracy when the token is secret; security degrades to that of the underlying biometric if the token is leaked (the "stolen token" scenario).
- **Fuzzy commitment** (Juels and Wattenberg, 1999): Binds a secret key to a biometric using error-correcting codes. Tolerates bounded noise. Well-suited to iris codes, which have a stable bit representation.

- **Fuzzy vault** (Juels and Sudan, 2002; Nandakumar, Jain, and Pankanti, 2007): Hides a secret in a polynomial whose evaluation points are mixed with chaff points; the genuine biometric reveals enough true points to reconstruct the polynomial.
- **Homomorphic encryption-based matching** (Bringer, Chabanne, and Patey, 2013; Boddeti, 2018): Computes match scores directly on encrypted templates. Strongest privacy guarantees, but computationally expensive — particularly for 1:N search.

ISO/IEC 24745:2022 codifies the requirements (irreversibility, unlinkability, renewability) that any BTP scheme should satisfy. ISO/IEC 30136:2018 specifies the methodology for evaluating these properties empirically. We rely on both standards as the basis for the framework's claims.

4.2 Digital Identity Models

The World Bank's ID4D Practitioner's Guide (World Bank, 2019; updated 2024) distinguishes centralized, federated, and user-centric (self-sovereign) identity models. The European Union's eIDAS Regulation, particularly the 2024 eIDAS 2.0 update introducing the European Digital Identity Wallet, represents the most mature federated implementation. WHO's SMART Guidelines and the Digital Public Goods Alliance's recommendations explicitly call for open-standards-based, federated approaches in health-sector identity (WHO, 2023). The Modular Open-Source Identity Platform (MOSIP), deployed in Morocco, the Philippines, Ethiopia, and several other countries, demonstrates that a federated foundational ID can be built and operated at national scale.

4.3 Health Information Exchange Standards

HL7 FHIR (versions R4 and R5) is the de facto standard for modern health data exchange. The IHE Patient Identifier Cross-Reference (PIXm) and Patient Demographics Query for Mobile (PDQm) profiles, both FHIR-based, define the patterns for patient matching across domains. The HL7 Identity Matching Implementation Guide (HL7 International, 2024) standardizes the demographic attributes used for probabilistic matching. The framework presented here treats these as the interoperability substrate rather than replacing them.

5. The Curely Identity Framework

CIF is a layered architecture with five components: (1) a multimodal biometric capture layer, (2) a template protection layer combining cancelable transforms and fuzzy commitment, (3) an indexing and matching layer supporting 1:1 verification and bounded 1:N identification, (4) an HL7 FHIR interoperability layer, and (5) a governance layer covering consent, audit, and federation.

5.1 Design Principles

- **Privacy by default.** No raw biometric image or unprotected feature vector is persisted after enrollment. Capture buffers are zeroed within the enrollment session.

- **Federation over centralization.** Each jurisdiction (country, health authority) operates its own protected template store. Cross-border matching is performed via bilateral or multilateral trust frameworks, not by replicating data centrally.
- **Standards over proprietary protocols.** HL7 FHIR R4/R5, ISO/IEC 24745, ISO/IEC 19794 feature formats, OAuth 2.0 / SMART-on-FHIR for access control.
- **Honest about tradeoffs.** Where a privacy-preserving primitive is too slow for the use case, the system uses a less private but faster primitive and documents the tradeoff explicitly in the threat model, rather than silently weakening guarantees.

5.2 The Role of Machine Learning

The framework uses machine learning in three specific, bounded roles, none of which involve learned representations of identity. First, capture-time quality assessment (a convolutional model derived from NIST NFIQ 2 for fingerprint and equivalent open models for iris and face) accepts or rejects a sample at enrollment, reducing FTE and downstream FMR/FNMR. Second, demographic-aware threshold calibration adapts decision thresholds based on capture conditions and demographic factors disclosed at enrollment, mitigating known accuracy disparities. Third, enrollment anomaly detection flags patterns consistent with synthetic identity fraud or coerced enrollment for human review. These are auxiliary services; the core identity claim is cryptographic, not learned.

6. Technical Architecture

6.1 Enrollment Workflow

Enrollment proceeds in six steps:

1. Biographic capture: Name, date of birth, sex, and other ISO/IEC 5218-compliant attributes are collected and persisted as a FHIR Patient resource.
2. Biometric capture: Multimodal samples are acquired — typically four fingerprints (right and left index, right and left thumb), both irises, and a face image. ML-based quality assessment runs in real time; rejected samples trigger recapture.
3. Feature extraction: Each modality is processed by its standardized feature extractor (NIST MINEX-compliant minutiae for fingerprint, ISO/IEC 19794-6 IrisCode for iris, ISO/IEC 19794-5 for face).
4. Cancelable transformation: A user-specific (or session-specific) random parameter is applied to each feature vector to produce a cancelable template. The transform parameter is itself protected, either by binding to a user-held credential or by storing in a hardware security module.
5. Fuzzy commitment: For iris codes specifically, a fuzzy commitment binds a randomly generated identifier (the C-GPI) to the cancelable template via BCH or Reed-Solomon error-correcting

codes. The commitment can be opened only by presenting an iris sample sufficiently close to the original.

6. Deduplication: The cancelable templates are queried against the existing protected gallery using the indexing scheme described in §6.3. If no match is found, the C-GPI is recorded as a new identity. If a match is found, the enrollment is rejected pending human review.

6.2 Authentication Workflow

At each subsequent encounter, the patient presents one or more biometric samples. The same feature extraction and cancelable transformation are applied (using the same transform parameter, retrieved via the patient's credential or a 1:N lookup over the cancelable index). The system attempts to open the fuzzy commitment; success yields the C-GPI, which is then used to retrieve the patient's FHIR record subject to the consent policies attached to that record.

6.3 The 1:N Matching Problem and Our Approach

Privacy-preserving 1:N identification at population scale is the framework's hardest technical problem and the one most likely to constrain real-world deployment. A naïve approach — running a secure comparison against every enrolled template — is O(N) in cryptographic operations per query and becomes infeasible above roughly 10^5 enrollees with current homomorphic encryption performance.

CIF uses a two-stage matching pipeline. Stage one builds an index over the cancelable templates using locality-sensitive hashing (LSH) tuned to the metric used by the underlying biometric (Hamming distance for iris codes, modified Jaccard for fingerprint minutiae). LSH is not privacy-preserving in the strong cryptographic sense — it reveals approximate proximity in template space — but combined with the cancelable transform it leaks bounded information about the underlying biometric. Stage one returns a candidate set, typically of size 10–50. Stage two runs a secure verification (fuzzy commitment opening, or homomorphic comparison if higher assurance is required) only against the candidate set, making the secure operation count O(1) in the enrollee population.

This is a deliberate privacy-performance tradeoff. The threat model assumes that an adversary who compromises the index learns approximate template proximity but not the underlying biometric, and that approximate proximity in the cancelable template space is significantly less sensitive than raw biometric data. We do not claim this is equivalent to fully homomorphic search; we claim it is the best currently-achievable point on the privacy-performance frontier for population-scale deployment, and we recommend periodic re-evaluation as homomorphic search performance improves.

6.4 Layer Summary

Table 2. CIF architectural layers.

Layer	Function	Key Standards
Capture	Multimodal biometric acquisition with real-time quality assessment.	ISO/IEC 19794, NIST NFIQ 2
Template Protection	Cancelable transform; fuzzy commitment for identifier binding.	ISO/IEC 24745, ISO/IEC 30136
Matching	LSH index for candidate generation; secure verification on candidates.	Application-specific
Interoperability	Patient lookup, consent retrieval, audit emission over FHIR.	HL7 FHIR R4/R5, IHE PIXm/PDQm
Trust & Governance	Federation agreements, access control, consent enforcement.	OAuth 2.0, SMART-on-FHIR, eIDAS-aligned trust lists

7. HL7 FHIR Interoperability Model

CIF treats HL7 FHIR R4/R5 as the wire-level protocol for all patient lookups, consent operations, and audit emissions. The specific resources used are:

- **Patient:** Stores biographic attributes and a system-issued identifier (the C-GPI carried as a Patient.identifier with a designated system URI). The C-GPI is the only persistent link to the protected biometric template; no biometric data appears in the Patient resource itself.
- **Consent:** Records the patient's explicit authorization for specific purposes and scopes of data sharing. Each cross-jurisdictional query checks for an applicable Consent resource before proceeding.
- **AuditEvent:** Emitted for every C-GPI lookup, including the requesting party, purpose-of-use, and outcome. Patients have read access to their own AuditEvent stream.
- **Provenance:** Attached to clinical records to indicate which enrollment produced the identifier and which authority vouched for it.
- **Endpoint and Organization:** Used to publish the federation directory — the list of recognized jurisdictional authorities and their FHIR endpoints.

Patient matching across jurisdictions uses IHE PIXm with the C-GPI as the master identifier. Where the C-GPI is unavailable (e.g., a jurisdiction has not yet joined the federation), the system falls back to IHE PDQm with demographic attributes, accepting the reduced match accuracy this entails.

8. Consent and Governance Model

Consent in CIF is purpose-bound, revocable, and machine-readable. The patient grants consent for specific data scopes (e.g., immunization history, chronic disease summaries) to specific classes of requesters (e.g.,

"any licensed clinician in jurisdiction X") for specific purposes (e.g., "emergency treatment"). Consent is captured at enrollment using an assisted consent process recommended by the WHO SMART Guidelines for low-literacy contexts, and is recorded as a FHIR Consent resource. Revocation propagates to all federation members within a defined latency target (we recommend 24 hours).

Table 3. Consent operations.

Operation	Description
Request	Present scope and purpose to patient via accessible medium (visual, audio, or assisted).
Capture	Record the consent as a FHIR Consent resource signed (cryptographically or via witnessed mark) by the patient.
Review	Patient (or authorized proxy) can view active consents and AuditEvent history.
Revoke	Immediate invalidation; propagated to federation members within target latency.
Expire	Consents have default expiration aligned with the purpose (e.g., 1 year for routine care).

On the opt-in versus opt-out question, the framework defaults to opt-in for cross-border data sharing, consistent with GDPR Article 9 (special categories) and the African Union's Malabo Convention. Within-jurisdiction sharing may follow each jurisdiction's prevailing default; the framework does not attempt to harmonize this.

9. Privacy, Security, and Threat Model

9.1 Threat Model

The framework's security claims are made against an adversary with the following capabilities:

- **External attacker:** Can intercept network traffic; defeated by TLS 1.3 with certificate pinning and mTLS between federation members.
- **Compromised gallery:** Adversary obtains the protected template database. Cancelable transforms and fuzzy commitment make raw biometric recovery infeasible. Cross-database linkage is bounded by the unlinkability property of the BTP scheme.
- **Compromised matching service:** Adversary observes the LSH index. Recovers approximate template proximity, which is sensitive but bounded; cannot recover raw biometrics. This is the framework's principal residual risk and is documented honestly.
- **Insider with query access:** Adversary can issue lookups but cannot bulk-export data. AuditEvent emission and rate limiting bound the damage; per-query consent enforcement makes patterns of misuse detectable.

- **Coerced enrollment:** Not solved technically. Mitigated through enrollment witness requirements, anomaly detection on enrollment patterns, and the legal frameworks of participating jurisdictions.

9.2 Privacy-by-Design Alignment

CIF aligns with Cavoukian's (2009) seven Privacy-by-Design principles. The most operationally consequential are end-to-end protection (no raw biometric persisted past the enrollment session), data minimization (only the C-GPI and required FHIR fields flow across jurisdictional boundaries), and full lifecycle protection (revocation and renewal procedures are first-class operations, not afterthoughts).

10. Implementation Roadmap

We propose a four-phase deployment trajectory. Phase boundaries are defined by milestones, not calendar dates.

7. Diagnostic phase: Conduct an ID4D Diagnostic and Identity Enabling Environment Assessment (IDEEA) in the candidate jurisdiction. Establish baseline metrics for existing identification coverage and health system fragmentation.
8. Single-jurisdiction pilot: Deploy CIF within one health authority (e.g., a national immunization registry, or a refugee health service). Target population: 50,000–500,000 enrollees. Validate accuracy targets and operational metrics; iterate on the consent UX with the affected community.
9. Bilateral federation: Extend to a second jurisdiction with a bilateral trust agreement covering data residency, mutual recognition, and audit-sharing. Validate cross-border matching workflows under real conditions.
10. Multilateral federation: Onboard additional jurisdictions under a common trust framework, drawing on the eIDAS pattern. Establish a governing body with representation from participating jurisdictions and civil society.

We do not propose a global rollout. Cross-border health identity at global scale is constrained by political and legal factors that no technical framework can resolve unilaterally. The federation model is designed to grow incrementally as those constraints are resolved bilaterally.

11. Evaluation Metrics and Target Performance

We specify concrete performance targets, informed by NIST FRVT and FpVTE benchmarks for unprotected systems and adjusted for the typical 1–3 percentage-point accuracy degradation introduced by template protection (Patel et al., 2015).

Table 4. Target performance metrics.

Metric	Target	Notes
False Match Rate (FMR)	≤ 10^{-5}	At population scale, this still produces multiple candidate matches; resolved by the two-stage pipeline.
False Non-Match Rate (FNMR)	≤ 2%	Measured per modality; multimodal fusion brings system-level FNMR below 0.5%.
Failure to Enroll (FTE)	≤ 1.5%	Multimodal fallback required to reach this target; single-modality FTE for manual laborers can exceed 5%.
Identification Latency	≤ 3 s p95	End-to-end at the FHIR API for a federated query. Local-only verification: ≤ 500 ms.
Demographic Parity (FMR/FNMR)	Δ ≤ 1.5×	Maximum ratio of FMR or FNMR between any two reported demographic subgroups; publicly reported per release.
Unlinkability (ISO 30136)	D ≤ 0.1	Average linkability metric across protected templates of the same subject in different deployments.

Each pilot deployment must publish these metrics, broken down by demographic subgroup, as a condition of joining the federation. Refusal or inability to report disqualifies a deployment.

12. Use Cases

12.1 Cross-Border Continuity of Care for Refugees

A refugee enrolled at a UNHCR site in one country presents at a health facility in another. The receiving facility queries the federation directory, locates the issuing authority via the C-GPI, retrieves immunization and chronic-care history under an applicable consent, and continues care without restarting baseline workups. The patient's biometrics never leave the issuing jurisdiction's protected gallery; only the C-GPI and consented clinical attributes cross borders.

12.2 Childhood Immunization Schedule Completion

Children enrolled with their guardians at birth registration receive a C-GPI bound at the time to face and (where reliable) iris. Subsequent vaccination encounters re-authenticate against the bound modality, with re-enrollment at defined age milestones (face is unstable in young children; iris matures earlier). Schedule completion can be tracked across providers without exposing children's biometric data to a central registry.

12.3 Fraud Reduction in Health Insurance

Insurance schemes use biometric deduplication to remove ghost beneficiaries from rolls. CIF supports this by exposing a 1:N "unique-or-not" verification API that returns only a boolean and the C-GPI, never the candidate match list, reducing the surface for downstream privacy violations.

13. Risks and Limitations

We enumerate the framework's principal limitations explicitly, separating those that are solved-with-tradeoffs from those that remain open problems.

13.1 Acknowledged Tradeoffs

- **LSH index leaks proximity.** Stage-one indexing reveals approximate template proximity to anyone with access to the index. This is mitigated by the cancelable transform but not eliminated.
- **Demographic accuracy disparities.** Face recognition in particular performs unevenly across skin tone, age, and gender (NIST FRVT Part 3). CIF mitigates by preferring fingerprint and iris where viable and by per-subgroup threshold calibration, but does not eliminate the disparity.
- **Re-enrollment on transform compromise.** Since CIF does not retain raw biometrics, a compromised transform parameter requires re-enrolling the affected population from the original biometric.

13.2 Open Problems

- **Fully secure 1:N search at population scale.** Currently relies on the LSH-based two-stage pipeline; a fully homomorphic alternative exists in the literature but is impractical above ~10^5 enrollees with current hardware.
- **Cross-modality re-enrollment.** Switching sensor families (e.g., capacitive to optical fingerprint) typically requires re-enrollment; cross-modality template translation is not a solved problem.
- **Coerced or fraudulent enrollment.** Technical mitigations (anomaly detection, witness requirements) are partial. The residual risk is governance-bound, not technically-bound.
- **Long-term cryptographic agility.** Post-quantum-safe variants of fuzzy commitment and homomorphic matching are an active research area; deployments should plan for cryptographic migration.

14. Future Research

- **Practical homomorphic 1:N search.** Advances in CKKS and TFHE schemes are pushing performance boundaries; periodic re-evaluation should determine when full homomorphic search displaces the LSH-based pipeline.
- **Decentralized identifier (DID) integration.** W3C DID and verifiable credentials may allow patient-held credentials to substitute for federation lookups in some workflows, reducing reliance on jurisdictional infrastructure.

- **Post-quantum BTP.** Adapting fuzzy commitment and cancelable schemes to lattice-based and code-based primitives compatible with NIST PQC selections.
- **Federated learning for quality assessment.** Improving capture-quality models without centralizing training data.

15. Conclusion

Closing the global identification gap is a precondition for universal health coverage, but it must be done without manufacturing the surveillance infrastructure that erodes the trust on which healthcare depends. The Curely Identity Framework attempts a middle path: a federated, standards-based architecture that produces a stable global patient identifier without persisting recoverable biometric data, and that interoperates with existing health information systems through HL7 FHIR. The framework is explicit about what it solves cryptographically, what it approximates for performance reasons, and what remains open research. None of the underlying primitives are novel; the contribution is integration, honest tradeoff documentation, and a deployment model — federated regional pilots, not a centralized global registry — that is technically and politically tractable.

We invite collaboration on pilot deployments, particularly with health authorities serving displaced and underdocumented populations, where the gap between identification need and identification supply is largest.

References

Boddeti, V. N. (2018). Secure face matching using fully homomorphic encryption. In IEEE 9th International Conference on Biometrics Theory, Applications and Systems (BTAS) (pp. 1–10).

Bringer, J., Chabanne, H., & Patey, A. (2013). Privacy-preserving biometric identification using secure multiparty computation. IEEE Signal Processing Magazine, 30(2), 42–52.

Cavoukian, A. (2009). Privacy by Design: The 7 Foundational Principles. Information and Privacy Commissioner of Ontario.

ECRI Institute. (2017). Patient Identification: Executive Summary. ECRI Institute PSO Deep Dive.

HL7 International. (2024). Identity Matching Implementation Guide. Ann Arbor, MI: HL7 International.

HL7 International. (2024). FHIR Release 5 Specification. Ann Arbor, MI: HL7 International.

ISO/IEC 19794-2:2011. Information technology — Biometric data interchange formats — Part 2: Finger minutiae data.

ISO/IEC 24745:2022. Information security, cybersecurity and privacy protection — Biometric information protection.

ISO/IEC 30136:2018. Information technology — Performance testing of biometric template protection schemes.

Jain, A. K., Nandakumar, K., & Nagar, A. (2008). Biometric template security. EURASIP Journal on Advances in Signal Processing, 2008, Article 579416.

Juels, A., & Sudan, M. (2002). A fuzzy vault scheme. In Proceedings of the IEEE International Symposium on Information Theory (p. 408).

Juels, A., & Wattenberg, M. (1999). A fuzzy commitment scheme. In Proceedings of the 6th ACM Conference on Computer and Communications Security (pp. 28–36).

Nandakumar, K., Jain, A. K., & Pankanti, S. (2007). Fingerprint-based fuzzy vault: Implementation and performance. IEEE Transactions on Information Forensics and Security, 2(4), 744–757.

NIST. (2024). Face Recognition Vendor Test (FRVT) Part 3: Demographic Effects. National Institute of Standards and Technology.

Patel, V. M., Ratha, N. K., & Chellappa, R. (2015). Cancelable biometrics: A review. IEEE Signal Processing Magazine, 32(5), 54–65.

Pew Charitable Trusts. (2018). Enhanced Patient Matching Is Critical to Achieving Full Promise of Digital Health Records.

Ratha, N. K., Connell, J. H., & Bolle, R. M. (2001). Enhancing security and privacy in biometrics-based authentication systems. IBM Systems Journal, 40(3), 614–634.

Teoh, A. B. J., Goh, A., & Ngo, D. C. L. (2006). Random multispace quantization as an analytic mechanism for BioHashing of biometric and random identity inputs. IEEE Transactions on Pattern Analysis and Machine Intelligence, 28(12), 1892–1901.

World Bank. (2019). ID4D Practitioner's Guide: Version 1.0. Washington, DC: World Bank Group.

World Bank. (2024). The Global Identification Challenge: Who Are the 850 Million People Without Proof of Identity? Washington, DC: World Bank ID4D Initiative.

World Health Organization. (2023). SMART Guidelines: Standards-based, Machine-readable, Adaptive, Requirements-based, Testable. Geneva: WHO.