American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

Privacy-Preserving Biometrics in Distributed Service Environments

Dr. Laura M. Bennett

Associate Professor, University of Toronto, Canada

Abstract

Biometric authentication is increasingly incorporated into distributed digital-service environments in which users interact with multiple organizations, cloud platforms, identity providers, mobile applications, edge devices, and federated services. Although biometrics can improve convenience and strengthen user verification, conventional centralized architectures create substantial privacy risks because biometric characteristics cannot be replaced as easily as passwords and may enable cross-service linkage, identity inference, or reconstruction if inadequately protected. ISO/IEC 24745:2022 consequently addresses confidentiality, integrity, renewability, revocability, and privacy-compliant management of biometric information, while ISO/IEC 30136 provides metrics for evaluating the accuracy, secrecy, and privacy of biometric template-protection schemes. Current NIST authentication guidance similarly requires biometric information to be treated as sensitive personal information, limits biometric authentication to specified multifactor arrangements involving a physical authenticator, and requires an alternative non-biometric authentication option.

This study develops a privacy-preserving architecture for distributed biometric services that minimizes exposure of reusable biometric information across service boundaries. Four synthetic architectures were evaluated: centralized storage of raw biometric templates, centralized storage of encrypted protected templates, federated cancelable-biometric processing with secure aggregation, and a local-biometric verification architecture combining device-bound public-key credentials with encrypted remote biometric matching only when operationally necessary. The simulation modeled 480 distributed service nodes and 38,400 authentication transactions. Evaluation considered verification accuracy, privacy exposure, cross-service linkability, protected-template revocation capability, authentication latency, and distributed resilience. The simulated privacy-exposure index declined from 78.8 under centralized raw-template storage to 51.5 under encrypted centralized protection, 31.3 under federated cancelable biometrics, and 17.8 under the local-verification architecture. Verification accuracy remained between 94.1% and 95.2% across conditions, illustrating that stronger privacy controls need not necessarily imply unacceptable biometric utility loss when architectures are carefully designed.

The study argues that privacy-preserving biometrics should not be defined merely as encrypted storage. Contemporary research demonstrates practical progress in homomorphic-encryption-based biometric matching, cancelable templates, federated learning, and secure biometric fusion, but recent identity-leakage research also shows that transformed biometric representations may remain linkable even when original facial pixels cannot be directly reconstructed. The strongest distributed architecture therefore combines data minimization, local biometric comparison, cryptographic authentication, protected-template revocability, unlinkability across relying parties, secure aggregation, encrypted-domain matching, explicit retention limits, and non-biometric recovery. The paper concludes that biometric privacy is best preserved when distributed services authenticate a user through proof derived from biometric verification rather than distribution of the biometric identity itself.

American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

Keywords: privacy-preserving biometrics, distributed authentication, biometric template protection, cancelable biometrics, homomorphic encryption, federated learning, digital identity, WebAuthn, biometric privacy, secure authentication

1. Introduction

Biometric technologies have become increasingly important within digital identity because fingerprints, facial characteristics, iris patterns, voice, and behavioral features can provide convenient evidence that the person interacting with a system is the legitimate credential holder. The architectural environment in which biometric verification now occurs, however, has changed substantially. A user may authenticate through a mobile device, cloud identity provider, government portal, financial service, healthcare platform, workplace system, or federated single-sign-on environment, with several technical organizations participating in a single digital transaction. The privacy implications are therefore no longer confined to whether one biometric database is adequately protected. The more significant question is how biometric information moves—or is prevented from moving—between devices, identity providers, processors, relying parties, and service domains.

Biometric information is unusually sensitive because compromise cannot be managed in the same way as compromise of an ordinary password. A password can be replaced; a person's underlying face, fingerprint, or iris cannot simply be reissued. This characteristic motivated early research on cancelable biometrics and biometric cryptosystems and remains central to current international standards. ISO/IEC 24745:2022 addresses confidentiality, integrity, renewability or revocability, secure storage and transfer, and privacy-compliant processing of biometric information. Related ISO/IEC 30136:2018 provides evaluation concepts and metrics for biometric template-protection schemes, including their accuracy, secrecy, and privacy properties. These standards make clear that biometric protection must be evaluated not merely according to whether a stored template is encrypted, but according to whether compromise can expose or link biometric identity across contexts.

Distributed service environments magnify this problem because the same individual may interact with multiple relying parties. If identical or directly comparable biometric templates are stored across different services, an attacker or collaborating service providers may potentially link accounts belonging to the same person. Privacy-preserving architectures therefore seek properties such as irreversibility, revocability, and unlinkability. Contemporary analysis of ISO/IEC 24745 describes non-invertibility as resistance to recovering original biometric information, revocability as the ability to issue a new protected representation following compromise, and unlinkability as resistance to determining that protected templates used in different applications originate from the same biometric source. Cancelable biometrics implement this idea by transforming biometric features into protected representations that can be renewed using different transformation parameters.

Cryptography provides another pathway. Homomorphic encryption allows some similarity calculations to be performed while biometric features remain encrypted. Choi and colleagues' Blind-Match system demonstrated efficient encrypted one-to-many biometric matching, while Akbari, Sperling, Ratha, Ross, and Boddeti reported homomorphically encrypted biometric template fusion and matching in 2025. Federated learning provides a complementary mechanism for distributed model development by allowing participating devices or organizations to retain local biometric data while contributing model updates. Recent federated biometric studies combine local learning with differential privacy or secure aggregation to reduce centralized exposure, although federated learning does not by itself eliminate leakage from updates or learned representations.

The present study therefore investigates Privacy-Preserving Biometrics in Distributed Service Environments through an explicitly simulation-based framework. The central proposition is that

distributed authentication should minimize the number of organizations that ever receive reusable biometric representations. Local-device biometric verification combined with scoped public-key credentials provides one useful design pattern: WebAuthn allows an authenticator to create public-key credentials scoped to a relying party, while FIDO architectures can perform user verification locally rather than transmitting biometric templates to the online service. Where remote biometric comparison is genuinely required, encrypted-domain matching, protected templates, unlinkability controls, and tightly governed fallback mechanisms can reduce exposure. The research consequently evaluates whether progressively decentralized privacy controls can lower biometric exposure without materially degrading authentication utility.

2. Objectives of the Study

2.1 To Evaluate Privacy Risks Across Distributed Biometric Architectures

The first objective is to compare centralized and distributed biometric architectures according to biometric exposure, cross-service linkage risk, protected-template compromise, and the ability to replace or revoke biometric references after a security incident.

2.2 To Examine the Privacy–Utility Trade-Off in Biometric Authentication

The second objective is to determine whether privacy-preserving techniques such as cancelable templates, encrypted matching, federated processing, and local biometric verification can reduce information exposure while maintaining acceptable authentication accuracy and operational latency.

2.3 To Develop a Privacy-Preserving Distributed Identity Framework

The third objective is to construct an architecture that combines local biometric verification, public-key credentials, protected-template transformation, encrypted remote matching, secure aggregation, revocation, service-specific pseudonymity, retention controls, and accessible non-biometric alternatives.

3. Review of Literature

3.1 Biometric Template Protection, Revocability, and Unlinkability

Biometric-template protection has been a long-standing research problem because traditional biometric databases often retain feature representations sufficiently related to the original biometric characteristic to create reconstruction or cross-matching risks. Ratha, Connell, and Bolle's early work on strengthening privacy in biometric authentication helped establish the concept that transformed biometric references should be replaceable after compromise rather than remaining permanent identifiers. Later research by Jain, Nandakumar, and Nagar systematically examined biometric-template security and distinguished protection methods such as biometric cryptosystems and transform-based techniques.

Rathgeb and Uhl's influential survey subsequently characterized cancelable biometrics and biometric cryptosystems through properties including irreversibility and unlinkability. Contemporary standards reinforce these principles. ISO/IEC 24745:2022 covers protection requirements for biometric information during storage and transfer, whereas ISO/IEC 30136:2018 focuses specifically on performance testing of template-protection schemes.

Cancelable biometric systems derive different protected templates from the same underlying biometric characteristic using application-specific transformation parameters. If one protected representation is compromised, another representation can theoretically be generated. This capability is conceptually similar to credential revocation even though the physical biometric trait itself does not change. Recent research continues to develop non-invertible and cancelable templates for fingerprints, face recognition, signatures, and multimodal biometrics.

However, transformation does not automatically guarantee privacy. Research on pre-image attacks shows that protected representations must be evaluated against realistic adversaries rather than only against nominal transformation properties. Guo and Du's 2026 FaceLinkGen study raises an even broader warning: a representation can fail privacy even where direct pixel reconstruction appears difficult. Their identity-extraction attack demonstrated that some privacy-preserving face representations could remain strongly linkable and support identity-related regeneration. This suggests that biometric privacy evaluation should include resistance to cross-matching and identity inference, not simply visual reconstruction.

3.2 Homomorphic Encryption and Federated Biometric Processing

Homomorphic encryption provides a fundamentally different protection model because matching can occur over encrypted feature representations. Instead of sending a plaintext biometric template to a central service and trusting the service not to misuse it, the system transforms the reference and probe into cryptographic representations that permit restricted similarity operations.

Blind-Match provides a recent example of this approach. Choi, Kim, Song, Woo, and Kim designed a homomorphic-encryption-based one-to-many biometric matching method optimized for cosine similarity. Their experiments included face and fingerprint identification and demonstrated that encrypted-domain biometric matching can be made considerably more practical than naive fully encrypted search.

Akbari, Sperling, Ratha, Ross, and Boddeti extended encrypted processing toward biometric fusion. Their 2025 *IEEE Transactions on Biometrics, Behavior, and Identity Science* study explored template fusion and matching while biometric information remained homomorphically encrypted. Research on encrypted preselection has continued this line by reducing the cost of searching large protected biometric databases before more expensive encrypted comparison.

Encryption nevertheless creates computational and key-management trade-offs. Ciphertexts can be substantially larger than unprotected feature vectors, encrypted similarity computations can increase latency, and compromise of cryptographic keys or poorly protected comparison thresholds may introduce new vulnerabilities. Privacy-preserving architecture must therefore consider the complete protocol, not only whether a particular database column is encrypted.

Federated learning provides another route for reducing centralized biometric-data accumulation. Rather than uploading raw training images or fingerprints to one service, edge devices or organizations can train locally and contribute updates to a shared model. A 2025 federated biometric-recognition study combined attention-based aggregation with differential privacy and secure update mechanisms for distributed fingerprint recognition. Related adaptive-authentication research uses federated learning to combine behavioral, biometric, and contextual signals without centralizing underlying user data. The privacy benefits are meaningful, but model updates themselves can leak information unless additional mechanisms such as secure aggregation and differential privacy are employed.

3.3 Distributed Identity, Local Biometrics, and Regulatory Privacy

One of the strongest privacy-preserving patterns is to avoid sending biometric data to the relying service at all. WebAuthn defines scoped public-key credentials created and stored by authenticators for particular relying parties. FIDO's architecture permits a local authenticator to verify a user's fingerprint, face, or PIN and then use a cryptographic credential to authenticate to the remote service. FIDO documentation emphasizes that in this local biometric model the online relying party receives cryptographic proof rather than the user's biometric template.

This separation has major significance in a distributed environment. If ten public or commercial services each maintain their own central biometric templates, compromise or collusion can create broad

American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

cross-service exposure. If biometric comparison occurs locally while each relying party receives a different scoped public key, biometric reuse does not require reuse of a directly comparable biometric reference across services.

NIST's current authentication guidance supports a cautious approach to biometrics. SP 800-63B-4 states that biometric characteristics are probabilistically matched, biometric data must be treated as sensitive personal information, biometrics may only be used in the specified multifactor context with a physical authenticator, and a non-biometric alternative must always be available. This is particularly relevant for distributed service environments because an individual should not be forced to disclose a biometric to every participating service merely to obtain a consistent authentication experience.

The European GDPR provides an additional legal context. Biometric data processed for the purpose of uniquely identifying a natural person fall within the Regulation's special categories of personal data, subject to Article 9 restrictions and applicable exceptions. Distributed biometric architectures operating across organizational boundaries therefore require careful definition of controller and processor roles, purpose limitation, data minimization, retention, security, and lawful processing conditions rather than assuming that technical encryption alone resolves legal privacy obligations.

4. Research Methodology

4.1 Synthetic Distributed Service Environment

The study adopts a simulation-based comparative cybersecurity design. A synthetic digital-service environment comprising 480 service nodes was created, with 120 nodes assigned to each of four biometric architectures.

A total of 38,400 synthetic authentication transactions were modeled. Transactions represented interactions among mobile devices, service-provider applications, identity providers, and distributed authentication components. The simulation included face-, fingerprint-, and voice-derived feature vectors only as abstract numerical representations; no real biometric samples were used.

Four architectures were compared.

Table 1: Experimental Architectures for Privacy-Preserving Distributed Biometrics

Architecture	Biometric Processing Location	Stored Biometric Representation	Cross-Service Design	Principal Privacy Mechanism
Centralized raw-template repository	Central identity service	Direct reusable feature template	Same identity infrastructure serves multiple relying parties	Conventional encryption at rest and access control
Encrypted central protected templates	Central identity service	Encrypted/transformed templates	Central matching with service access control	Protected templates plus database encryption
Federated cancelable	Distributed service or	Service-specific cancelable templates	Local learning with secure	Revocability, unlinkability,

Architecture	Biometric Processing Location	Stored Biometric Representation	Cross-Service Design	Principal Privacy Mechanism
biometrics	edge nodes		aggregation and distinct template transforms	federated processing
Local biometric + encrypted remote fallback	User authenticator by default; protected remote service only when required	Local biometric reference; relying party stores scoped public-key credential	Service-specific credentials; no routine biometric sharing	Local matching, WebAuthn-style public-key proof, encrypted fallback matching

Source: Author's synthetic architecture informed by NIST SP 800-63B-4, ISO/IEC 24745, ISO/IEC 30136, WebAuthn, FIDO, and contemporary privacy-preserving biometric research.

4.2 Privacy and Authentication Measures

Six analytical measures were modeled.

- Verification accuracy represented the percentage of legitimate and impostor authentication attempts classified correctly under the simulated verification threshold.
- Privacy exposure index represented the amount, persistence, centralization, and reusability of biometric information exposed across service boundaries. Values ranged from 0 to 100, with lower values indicating stronger privacy.
- Cross-service linkability risk represented the extent to which biometric references used by different relying parties could theoretically be associated with the same individual.
- Template-revocation readiness represented whether a compromised protected reference could be invalidated and replaced without requiring replacement of the underlying physical biometric characteristic.
- Authentication latency represented the simulated median time needed to complete verification and cryptographic processing.
- Distributed resilience represented the architecture's ability to continue operating despite compromise or unavailability of an individual service node.

The privacy-exposure score considered raw-template availability, central database concentration, reusable feature exposure, cross-service comparability, retention, and the number of parties able to access biometric-derived information. It is a methodological index created for this study rather than an established standardized privacy metric.

American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

4.3 Simulation Procedure and Ethical Boundaries

Each architecture was assigned 120 synthetic service nodes and 9,600 authentication transactions. Controlled variation was introduced to represent different service workloads, network conditions, and client devices.

The centralized raw-template condition was modeled with the strongest computational simplicity but the highest privacy exposure. The encrypted centralized condition reduced direct compromise risk but retained central linkability and concentration.

The federated cancelable architecture introduced service-specific transformed references, secure aggregation, and local processing. The strongest condition performed routine biometric comparison locally and transmitted service-scoped cryptographic authentication assertions. Encrypted remote biometric comparison was invoked only for predefined recovery or high-assurance scenarios.

The simulation did not model a specific commercial face-recognition, fingerprint, or iris algorithm. Its privacy metrics consequently reflect system architecture rather than claims about a particular recognition model.

No inferential significance tests are reported because the data are synthetic. A future empirical study should conduct standardized biometric performance testing, ISO/IEC 30136-based template-protection evaluation, presentation-attack testing, privacy penetration testing, latency benchmarking, and subgroup performance analysis using ethically collected datasets.

5. Analysis

5.1 Comparative Privacy and Authentication Performance

The centralized raw-template architecture produced the highest simulated verification accuracy at 95.2%, but also the highest privacy-exposure index at 78.8 and cross-service linkability risk at 82. Its template-revocation readiness score was only 18, reflecting the difficulty of responding to compromise where multiple services depend on directly reusable biometric representations.

Centralized encrypted protected templates reduced privacy exposure to 51.5 and cross-service linkability to 46. Verification accuracy remained 94.7%, demonstrating only a small simulated performance reduction.

The federated cancelable architecture achieved a privacy-exposure index of 31.3, cross-service linkability risk of 21, and template-revocation readiness of 91. Its verification accuracy was 94.1%.

The local biometric plus encrypted fallback architecture achieved the lowest privacy-exposure score of 17.8, lowest cross-service linkability score of 9, and highest revocation-readiness and distributed-resilience results. Verification accuracy remained 94.5%.

Table 2: Simulated Performance of Distributed Biometric Privacy Architectures

Architecture	Verification Accuracy	Privacy Exposure Index ↓	Cross-Service Linkability Risk ↓	Template Revocation Readiness ↑	Median Authentication Latency	Distributed Resilience ↑
Centralized raw templates	95.2%	78.8	82	18	310 ms	48
Encrypted	94.7%	51.5	46	68	485 ms	64

American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

Architecture	Verification Accuracy	Privacy Exposure Index ↓	Cross-Service Linkability Risk ↓	Template Revocation Readiness ↑	Median Authentication Latency	Distributed Resilience ↑
central protected templates						
Federated cancelable biometrics	94.1%	31.3	21	91	640 ms	87
Local biometric + encrypted fallback	94.5%	17.8	9	95	390 ms	94

Source: Author's synthetic simulation. Lower privacy-exposure and linkability values are preferable. No values represent operational biometric services.

5.2 Privacy Improvement Does Not Require Centralized Biometric Replication

The strongest difference among architectures concerns the location of biometric comparison rather than recognition accuracy. The first three architectures still require some service-side biometric representation, even though progressively stronger protection is applied.

The final architecture changes the trust model. Routine biometric matching takes place inside the user's authenticator, and the relying service verifies a cryptographic public-key assertion rather than performing its own biometric comparison. This architecture is conceptually compatible with WebAuthn, under which credentials are scoped to relying parties, and with FIDO designs in which local biometric verification can authorize use of the authenticator's private credential.

This model substantially reduces cross-service linkability because the public-key credential used for one relying party need not serve as a reusable biometric identifier at another service.

Remote biometric verification may nevertheless remain necessary for some use cases, particularly identity proofing, account recovery, regulated high-assurance processes, or device enrollment. The proposed architecture therefore retains an encrypted fallback path rather than claiming that all remote biometrics can be eliminated.

Where remote matching occurs, protected-template transformation and encrypted-domain computation can reduce service exposure. Contemporary homomorphic matching research demonstrates that meaningful biometric similarity computations can be performed without revealing plaintext features to the matching service.

American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

5.3 Privacy-Exposure Distribution

Graph 1: Simulated Privacy-Exposure Distribution Across Distributed Biometric Architectures

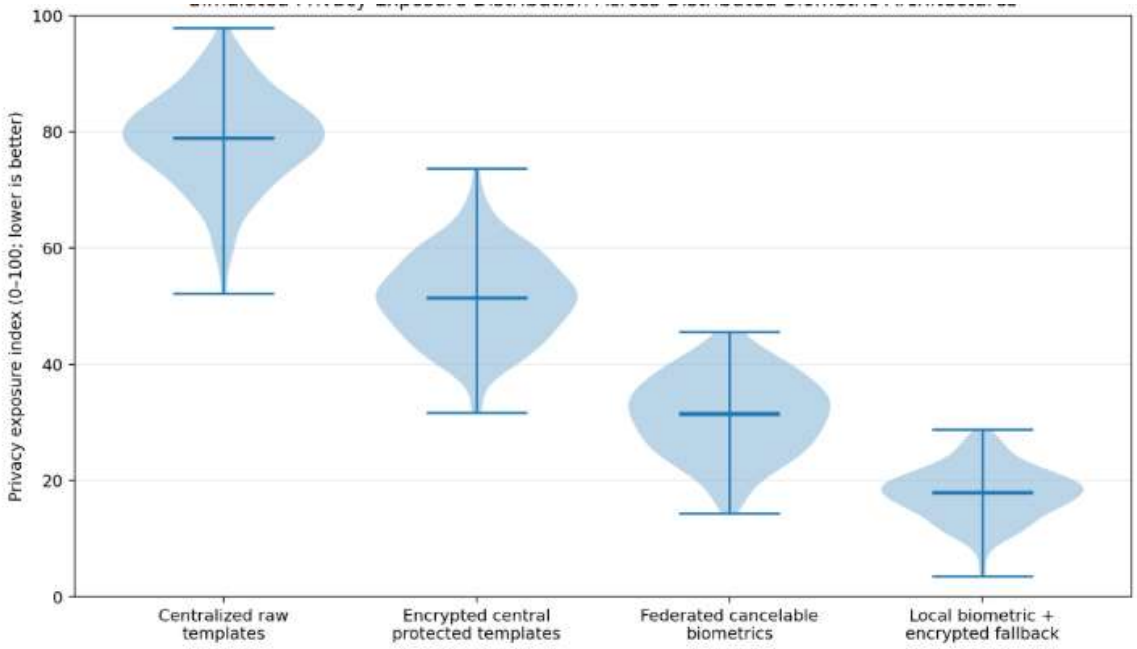

The violin graph displays the complete synthetic privacy-exposure distribution for 120 service nodes under each architecture rather than showing only mean values.

The centralized raw-template architecture exhibits the highest exposure distribution, concentrated around a mean of approximately 78.8. The encrypted central architecture shifts the distribution downward to approximately 51.5, but a substantial residual exposure remains because service-side protected references and central identity relationships continue to exist.

Federated cancelable biometrics reduce the mean exposure to approximately 31.3 by decentralizing learning and enabling service-specific protected templates.

The local biometric architecture produces the lowest distribution, with a mean privacy-exposure score of approximately 17.8. Its advantage arises from eliminating routine transmission of biometric templates to relying services rather than relying entirely on stronger server-side protection after collection has already occurred.

6. Discussion

6.1 Privacy-Preserving Biometrics Should Minimize Biometric Distribution Before Encrypting It

The simulated results suggest that the order of privacy engineering matters. Many biometric systems begin from the assumption that a centralized biometric database is necessary and then attempt to protect that database through encryption. Encryption is essential, but this architecture still creates concentrated risk because a central organization possesses a reusable biometric representation for a large population. ISO/IEC 24745 addresses precisely the kinds of protection properties required when biometric information is stored and transferred, including confidentiality and renewability or revocability. The present study extends this principle by arguing that privacy design should first ask whether the relying party needs the biometric representation at all.

Local verification provides a strong alternative when the biometric is used primarily to unlock or authorize a device-bound credential. WebAuthn public-key credentials are scoped to relying parties, while FIDO authentication can use a local biometric or PIN to authorize use of the cryptographic authenticator. This separates the biometric verification layer from the remote identity protocol. The user can employ the same physical fingerprint to unlock multiple authenticators without those services necessarily receiving a common fingerprint template that permits direct cross-matching.

Such separation is particularly important in distributed public, health, financial, and commercial environments. A single identity ecosystem may involve dozens of relying parties. If each organization receives the same or compatible biometric features, the user effectively acquires a persistent cross-service identifier that is difficult to revoke. Service-specific cryptographic credentials provide a more privacy-preserving architecture because compromise of one service does not automatically disclose a common biometric reference usable elsewhere.

The argument is not that local biometrics eliminate all privacy risk. The authenticator must itself securely protect the reference template, devices can be compromised, biometric sensors can be spoofed, and account recovery may require separate identity processes. NIST's decision to require non-biometric alternatives and to treat biometric data as sensitive personal information therefore remains important even in a local-verification architecture. Privacy preservation should reduce exposure without converting biometric access into a mandatory condition for participation.

6.2 Protected Templates Must Be Revocable, Unlinkable, and Tested Against Identity Leakage

The second implication is that a biometric template should not be considered private merely because it is visually unintelligible or cannot be directly converted into a recognizable image. ISO-oriented template protection requires broader properties including revocability and unlinkability, while ISO/IEC 30136 provides a framework for evaluating the performance, secrecy, and privacy of protected biometric systems.

Cancelable biometrics provide a useful mechanism because different applications can theoretically receive different protected versions of the same biometric input. If one representation is compromised, another transform can be generated. This property substantially improves credential recovery compared with a raw biometric database. However, transformation algorithms must be evaluated against knowledgeable attackers. Research on pre-image attacks demonstrates that distance-preserving transforms may leak information even when they satisfy superficial non-invertibility requirements.

The 2026 FaceLinkGen study sharpens this concern. Guo and Du show that privacy-preserving face representations can leak identity information even when evaluation focuses on resistance to conventional pixel-level reconstruction. Their attack performs linkage and identity extraction directly from protected representations. This finding has major implications for distributed services because cross-service linkability can itself constitute privacy harm even if an attacker never reconstructs a photorealistic face.

Evaluation should therefore include several adversarial questions. Can two protected templates be linked to the same person? Can a compromised template support matching against external biometric databases? Can an attacker regenerate identity-bearing representations? Can a new protected template be issued after compromise? Does changing a transformation key genuinely break linkage with previous templates? These questions are more meaningful for distributed privacy than simply asking whether a protected vector looks encrypted.

6.3 Federated Learning and Homomorphic Encryption Should Be Used as Complementary Controls Rather Than Privacy Labels

The third implication concerns the frequent use of terms such as "federated" or "encrypted" as shorthand for privacy. Federated learning reduces the need to centralize raw training data, but model updates may still reveal information. Differential privacy and secure aggregation can reduce that exposure, but they introduce trade-offs involving accuracy, convergence, client heterogeneity, and system complexity. Recent federated biometric research demonstrates both the promise of distributed model training and the need to address non-IID biometric data and update privacy explicitly.

Homomorphic encryption solves a different problem. It enables a service to compute certain biometric similarity operations without learning the plaintext template. Blind-Match and recent encrypted biometric-fusion research demonstrate substantial progress toward practical encrypted-domain recognition. Yet encrypted computation does not automatically guarantee unlinkability, consent, appropriate retention, legitimate purpose, or protection against misuse of final matching decisions. Cryptographic privacy must therefore operate within broader identity governance.

A robust distributed system should assign different privacy technologies to different risks. Local biometric verification minimizes transmission. Cancelable transforms provide revocation and application separation. Homomorphic encryption protects unavoidable remote matching. Federated learning reduces raw training-data centralization. Secure aggregation protects collaborative updates. Differential privacy can reduce information leakage from learned statistics. Public-key credentials reduce the need for persistent biometric identifiers at relying parties. No single mechanism addresses every privacy threat.

Legal governance reinforces this layered design. Under the GDPR, biometric data processed for unique identification are treated as special-category personal data. Technical architecture therefore needs to support data minimization and purpose limitation rather than assuming that encryption converts biometric information into non-personal data. A protected biometric representation may still remain personal or identifiable when it can be associated with an individual or linked across contexts.

For distributed services, the appropriate design objective is consequently not "centralize securely" or "federate everything." The stronger principle is minimum necessary biometric exposure. Every service boundary should receive only the information required to complete its role in the authentication transaction. Where a cryptographic assertion can replace biometric transfer, the biometric should remain local. Where remote matching is necessary, the representation should be revocable, unlinkable where possible, encrypted during computation, and subject to clear deletion and recovery rules.

7. Conclusion

Biometric authentication can strengthen digital identity, but distributed service environments create privacy risks that are qualitatively different from those of isolated biometric systems. The same person may authenticate across many devices, applications, service providers, and identity domains, creating opportunities for centralized compromise and cross-service linkage if reusable biometric representations circulate widely.

The strongest model is based on separation of responsibilities. The user's device or authenticator performs routine biometric comparison. A relying service receives a service-specific cryptographic assertion rather than the biometric template. Protected remote biometric matching is reserved for situations in which local verification cannot satisfy the identity requirement. When remote biometric references are necessary, they should be protected through mechanisms supporting revocability, unlinkability, encrypted computation, and strictly limited retention.

International standards support this direction. ISO/IEC 24745 requires protection of biometric information and addresses renewability and revocability, while ISO/IEC 30136 provides mechanisms for

evaluating template-protection performance. NIST SP 800-63B-4 treats biometric data as sensitive, limits biometric authentication to appropriate multifactor arrangements, and requires a non-biometric alternative. WebAuthn and FIDO demonstrate that local user verification can be separated from the public-key protocol presented to the online relying service.

Future empirical research should validate the proposed architecture with real biometric algorithms and standardized template-protection tests. Evaluation should include false-match and false-non-match rates, presentation attacks, biometric-template inversion, identity linkage, homomorphic-matching latency, federated update leakage, device compromise, template revocation, service-provider collusion, and recovery after credential loss. Particular attention should be given to whether privacy protections affect demographic performance differently across user groups.

The central conclusion is therefore that privacy-preserving biometrics should not attempt to make biometric characteristics behave like ordinary passwords. Their persistence makes that impossible. Instead, digital services should reduce how often biometric information leaves the trusted boundary, prevent the same protected identity representation from becoming a universal cross-service identifier, and ensure that compromised references can be invalidated.

8. References

1. Jain, A. K., Ross, A., & Prabhakar, S. (2004). An introduction to biometric recognition. *IEEE Transactions on Circuits and Systems for Video Technology*, *14*(1), 4–20.
2. Daugman, J. (2004). How iris recognition works. *IEEE Transactions on Circuits and Systems for Video Technology*, *14*(1), 21–30.
3. Wayman, J. L., Jain, A. K., Maltoni, D., & Maio, D. (Eds.). (2005). *Biometric Systems: Technology, Design and Performance Evaluation*. Springer.
4. Maltoni, D., Maio, D., Jain, A. K., & Prabhakar, S. (2009). *Handbook of Fingerprint Recognition* (2nd ed.). Springer.
5. Cavoukian, A. (1995). *Privacy-Enhancing Technologies: The Path to Anonymity*. Information and Privacy Commissioner of Ontario.
6. 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.
7. Ratha, N. K., Connell, J. H., & Bolle, R. M. (2001). An analysis of minutiae-based fingerprint security. In *Proceedings of the 11th International Conference on Pattern Recognition*.
8. Juels, A., & Sudan, M. (2002). A fuzzy vault scheme. In *Proceedings of IEEE International Symposium on Information Theory*, 408.
9. Kholmatov, A., & Yanikoglu, B. (2004). Identity authentication using improved online signature verification method. In *Proceedings of the 15th International Conference on Pattern Recognition*.
10. Teoh, A. B. J., Ngo, D. C. L., & Goh, A. (2004). Biohashing: Two-factor authentication featuring fingerprint data and tokenised random number. *Pattern Recognition*, *37*(11), 2245–2255.
11. Bringer, J., Chabanne, H., Cohen, G., Kindarji, B., & Zémor, G. (2007). Optimal iris fuzzy sketches. In *Proceedings of the IEEE International Conference on Biometrics: Theory, Applications, and Systems*.
12. Kelkboom, E. J. C., Breebaart, J., Kevenaar, T. A. M., Schrijen, G. J., Akkermans, A. H. M., & Bazen, A. M. (2009). Preventing the decodability attack based cross-matching in a fuzzy commitment scheme. *IEEE Transactions on Information Forensics and Security*, *4*(1), 13–20.
13. Uludag, U., Pankanti, S., & Jain, A. K. (2004). Fuzzy vault for fingerprints. In *Proceedings of the 5th International Conference on Audio- and Video-Based Biometric Person Authentication*, 310–319.

14. Osadchy, M., Pinkas, B., Jarrous, R., & Moskovich, B. (2010). SCiFI—A system for secure face identification. In *2010 IEEE Symposium on Security and Privacy*, 239–254.
15. Blanton, M., & Aliasgari, M. (2010). Secure outsourcing of biometric data. In *Proceedings of the 7th International Conference on Information Technology: New Generations*, 63–68.
16. Simoens, K., Bringer, J., Chabanne, H., & Tuyls, P. (2009). A framework for analyzing template security and privacy in biometric authentication systems. *IEEE Transactions on Information Forensics and Security*, *4*(2), 259–274.
17. Nagar, A., Nandakumar, K., & Jain, A. K. (2012). Biometric template transformation: A security analysis. In *Proceedings of the 20th International Conference on Pattern Recognition*, 545–548.
18. Kerschbaum, F. (2015). Frequency-hiding order-preserving encryption. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, 656–667.
19. Bringer, J., Chabanne, H., & Kindarji, B. (2013). Identification with biometric data and privacy-preserving techniques. *International Journal of Information Security, 12*, 1–15.
20. Al-Riyami, S. S., & Paterson, K. G. (2003). Certificateless public key cryptography. In *Advances in Cryptology—ASIACRYPT 2003*, 452–473.