authID

QUANTUM RESISTANT BIOMETRIC AUTHENTICATION

MPC and Post-Quantum Biometric Digital Signatures

www.authid.ai | (516) 778-5639

CONTENTS

1. WHY QUANTUM RESISTANCE IN BIOMETRICS IS IMPORTANT

Unlike a password, you can't change your face. If an attacker compromises an encrypted biometric database today and decrypts it later using a quantum computer, the damage is irreversible. This is the "harvest now, decrypt later" threat; adversaries are already collecting encrypted biometric data with the intention of breaking it once quantum hardware matures. The irreversibility of biometric data means the transition to quantum-resistant biometrics needs to happen before quantum hardware matures, not after.

1.1 How current systems are vulnerable

Most biometric systems protect stored facial templates and data in transit using RSA or elliptic curve cryptography (ECC). Shor's algorithm, running on a sufficiently powerful quantum computer, can break both of these in polynomial time, making them effectively useless against a quantum-capable attacker. The facial feature vectors (numerical representations of your face) stored in databases would be exposed.

More advanced solutions use Sharding + MPC, a genuinely strong architecture for protecting these biometric templates against classical adversaries. But it creates a false ceiling of security against quantum adversaries unless every cryptographic primitive surrounding the scheme, encryption, key exchange, authentication, and commitment, is replaced with post-quantum alternatives.

2. QUANTUM-RESISTANT BIOMETRICS IS HERE

authID's PrivacyKey stores no biometric information, and it uses cryptographic key-regeneration technology underpinning the processing of authentication and identification transactions. It utilizes a PrivacyKeyMap, which is an encrypted file that contains zero-knowledge guidance data, a description of *where* to sample biometric input at subsequent authentication events. PrivacyKey was built around cryptographic primitives that are secure against a cryptographically relevant quantum computer. NIST recognizes two primitives central to the authID PrivacyKey architecture as quantum-resistant:

- **SHA-256**, which produces the 32-byte seed that drives PrivacyKey's ephemeral keypair generation at every authentication event, retaining approximately 128 bits of post-quantum preimage security under Grover's algorithm (the tightest known quantum speedup for hash-function preimage search). This meets NIST's Category 1 bar for post-quantum symmetric strength.
- **AES-256-GCM**, which protects every PrivacyKeyMap at rest, retaining 128 bits of post-quantum security under Grover. This meets NIST's Category 1 bar for post-quantum symmetric encryption. No published cryptanalytic result as of 2026 threatens this posture.

Building on this post-quantum-safe symmetric foundation, authID has added two major advances in quantum security:

Sharding/MPC regeneration of per-PrivacyKeyMap AES-256 keys: A t-of-n threshold Oblivious Pseudorandom Function (OPRF), distributed across authID-operated nodes in distinct trust and failure domains, actively gates the regeneration of the unique AES-256 key that protects each PrivacyKeyMap. Every encrypt and decrypt operation on a PrivacyKeyMap traverses a distributed, policy-enforced regeneration boundary enabling per-transaction authorization, immediate revocation on the next operation, rate limiting, and node-tier operational telemetry over biometric-gated key access.

Support for post-quantum cryptography (PQC) biometric digital signatures: PrivacyKey's ephemeral-key architecture produces an algorithm-agnostic SHA-256 seed that drives keypair generation at every authentication event. PrivacyKey's existing per-operation algorithm-selection API is extended with three NIST-standardized post-quantum signature algorithms ML-DSA-65 (FIPS 204), SLH-DSA-128s (FIPS 205), and SLH-DSA-256s (FIPS 205) available as primitives alongside the existing classical elliptic-curve options. Customers select the algorithm per operation through the application layer.

Both advances result in:

- **No change to the 1:1 Bn match accuracy.** Neither advance modifies PrivacyKey's biometric matching pipeline or the PrivacyKeyMap construction itself. The threshold MPC layer operates on the per-PrivacyKeyMap AES-256 key upstream of any biometric matching operation, and transparent to it once the PrivacyKeyMap is decrypted. The post-quantum signature primitives replace only the final keypair-generation step downstream of all bitwise conversions of biometric data. Both insertion points are architecturally distinct from the matching pipeline.
- **No change to ISO 30136 certified core security claims.** The structural properties of PrivacyKey — no stored biometric data, ephemeral private keys, PrivacyKeyMap non-invertibility, cross-relying-party key unlinkability are preserved exactly. Both advances are defense-in-depth additions to a validated architecture.
- **IMPERCEPTIBLE change to processing performance.** Less than 50ms impact on operational latency.

PrivacyKey-derived digital signatures remain what they have always been: deterministic cryptographic proof that a specific individual, not merely a registered device or authenticator, was present during the signing event.

3. PRIVACYKEY ARCHITECTURE

The following architectural properties of PrivacyKey are load-bearing for the discussion that follows. Together, they constitute the baseline to which the quantum security advances described are added.

3.1 Every authentication event is a digital signature

PrivacyKey is not a probabilistic pass/fail matcher and is not a binary pass/fail verifier. Every successful authentication event produces a **digital signature** over a relying-party-provided nonce, computed by a keypair that exists only in memory during that event.

The event unfolds as a challenge-response ceremony: the relying party supplies a nonce; the individual presents a live, liveness-verified biometric capture; PrivacyKey combines the capture with the enrolled PrivacyKeyMap to regenerate the ephemeral keypair, signs the nonce with the private key, returns the signature and public key to the relying party, and zeroizes the private key. A successful authentication is therefore not "the biometric matched with sufficient confidence" and not a Boolean "accept" token it is a cryptographic signature that proves a specific face was present at the moment the relying party's nonce was signed. Each event produces a unique signature that cannot be replayed across events, across relying parties, or across nonces, provided the relying party enforces nonce uniqueness as required in any challenge-response protocol. This pattern enables PrivacyKey's integration with any challenge-response authentication protocol.

3.2 Digital signing keys are ephemeral and robustly regenerated

Private keys exist only in memory during an authentication transaction. They are regenerated deterministically at each event from a live, liveness-verified biometric presentation and the corresponding PrivacyKeyMap, and are zeroized at transaction completion. No private key material is persisted, transmitted, or accessible to any external process. The public key and a signature over caller-provided data are the only cryptographic outputs observable outside PrivacyKey.

The cryptographic robustness of PrivacyKey's PrivacyKeyMap construction and its key regeneration is supported by BixeLab's 2024 findings against ISO/IEC 30136:2018 [8], which include system characterizations important from both privacy and post-quantum-cryptographic perspectives:

Property	Measurement	Threshold	Result
Shannon entropy of template bits	7.89 bits/byte	>7.5	Pass
Mutual information I(B;T), biometric to template	0.000	<0.01	Pass
Reconstruction attack, 10,000-template corpus	No successful reconstruction	—	Pass
Cross-subject template distribution overlap	Complete overlap	—	Pass
Template decidability index (d-prime)	0.008	<0.05	Pass
Cross-relying-party key correlation	None detectable	—	Pass
Key decidability index (d-prime, cross-RP)	0.003	<0.05	Pass

3.3 The PrivacyKeyMap

At each biometric enrollment event, PrivacyKey produces a persisted artifact called a **PrivacyKeyMap**. The PrivacyKeyMap is zero-knowledge guidance data, a description of *where* to sample biometric input at subsequent authentication events. It does not contain biometric templates, embeddings, feature vectors, compressed facial images, error-correction helper data, or any biometrically reconstructible representation of the enrolled biometric. This absence of biometric content is a structural property of the architecture, protected by issued patents, which require that the persisted artifact "does not include any of the biometric data and wherein the user cannot be identified."

At each subsequent authentication event, the PrivacyKeyMap is combined with a live, liveness-verified biometric presentation from the enrolled individual to regenerate the ephemeral cryptographic keypair bound to that enrollment. The PrivacyKeyMap alone, without a live biometric presentation, cannot produce key material. A stolen or exfiltrated PrivacyKeyMap yields sampling geometry with no biometric content; it cannot be used to reconstruct the enrolled biometric, impersonate the enrolled individual, or regenerate the associated key.

At rest, each PrivacyKeyMap is encrypted with AES-256-GCM using a **unique per-PrivacyKeyMap AES-256 key**. This uniqueness is structural: the AES-256 key is derived from a combination of

- a customer-scoped secret component (`customerKey`),
- a per-user identifier (`customerUuid`), and
- the per-PrivacyKeyMap cleartext identifier carried in the PrivacyKeyMap header (`privacykeyMapUuid`, a UUIDv4 generated at enrollment from a cryptographically secure pseudorandom number generator).

4. WHERE SHARDING/MPC IS APPLIED IN THE PRIVACYKEY ARCHITECTURE

Multi-party computation is a family of techniques that can be integrated at multiple points in a biometric authentication pipeline. Some architectures distribute biometrically sensitive data across multiple parties and compute matching operations across the shares. This creates unnecessary latency in the authentication processing, resulting in suboptimal user experiences.

PrivacyKey does not retain biometric templates, embeddings, or any biometrically reconstructible data. The protection that distributed matching constructions provide to stored biometric material is, in PrivacyKey's architecture, achieved at a more fundamental layer by not storing that material at all. The MPC construction offered is applied where it yields the greatest incremental benefit: gating per-PrivacyKeyMap AES-256 key regeneration, hence gating the authentication process.

The placement of Sharding/MPC in this architecture, therefore, optimizes for three properties simultaneously: the greatest incremental security benefit, maximum operational control flow over cryptographic key access, and the lowest network and compute overhead. The selected insertion point is the regeneration of the unique AES-256 key that protects each PrivacyKeyMap. At this point in the flow, a threshold MPC construction actively gates every PrivacyKeyMap encrypt and decrypt operation behind a distributed, policy-enforced regeneration event providing per-transaction authorization, immediate revocation, rate limiting, and comprehensive operational telemetry.

5. MPC REGENERATION OF PER-PRIVACYKEYMAP AES-256 KEYS

5.1 Summary of the change

Today, the per-PrivacyKeyMap AES-256 key is derived locally from `customerKey`, `customerUuid`, and `privacykeyMapUuid`. In the proposed construction, `customerKey` is no longer used directly as a key-derivation input. Instead, it becomes the input to a threshold OPRF evaluated across n authID-operated nodes, and the OPRF's output replaces `customerKey` in the regeneration of each PrivacyKeyMap's unique AES-256 key. The per-PrivacyKeyMap uniqueness of the AES-256 key is preserved exactly `privacykeyMapUuid` continues to provide unique domain separation per PrivacyKeyMap.

Because `privacykeyMapUuid` is unique per PrivacyKeyMap and is used as the HKDF salt, every PrivacyKeyMap's AES-256 key is cryptographically distinct from every other PrivacyKeyMap's AES-256 key, including PrivacyKeyMaps belonging to the same user or resulting from successive rotations of the same enrollment.

5.2 MPC participation gating

Participation in the threshold MPC ceremony is itself subject to integrity enforcement at the node tier. Before any scalar multiplication is performed against a Shamir share, each incoming request is evaluated against a configurable set of integrity conditions. authID exposes the following classes of signal for customers to compose under their own deployment policy:

- **Request binding.** A cryptographically-bound per-request value, for example, the authentication nonce that will subsequently be signed, tying the OPRF invocation to a specific authentication event. This prevents speculative invocation, replay, and relay of valid requests against unrelated sessions.
- **Transport-layer client authentication.** Mutual TLS with a client certificate issued by the customer's or authID's certificate authority and validated at the node against the current revocation state.
- **Customer risk signals.** Application-layer context geolocation, device reputation, session risk indicators, tenant-specific policy assertions carried in the request, and evaluated at the node under rules the customer has defined.

- **Platform integrity attestations.** Native platform signals of the kind routinely employed in FIDO2 ceremonies, iOS App Attest, Android Key Attestation, and Play Integrity, and equivalent mechanisms on other platforms.
- **Hardware remote attestation.** Per-request attestations originating from TPM, Secure Enclave, or comparable hardware roots of trust, for deployments that require the OPRF invocation to be bound to verified platform state.

Which combination of signals is required, and at what strength, is a customer policy decision; PrivacyKey does not impose defaults. A request that fails any required condition is rejected before it reaches the Shamir-share evaluation stage. A rejected request consumes no cryptographic work on the node, contributes no pressure to legitimate tenants' rate-limit budgets, and produces no telemetry correlatable to valid OPRF outputs; it simply never enters the ceremony.

5.3 Choice of OPRF construction

The proposed construction is the **2HashDH OPRF**, standardized in IETF RFC 9497 (*Oblivious Pseudorandom Functions Using Prime-Order Groups*). The threshold variant is a straightforward extension using Shamir secret-sharing of the OPRF key over the scalar field of the chosen group.

The recommended group is **ristretto255** (decaf-encoded Curve25519). Ristretto provides a prime-order group with high-performance scalar arithmetic, well-vetted implementations across multiple languages, and resistance to common implementation pitfalls of raw Edwards curves.

5.4 Per-operation latency

Component	Typical latency
Client-side blinding (1 scalar multiplication)	~50 µs
Network RTT to t = 2 nodes in parallel, same region	~2–10 ms (same DC) to ~20–50 ms (cross-AZ)
Per-node scalar multiplication	~50 µs
Client-side unblind + Lagrange combine	~100 µs
HKDF + AES-GCM (per-PrivacyKeyMap key derivation)	< 1 ms
Total added latency per operation	**~3–50 ms** (modeled; actual latency varies with deployment topology, client-to-node RTT, and concurrent load)

For human-driven endpoint operations, **this latency is imperceptible**. For high-throughput server-side workloads, co-locating OPRF nodes in the same region as the high-volume clients keeps added latency in the low-single-digit-millisecond range.

6. POST-QUANTUM SIGNATURE PRIMITIVES

6.1 Seed-based keypair generation in PrivacyKey

PrivacyKey regenerates the signing keypair at every authentication event from a combination of biometric input, PrivacyKeyMap data, and an application-supplied `key_scope` a relying-party-assigned parameter that produces statistically independent derivative keypairs from the same biometric root, ensuring key unlinkability across relying parties. The final stage of the pipeline produces a 32-byte SHA-256 seed that drives keypair generation. This seed is algorithm-agnostic: it is not specific to elliptic curve cryptography.

6.2 Available signature primitives

PrivacyKey's existing `AttestationData.cose_alg_id` API parameter selects the signature algorithm on a per-operation basis. The enumeration is extended with three post-quantum primitives:

Algorithm	Standard	NIST Category	Basis	Public key	Signature
ES256 (ECDSA P-256)	SEC1 / FIPS 186-5	—	ECC	65 bytes (uncompressed)	~64 bytes
ES384 (ECDSA P-384)	SEC1 / FIPS 186-5	—	ECC	97 bytes (uncompressed)	~96 bytes
ES512 (ECDSA P-521)	SEC1 / FIPS 186-5	—	ECC	133 bytes (uncompressed)	~132 bytes
EdDSA (Ed25519)	RFC 8032	—	EdDSA	32 bytes	64 bytes
EdDSA (Ed448)	RFC 8032	—	EdDSA	57 bytes	114 bytes
ML-DSA-65	**FIPS 204 [7]**	**3**	**Lattice (module-LWE)**	**1,952 bytes**	**3,309 bytes**
SLH-DSA-128s	**FIPS 205 [9]**	**1**	**Hash-only**	**32 bytes**	**7,856 bytes**
SLH-DSA-256s	**FIPS 205 [9]**	**5**	**Hash-only**	**64 bytes**	**29,792 bytes**

Classical ECC and EdDSA algorithms remain supported without change. The three post-quantum primitives are additions to the menu, not replacements.

6.3 Algorithm family diversity

The three post-quantum primitives span two distinct mathematical foundations:

- **ML-DSA-65** (FIPS 204) is a lattice-based construction. Its security rests on the hardness of module learning-with-errors.
- **SLH-DSA-128s and SLH-DSA-256s** (FIPS 205, different NIST security categories) are stateless hash-based constructions. Their security rests only on the properties of SHA-256 / SHAKE.

These families share no hard mathematical assumption. Note that SLH-DSA-128s and SLH-DSA-256s are both members of the hash-only family at different NIST security levels, not separate families; genuine cross-family diversity therefore requires pairing ML-DSA-65 with one of the SLH-DSA parameter sets. Customers for whom structural diversity of cryptographic assumptions is a requirement may select algorithms from both families for different operations according to their own policy, while customers who require only a single post-quantum assurance posture may select one.

6.4 What does not change

PrivacyKey's biometric pipeline, PrivacyKeyMap structure, and `key_scope` mechanism are unchanged by the introduction of post-quantum primitives. The validated properties of the SHA-256 seed per-bit independence, cross-identity independence, cross-enrollment independence, identical key reconstruction across captures, rotation-chain independence, are empirical properties of the seed itself; they carry forward through HKDF-Expand (a pseudorandom function) to the expanded seeds used for SLH-DSA keygen, and they carry forward directly to the 32-byte consumed by ML-DSA-65 keygen. The post-quantum algorithms named in this paper are compatible with the PrivacyKey pipeline at the keygen stage: ML-DSA-65 directly, and SLH-DSA-128s and SLH-DSA-256s via seed-expansion construction.

6.5 Customer options

Algorithm selection is performed by the customer's application layer on a per-operation basis. PrivacyKey does not impose defaults, heuristics, or scope-based selection policy. The customer decides which algorithm is appropriate for each class of attestation based on their own risk model, compliance requirements, and operational constraints.

7. THREAT MODEL AND PROPERTIES

7.1 Threat actors considered

Actor	Capability	Outcome
Network attacker, client-to-node link	Passive observation, active MitM up to TLS limits	No leakage. Blinded values are information-theoretically unlinkable to `customerKey`.
Attacker compromising a single OPRF node	Full read/write of node, including its Shamir share	No leakage of `k`, no regeneration capability. Cannot produce OPRF outputs without t−1 additional shares.
Attacker compromising t−1 OPRF nodes	All shares from compromised set	Still no regeneration capability. Threshold not met.
Attacker compromising t OPRF nodes	Threshold met	Can produce OPRF outputs for inputs they choose to query. Still requires `customerKey` to regenerate any specific PrivacyKeyMap's AES-256 key, and the decrypted PrivacyKeyMap is not usable as key material key regeneration requires a live, liveness-verified biometric presentation.
Attacker stealing client-side material (`customerKey`, `customerUuid`, `privacykeyMapUuid`)	Complete client-side secret set	Cannot regenerate any per-PrivacyKeyMap AES-256 key without successfully invoking the OPRF threshold. Subject to node-tier rate limiting, authentication, anomaly detection, and revocation.
Compromised endpoint device (rooted / jailbroken device, tampered application binary, or untrustworthy platform state)	Full endpoint control; presents valid client-side material and attempts to invoke the MPC ceremony	Where platform attestation is configured as a required gating signal iOS App Attest, Android Key Attestation / Play Integrity, or equivalent the attestation check fails and MPC ceremony participation is denied before any Shamir-share processing. No OPRF work is performed and no output is produced. Defeating this layer requires compromising the platform attestation chain itself.

8. PRIVACYKEY SECURITY PROPERTIES

8.1 Security properties of the core product

- No biometric templates, embeddings, or biometrically reconstructible data are stored. BixeLab's 2024 ISO/IEC 30136:2018 compliance findings apply to the underlying PrivacyKey pipeline, which is unchanged by the advances described in this paper.
- Private keys remain ephemeral, existing only in memory during the transaction in which they are regenerated.
- Signing keys remain bound to the combination of the enrolled PrivacyKeyMap and the live biometric presentation.
- Keys remain unlinkable across relying parties.
- Every PrivacyKeyMap continues to be protected with its own unique AES-256 key.

8.2 Security properties provided by the threshold MPC layer

The threshold MPC layer actively gates per-PrivacyKeyMap AES-256 key regeneration and, through that gate, provides:

- **Per-transaction authorization.** Every PrivacyKeyMap encrypt and decrypt operation requires a successful threshold OPRF invocation, conditioned on the integrity checks. There is no cached regeneration path to wait out.
- **Unconditional revocation on the next operation.** A policy change at the OPRF node tier takes effect on the next operation. Revocation does not depend on client cooperation, device state, or eventual-consistency propagation.
- **Per-tenant rate limiting.** Node-tier enforcement caps operation rates per tenant, per client, or per arbitrary policy dimension.
- **Comprehensive operation telemetry.** Every encrypt and decrypt event produces a record at the node tier, enabling anomaly detection, geographic and temporal access policies, and compliance audit trails.
- **Single-host-compromise resistance within authID's infrastructure.** No single host, single administrator, or single secret-exfiltration event is sufficient to regenerate any PrivacyKeyMap's AES-256 key.
- **Global suspension capability.** In an incident response scenario, a single policy change at the node tier can suspend MPC service to a tenant globally, immediately preventing new PrivacyKeyMap operations for that tenant.

8.3 Security properties provided by post-quantum primitives

- **Per-operation algorithm agility.** The customer's application layer selects the signature algorithm at sign time, for each operation, from a menu that includes both classical and post-quantum primitives.
- **Standards-aligned post-quantum signing.** Three NIST-standardized post-quantum signature algorithms are available as first-class primitives within PrivacyKey's existing API.
- **Mathematical-foundation diversity.** The three post-quantum primitives span two non-overlapping assumption classes (lattice and hash-only), enabling customer policy that combines algorithms from distinct foundations where structural diversity is required.

9. SUMMARY

authID's PrivacyKey privacy-preserving, quantum-resistant architecture provides the data privacy, biometric accuracy, scalable performance, regulatory compliance, and quantum-resistant security needed in facial biometric solutions.

Customer policy-driven algorithm, key scope string, threshold parameter, and access policy selection empower customers to choose the right level of security for each of their use cases. authID PrivacyKey provides cryptographic primitives; customers compose a policy around them.

The quantum era isn't theoretical. Legacy biometric systems are vulnerable. NIST has finalized its post-quantum standards. The question isn't whether you need quantum-resistant biometric security, it's whether you'll have it in time.

APPENDIX A: GLOSSARY

- **AES-256-GCM:** Advanced Encryption Standard with 256-bit keys in Galois/Counter Mode. Authenticated encryption providing confidentiality and integrity.
- **COSE:** CBOR Object Signing and Encryption (RFC 9052).
- **CSPRNG:** Cryptographically Secure Pseudorandom Number Generator.
- **DKG (Distributed Key Generation):** Protocol allowing n parties to jointly generate a secret such that no single party ever knows it.
- **HKDF:** HMAC-based Key Derivation Function (RFC 5869).
- **ML-DSA:** Module-Lattice-Based Digital Signature Algorithm (FIPS 204).
- **ML-KEM:** Module-Lattice-Based Key-Encapsulation Mechanism (FIPS 203).
- **MPC (Multi-Party Computation):** Cryptographic protocols allowing multiple parties to jointly compute a function over their inputs without revealing those inputs.
- **OPRF (Oblivious Pseudorandom Function):** A protocol in which a client obtains the value of a pseudorandom function on an input while the server learns nothing about the input.
- **Proactive Secret Sharing:** A protocol allowing Shamir shareholders to refresh share values without changing the underlying secret.
- **ristretto255:** Prime-order group constructed from Curve25519, providing safe arithmetic for cryptographic protocols requiring a prime-order group.
- **Shamir Secret Sharing:** A scheme for splitting a secret into n shares such that any t reconstructs it, but t−1 reveals nothing.
- **SLH-DSA:** Stateless Hash-Based Digital Signature Algorithm (FIPS 205).
- **PrivacyKeyMap:** PrivacyKey's zero-knowledge guidance artifact: encrypted calibration geometry and validity mask describing where biometric samples are taken. Contains no biometric entropy.
- **`customerKey`:** Organization-scoped secret component held client-side. Under the MPC regeneration construction, `customerKey` is no longer used directly as key material: it becomes the input to the threshold OPRF, and the OPRF output replaces it in the HKDF step. Its confidentiality remains required an attacker who obtains `customerKey` and successfully invokes the OPRF threshold can regenerate PrivacyKeyMap AES-256 keys.
- **`customerUuid`:** Per-user identifier used in PrivacyKeyMap AES-256 key regeneration.
- **`privacykeyMapUuid`:** Per-PrivacyKeyMap cleartext identifier carried in the PrivacyKeyMap header; used as HKDF salt for per-PrivacyKeyMap AES-256 key uniqueness.
- **`key_scope`:** Application-supplied parameter producing derivative key pairs from the biometric root, ensuring key unlinkability across relying parties.

APPENDIX B: REFERENCES

1. RFC 5869: *HMAC-based Extract-and-Expand Key Derivation Function (HKDF)*.
2. RFC 9497: *Oblivious Pseudorandom Functions (OPRFs) Using Prime-Order Groups*.
3. NIST FIPS 197: *Advanced Encryption Standard (AES)*.
4. NIST SP 800-38D: *Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM)*.
5. NIST FIPS 203: *Module-Lattice-Based Key-Encapsulation Mechanism Standard (ML-KEM)*.
6. NIST FIPS 204: *Module-Lattice-Based Digital Signature Standard (ML-DSA)*.
7. ISO/IEC 30136:2018. *Information technology: Performance testing of biometric template protection schemes*. International Organization for Standardization, 2018.
8. NIST FIPS 205: *Stateless Hash-Based Digital Signature Standard (SLH-DSA)*.
9. NIST IR 8214C: *NIST First Call for Multi-Party Threshold Schemes*.
10. BixeLab. *Biometric Performance Validation* (Letter of Confirmation, test report 25_BXL039_TR_01). BixeLab Pty Ltd, August 2025.
11. ISO/IEC 19795-1 and ISO/IEC 19795-2. *Information technology: Biometric performance testing and reporting*. International Organization for Standardization.