REACHTECH IDENTITY ARCHITECTURE

Technical Specification — Document B

Human-Centric Authentication (HCA) Framework

ReachTech Inc. — Wichita, Kansas
Justin Wieland, Founder & CEO

April 2026 — Version 0.2

CONFIDENTIAL — For Internal Review and Peer Evaluation

Revision History

v0.1 (April 3, 2026): Initial draft. Name-first resolution flow with five-step sequence, confidence scoring, threat model, account lifecycle, accessibility, FIDO2 compatibility, and implementation roadmap.

v0.2 (April 3, 2026): Major architectural revision based on peer review from six independent AI reviewers. Key changes: (1) Inverted the authentication flow from name-first to face-first, eliminating the template distribution problem identified in v0.1. (2) Replaced the name-entry step with a spoken name step, adding voice biometric as a second independent signal. (3) Integrated duress detection into the normal authentication flow rather than as a separate Panic Mode activation. (4) Introduced three-tier security model (Community, Standard, Protected) implementing proportional security. (5) Replaced static confidence weights with context-aware dynamic scoring. (6) Added session integrity section for continuous validation beyond login. (7) Added proportional security framework establishing that authentication rigor should scale with value at risk.

1. Purpose and Scope

This document specifies the technical architecture for ReachTech's Human-Centric Authentication (HCA) framework. It serves as the engineering companion to "Your Name Is Your Name" (Document A), which presents the philosophical and economic case for contextual identity.

HCA is designed for deployment across three initial ReachTech products: EnRoute (corridor-based delivery and rideshare), Jackson AI (personal AI companion), and Ghost Mobile (privacy-first mobile phone service). The architecture is product-agnostic and extensible to any service requiring user authentication.

This specification acknowledges that risk-based authentication (RBA), adaptive MFA, behavioral biometrics, and on-device biometric processing are established fields with significant prior art. ReachTech's contribution is not the invention of these components but their integration into a cohesive, human-centric system that prioritizes user dignity, transparency, and zero data extraction.

1.1 Design Principles

Principle 1: The system recognizes the human. The human does not authenticate to the system.

Principle 2: Your face opens the door. Your voice confirms you walked through it willingly.

Principle 3: Biometric data never leaves the device. This is an architectural constraint, not a policy choice.

Principle 4: Every signal collected is disclosed to the user. There is no covert data collection.

Principle 5: Security rigor scales with the value at risk. A delivery ride does not require the same posture as a wire transfer.

Principle 6: The system must work for a 72-year-old who has never heard the word "keychain" and a 22-year-old who has never carried cash.

1.2 Proportional Security

A fundamental critique of modern authentication is that it applies uniform security posture regardless of what is being protected. A login to check a delivery ETA requires the same password complexity as a login to initiate a wire transfer. This is disproportionate and creates unnecessary friction for low-risk interactions.

HCA implements proportional security: authentication rigor scales with the value at risk. An EnRoute account with a cash-loaded credit balance does not contain credit card numbers, bank account details, Social Security numbers, or advertising profiles. The maximum financial exposure is the credit balance, which was loaded with cash. A compromised EnRoute account gives the attacker access to a ride they must pay for in person. The economics of attacking this system do not justify sophisticated methods.

Furthermore, the data ReachTech collects has no extraction value. The system may know that a customer prefers Coke Zero, that their father recently passed away and conversations should be handled gently, or that they typically travel the Wichita-to-Salina corridor on Thursdays. This is relational context — data that exists to make the service more human. It has no market value. A hacker who steals it has stolen nothing worth selling. The only people who value this data are the people it belongs to.

> **DESIGN PRINCIPLE:** We do not need a fortress because there is nothing in the vault that a thief would want. Security posture is calibrated to actual risk, not theoretical worst-case scenarios.

2. The Face-First Authentication Flow

Version 0.1 of this specification used a name-first flow: the user typed their name, the system searched for candidates, and the face resolved among them. Peer review identified a fundamental flaw in this approach: how do the candidate face embeddings reach the device without a privacy violation? Sending other users' biometric templates to a device is unacceptable. Performing server-side face matching violates the on-device principle.

Version 0.2 inverts the flow. The face comes first. The device already knows who you are because your face was enrolled on this specific device. No candidate set is needed. No server is involved. No other user's biometric data touches your phone.

2.1 The Three-Step Authentication Sequence

Step 1 — Face: The user opens the app or device. The camera activates and performs a face match against the enrolled profile(s) stored on-device. This is a 1:1 match (single user) or 1:N where N is the number of profiles enrolled on this specific device (household scenarios, typically 2-5). No data leaves the device. No server is contacted. The system now knows who is holding the device.

Step 2 — Voice ("Name, please"): The system prompts: "Name, please." The user speaks their name. This step serves three simultaneous purposes: (a) Confirmation — a thief holding the phone to an unconscious person's face cannot pass this step. (b) Voice biometric — the spoken name is compared against the stored voice embedding, providing a second independent biometric signal. (c) Duress detection — the system analyzes how the name is spoken, comparing cadence, pitch, rhythm, and stress markers against the user's baseline. Deviation from baseline triggers configurable responses based on the user's security tier.

Step 3 — Contextual Scoring: In the background (under 500ms), the system evaluates device fingerprint, geolocation, time of day, and behavioral consistency. These signals combine with the face and voice scores to produce a composite confidence score. Access is granted, challenged, or denied based on the score and the user's security tier.

> **USER EXPERIENCE:** Open the app. Look at the screen. Say your name. You are in. The entire process takes under three seconds. The system has performed two independent biometric checks, a duress analysis, and a contextual evaluation without the user doing anything unnatural.

2.2 Per-Device Enrollment

Each device is enrolled separately. When a user accesses HCA on a new device for the first time, the enrollment process is:

For EnRoute (Community Tier): A pace car driver or verified user confirms the new user's identity in person. Face and voice enrollment occur on the user's device during this interaction. This is human-to-human vouching — the same model communities have used forever.

For Jackson AI (Standard Tier): The user authenticates on the new device using their password (or via push notification to an existing trusted device). Upon successful authentication, face and voice enrollment occur on the new device.

For Ghost Mobile (Protected Tier): Enrollment occurs in person at a Ghost Mobile activation point or via video verification with a ReachTech team member. Government ID is checked by a human (not scanned or stored). Face and voice enrollment occur on the Ghost Mobile device.

After enrollment, the device knows exactly whose face and voice to expect. The system does not search a database. It does not download templates. It simply recognizes the person it has already met — the same way a friend recognizes you when you walk into a room.

3. Tiered Security Model

HCA implements three security tiers. Users select their tier during enrollment and can change it at any time. Each tier adjusts the sensitivity of duress detection, the strictness of confidence thresholds, and the available response options.

	Community	Standard	Protected
Intended for	Everyday users (Regi). Low-risk accounts. Delivery, rideshare, basic services.	Most users. Moderate-risk accounts. Financial transactions, personal AI.	Users facing legal, political, or personal threats. Journalists, activists, anyone who needs maximum control.
Authentication	Face + spoken name. System learns voice baseline over time.	Face + voice biometric match. Active duress analysis on every login.	Face + voice + configurable duress responses including kill phrases, data wipe, and decoy modes.
Duress response	Passive monitoring. System learns baseline and would flag anomalies but takes no automatic action.	Anomalous voice triggers soft challenge (password prompt). Pre-configured duress phrase grants restricted access and silently alerts trusted contact.	Kill phrase locks device, wipes specified data, or presents a decoy environment. Silent alert to trusted contact and/or legal counsel. Configurable per-user.
Password	Optional. Advisory only.	Recommended. Used as fallback and for Panic Mode.	Required. Serves as Fifth Amendment-protected fallback. Separate duress password available.
Example product	EnRoute passenger	Jackson AI, EnRoute driver	Ghost Mobile

3.1 Community Tier — "Show Me Your Face, Tell Me Your Name"

The Community Tier is designed for users like Regi — a 72-year-old who has had the same email for ten years and does not know what a keychain password is. The authentication experience is: open the app, the system sees your face, it asks you to say your name, you are in.

The system begins learning the user's voice baseline from the first login. Over weeks and months of normal use, the voice model becomes increasingly precise. Duress detection is passive — the system monitors for anomalies but does not take automatic action at this tier. If a significant voice anomaly is detected, the system may ask a gentle follow-up question ("Everything okay today?") but does not lock the account or alert contacts without the user's prior configuration.

The Community Tier reflects the proportional security principle: an EnRoute passenger account with a $20 credit balance does not warrant the same friction as a bank account. The face confirms identity. The voice confirms presence and willingness. That is sufficient for the risk level.

3.2 Standard Tier — Active Verification

The Standard Tier adds active voice biometric matching and real-time duress analysis. The voice embedding is compared against the stored profile on every login. Stress markers — elevated pitch, compressed cadence, micro-tremors, unusual pauses — are compared against the user's established baseline.

Users at this tier may configure a duress phrase: a specific way of saying their name (using a middle name, a nickname, or a pre-arranged variation) that the system recognizes as a distress signal. When the duress phrase is spoken, the system grants access to a restricted view of the account while silently alerting a designated trusted contact with the device's location.

The restricted view is visually indistinguishable from a normal login. An attacker watching the screen cannot tell that the user has triggered a duress alert. The trusted contact alert is sent via a background channel that does not produce a visible notification on the device.

3.3 Protected Tier — Maximum Control

The Protected Tier is designed for users who face active legal, political, or personal threats to their privacy. This includes individuals under law enforcement investigation, journalists protecting sources, domestic abuse survivors, activists, and anyone who requires the highest level of control over their digital identity.

At this tier, the user configures one or more kill phrases — spoken words or phrases that trigger immediate defensive actions. These actions are configurable and may include:

Lock: The device immediately locks and requires the password (Fifth Amendment-protected testimonial evidence) to reopen. Biometric authentication is disabled until the password is entered.

Wipe: Specified data is securely erased from the device. The user pre-configures which data categories are wiped (messages, files, browsing history, specific apps). The wipe is silent and instantaneous.

Decoy: The device presents a decoy environment — a clean home screen with innocuous apps, a separate contact list, and no access to protected data. To an observer, the device appears to be functioning normally. The real environment is encrypted and inaccessible until the correct password is entered.

The kill phrase is processed entirely on-device using local speech recognition. It is never transmitted. The phrase itself is stored as a hashed embedding, not as plaintext. Even if the device's storage is forensically examined, the kill phrase cannot be recovered.

> **LEGAL NOTE:** The Protected Tier is designed to support users who wish to invoke constitutional protections against compelled device access. Courts have generally

> held that biometrics (face, fingerprint) can be compelled, while passwords are protected as testimonial evidence under the Fifth Amendment. Case law is evolving and jurisdictions are split. ReachTech does not represent these features as guaranteed legal protections. Users facing legal proceedings should consult counsel.

4. Dynamic Confidence Scoring

Version 0.1 used static weights for confidence scoring. Peer review identified this as a predictable attack surface: if an attacker knows the weights never change, they can optimize their approach. Version 0.2 implements context-aware dynamic scoring where weights shift based on environmental conditions.

4.1 Base Weights

Signal	Base	Dynamic Adjustment
Face Match	0-35	Reduced in low light, partially occluded face, or degraded camera conditions. System reports confidence level of the match, not binary yes/no.
Voice Match	0-25	Reduced in noisy environments. Increased if voice embedding matches with high confidence. New signal in v0.2.
Device	0-20	Known device = 20. Recently added = 10. Unknown = 0. Jailbroken/rooted device = reduced by 10 (requires hardware attestation).
Location	0-10	Home geofence = 10. Usual city = 7. Traveling = 3. New country = 0. Offline/no GPS = weighted at 0 (not penalized, just absent).
Behavior	0-5	Consistent time/cadence = 5. Anomalous = 0. Weight intentionally low — supplementary signal only.
Password	Override	Not a weighted signal. Correct password serves as a threshold override in low-confidence scenarios. Grants access at 50+ with all other signals, regardless of individual signal weakness.

> **DESIGN CHANGE v0.2:** Password is no longer a weighted signal (10 points in v0.1). It is now a threshold override. This addresses peer review feedback that a correct password on a stolen device with a mismatched face should not contribute additive points. Instead, password serves as a bypass when the system cannot achieve high confidence through biometrics alone.

4.2 Threshold Actions by Tier

Score	Community	Standard	Protected
70-95	Grant access	Grant access	Grant access
45-69	Grant with advisory ("We didn't hear you clearly — could you say your name again?")	Soft challenge: password prompt	Soft challenge: password required
20-44	Password required	Password + trusted device confirmation	Password + 15-minute delay + trusted contact notification
0-19	Deny. Retry available.	Deny. Account holder notified.	Deny. Account frozen. Trusted contact and/or

			counsel alerted.

4.3 Anomaly Detection

The system performs anomaly detection on failed and successful authentication attempts. This includes analysis of input patterns (for password-based fallback scenarios), voice stress analysis (for spoken name step), and contextual anomalies (login from new country, unusual time, device change). ReachTech acknowledges that behavioral biometrics and anomaly detection are established fields with extensive prior art (BioCatch, BehavioSec/LexisNexis, Sift, Riskified). A freedom-to-operate analysis is required before commercialization of anomaly detection features.

5. Threat Model

This section defines the adversaries HCA is designed to defend against, the attacks it mitigates, and the attacks it does not claim to prevent.

5.1 Adversary Tiers

Tier 1 — Opportunistic: Casual thief, nosy acquaintance. Physical access to device, possibly knows target's name. No specialized tools. HCA defends fully.

Tier 2 — Targeted: Ex-partner, identity thief, social engineer. Knows target's name, location, patterns. May have photos. HCA defends substantially through dual biometric requirement and duress detection.

Tier 3 — Sophisticated: Organized crime, state actor. Deepfake capability, device emulation, zero-day exploits. HCA raises the cost of attack significantly but does not claim to fully defend. High-value accounts should use hardware security keys in addition to HCA.

5.2 Attack Vectors and Mitigations

5.2.1 Stolen Device

> **THREAT:** Attacker possesses target's device.

> **MITIGATION:** Face match fails (wrong face). Even if attacker holds phone to unconscious target's face, Step 2 (voice) fails — attacker cannot produce target's voice with correct biometric signature. Dual biometric requirement (face + voice) is significantly harder to defeat than either alone.

5.2.2 Presentation Attack (Deepfake / Photo)

> **THREAT:** Attacker presents photo, video, or deepfake of target's face. Attacker plays recorded audio of target saying their name.

> **MITIGATION:** Presentation Attack Detection (PAD) for face: liveness checks (blink, micro-movement, depth sensing where available). Voice replay detection: the system introduces subtle variation in the prompt cadence or adds a randomized follow-up ("And your city?") that a pre-recorded audio cannot answer. On Ghost Mobile devices, LiDAR/structured light is required for face PAD.

> **RESIDUAL RISK:** Real-time deepfake video with synchronized voice synthesis may defeat 2D liveness on lower-end devices. Ghost Mobile hardware mitigates this. Third-party devices at Community Tier accept this residual risk as proportional to the low value at stake.

5.2.3 Coercion

> **THREAT:** Attacker physically forces target to authenticate.

> **MITIGATION:** Duress detection is built into the normal authentication flow. The target's voice under coercion will exhibit stress markers detectable by the system. At Standard Tier, a pre-configured duress phrase triggers restricted access + silent alert. At Protected Tier, a kill phrase can lock, wipe, or present a decoy

> environment. The attacker cannot see or hear the difference between a normal login and a duress login.

5.2.4 SIM Swap

> **THREAT:** Attacker transfers target's phone number to attacker's SIM.

> **MITIGATION:** HCA does not use SMS-based verification at any tier. There is no phone number in the authentication chain. SIM swaps are irrelevant to HCA.

5.2.5 Biometric Model Poisoning

> **THREAT:** Attacker gains device access and repeatedly authenticates, gradually training the biometric model.

> **MITIGATION:** Model updates only occur on sessions scoring 70+. Blend rate limits any single session to less than 2% influence. Face embedding drift exceeding a configurable threshold from original enrollment triggers account freeze and password-only access. Voice embedding follows the same anti-drift protocol. 90-day rollback checkpoints for both face and voice models.

5.2.6 Session Hijacking

> **THREAT:** Attacker intercepts session token after successful authentication.

> **MITIGATION:** Session tokens are bound to device fingerprint. Continuous session validation checks device and behavioral signals at configurable intervals (default: every 5 minutes). Anomalous device change during an active session triggers immediate session termination and re-authentication. Financial transactions above a configurable threshold require step-up re-authentication (face + voice) regardless of session state.

5.3 Impact on Identity Theft

The overwhelming majority of identity theft is remote credential abuse: stolen passwords, leaked SSNs, and compromised email accounts used to open fraudulent accounts or authorize transactions. HCA substantially increases the cost and reduces the scalability of these attacks by requiring a live face and voice match for every interaction. An attacker cannot use a stolen SSN to generate a matching face and voice.

HCA does not eliminate all identity theft. Enrollment fraud (fake ID during verification), recovery path social engineering, session hijacking, and high-quality real-time deepfakes remain theoretical vectors. These are addressed by the threat model mitigations above and by the proportional security framework: the residual risk is calibrated to the value at stake.

5.3.1 The Farmer State Bank Model

On April 3, 2026, a ReachTech founder called Farmer State Bank in Oakley, Kansas, to request a $900 wire transfer. The teller processed the wire based on contextual signals: the caller's name matched an account; she recognized his voice; his father had deposited a check from a known source minutes earlier; the wire recipient was the same person who received wires in the preceding months; and the request was consistent with the account's history.

She did not request a password, a verification code, or a second device. She composited five human signals and made a confident judgment. She was correct. HCA digitizes this process. Small-town banks have been running human-centric authentication for a century. The problem is not that nobody knows how to do this. The problem is that when banking went digital, the teller's judgment was replaced with a password field and called progress.

5.4 Security Architecture vs. Security Theater

This analogy is offered as a structural metaphor, not an empirical comparison. Prior to September 11, 2001, aviation security relied on a single checkpoint. After September 11, security became structural: hardened cockpit doors, crew protocols, and passengers who will act against threats. The checkpoint (TSA screening) still exists, but the real security is architectural.

Current digital authentication is a single checkpoint: the password. HCA implements structural security: the face is the hardened door (physical presence required), the voice is the crew protocol (second independent barrier), duress detection is the passenger awareness (continuous monitoring for threats), and contextual signals provide environmental intelligence. Password complexity requirements are the shoe removal: theater performed for comfort, not security.

6. Continuous Biometric Evolution

Both face and voice embeddings evolve with every high-confidence authentication (score 70+). The blend rate is conservative: less than 2% influence per session. The system tracks gradual changes — aging, weight change, vocal changes from illness or aging — without hitting a recognition cliff. Checkpoint snapshots are retained for 90 days, enabling rollback if drift is detected.

The user is informed that biometric models are being updated. A Biometric Health dashboard shows last update date and model confidence. Through Jackson AI, the system may optionally acknowledge visible changes ("New haircut — looks good."). This is configurable and can be disabled.

Anti-drift protections: updates only on 70+ sessions; 2% blend rate cap; drift exceeding threshold from enrollment triggers password-only access; all checkpoints are signed and tamper-evident.

7. Account Lifecycle

7.1 Enrollment

Enrollment is per-device and tiered by product (see Section 2.2). Face and voice enrollment occur simultaneously during initial setup. The system captures the user's face and records them speaking their name to establish both biometric baselines. Community Tier enrollment requires in-person vouching. Standard Tier requires password or trusted device push. Protected Tier requires in-person verification with government ID checked by a human.

7.2 Account Recovery

Path 1 — Password: If biometrics fail (camera broken, voice lost to illness, device change), the user enters their password.

Path 2 — Trusted Contact: 1-3 designated contacts can initiate recovery. Requires the trusted contact's own authentication plus a mandatory waiting period (24-72 hours, configurable by tier). The waiting period prevents impulsive or coerced recovery attempts. The account holder is notified immediately when a recovery is initiated and can cancel it during the waiting period.

Path 3 — In-Person: User visits a ReachTech-affiliated location and re-verifies identity with a human. Fallback of last resort.

Path 4 — Physical Key: A printed one-time recovery code generated during enrollment. ReachTech does not retain a copy.

> **DESIGN DECISION:** There is no email-based recovery. Email accounts can be compromised. Recovery flows that depend on another digital account move the problem rather than solving it.

7.3 Device Transfer (The Regi Scenario)

When a user acquires a new device: (a) open the app on the new device; (b) authenticate via password (face and voice unavailable on new device); (c) if password is unknown, use Trusted Contact or In-Person recovery; (d) upon successful authentication, enroll face and voice on the new device; (e) old device enrollment is deprecated after 30 days or upon user confirmation.

At no point does the system send a verification code to a phone number that may have changed.

7.4 Legacy Contact

Users may designate a Legacy Contact who can access the account in the event of the user's death or incapacitation. Legacy access requires: (a) the Legacy Contact's own authentication; (b) a waiting period of 7 days; (c) the account holder does not cancel during the waiting period. Legacy access grants read-only access to account data and the ability to close the account. It does not grant the ability to make transactions.

8. Accessibility and Shared Devices

HCA must serve all users. A biometric-first system that excludes non-biometric users is technology-centric, not human-centric.

Visual impairment or blindness: Voice becomes the primary biometric. Face enrollment is optional. Authentication path: voice + device + location + password if needed. The spoken name step is accessible to all sighted and non-sighted users.

Facial differences or progressive conditions: Continuous biometric evolution tracks gradual changes. For acute changes (surgery, accident), Trusted Contact or In-Person recovery enables re-enrollment.

Cultural or religious objections: Users who decline facial or voice capture authenticate via password + device + location. No penalty, no reduced functionality.

Shared devices: Multiple enrolled profiles per device. Face match at Step 1 identifies which user is present. Voice confirms. Household mode (single account, PIN toggle between users) available at reduced security for Community Tier only.

Camera or microphone failure: If hardware is unavailable, authentication falls back to password + device + location. The system informs the user: "We can't see or hear you right now. Please enter your password."

9. Regulatory Compliance

9.1 BIPA (Illinois)

ReachTech will obtain explicit written consent for biometric collection from all users, regardless of state, as a matter of policy. Consent specifies: what is collected (face embedding, voice embedding), the purpose (authentication and duress detection), and the retention period (deleted on device wipe, account deletion, or user opt-out). Biometric data never leaves the device and is not transmitted to ReachTech servers. BIPA statutory damages of $1,000-$5,000 per negligent violation apply. On-device processing significantly reduces exposure but does not create a safe harbor. Legal counsel review is required.

9.2 GDPR and CCPA/CPRA

Face and voice data are special category / sensitive personal information. Legal basis: explicit consent (GDPR Art. 9(2)(a)). Users may opt out at any time without loss of service access. Biometric-free authentication path is available at all tiers.

9.3 Password Policy and Duty of Care

Passwords are never the sole authentication factor. They serve as a fallback within a multi-signal system. During enrollment, the system advises users who select very short or common passwords. This is advisory, not enforced. When a password is used outside a fully multi-signal context (e.g., Panic Mode on a new device with no biometrics available), the system recommends a stronger password for that specific scenario. ReachTech maintains cyber liability insurance.

10. FIDO2 and Passkey Compatibility

HCA is built on top of FIDO2/WebAuthn, not in opposition to it. The device signal is implemented as a FIDO2 passkey stored in the device's secure enclave. The passkey provides the cryptographic proof that the device is recognized. The face and voice provide the human proof that the correct person is holding the device. The user never sees or manages the passkey.

On devices without camera or microphone (e.g., library computers), authentication falls back to passkey + password, which is functionally standard WebAuthn with an additional identity layer.

11. Limitations and Non-Goals

HCA explicitly does not claim to solve the following:

Nation-state adversaries: Tier 3 actors with unlimited resources and zero-day exploits are beyond the scope of HCA. High-value targets should supplement HCA with hardware security keys.

Server-side breaches: HCA secures the authentication layer. Database compromise, session token theft at the server level, and infrastructure attacks require separate security controls.

Enrollment fraud: A person who presents a convincing fake ID during in-person enrollment can create a fraudulent account. HCA does not solve identity proofing beyond the human judgment of the enrollment verifier.

Universal biometric accuracy: Face and voice recognition have non-zero false rejection and false acceptance rates. These rates vary by hardware, lighting, environmental noise, and user physiology. Phase 1 calibration will establish real-world performance baselines.

Devices without cameras or microphones: HCA degrades gracefully to password + device + location but cannot provide its full security posture on hardware that lacks biometric sensors.

12. Implementation Roadmap

Phase 1 — EnRoute MVP (Q2-Q3 2026)

Community Tier only. Face + spoken name authentication for driver and passenger accounts. Voice baseline learning begins. Confidence scoring active but logging-only for 90 days (no automated denials) to calibrate thresholds, measure false rejection rates, and identify real-world edge cases. All data used to tune the system before enforcement begins.

Phase 2 — Jackson AI + Standard Tier (Q3-Q4 2026)

Standard Tier with active voice biometric matching and duress detection. Continuous biometric evolution launched. Biometric Health dashboard. Behavioral signal collection with explicit consent and transparency.

Phase 3 — Ghost Mobile + Protected Tier (2027)

Full HCA at OS level on Ghost Mobile. LiDAR/structured light for face PAD. Kill phrase with lock/wipe/decoy. FIDO2 passkey as cryptographic layer. BIPA/GDPR compliance review with legal counsel completed.

Phase 4 — Open Protocol (2028+)

Publish HCA specification as open protocol. Invite third-party adoption. Explore decentralized identity (DID) compatibility. Goal: HCA becomes a standard, not a proprietary feature.

———

End of Document B — ReachTech Identity Architecture v0.2