Proof of Human *Biological Cryptographic Identity* in the Age of AI Impersonation

A protocol for distinguishing humans from machines when every external credential has fallen.

SOMA RESEARCH / OPENCREW • APRIL 2026 • PROOFOFHUMAN.WORLD

ABSTRACT

Artificial intelligence can now clone a human voice from three seconds of audio, generate photorealistic deepfake video from a single photograph, and pass identity verification checks with higher accuracy than legitimate users. Every external authentication method, from passwords and biometrics to CAPTCHA and liveness detection, has been defeated or is failing. This paper examines the architectural collapse of digital trust, quantifies the scale of AI-driven identity fraud, and proposes that **cryptographic identity derived from the continuous biological signals of a living human body** represents the only structurally undefeatable solution. Soma is building this infrastructure as an open protocol layer for the internet.

KEYWORDS cryptographic identity · biometric fuzzy extractors · deepfake fraud · proof-of-personhood · zero-knowledge authentication · biosignal cryptography · decentralized verification

§ 01 The Identity Verification Crisis

In July 2025, a mother in Dover, Florida received a phone call from her daughter. The voice was crying, desperate, claiming she had been in a car accident and lost her unborn child. The voice pleaded for $15,000 to avoid criminal charges. The mother sent the money. She later discovered she had been speaking to an AI-generated clone of her daughter's voice the entire time.[10]

This is not an isolated case. In February 2024, a finance worker at UK engineering firm Arup was tricked into wiring $25 million after participating in a video conference where every other participant, including the company's CFO, was an AI-generated deepfake. The worker reported that the synthesized executives looked, sounded, and behaved exactly like their real counterparts.[6, 11]

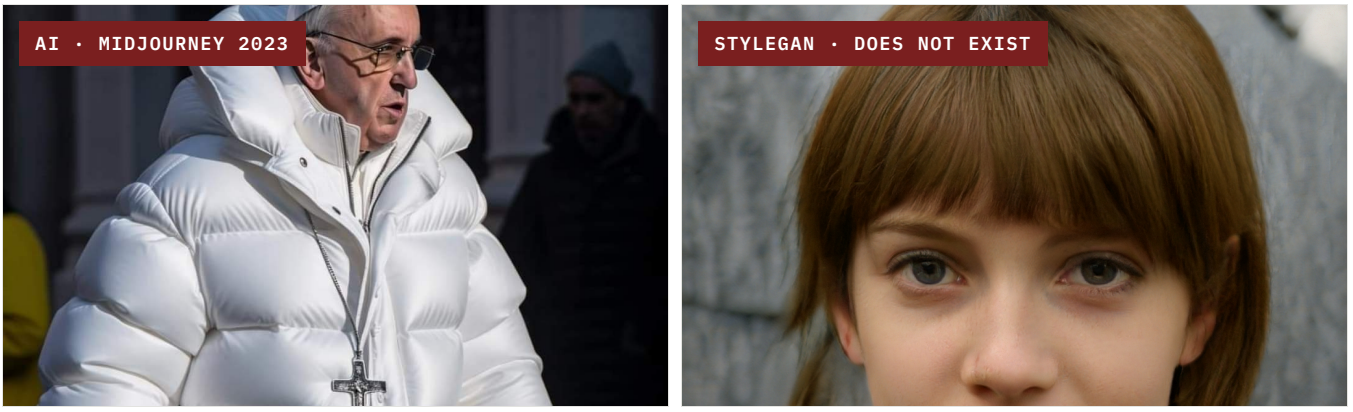

FIGURE 1 **Photorealistic synthetic humans.** Left: the viral 2023 deepfake of Pope Francis in a designer puffer jacket, generated by Midjourney, fooled millions before being identified as fake. Right: a StyleGAN-generated portrait of a woman who has never existed. No birth certificate, no photographs, no body. By 2025, the same neural architectures produce results that human observers cannot distinguish from real photographs; detection accuracy has fallen to just **24.5%** for high-quality fakes. [7,16] Source: Wikimedia Commons (public domain).

$1.1B	3SEC	1,600%	1/4
DEEPFAKE FRAUD LOSSES · US 2025	AUDIO SAMPLE TO CLONE ANY VOICE	SURGE IN AI VISHING · Q1 2025	ADULTS HIT BY AI VOICE SCAM
Keepnet Labs [5]	McAfee [1]	Right-Hand [8]	McAfee [1]

The scale of the problem is accelerating. Deepfake fraud losses in the United States reached $1.1 billion in 2025, tripling from $360 million the year before. Deepfake-enabled vishing surged over 1,600% in Q1 2025 compared to the end of 2024. One in four adults have experienced an AI voice scam, with 77% of targeted victims reporting financial losses.[4, 5, 8]

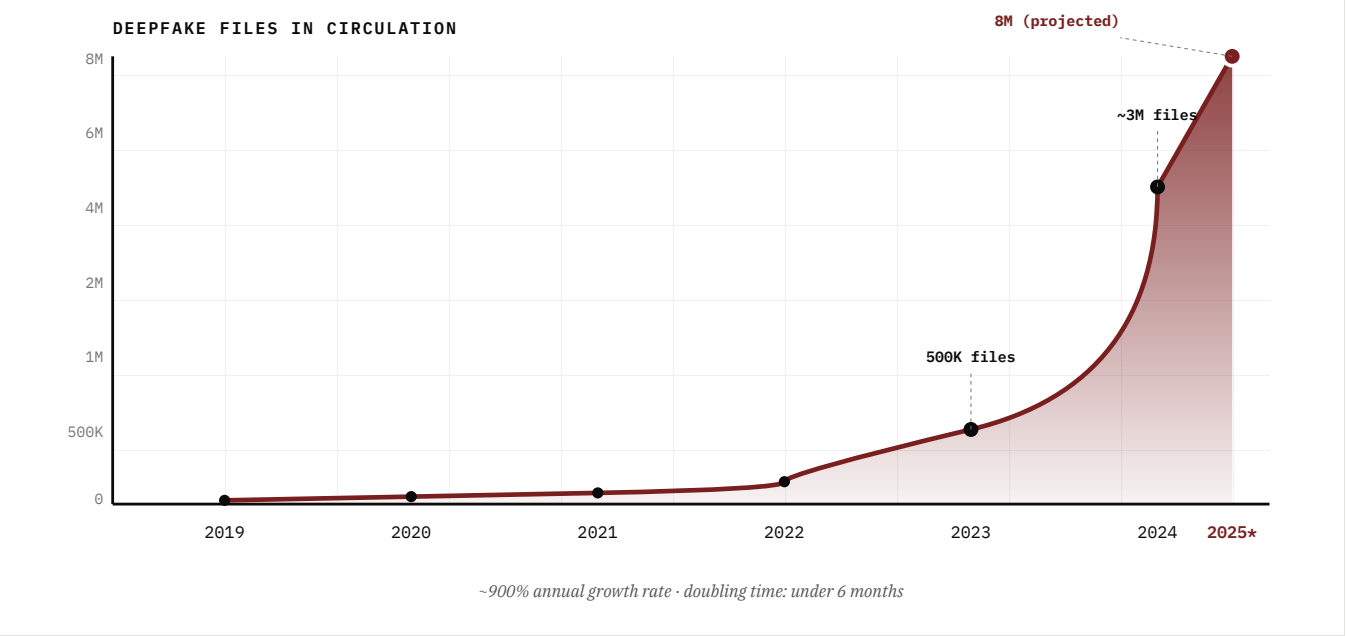

FIGURE 2 **Exponential growth in deepfake content.** Circulating deepfake files grew from approximately 500,000 in 2023 to a projected 8 million by 2025, a 900% annual growth rate. The doubling time has fallen below six months, outpacing every detection effort deployed against it. Sources: Sumsub Identity Fraud Report 2025, Deepstrike.io.[4, 7]

1.1 Voice Cloning: The Most Accessible Attack Vector

Modern AI tools can clone a voice using just three seconds of clear audio, achieving an 85% voice match. Source audio is scraped from social media, podcasts, and webinars. The deepfake robocall of President Biden used to disrupt the 2024 New Hampshire primary cost $1 to create and took less than 20 minutes. CEO fraud using cloned voices now targets at least 400 companies per day. Voice cloning fraud rose 680% in the past year.[1, 6, 7]

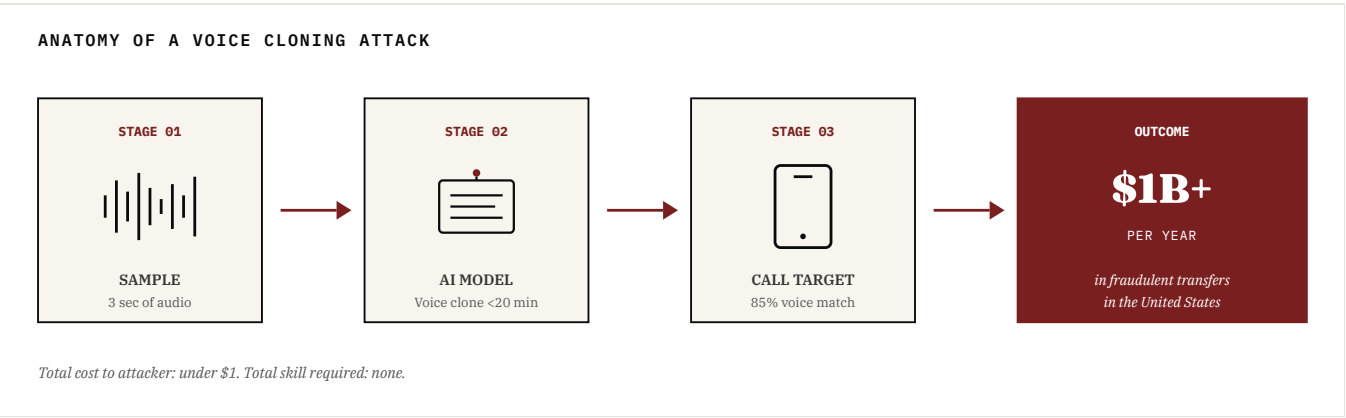

FIGURE 3 **Anatomy of a voice cloning attack.** The entire pipeline, from scraping social media audio to executing a fraudulent call, can be completed in under twenty minutes for less than one US dollar. The 2024 deepfake robocall of President Biden used to disrupt the New Hampshire primary cost exactly $1 to produce.[1, 7]

1.2 The Geographic Spread

The crisis is not regional. North America saw a **1,740%** increase in deepfake fraud between 2022 and 2023, with losses exceeding $200 million in Q1 2025 alone. Asia Pacific recorded a **1,530%** increase. The

Maldives saw a **2,100%** year-over-year surge. Sumsub's 2025 Identity Fraud Report finds that 11% of all global fraud is now deepfake-driven.[7]

§ 02 The Authentication Defeat Timeline

Every verification method humans have built relies on a credential that exists outside the human body. This architectural flaw means any sufficiently advanced AI can copy, intercept, or synthesize the credential. The defeat is not a future risk. It has already happened.

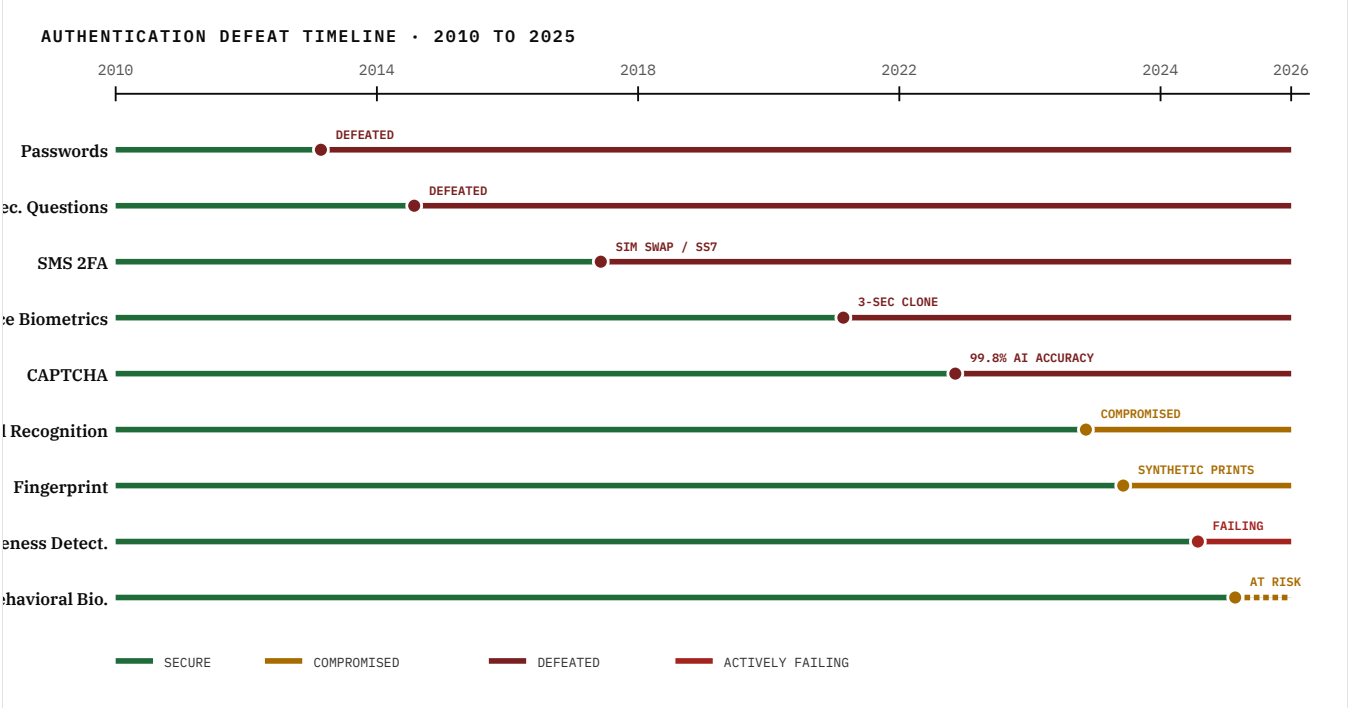

FIGURE 4 **Authentication defeat timeline.** Every external authentication method has either fallen or is actively failing against AI-powered attacks. The sequence is accelerating: methods that took a decade to break in the 2010s now fall within months of deployment. Sources: synthesized from McAfee, ETH Zurich, UC Irvine.[1, 2, 3]

TABLE 1 **Status of major authentication methods against AI-powered attacks.**

METHOD	PRIMARY ATTACK VECTOR	STATUS
Passwords	Brute force, phishing, credential stuffing	DEFEATED
Security Questions	Social engineering, public data mining	DEFEATED
SMS Two-Factor	SIM swapping, SS7 protocol interception	DEFEATED
Voice Biometrics	Real-time voice cloning from 3-sec sample	DEFEATED
CAPTCHA	AI solves all variants: 99.8% vs human 50–84%	DEFEATED
Facial Recognition	AI-generated faces, deepfake video injection	COMPROMISED
Fingerprint Scanners	Synthetic fingerprints from photographs	COMPROMISED
Liveness Detection	Adversarial ML, API-level injection	FAILING
Behavioral Biometrics	AI behavioral modeling from minimal data	AT RISK

The CAPTCHA case is illustrative. Researchers at UC Irvine found that AI bots solve CAPTCHAs with 99.8% accuracy, while human accuracy ranges from 50–84%. In September 2024, ETH Zurich researchers demonstrated 100% success against Google's reCAPTCHA v2.[2, 3, 13] The system designed to distinguish humans from machines now performs that task in reverse: it identifies humans by their inferiority at the task.

> ***The common vulnerability is architectural: the credential exists outside the human body. Anything external to the body can be replicated by a sufficiently capable adversary. AI is now that adversary.***

§ 03 The Soma Thesis

If every external credential can be forged by AI, the only remaining option is a credential that *is not external*, one that is generated by the living body itself, in real time, and cannot exist independently of the person who produces it.

The continuous biological signals of a living human body possess three properties that make them uniquely resistant to AI synthesis:

THE THREE PROPERTIES

Irreducible complexity. These signals emerge from billions of cells operating in concert. They are the output of a living system whose full state cannot be computationally modeled in real time.

Continuous variability. Unlike a fingerprint or iris pattern, biological signals change moment to moment while maintaining individual-specific invariant features. A recorded sample is immediately stale.

Physical inseparability. The signals cannot be separated from the body. There is no database to breach, no token to steal. If the body is not present, the key does not exist.

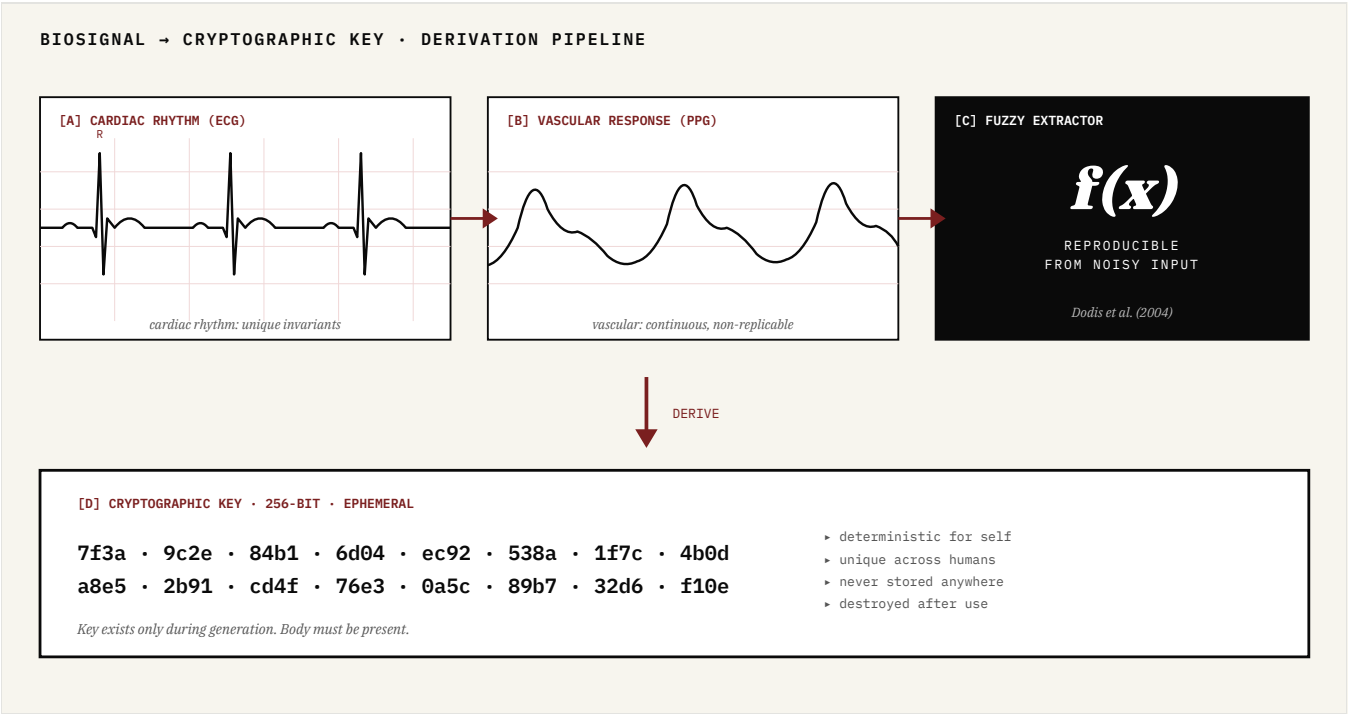

FIGURE 5 **Biological signal → cryptographic key derivation.** Soma combines multiple continuous biological streams (cardiac rhythm, vascular response) and feeds them through a fuzzy extractor, a cryptographic primitive that derives stable keys from noisy biometric input. The same person reliably produces the same key; a different person produces a different key. The key exists only during generation and is destroyed immediately after use. No signal, no key.[16]

Soma derives cryptographic identity from these signals using **fuzzy extractors**, cryptographic primitives designed to derive stable keys from noisy biometric data. The same person reliably produces the same key. A different person produces a different key. The key exists only during generation and is destroyed immediately after use.

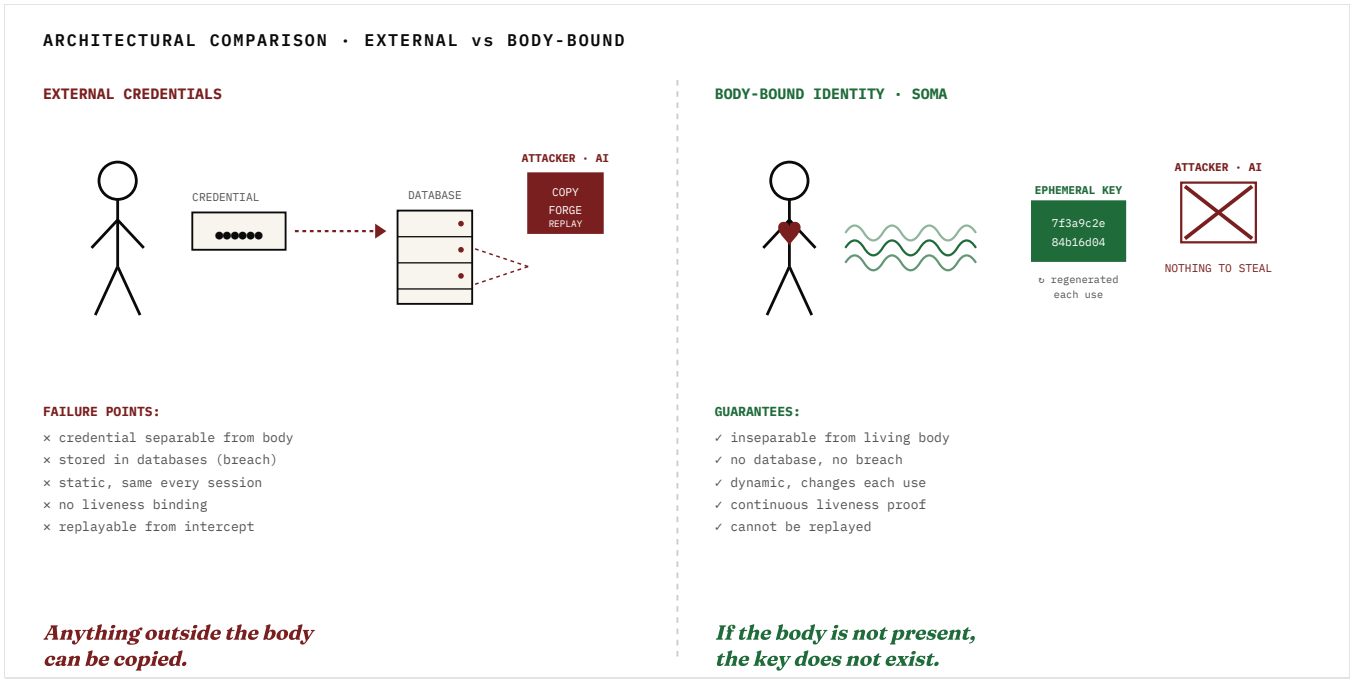

FIGURE 6 **Architectural comparison.** Left: every external authentication scheme externalizes the credential, creating an attack surface that AI now exploits at industrial scale. Right: Soma's body-bound architecture removes the credential entirely. There is nothing in any database to breach, nothing on any device to steal. The key is regenerated from continuous biological signal each session and destroyed immediately after.

EXTERNAL CREDENTIALS	BODY-BOUND IDENTITY (SOMA)
× Exist outside the body	✓ Inseparable from the living body
× Can be copied or intercepted	✓ Cannot be copied or intercepted
× Static, same every session	✓ Dynamic, changes moment to moment
× Stored in databases (breach risk)	✓ Never stored anywhere
× One-time check, no liveness	✓ Continuous liveness proof
× Separable from the person	✓ Key exists only during generation

§ 04 System Overview

Soma transforms biological signals into cryptographic proof of identity in three stages. No biometric data is ever stored, transmitted, or accessible to any external party.

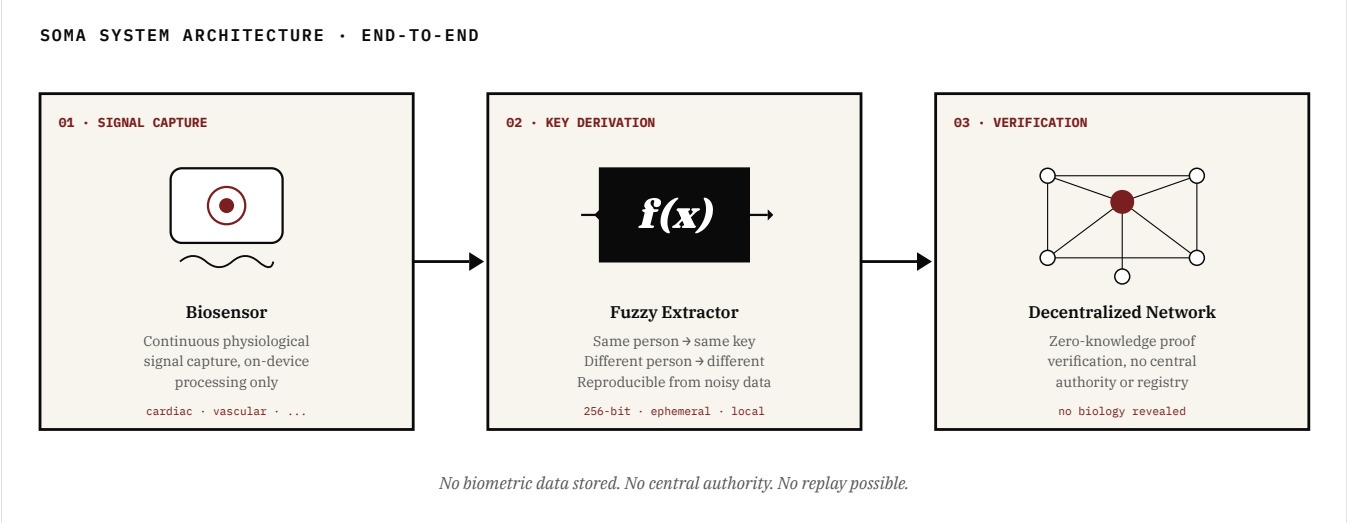

FIGURE 7 **The Soma protocol stack.** A miniaturized biosensor captures continuous physiological signals, processes them entirely on-device, derives a session key via fuzzy extractor, and broadcasts a zero-knowledge proof of humanness to a decentralized verification network. The biological signal never leaves the device. The key never persists. The architecture has no central authority and no central registry to compromise.

TABLE 2 **Cryptographic properties of the Soma key derivation.**

PROPERTY	DESCRIPTION
DETERMINISTIC	Same person reliably produces the same cryptographic key across sessions
UNIQUE	Different person produces a different key; collision probability is negligible
EPHEMERAL	Key exists only during generation, destroyed after use, never stored
NON-EXTRACTABLE	Biological signal cannot be reverse-engineered from the public key
LIVENESS-BOUND	Key generation requires continuous biological signal; stops when body separates
ZERO-KNOWLEDGE	Verification proves humanness without revealing any biological data

§ 05 Applications

5.1 Identity Verification

Prove humanness and identity without passwords, biometric databases, or central authority. Replaces the entire authentication stack with a single biological proof. Applicable to platform onboarding, KYC compliance, and age verification.

5.2 Authenticated Communication

Every call, message, and video carries cryptographic proof that a verified human created it. Deepfake calls and AI-generated messages become instantly identifiable as unsigned. Applicable to voice, video, messaging, and email.

5.3 Financial Authentication

Authorize transactions with proof that a living human initiated the action. No amount of stolen data can replicate the biological key. Applicable to payments, banking, trading, and contracts.

5.4 Platform Trust

Verified human status anchored to cryptographic biological proof. Bots, farms, and synthetic identities become structurally impossible. Applicable to social media, marketplaces, voting, and reviews.

§ 06 Why Now

6.1 The AI Capability Explosion

Synthetic humans are indistinguishable from real ones in digital contexts. Deepfake content grows at 900% annually. Fraud attempts increased 3,000% in a single year. The window between “AI can fool some systems” and “AI can fool all systems” is closing.

6.2 The Digital Identity Crisis

Institutions are responding to fraud by mandating centralized digital identity systems: databases linking biometric data to government credentials. These create single points of failure and concentrate power over individual identity. *Centralized identity creates centralized vulnerability.*

6.3 The Biological Authentication Window

Biosensor miniaturization, low-power cryptographic processing, and body-area networking have made biologically-derived cryptography technically feasible for the first time. The components exist. The science is sound. What has been missing is the architecture. That architecture is Soma.

§ 07 Design Principles

BODY-BOUND	Identity is inseparable from the living body. If the body is not present, the key does not exist.
ZERO-KNOWLEDGE	Verification proves humanness without revealing biological data. No biometric database. No central store to breach.
DECENTRALIZED	No single entity holds unilateral control over identity issuance, verification, or revocation.
MINIMALIST	The system proves exactly what needs to be proven and nothing more.
CONTINUOUS	Ongoing proof that the authenticated human remains present. Body separates → authentication ceases.
OPEN PROTOCOL	Any platform, application, or institution can integrate verified human status. Global standard, not proprietary product.

§ 08 Conclusion

The infrastructure for verified human identity does not exist yet. No system currently in operation can provide cryptographic proof that a digital interaction originates from a living human being. The threat is measured in billions of dollars of fraud, voice cloning from seconds of audio, and video conferences where every participant is synthetic.

Soma is building the replacement: a protocol layer that allows the internet to distinguish human from machine. The approach is grounded in one insight: **the living human body is the only source of identity that AI cannot forge.**

> ***For the first time since the internet was built, there is a credential that AI cannot copy, because it is not a credential at all. It is the living body itself.***

SOMA

PROOF OF HUMAN

devs@opencrewai.com

proofofhuman.world

opencrewai.com

§ References

[1] McAfee. *The Artificial Imposter: AI Voice Cloning Consumer Survey.* 2024.

[2] Tsudik, G., Searles, A. et al. *An Empirical Study and Evaluation of Modern CAPTCHAs.* University of California, Irvine. 2023.

[3] ETH Zurich. *Breaking reCAPTCHAv2 with 100% Accuracy Using AI.* September 2024.

[4] Resemble AI. *Q1 2025 Deepfake Incident Report.* 2025.

[5] Keepnet Labs. *Deepfake Fraud Losses in the United States.* 2025.

[6] Brightside AI. *Deepfake CEO Fraud: $50M Voice Cloning Threat to CFOs.* October 2025.

[7] Deepstrike.io. *Deepfake Statistics 2025: AI Fraud Data and Trends.* September 2025.

[8] Right-Hand Cybersecurity. *The State of Deep Fake Vishing Attacks in 2025.* October 2025.

[9] European Parliament / EPRS. *Scam Calls in Times of Generative AI.* 2025.

[10] American Bar Association. *The Rise of the AI-Cloned Voice Scam.* September 2025.

[11] CNN. *Finance worker pays out $25 million after video call with deepfake CFO.* February 2024.

[12] Deloitte Center for Financial Services. *Generative AI Fraud Projections.* 2024.

[13] CoinDesk. *Kill the Captcha: They Don't Work, Here's What Does.* September 2025.

[14] Norton. *Top 5 Ways Scammers Have Used AI and Deepfakes in 2025.* October 2025.

[15] Sumsub. *Identity Fraud Report 2025-2026.* 2025.

[16] Dodis, Y., Reyzin, L., Smith, A. *Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data.* EUROCRYPT 2004.

[17] Rossler, A. et al. *FaceForensics++: Learning to Detect Manipulated Facial Images.* ICCV 2019.

[18] SoSafe. *How to Spot a Deepfake.* 2024.