Stateless Device Pseudonymity and Secure Messaging in Cognition-Native Systems

by Nick Clark | Published May 25, 2025 | Modified January 19, 2026

Introduction: The Limits of Key-Based Identity

Continuity-based biological identity treats identity as an evolving trajectory rather than a static credential. If identity and accountability belong to a continuously validated human, then devices must not become permanent identity anchors. Persistent device identifiers and long-lived keys create correlation surfaces that undermine continuity-based privacy and allow captured hardware to impersonate, surveil, or outlive the authority that originally used it.

Conventional digital authentication relies on long-lived public–private keypairs, centralized registries, and revocation infrastructure. These systems introduce systemic risk: keys can be copied, stolen, correlated, or rendered obsolete by advances in cryptanalysis, including quantum algorithms.

In stateless, distributed, or cognition-native environments—such as autonomous agents, edge devices, or intermittently connected systems—the assumption that persistent secrets can be safely stored and managed is increasingly untenable.

The architecture below therefore treats devices as revocable delegates. Trust is accumulated through continuity of behavior over time, not through possession of static secrets.

References to fragility, surveillance risk, or hardware capture describe architectural failure modes rather than claims about specific systems, actors, or threat environments. They are included to motivate structural design choices, not to assert inevitability or operational conclusions.

1. Memory-Native Identity and Trust Slopes

The disclosed architecture replaces static credentials with memory-native identity. Devices and agents express identity as a trust slope: a monotonically advancing sequence of Dynamic Device Hashes (DDHs) or Dynamic Agent Hashes (DAHs), each derived from a prior trusted state and fresh, non-exportable unpredictability.

A verifier does not ask whether an identifier is valid in isolation. It asks whether the presented identity is a valid successor of a previously trusted state under policy-bounded continuity rules. Identity is therefore evaluated as behavior over time, not possession of a secret.

This makes device participation compatible with continuity-based human identity: the device contributes continuity evidence without becoming a stable identifier that can be tracked across contexts.

2. Dynamic Device Hashes and Dynamic Agent Hashes

A Dynamic Device Hash is generated by a host device using locally available unpredictability, such as a hardware anchor combined with volatile salts, a stability-tuned local state vector processed by a strong extractor, or a hybrid of both. A Dynamic Agent Hash is generated analogously by a semantic agent and may be entangled to the host on which it executes.

Each new DDH or DAH is computed as a successor of the prior value, producing a cryptographically verifiable lineage. Observing a hash does not enable impersonation; generating a valid successor requires access to the device or agent's local entropy and memory.

Because the identifier advances, correlation is minimized and capture resistance improves: prior observations do not enable durable tracking or forward impersonation.

3. Two-Stage, Memory-Resolved Authentication

Secure messaging is achieved through a two-stage authentication process. First, a sender places

its current dynamic hash in the transport header, allowing receivers to perform fast continuity screening before any decryption occurs.

Second, payload encryption keys are derived from the recipient's current dynamic identity, and an embedded copy of the sender's identity is included within the encrypted payload. A message is accepted only if both header-level and payload-level validations succeed.

This design binds routing integrity, confidentiality, and semantic authenticity to the same memory-native identity substrate, while keeping identity rotation and revocability native to the protocol rather than dependent on external registries.

4. Resistance to Spoofing, Replay, and Key Compromise

Spoofing is prevented because an attacker cannot synthesize valid successors without access to local unpredictability. Replay is prevented by enforcing monotonic progression along the trust slope and rejecting reused or regressive identities.

Because no persistent private keys exist, there is nothing to exfiltrate, rotate, or revoke. Compromise of a single identity state does not enable future impersonation.

This is particularly important in delegation settings: a captured device does not remain a durable identity, and its ability to participate can be degraded through continuity failure rather than brittle revocation mechanisms.

5. Post-Quantum Alignment

The security of the system does not depend on discrete logarithms, elliptic curves, or factorization. Instead, it relies on hash preimage resistance and the unpredictability of locally derived entropy.

Even under quantum search acceleration, the probability of forging a valid successor remains negligible with appropriate entropy parameters, making the architecture suitable for long-lived

defense and infrastructure deployments.

Post-quantum suitability is discussed in terms of cryptographic primitives and entropy assumptions, not as a guarantee against all future cryptanalytic advances or implementation flaws.

6. Deployment in Defense, Infrastructure, and Autonomous Systems

Memory-native authentication is well suited for military, intelligence, and critical infrastructure environments where devices may be disconnected, captured, or operating without centralized trust services.

The architecture supports stateless operation, delayed verification, quorum-based recovery, and compatibility with legacy PKI systems through isolated adapters—without contaminating the core identity substrate.

Because devices remain pseudonymous and revocable, the system can support large fleets and autonomous deployments without creating a single persistent identity surface that adversaries can inventory and exploit.

Deployment contexts are illustrative of structural applicability rather than claims of readiness, authorization, or adoption by any specific defense, intelligence, or infrastructure organization.

Conclusion

Dynamic Device Hashes, Dynamic Agent Hashes, and trust-slope continuity define a fundamentally different approach to identity and authentication. By deriving trust from memory and behavior rather than static secrets, this architecture defines conditions under which secure, scalable, and post-quantum-aligned authentication can be computed in cognition-native systems, without asserting deployment completeness or outcome guarantees.

Crucially, this model complements continuity-based biological identity by keeping devices in a subordinate role: devices contribute continuity evidence and messaging security without becoming permanent identity anchors.