



WHITEPAPER

The Architecture of Accumulation

How Attackers Cultivate Identity in a Transactional Era

By Jen Baldwin, CFE

© 2026 myNetWatchman, LLC. All rights reserved.

Abstract

This white paper examines the structural shift in digital fraud from immediate transactional attacks to the long-term cultivation of identity assets. By exploiting the inbox as a universal digital passport and continuity layer, adversaries bypass traditional point-in-time defenses. We explore the mechanics of trust accumulation—including inbox aging and alias fanout—and argue for a transition toward longitudinal, predictive identity verification systems to secure modern digital infrastructure.

1. Introduction: The Institutional Blind Spot

The architecture of digital trust is suffering from a persistent structural lag. Over the past two decades, the fraud and risk industry optimized its systems for a highly specific set of problems: stolen credit cards, automated password guessing, and transactional friction. To combat these threats, organizations invested heavily in securing the application perimeter. They deployed advanced device fingerprinting, behavioral biometrics, and multi-factor authentication, operating under the assumption that identity risk could be evaluated at the moment a user attempts to log in or initiate a transfer.

Meanwhile, attackers evolved. Recognizing that the front door had been heavily fortified, adversaries shifted their focus upstream, moving away from immediate, transactional fraud and toward the long-term cultivation of identity assets.

Modern identity systems increasingly depend on email infrastructure while inheriting many of its weaknesses. The email address was engineered in the 1970s as a pragmatic routing mechanism for networked mainframes, not as a cryptographic proof of personhood.¹ Yet, as the digital economy expanded, platforms required a persistent identifier that could follow a user across sessions and devices. The inbox quietly assumed the role of a universal digital passport.

By designing trust systems around an identifier that was never intended to serve as a secure, lifelong credential, organizations inadvertently granted attackers a centralized target.

The Continuity Layer and the Recovery Fallback

The vulnerability of this dependency becomes acutely apparent in the context of passwordless authentication. The security industry has aggressively promoted passkeys, magic links, and biometric authenticators as a decisive break from the vulnerabilities of legacy credentials.²

When a user loses their mobile device, replaces a hardware token, or simply forgets a PIN, the system must provide a mechanism to restore access. That mechanism is overwhelmingly a secure link or one-time passcode routed directly to the user's registered email address.³ Through this structural reliance, the inbox functions operationally as a master credential.

An adversary who compromises an email account can often bypass downstream authentication controls by exploiting recovery workflows. While some highly mature platforms utilize device-bound hardware constraints, the vast majority of the digital economy still treats the inbox as the ultimate root of trust. Attackers do not necessarily need to clone a fingerprint or steal a physical security key; for most systems, they simply initiate a recovery request. Once inside the inbox, the attacker intercepts the multi-factor authentication codes, monitors financial correspondence to understand the victim's habits, and executes password resets across disparate systems like banking, healthcare, workplace software; systems that were never designed to share a security boundary. The scale of this exposure is massive; in a single recent year, infostealer malware compromised over 183 million accounts, demonstrating that the inbox is frequently the most fragile link in the trust chain.⁴

The Failure of Transactional Evaluation

Even as the inbox has become the continuity layer for identity, defenders continue to evaluate identity transactionally while attackers build identity longitudinally.

When a user registers for a new service, the platform's risk engine typically executes a point-in-time snapshot. It queries whether the email syntax is valid, whether the domain resolves to an active mail server, and whether the address originates from a known disposable provider. If the address clears these rudimentary hurdles, it is generally granted a baseline of trust that serves as the foundation for the ongoing customer relationship.

This transactional evaluation is structurally blind to the lifecycle of malicious assets. Identity is no longer created at onboarding. It is cultivated over time, often long before it is observed by the systems that rely on it. Industrialized fraud rings operate with the patience of legitimate consumers. A single onboarding check cannot reliably distinguish between a genuine human, a synthetic identity that has been deliberately aged, and a compromised legitimate inbox carrying years of authentic history. At the exact moment a risk engine parses the registration payload, all three present virtually identical static characteristics.

2. The Mechanics of Cultivation

To understand how attackers bypass transactional defenses, one must examine the concrete behavioral mechanics of how they cultivate identity infrastructure over time. Attackers understand that credibility is accumulated, not declared. They simulate that accumulation through several specific operational patterns.

Inbox Warm-Up and Aging

Newly minted email addresses are highly scrutinized by anti-spam and fraud filters. To launder a synthetic identity, operators deploy automated scripts to simulate normal human administration. A newly created address is systematically subscribed to dozens of benign, mainstream newsletters and retail promotions. The automated scripts may occasionally

open messages or click links to generate engagement metrics. The address then sits largely dormant for six to eight months. By the time it is directed at a target financial institution, the address possesses a deep history of incoming B2C communications and a mature creation date, allowing it to seamlessly bypass basic velocity and age-based filters.

Alias Fanout and Routing Exploitation

Adversaries exploit the technical disconnect between how email providers route messages and how platforms parse identities. While SMTP standards dictate that the local part of an email address should be case-sensitive, every major email provider silently enforces case-insensitive handling.⁵ Furthermore, providers offer aliasing features like appending a plus sign or inserting dots, that route infinite variations of an address to a single underlying inbox.

Because online platforms typically treat these aliases as distinct users, attackers can execute massive registration fanouts without the overhead of managing thousands of actual inboxes. In one extensive campaign targeting the NPM software registry, a single base Gmail address was used to spawn 139 unique alias accounts.⁶ These aliases, which utilized a combination of plus-suffixing and case variation, bypassed registration limits and subsequently published nearly 4,000 spam packages.⁷ The attackers successfully infinite-scaled their access because the defending platform evaluated the identifier structurally rather than resolving its routing behavior.

Dormant-to-Active Transitions and Compromised Blending

The most difficult threat to detect is the weaponization of accumulated trust. When attackers compromise a legitimate, ten-year-old inbox via credential stuffing or infostealers, they inherit a pristine behavioral history. A sophisticated operator will not immediately trigger fraudulent transactions. They will observe the account, studying the timing of the user's interactions and the nature of their regular purchases. When the attacker eventually initiates a fraudulent password reset or intercepts an authorization code, the anomalous activity is masked by the overwhelming volume of historically legitimate traffic. To a static risk engine, the compromised identity appears unimpeachable precisely because its history is authentic.

This dependency becomes most visible in how modern systems recover identity. See Figure 1.

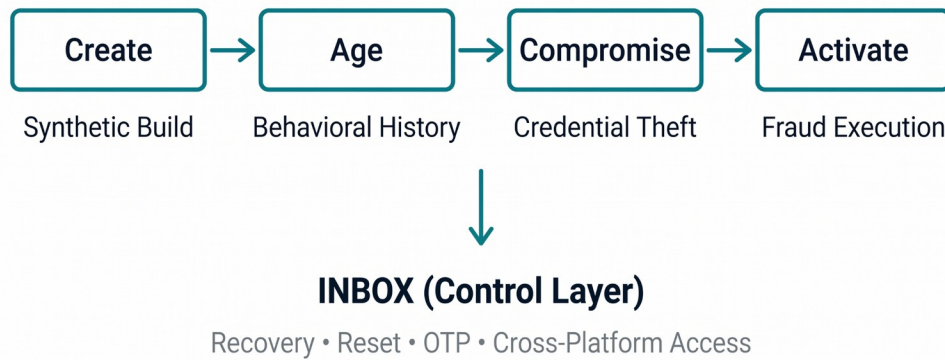


Fig 1. Identity Lifecycle and Control Model

3. Realigning the Defense

Defending modern digital infrastructure requires a shift toward observing behavioral accumulation. Organizations must deploy continuous risk assessment that tracks how identities behave across their lifecycle. Does the address exhibit the organic, irregular cross-platform reuse typical of a human being, or does it display the tight, clustered registration timing of an automated paper mill? Does the routing logic of the address point to a known alias network?

By analyzing the operational mechanics of the continuity layer, defenders can intercept threats during the infrastructure preparation phase. If the industry continues to evaluate trust one transaction at a time, it will remain permanently outmaneuvered by adversaries who understand that identity is something you cultivate, not something you create.

Operationalizing Longitudinal Models

Executing this shift at scale requires replacing static rule engines with models explicitly designed to evaluate longitudinal feature sets. Instead of parsing isolated payloads, these

models ingest streams of historical metadata to map an identity's behavioral baseline against known adversarial typologies. This allows systems to distinguish between organic behavioral variance and coordinated identity construction.

In practice, a predictive model evaluates the delta between inbox age and observable activity. An address created two years ago but demonstrating zero inbound B2C communication history before a sudden burst of financial platform registrations represents a high-confidence anomaly. Similarly, the models track alias expansion patterns, mapping seemingly distinct identities—such as those utilizing sub-addressing or domain manipulation—back to a single, clustered routing node. Credential exposure timing serves as another critical feature; a pristine historical baseline followed immediately by a confirmed appearance in a credential stuffing combo-list dramatically alters the risk score of subsequent authentication attempts.

By scoring these accumulated signals, systems transition away from static evaluations. Low-risk identities that demonstrate stable, organic behavior over time move through systems with minimal resistance. For business leaders, this means predictive identity verification not only isolates synthetic infrastructure, but actively reduces false-positive friction, driving higher conversion rates for legitimate customers. Conversely, identities exhibiting rapid synthetic construction, alias amplification, or compromised blending trigger progressively stronger verification requirements. Trust becomes a dynamic state, continuously evaluated rather than permanently granted.

Email as Identity Infrastructure

This evolution reinforces a broader implication: securing email is no longer a secondary concern. As the inbox continues to function as the continuity layer for digital identity, its security becomes foundational to the integrity of the systems built upon it. The distinction between communication infrastructure and identity infrastructure is no longer meaningful in practice.

Identity systems will not fail because authentication mechanisms are insufficient. They will fail if they continue to evaluate identity only at the moment it appears, rather than understanding how it was constructed.

myNetWatchman provides identity intelligence derived from continuous observation of live attacker activity. Since 2003, the company has monitored attacker-controlled infrastructure to help organizations prevent account takeover, fake account creation, and online fraud before malicious identities reach onboarding, login, or recovery workflows.

References

Lemelson-MIT Program. "Ray Tomlinson." Accessed May 18, 2026.
<https://lemelson.mit.edu/ray-tomlinson>

²⁻³ Authgear. "Passwordless Authentication: Magic Links vs Passkeys vs OTP." February 24, 2026. Authgear Blog

⁴ Don Bush, "The Digital Identity Paradox: Why Email Verification is the New Security Frontier," myNetWatchman Blog, February 10, 2026.

⁵⁻⁷ Wu, Mengying, Geng Hong, Jiatao Chen, Baojun Liu, Mingxuan Liu, and Min Yang. "One Email, Many Faces: A Deep Dive into Identity Confusion in Email Aliases." Proceedings of the Network and Distributed System Security Symposium (NDSS), Internet Society, 2026.