NextGen Scientific Publication
American International Journal of Computer Science and Technology

Volume 3 Issue 2, Pg. No. 53-62, AIJCST-V3I2P106, 2021
https://doi.org/10.63282/3117-5481/AIJCST-V3I2P106

Original Article

Account Recovery as a Security Boundary: A Risk-Adaptive, Layered Architecture for Financial Services

Sruchit Kumar Bommineni[1], Brian Campbell[2]
[1]Manager, Software Development and Engineering, Charles Schwab & Co., Inc., USA.
[2]Distinguished Engineer, Ping Identity, USA.

Abstract:

Financial institutions have invested heavily in hardening primary authentication through multifactor authentication (MFA), risk-based authentication (RBA), device intelligence, behavioral analytics and phishing-resistant cryptographic authenticators. Account recovery frequently operates at a lower level of assurance than the primary authentication it replaces. The gap is structural: recovery must restore access precisely when one or more enrolled authenticators are unavailable, which creates a persistent tension between security and accessibility. Adversaries who cannot defeat primary authentication increasingly target recovery pathways through credential theft, social engineering, SIM swap and number-port-out fraud, email account compromise and automated abuse of recovery endpoints. This paper argues that account recovery should be modeled as an independent authentication and identity-assurance decision rather than as a customer-service or password-reset function. It formalizes the relationship between recovery assurance and effective account assurance, presents a threat model specific to financial-services recovery, and proposes a layered, risk-adaptive recovery architecture in which identity evidence, authentication assurance, contextual risk and authorization scope are evaluated as separate decisions. The architecture incorporates channel independence, graduated post-recovery authorization, protections for authenticator lifecycle changes, post-recovery fraud monitoring, and resistance to enumeration and automated abuse. The paper also proposes operational metrics for evaluating recovery security and usability. The approach is intended to reduce account-takeover (ATO) exposure without imposing maximum friction on legitimate customers.

Keywords:

Identity and Access Management, Account Recovery, Authentication Assurance, Multifactor Authentication, Risk-Based Authentication, Account Takeover, Fraud Prevention, Financial Services Security, Identity Orchestration.

Article History:

Received: 23.01.2021

Revised: 26.02.2021

Accepted: 07.03.2021

Published: 12.03.2021

1. Introduction

Authentication is a foundational control in financial-services applications. Institutions combine knowledge factors, possession factors, device recognition, behavioral analytics, biometrics, passkeys and transaction monitoring to protect customer accounts. Because a recognized, previously verified device commonly substitutes for a fresh MFA challenge, a representative primary authentication decision may be expressed as:

$$\text{Access} = \text{Password} \wedge (\text{MFA} \vee \text{TrustedDevice}) \wedge (R \leq \theta) \quad (1)$$

Where R is a contextual risk score and θ is a policy threshold. When a customer forgets a password, replaces a device, changes a phone number or is locked out, the institution must provide an alternative means of restoring access. Recovery is therefore unavoidable, and its design determines whether the institution's primary controls hold in practice.

Copyright by the Author(s). This work is licensed under a Creative Commons Attribution-ShareAlike 4.0 International (CC BY-NC-ND 4.0) License (https://creativecommons.org/licenses/by-sa/4.0/)

The central problem is that *a recovery process must be accessible enough for legitimate customers who have lost an authenticator, yet resistant enough to adversaries who deliberately seek a path around primary authentication.* If recovery provides materially lower assurance than primary authentication, an adversary does not need to defeat the institution's strongest controls; it needs only to defeat the recovery process.

This paper draws on the author's 2020–2022 implementation work on authentication and access-recovery services at a U.S. brokerage and wealth-management firm. That work included integrating device recognition directly into forgot-password and forgot-login-ID flows, so that recovery decisions could rely on previously established device trust, and onboarding internal business lines to forced password and login-ID change policies following recovery events. The paper focuses on the architecture of digital recovery: risk-adaptive decisioning, fallback assurance, abuse resistance and post-recovery authorization. It generalizes from that work and reports no proprietary data.

This paper makes five contributions:

- It frames account recovery as an authentication decision and formalizes the effect of recovery assurance on effective account assurance (Sections 2 and 4).
- It presents a threat model of recovery-targeted attacks relevant to financial services (Section 3).
- It proposes a layered, risk-adaptive recovery architecture that separates identity evidence, authentication assurance, contextual risk and authorization (Sections 6 through 12).
- It specifies controls for abuse resistance, information disclosure and notification (Sections 13 and 14).
- It proposes measurable indicators for evaluating recovery security and usability (Section 17).

2. Background and Terminology

Account recovery refers to any process that restores or re-establishes a subject's ability to authenticate when one or more previously bound authenticators are unavailable. This includes password reset, login-identifier recovery, account unlock, authenticator replacement, contact-channel changes and credential reactivation.

Identity assurance refers to confidence that the claimant is the person to whom the account belongs. Authentication assurance refers to confidence that the claimant controls authenticators bound to that account. NIST SP 800-63B-3 expresses these concepts as Identity Assurance Levels (IAL) and Authentication Assurance Levels (AAL) [1]. The fourth revision of SP 800-63B treats account recovery as a defined part of authenticator management: it specifies recovery methods (saved and issued recovery codes, recovery contacts and repeated identity proofing [2]), sets the combinations of those methods required at each assurance level, and requires that every recovery event generate a notification to the subscriber [3]. The FIDO Alliance reaches the same position from the relying-party side, describing recovery methods as forms of authentication that must be counted when judging the strength of an account [4].

Effective assurance. An account can be reached through a set of access paths P, which includes primary authentication and every supported recovery path. Because an adversary will select the weakest available path, the effective assurance of the account is bounded by the least-assured path:

$$A_{eff}(\text{account}) = \min \{ A(p) : p \in P \} \qquad (2)$$

Equation (2) is the formal basis for the weakest-link problem discussed in Section 4. Adding controls to primary authentication does not raise A_{eff} if any recovery path remains at a lower level.

3. Threat Model

The adversary's objective is persistent, transactable access to a customer's financial account. The adversary may possess one or more of the following capabilities: knowledge of the victim's personal data from breaches or public sources; valid or partially valid credentials obtained through phishing or credential stuffing; control of the victim's phone number through SIM swap or port-out fraud [10]; control of the victim's email account; the ability to socially engineer contact-center staff; and access to automation infrastructure distributed across residential proxies and device farms. Table 1 maps these capabilities to the recovery paths they target.

Table 1. Recovery-Targeted Attack Vectors and Primary Mitigations

Attack Vector	Recovery Path Targeted	Primary Mitigating Controls
SIM swap / port-out	SMS OTP reset; phone-number change	Carrier SIM-change and port-out signals [11]; channel independence; cooling-off periods
Email account compromise	Email reset links	Existing-trust factors; device binding; notification to all channels
Social engineering of agents	Assisted recovery	Scripted verification; agent-side risk signals; restricted agent privileges
Knowledge-based answer harvesting	KBA fallback	Deprecating static KBA; multi-signal corroboration
Automated recovery abuse	Reset and lookup endpoints	Velocity controls; bot detection; uniform responses
Account enumeration	Lookup and reset endpoints	Non-disclosing responses; timing uniformity
Authenticator replacement	MFA reset and enrollment	Elevated assurance; delayed activation; old-authenticator notification

Out of scope are full compromise of the institution's infrastructure and fully compromised customer endpoints with persistent malware, although post-recovery monitoring (Section 12) provides partial mitigation for the latter.

4. Recovery as an Authentication Decision

Recovery is commonly implemented as a supporting feature of login. From a security perspective, each recovery operation is a state transition on the account's identity and authentication configuration, and transitions differ in impact (Table 2).

Table 2. Recovery Operations, the State Change Each Produces, and Impact Tier

Recovery Operation	State Change	Impact Tier
Password reset	New knowledge credential	High
Login-ID recovery	Disclosure of identifier	Moderate
Account unlock	Restores authentication attempts	Moderate
MFA device replacement	New possession factor bound; old factor possibly removed	Critical
Phone-number or email change	Recovery and notification channel altered	Critical
New trusted-device registration	New device trust relationship	High

These operations sit within an interdependent lifecycle that runs from identity proofing through authentication, recovery, credential enrollment and device trust to authorization. A weakness at any stage can nullify the protections established at another. Figure 1 shows the resulting problem: the recovery path functions as a second authentication system with a lower assurance requirement, and by Equation (2) it sets the account's effective assurance.

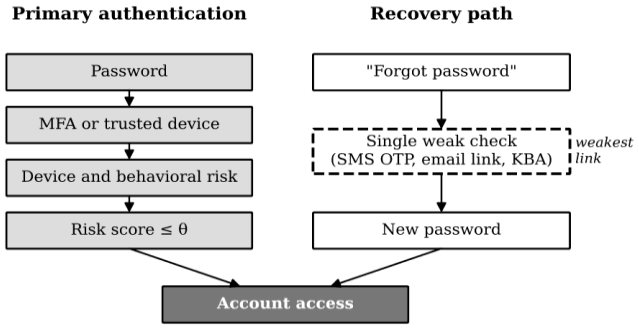


Figure 1. Primary Authentication and A Typical Recovery Path Converge on the Same Account Access. A Single Weak Recovery Check (Dashed) Bounds the Effective Assurance of the Account

The architecture proposed here treats recovery as a distinct authentication decision, with explicitly defined assurance requirements tied to the operation being performed.

5. Why Financial Services Present a Distinct Recovery Challenge

Account recovery matters for all digital services, but financial services raise the stakes in two ways. First, the consequences of compromise are direct and often irreversible: unauthorized transfers, payment redirection, and beneficiary modification, exposure of financial data, unauthorized profile changes, persistent adversary access and use of the compromised account to defraud third parties. Second, the rate of legitimate recovery is high. Customers lose authenticators for ordinary reasons such as device loss or replacement, phone-number changes, loss of email access, hardware-token replacement and travel. An architecture that rejects these customers imposes significant cost and pushes them toward less secure assisted channels. Institutions must therefore prevent unauthorized recovery and enable efficient legitimate recovery at the same time. The design goal is *risk-adaptive* recovery, in which required assurance scales with contextual risk and with the impact of the requested operation, rather than uniformly stricter recovery. Regulatory guidance reinforces the expectation that controls be commensurate with risk, including the FFIEC's 2021 authentication guidance [5] and the strong customer authentication requirements under the EU's revised Payment Services Directive [6].

6. Layered Recovery Architecture

The proposed architecture decomposes recovery into discrete stages, each with its own control objective (Figure 2). Rather than asking only whether a customer may reset a password, the system resolves seven questions:

- Which account is the subject of the request?
- What evidence supports the claimant's association with that account?
- Which previously bound authenticators remain available?
- What contextual risk accompanies the request?
- What state transition is being requested, and what is its impact tier?
- What authorization scope should be restored?
- What controls should apply during the post-recovery period?

Separating these questions prevents a single successful check from implicitly authorizing every downstream action. The final stage feeds confirmed outcomes (fraud labels, customer reports, false rejects) back into risk models and policy thresholds, so the architecture improves over time without requiring every customer to experience maximum friction.

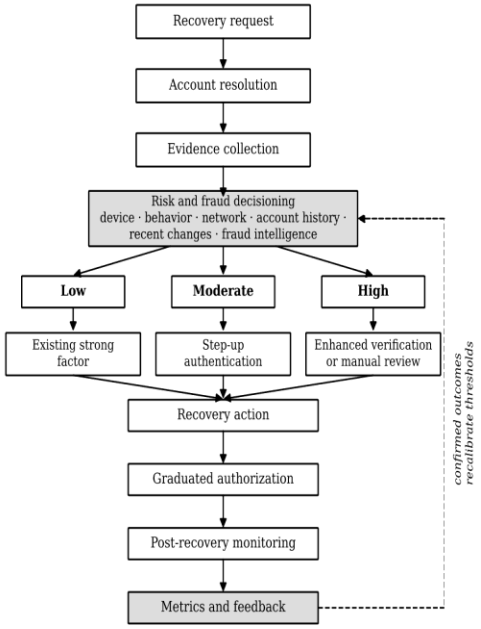


Figure 2. Layered, Risk-Adaptive Recovery Architecture. Contextual Risk Selects the Verification Path; Confirmed Outcomes Recalibrate the Risk Thresholds (Dashed)

7. Recovery Authentication Patterns

7.1. Preserve Existing Trust

Recovery should prefer authenticators and trust relationships established before the recovery event: synced or device-bound passkeys, registered platform or roaming authenticators, hardware security keys, previously trusted devices, saved recovery codes and existing authenticated sessions. Evidence bound before the event is more resistant to manipulation than evidence supplied during it. A password reset initiated from a device with a long, consistent history presents a different risk profile from the same request made from a previously unseen device. The FIDO Alliance recommends encouraging customers to register multiple authenticators so that recovery can rely on an existing bound credential [4], and NIST SP 800-63B-4 sets minimum requirements for saved recovery codes, including entropy, one-way storage, throttling and single use [3].

7.2. Combine Multiple Independent Signals

Recovery decisions should not depend on a single attribute. Relevant signals include possession of an existing authenticator, device identity and reputation, account tenure and history, behavioral biometrics, network characteristics, geographic consistency, historical authentication patterns, recent changes to security attributes, and consortium or internal fraud intelligence. No single signal should be dispositive. Signals are combined into a composite assessment, with care taken that correlated signals (for example, IP geolocation and network ASN) are not double-counted as independent evidence.

7.3. Enforce Channel Independence

A channel should not be used to authorize a change to itself. A one-time code sent to a phone number should not be sufficient to change that phone number, and a reset link sent to an email address should not be sufficient to change that address. Where a channel change is requested, verification should come from an independent bound authenticator, and notification should be sent to the channel being replaced.

8. Risk-Based Recovery Decisioning

The same stated intent, "I forgot my password," can carry very different risk. Table 3 contrasts two such requests.

Table 3. Two Password-Reset Requests with the Same Intent and Different Risk

Attribute	Request A	Request B
Device	Established, recognized	Previously unseen
Network	Consistent with history	Anonymizing proxy or unusual ASN
Prior attempts	None	Multiple failed recovery attempts
Recent changes	None	Phone number changed within 48 hours
Geography	Consistent	Inconsistent with recent activity
Existing authenticator	Available	Unavailable

A recovery risk engine aggregates these dimensions:

$$R = f(R_{identity}, R_{device}, R_{behavior}, R_{network}, R_{change}, R_{fraud}) \quad (3)$$

Required assurance is then a function of both the risk score and the impact tier of the requested operation (Table 2). Listing 1 gives a simplified decision procedure, and Table 4 summarizes representative responses by risk level.

Listing 1. Simplified risk-adaptive recovery decision.

```
function decide_recovery(request):
    ctx      = collect_context(request)
    evidence = evaluate_evidence(request.factors)
    risk     = score(ctx)
    required = policy.required_assurance(request.action, risk)

    if evidence.assurance >= required:
        grant(request.action, scope = policy.scope(risk))
```

```
    start_monitoring(request.account, policy.window(risk))
    notify_all_channels(request.account, request.action)
  else if policy.step_up_available(evidence, required):
    challenge(policy.next_factor(evidence, required))
  else:
    route_to_enhanced_verification(request)
```

Table 4. Representative recovery responses by risk level.

Risk Level	Representative Response
Low	Recovery via existing strong factor with streamlined flow
Moderate	Step-up authentication or corroborating evidence
High	Enhanced identity verification, restricted scope, delayed activation or manual review

Thresholds should be calibrated to the institution's risk appetite, regulatory obligations, observed threat activity and account characteristics, and validated against labeled fraud outcomes.

9. Fallback Assurance Hierarchy

Fallback mechanisms are necessary but become the weakest factor when treated as equivalent to strong authentication. Each fallback should carry an explicitly assigned assurance level, and each level should map to a defined set of permitted operations. Figure 3 gives ordering.

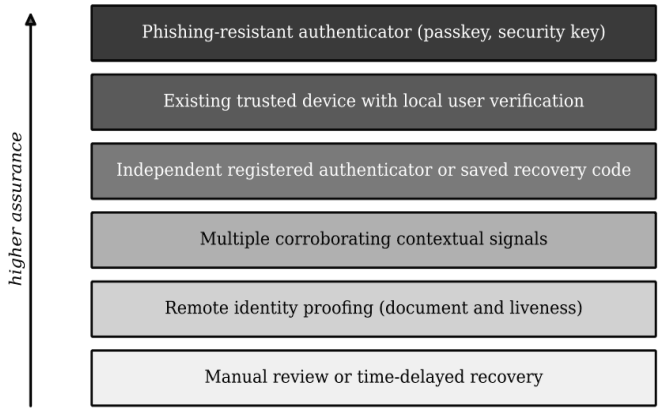


Figure 3. Fallback Assurance Hierarchy, from Highest (Top) to Lowest (Bottom) Assurance. Exact Ordering Depends on Implementation Quality

NIST SP 800-63B classifies out-of-band authentication over the public switched telephone network (SMS or voice OTP) as a restricted authenticator because of SIM swap and interception risk [3], and CISA identifies phishing-resistant authenticators as the target state for high-value accounts [9]. Static knowledge-based questions provide weak assurance given the volume of breached personal data. Institutions that rely on these mechanisms should compensate with additional signals and restricted post-recovery scope.

10. Protecting Authenticator Lifecycle Changes

Password reset and MFA replacement are not equivalent operations. Consider an adversary who controls a victim's phone number or email account and can:

- initiate recovery,
- reset the password,
- remove the existing MFA authenticator, and
- Bind a new authenticator under the adversary's control.

At step 4 the adversary achieves persistence, and the legitimate customer may lose the ability to recover the account. MFA replacement should therefore be treated as a critical state transition, with the following controls:

- Higher required assurance than password reset, preferably verified through an existing bound authenticator or trusted device.
- Out-of-band notification to the authenticator and channels being replaced, with a clear revocation path.
- Delayed activation (cooling-off period) during which the new authenticator cannot authorize high-risk operations and the prior authenticator can cancel the change.
- Session and token revocation for all existing sessions and refresh tokens once the change is committed.
- Enhanced fraud review when replacement coincides with other high-risk signals.

These controls matter most when the existing authenticator is being *removed* rather than supplemented.

11. Graduated Authorization after Recovery

Successful recovery need not restore every capability immediately. Authentication establishes a level of confidence; authorization determines what that confidence permits. Treating these as separate decisions allows institutions to restore access quickly while limiting the damage of an erroneous recovery (Table 5).

Table 5. Post-recovery treatment of account capabilities

Capability	Post-Recovery Treatment
View balances and transactions	Permitted immediately
Update password	Permitted immediately
Register additional authenticator	Permitted with notification
Change contact channels	Step-up required; cooling-off period
Add beneficiary or external account	Step-up required; holding period
High-value or first-time-payee transfer	Step-up required; holding period and enhanced monitoring

Holding periods and step-up requirements can be shortened when recovery was completed with high-assurance factors and low contextual risk.

12. Post-Recovery Monitoring and Fraud Integration

A recovery event is itself a risk indicator. Following recovery, the account can enter a time-bounded heightened-monitoring state in which new-device registration, authenticator replacement, contact-channel changes, beneficiary or external-account additions, anomalous transaction behavior and rapid sequential changes to multiple security attributes receive elevated scrutiny. A sensitive change detected in this window triggers transaction risk evaluation and, where warranted, a step-up challenge or a hold. Monitoring provides a second line of defense if the recovery decision itself was wrong.

Effective monitoring requires integration between identity and fraud systems, which many institutions operate independently. Recovery events should be published as features to fraud models, and fraud-model outputs should be available to authentication and recovery policy. Section 16 places this integration within a centralized orchestration layer (Figure 4).

13. Abuse Resistance: Automation and Enumeration

13.1. Automated Recovery Abuse

Recovery endpoints are attractive targets for automation. Controls should include per-account and per-identifier velocity limits, bot and automation detection, device reputation, network risk scoring, failed-attempt thresholds, behavioral analysis and correlation of recovery attempts across accounts. Controls should not rely solely on IP address, because adversaries distribute traffic across residential proxies, mobile networks and device farms. Correlation on device, behavioral and request-pattern features is more robust.

Recovery artifacts should meet baseline properties consistent with OWASP guidance [7], [8]: tokens and codes generated from a cryptographically secure random source, single-use, short-lived, stored only as hashes, bound to the initiating session or device where feasible, and invalidated on use or on any subsequent recovery request. Password reset should revoke existing sessions.

13.2. Enumeration and Information Disclosure

A message such as "The email address you entered is not registered" confirms whether a person has a relationship with the institution. In financial services this is sensitive information and aids targeted attacks. Recovery flows should return responses that are uniform in content, status code and timing regardless of whether an account exists, for example: "If an account matches the information provided, we have sent instructions." The same analysis applies to login-ID recovery, authenticator enrollment, account lookup and registration. Usability and disclosure analysis should be performed together, since non-disclosing responses must still guide legitimate customers.

14. Independent Customer Notifications

Security-relevant recovery operations should trigger notifications through all registered channels, including the channel being replaced; NIST SP 800-63B-4 makes notification of every account-recovery event a requirement [3]. Notifiable events include password reset, authenticator replacement, new-device registration and changes to phone number, email address or security settings. Notifications should identify the event, its approximate time and context, and a direct, verifiable way to report unauthorized activity, without links or content that could be imitated in phishing. Customer reports of unrecognized events are high-quality fraud labels, so notifications serve fraud detection as well as customer awareness.

15. Usability as a Security Control

In recovery, usability and security are closely coupled. Excessively complex recovery drives customers to retry repeatedly, seek help through less secure channels, share one-time codes, disable security features, reuse credentials or follow fraudulent "support" instructions. A well-designed recovery experience therefore reduces risk. It should explain why additional verification is required without exposing detection logic; follow a predictable sequence (identify account, verify identity, restore access, secure account, resume normal use); offer legitimate alternatives when one factor is lost; and behave consistently across applications and channels, including assisted channels.

16. Centralized Identity Orchestration

Large institutions operate many applications and business lines. Implementing recovery independently in each creates inconsistent policy and multiplies attack surface. A centralized orchestration layer provides common recovery capabilities alongside login, single sign-on and enrollment, and connects them to shared authentication, risk and fraud services (Figure 4).

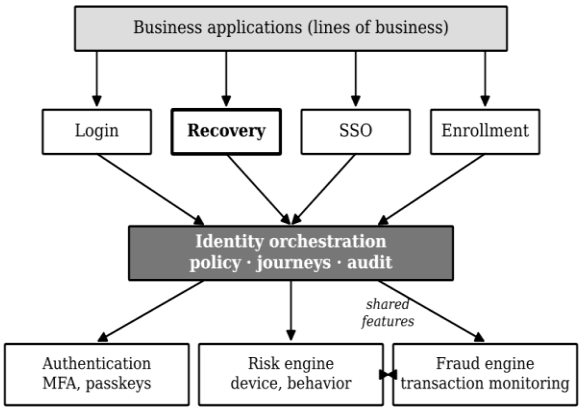


Figure 4. Recovery as One Journey within a Centralized Identity Orchestration Layer Shared across Business Lines, with Risk and Fraud Engines Exchanging Features

Benefits include consistent policy enforcement, reusable controls, standardized journeys, centralized risk evaluation, unified audit trails, faster application onboarding and stronger governance. Centralization should still expose policy hooks so that business lines can apply transaction-specific requirements, since the sensitivity of a brokerage account, a commercial treasury platform and a retail card account differs.

17. Evaluation Metrics

Recovery programs should be measured on both security and usability. Suggested indicators:

- Recovery-induced ATO rate: confirmed takeovers in which a recovery event preceded the fraud, per 10,000 recoveries.
- Recovery false-reject rate: legitimate recovery attempts that failed or were abandoned, estimated through follow-up or assisted-channel escalation.
- Recovery completion rate and time to recover, segmented by risk tier.
- Step-up and manual-review rates by risk tier, used to validate threshold calibration.
- Assisted-channel escalation rate, as a proxy for digital recovery friction.
- Post-recovery intervention rate: holds or step-ups triggered during the monitoring window and their confirmed-fraud precision.
- Time to customer report following unauthorized-change notifications.

Tracking these together prevents optimizing one dimension, such as fraud loss, at the unmeasured expense of another, such as customer lockout.

18. Design Principles

- Treat recovery as authentication. Recovery paths require assurance commensurate with the access they restore.
- Preserve existing trust. Prefer authenticators and devices bound before the recovery event.
- Avoid single-factor recovery. A single weak factor should not restore unrestricted access to a high-value account.
- Make recovery risk-adaptive. Required assurance should scale with contextual risk and operation impact.
- Enforce channel independence. No channel should authorize changes to itself.
- Protect authenticator lifecycle changes. Replacing or removing MFA warrants heightened assurance, notification and delayed activation.
- Separate authentication from authorization. Recovery confidence should govern the scope of restored capability.
- Monitor after recovery. The post-recovery window warrants elevated scrutiny.
- Integrate identity and fraud controls. Each system should inform the other.
- Minimize information disclosure. Responses should not reveal account existence.
- Make security usable. Legitimate customers need clear, practical recovery paths.
- Measure outcomes. Evaluate recovery on both fraud and customer-access metrics.

19. Limitations and Future Work

The architecture depends on the quality of available risk signals; institutions with limited device or behavioral telemetry will have less ability to differentiate risk. Holding periods and cooling-off windows introduce friction that must be balanced against customer expectations, particularly for time-sensitive payments. Assisted recovery through contact centers remains difficult to bring under the same policy engine and warrants further study. The paper proposes metrics but does not report measured outcomes. Future work includes empirical evaluation of threshold calibration, the effect of synced passkeys on recovery demand, and privacy-preserving sharing of recovery-abuse signals across institutions.

20. Conclusion

Account recovery operates under a distinctive condition: the customer seeks access precisely when established authenticators are unavailable. This makes recovery both a usability problem and a security boundary. Recovery mechanisms weaker than primary authentication create an alternative path to account takeover, while uniformly high-friction recovery burdens legitimate customers and pushes them toward less secure channels.

A risk-adaptive, layered architecture addresses both concerns by combining previously established trust, multiple independent signals, channel independence, explicitly graded fallback assurance, contextual risk and fraud intelligence, abuse resistance, graduated authorization, post-recovery monitoring and a clear customer experience. By Equation (2), an account's authentication is only as strong as the weakest mechanism through which an adversary can regain access to it. For financial institutions, strengthening account recovery is therefore a direct way to raise the effective assurance of every customer account.

References

[1] National Institute of Standards and Technology, *Digital Identity Guidelines*, NIST Special Publication 800-63B-3.
[2] National Institute of Standards and Technology, *Digital Identity Guidelines: Identity Proofing and Enrollment*, NIST SP 800-63A-4.
[3] National Institute of Standards and Technology, *Digital Identity Guidelines: Authentication and Authenticator Management*, NIST SP 800-63B-4, sections on account recovery and restricted authenticators.
[4] FIDO Alliance, *Recommended Account Recovery Practices for FIDO Relying Parties*, white paper, Feb. 2019.
[5] Federal Financial Institutions Examination Council, *Authentication and Access to Financial Institution Services and Systems*, Aug. 2021.
[6] European Commission, Commission Delegated Regulation (EU) 2018/389 of 27 November 2017 supplementing Directive (EU) 2015/2366 with regard to regulatory technical standards for strong customer authentication and common and secure open standards of communication, *Official Journal of the European Union*, L 69, 13 Mar. 2018.
[7] OWASP Foundation, *Forgot Password Cheat Sheet*, OWASP Cheat Sheet Series.
[8] OWASP Foundation, *Application Security Verification Standard*, version 4.0.0, 2019.
[9] Cybersecurity and Infrastructure Security Agency, *Implementing Phishing-Resistant MFA*, fact sheet, Oct. 2021.
[10] Jordaan, L., & Von Solms, B. (2011). A biometrics-based solution to combat SIM swap fraud. In Open Research Problems in Network Security: IFIP WG 11.4 International Workshop, iNetSec 2010, Revised Selected Papers (pp. 70–87). Springer. https://doi.org/10.1007/978-3-642-19228-9_7
[11] Federal Communications Commission, *Protecting Consumers from SIM-Swap and Port-Out Fraud*, Report and Order, Sept 2021 NPRM (FCC 21-102).