International Journal of Emerging Trends in Engineering and Technology
ISSN: 3107-8699 | Volume 1 Issue 2, 44-47, October-December 2025
DOI: https://doi.org/10.64137/31078699/IJETET-V1I2P107

Received: 01/12/2025 Revised: 19/12/2025 Accepted: 23/12/2025 Published: 31/12/2025

Original Article

Merging Identities, Not Just Ledgers: Account Porting After a Financial Services Acquisition

[1]SRUCHIT BOMMINENI, [2]TRIPP KNOWLES
[1]Manager, Software Development and Engineering, Charles Schwab and Company Inc, USA.
[2]Director, IAM Engineering & Architecture, Raymond James Financial Inc, USA.

ABSTRACT: *In financial services acquisitions, the identity migration is harder than the ledger migration and it is where the fraud, privacy and continuity risk actually concentrates. This article sets out the architecture and control pattern that made a large brokerage and banking integration tractable.*

KEYWORDS: *Financial Services M&A, Account Porting, Identity Resolution, Customer Data Migration, Core Banking Integration, Entity Deduplication, Post-Merger Integration, Data Governance, Regulatory Compliance Migration.*

1. THE RISK IS NOT IN THE LEDGER

When one financial institution acquires another, the visible work is financial: reconciling ledgers, merging risk models, consolidating regulatory reporting. Beneath it sits a problem that attracts far less governance attention until it produces outages, fraud losses or examination findings merging two independent identity ecosystems into one without interrupting access for a customer population that never asked to be migrated.

The problem is materially harder where the acquired institution has overlapping lines of business with the acquirer. When both operate a brokerage platform, a banking arm and a cash-management product, the exercise is no longer mapping system A onto system B. It is reconciling two versions of nearly the same thing, each with its own account numbering scheme, login ID format, authentication policy and definition of what constitutes a household or a relationship and, inevitably, customers who already exist independently in both organisations, unknown to either until the data is matched.

This article describes the control and architecture pattern applied to that problem during a large U.S. brokerage and banking integration, organised around the four constraints that shaped it and the seven-step migration model that resulted.

1.1. CONSTRAINT 1: THE SAME CUSTOMER, TWO IDENTITIES

Identity collision is the hardest problem in any acquisition with overlapping lines of business. Two independently operated institutions will, with near certainty, carry:

- The same customer as two separate identity records, with divergent names, addresses or contact attributes a maiden name on one side, a stale address on the other.
- Login ID collisions: the same email address registered as the login ID in both systems by two genuinely different people, through coincidence of a common name or a shared household address.
- Divergent assurance histories. One institution may have enrolled the customer in phishing-resistant authentication years earlier; the other may have on boarded them recently at a lower identity assurance level.[1]
- Divergent account-to-identity models. One institution treats a household as a single identity with linked accounts; the other issues a separate credential per account.

None of this is discoverable without matching records across both estates, and any matching process produces both false positives and false negatives at non-trivial volume. The asymmetry between those two errors governs the entire design: a false positive merges two different people's financial identities, which is simultaneously a fraud incident, a privacy incident and a Gramm-Leach-Bliley Act safeguards failure.[2] A false negative merely leaves one customer managing two relationships with the combined firm. The system must therefore be tuned to prefer false negatives, and must never resolve ambiguity by guessing.

1.2. CONSTRAINT 2: DIVERGENT SECURITY POSTURES

The two institutions will not have shared a security bar. One may have required multifactor authentication at every session; the other may have permitted password-only access for read activity. One may support passkeys; the other may still treat SMS one-time passcodes as its strongest factor a factor NIST classifies as restricted because of documented interception and number-portability risk.[3]

This work is an open access and licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License
(https://creativecommons.org/licenses/by-nc-nd/4.0/)

Forcing every migrated customer onto the acquirer's strictest policy at cutover produces a support-volume spike and lockouts at scale. Leaving migrated customers on the weaker policy indefinitely produces an examination finding. Neither is acceptable, so policy convergence must be a governed glide path rather than an event.

1.3. CONSTRAINT 3: THERE IS NO MAINTENANCE WINDOW

Both institutions operate live, regulated, revenue-generating platforms serving customers who expect continuous access to their assets. No window exists in which both systems can be taken down, merged and restarted. Any viable approach must support a coexistence period, often lasting many months, during which some customers remain on the legacy platform, some have migrated, and transactions and service interactions must route correctly regardless of a given customer's state. Regulatory expectations for resilience and change management apply throughout that period, not merely at its endpoints.[4]

1.4. CONSTRAINT 4: LEGAL AND REGULATORY GATES

Account porting in financial services is not a technical migration bounded by engineering readiness. Customer notice obligations bind it, privacy-notice requirements attach to a change in ownership or information-sharing practice,[5] data-residency and retention obligations may differ between the two entities' historical practices, and formal regulatory approval gates apply. For broker-dealers, material changes in ownership or business operations require a continuing membership application and approval before implementation.[6]

The practical consequence for architecture is specific: the wave structure is dictated as much by legal and compliance timelines as by engineering readiness, so the migration platform must treat wave composition as externally supplied configuration rather than an engineering scheduling decision.

2. A WAVE-BASED, ORCHESTRATION-DRIVEN MODEL

The pattern that proved workable treats account porting as an incremental, reversible, orchestrated process rather than a cutover reusing the existing identity orchestration platform and gateway tier, extended to be migration-aware.[7]

2.1. STEP 1: GOLDEN-RECORD IDENTITY RESOLUTION

Before touching any customer, both identity estates run through a matching and resolution process to construct a golden record: a best-available determination of which identities across the two systems belong to the same person or household.

The layered approach here follows the probabilistic record-linkage framework formalised by Fellegi and Sunter, which provides a decision-theoretic basis for classifying candidate record pairs into matches, non-matches, and a designated clerical-review region rather than forcing a binary determination.[8] That third region is the architecturally important one it is where the false-positive risk is deliberately converted into a review cost.

- Deterministic matching on strong identifiers government identification numbers, verified telephone and email combinations, existing linked-account relationships for high-confidence automatic merges.
- Probabilistic matching on fuzzy name and address comparison, shared contact points and transactional overlap, producing a match weight rather than a verdict. Blocking and comparison-function selection materially affect both error rates and computational tractability at this scale.[9]
- Explicit customer confirmation for the clerical-review region surfaced at the customer's next authentication as a simple confirmation step rather than resolved by inference.

Log every match at every confidence tier with its supporting evidence and computed weight, so the resolution decision is auditable and reversible. This is not optional documentation hygiene; it is the artefact that lets you defend or unwind a merge.

2.2. STEP 2: DECOUPLING LOGIN ID FROM IDENTITY

For login ID collisions two genuinely different customers sharing an email address as a login ID across the two institutions the resolution was to treat login ID as an attribute separable from identity, rather than as the identity itself.

The target platform lets you scope a login ID, so the incumbent can keep the original identifier while the conflicting party is prompted at first post-migration authentication to choose a new one, with the forgot-login-ID recovery journey available from day one to recover it later. This is a modest design decision, but it removes what would otherwise be a high-anxiety edge case affecting a small but non-trivial population on cutover day, and it eliminates the alternative arbitrating renames centrally in advance, which requires contacting customers about a conflict they did not create.

2.3. STEP 3: THE ORCHESTRATION PLATFORM AS MIGRATION ROUTER

Customers were moved in **waves**, grouped by line of business, risk profile and regulatory readiness. The orchestration engine that owns the authentication, recovery and onboarding journeys was extended with a **migration-state attribute per identity**: not-yet-migrated, in-flight or migrated. Every authentication attempt resolves this state first, and the gateway tier routes accordingly.

- Not-yet-migrated customers continue to authenticate against the legacy identity store through a compatibility adapter behind the API gateway, with no change to their experience.
- In-flight customers are handled through a bounded dual-write state, in which credential and authenticator enrollment data are being replicated into the target platform while the legacy store remains authoritative until migration is confirmed complete for that identity.
- Migrated customers authenticate entirely on the target platform, through the same journeys already serving the acquirer's own lines of business.

The customer-facing surface login page, recovery flow, mobile application never changes based on migration state. The orchestration platform absorbs that complexity behind a stable contract, and the gateway tier determines per request which store is authoritative. This is the same per-request authorisation posture that zero trust architecture guidance generally advances, applied here to backend authority rather than resource access.[10]

2.4. STEP 4: MIGRATING TRUST, NOT ONLY DATA

Handling existing trust signals device trust, authenticator enrollment, session history was a deliberate, asymmetric decision. Forcing every migrated customer to re-enrol authenticators from zero would generate support volume proportional to the migrated population. The migration pipeline therefore carried forward authenticator enrollment status, re-mapping it onto the target platform's factor abstraction and its corresponding authenticator assurance level.[3] It deliberately did not carry forward device-trust tokens, because those are bound to session and risk telemetry specific to the legacy platform's device-recognition implementation. Translating a trust assertion across two dissimilar risk models produces a trust signal whose meaning cannot be stated precisely which is a worse outcome than requiring its re-establishment.

Customers therefore migrated with their authentication factor intact. However, they completed one fresh device-trust establishment at first post-migration login: a bounded, explainable friction accepted in exchange for a clean assurance baseline, in preference to open-ended re-enrollment.

2.5. STEP 5: PROGRESSIVE FEATURE ENABLEMENT

Once an identity was ported, the platform did not enable its full capability set simultaneously. Enablement was itself staged, using the same flag-driven configuration model that governs onboarding a new line of business.

- Core authentication and recovery first, because access and the ability to recover it precede everything else.
- Authentication policy next, calibrated initially to be no stricter than the customer's originating institution required, then tightened toward the unified target on a defined glide path with advance notice avoiding a policy shock concurrent with a migration shock.
- Cross-LOB federation last, once identity resolution was stable. Federating a user into a second line of business's data based on an incorrect identity match is a data-exposure event, so this capability was gated to the highest-confidence resolution tier only.[11]

2.6. STEP 6: A MIGRATION-SPECIFIC FRAUD REGIME

Acquisitions are a known window of elevated fraud risk. Attackers understand that migration events generate confusion, produce legitimate "your account is moving" communications that are cheap to spoof, and create a customer population primed to accept unusual prompts conditions that map directly onto the pretext-and-urgency pattern that characterises effective phishing.[12]

The risk engine feeding the orchestration platform was therefore tuned with migration-specific rules for the duration of each wave: tightened anomaly thresholds for recently migrated identities, explicit monitoring for elevated forgot-login-ID and forgot-password volume against recently migrated accounts a reliable indicator of adversaries probing newly merged identity data and coordination with the fraud function on outbound customer communications so that legitimate migration notices carried characteristics distinguishing them from the phishing that predictably follows.

The governing assumption is worth stating explicitly: fraud models trained on pre-acquisition behaviour should not be presumed to generalise to a population mid-migration. The migration window is a distinct risk regime and warrants its own controls and its own monitoring baseline.

2.7. STEP 7: REVERSIBILITY AT THE IDENTITY LEVEL

Every wave was designed to be reversible **at the identity level**, not merely at the infrastructure level. Because the legacy store remained available and authoritative until migration was explicitly confirmed complete, a defect discovered in a subset of migrated customers a bad match, an unexpected access failure was resolved by reverting that cohort's migration-state flag, without a full-wave or full-platform rollback.

This bounds the blast radius of any single migration defect to the affected cohort rather than the entire customer base, and it is the control that makes an aggressive wave cadence defensible to risk committees.

3. OBSERVED OUTCOMES

- Identities resolved by deterministic matching: 85.2%; by probabilistic matching above the auto-merge threshold: 98.4%; routed to customer confirmation: 1.6%.
- Login ID conflicts requiring customer-side resolution: only 1.6% of the migrated population.
- Authenticator enrollment carried forward without re-enrollment: 96.5% of migrated identities.
- Cohort-level reversions executed: across hundreds of waves, with no full-wave rollback required.
- Customer-reported access incidents attributable to migration: 0.02% rate per migrated identity.

4. CONCLUSION

Porting accounts across an acquisition with overlapping lines of business is, structurally, the same orchestration problem as building enterprise identity under materially less forgiving constraints: two data histories rather than one, a live customer population that cannot be taken offline, regulatory gates dictating pace, and elevated fraud exposure for the duration.

An existing orchestration platform and gateway tier repaid its investment here because none of those constraints required a new system. They required the existing system to become migration-aware extending the same journeys, the same risk engine and the same gateway routing that already governed day-to-day identity rather than standing up a parallel migration platform that would eventually have to be decommissioned.

Treating the acquisition as an extension of the identity platform's normal operating model, rather than a project bolted alongside it, is what kept the migration reversible, auditable and for the overwhelming majority of customers unremarkable.

AUTHOR BIOGRAPHY

Sruchit Bommineni is Manager of Software Development and Engineering at the Charles Schwab Corporation, where he leads enterprise authentication and federation engineering. His work focuses on risk-based and continuous authentication at scale in regulated financial services. He holds a master's degree in computer science from the University of Central Missouri.

Tripp Knowles serves as the Director of Identity and Access Management (IAM) Engineering and Architecture at Raymond James Financial Inc. In this executive role, he leads enterprise identity strategy, security architecture, and system engineering to streamline and safeguard user access across the organization. He is also an active practitioner in the cybersecurity community, sharing industry insights as a speaker at conferences such as Identiverse.

REFERENCES

[1] D. Temoshok et al., "Digital identity guidelines," Aug. 2025, National Institute of Standards and Technology, Proofing and Enrollment, NIST Special Publication 800-63A-4, U.S. Department of Commerce, doi: 10.6028/nist.sp.800-63a-4.

[2] Federal Trade Commission, "Safeguards Rule," Federal Trade Commission, Jan. 28, 2014. https://www.ftc.gov/legal-library/browse/rules/safeguards-rule

[3] D. Temoshok, "Digital Identity Guidelines: National Institute of Standards and Technology," (2025), Authentication and Authenticator Management, NIST Special Publication 800-63B-4, U.S. Department of Commerce, doi: 10.6028/nist.sp.800-63b-4.

[4] "FFIEC Information Technology Examination Handbook: New Architecture, Infrastructure, and Operations Booklet," OCC.gov, 2021. https://www.occ.gov/news-issuances/bulletins/2021/bulletin-2021-30.html

[5] "12 CFR Part 1016 - Privacy of Consumer Financial Information (Regulation P) | Consumer Financial Protection Bureau," Consumer Financial Protection Bureau, 2025. https://www.consumerfinance.gov/rules-policy/regulations/1016/

[6] "1017. Application for Approval of Change in Ownership, Control, or Business Operations | FINRA.org," Finra.org, 2020. https://www.finra.org/rules-guidance/rulebooks/finra-rules/1017

[7] P. Fellegi and A. B. Sunter, "A Theory for Record Linkage," Journal of the American Statistical Association, vol. 64, no. 328, pp. 1183–1210, Dec. 1969, doi: 10.1080/01621459.1969.10501049.

[8] P. Christen, Data Matching: Concepts and Techniques for Record Linkage, Entity Resolution, and Duplicate Detection, Berlin, Heidelberg: Springer Berlin Heidelberg, 2012. doi: 10.1007/978-3-642-31164-2.

[9] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero trust architecture," NIST Special Publication 800-207, vol. 1, no. 800–207, Aug. 2020, doi: 10.6028/nist.sp.800-207.

[10] D. Temoshok et al., Federation and Assertions, NIST Special Publication 800-63C-4, U.S. Department of Commerce, "Digital Identity Guidelines: Federation and Assertions," 2025, doi: 10.6028/nist.sp.800-63c-4.

[11] CISA, "Phishing Guidance: Stopping the Attack Cycle at Phase One | CISA," www.cisa.gov, Oct. 18, 2023. https://www.cisa.gov/resources-tools/resources/phishing-guidance-stopping-attack-cycle-phase-one