International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(21s) 2026 pp. 1531-1544
https://doi.org/10.70917/ijcisim-2026-5874
Article

Continuity Of Regulatory Reporting Through Bank Resolution, Acquisition, And Entity Restructuring: Data Migration, Control Re-Establishment, And Supervisory Communication

Laxmi Pandrapragada

Independent Researcher,India. Email: pndlaxmi@gmail.com

Abstract: When a bank changes hands or changes shape, its regulatory reporting obligations do not pause. Required submissions continue on their established cadence while data, source systems, legal-entity hierarchies, controls, ownership and approval structures are simultaneously transferred or redesigned. The institution must determine which entity reports each balance, preserve historical comparability, prevent both omission and double counting, re-establish controls that the transition has invalidated, and confirm who is authorized to approve each filing. This paper documents an execution methodology for maintaining that continuity, drawn from thirteen years of regulatory reporting technology practice across three classes of transition: foundational capital-planning implementation, legal-entity creation and branch-to-subsidiary conversion, and post-acquisition integration under concurrent supervisory remediation. Four components are described. Data migration is treated as the preservation of reporting logic rather than the transfer of balances, on the argument that the fragile elements are not the balances but the contextual artifacts surrounding them: historical mappings, manual adjustments, undocumented transformation logic and definitions held locally or only by individuals. Control re-establishment is presented as a fourteen-step sequence in which a control is not restored when it runs again but when its execution can be evidenced and signed for. Legal-entity remapping is analysed through its hardest case, in which one exposure is simultaneously live in a legacy feed and migrated into a consolidated target, where including both feeds double counts and retiring the legacy feed early omits. Supervisory communication is treated as an evidence-production discipline rather than correspondence. The contribution is methodological. The paper reports no measurements, proposes a structured framework and an accompanying analytical model of the dual-feed overlap problem, and states its limitations directly.
Keywords: Regulatory reporting, capital planning, data migration, data lineage, legal-entity restructuring, internal control, supervisory remediation, RegTech

1. INTRODUCTION

Bank supervision has settled on a clear expectation. Under BCBS 239 [9] an institution must be able to aggregate risk data accurately and completely, and report it on time, with governance and architecture sufficient to sustain that under stress. The European Central Bank's 2024 guide on risk data aggregation and risk reporting [10] sharpened this into an expectation of complete and current lineage at the level of individual data attributes, traced from capture through to the final reported figure.

What these instruments describe is a required steady state. They are close to silent on how an institution holds that state while its own structure is being replaced underneath it.

That silence matters because structural change is common and because it attacks precisely the properties the regulation demands. An acquisition expands the reporting population. A subsidiarisation redraws the reporting perimeter. Standing up a new legal entity creates reporting obligations where none existed. In each case source systems are retired, data begins arriving from different platforms, entity hierarchies are redrawn, control owners change roles,


Copyright: © 2026 by the authors

and approval chains that represented the old organisation no longer represent the new one. Every one of those movements is a threat to completeness, accuracy, traceability or the authority to attest.

The obligations, meanwhile, continue. Exposure-level submissions remain monthly. Capital-planning schedules, holding-company reports and Call Report filings remain quarterly [11][12][13]. Public disclosures supported by the same underlying data continue on their own calendar. None of it is suspended because an institution is mid-integration.

The evidence that this is genuinely unsolved is available in the supervisors' own findings. The ECB's thematic review reported that no significant institution in its sample, including global systemically important institutions, had fully implemented the BCBS 239 principles [10], and that was an assessment of institutions under ordinary operating conditions rather than mid-transition. Industry survey evidence points the same way, with a substantial majority of banks still planning rather than operating end-to-end lineage [16].

This paper documents how the problem is actually solved in practice. The material derives from thirteen years in regulatory reporting technology spanning three phases of work, summarised in Table 1, and covers four components: what has to keep filing and why the consequences of failure escalate quickly, how data moves without destroying the reproducibility of prior periods, how controls are re-established and re-attested, and how legal-entity structures are remapped without double counting or omitting exposures.

Table 1. Transition classes and the boundary between owned and supported work.

Period	**Transition class**	**Led or owned**	**Supported, not led**
Phase one	Foundational capital-planning and exposure-level reporting implementation	Business and functional analysis, business-rule definition, source-to-target mapping review, reporting validation and reconciliation, control verification, defect management, user-acceptance testing	Overall capital-planning programme, formal regulatory interpretation, executive attestation, filing authorization, supervisory communication
Phase two	New legal entities, international expansion, branch-to-subsidiary conversion, acquisition integration, platform modernization	Reporting impact assessments, reporting architecture design, data-sourcing strategy, source-to-target mappings, reporting-rule definition, legal-entity and product mappings, validation and reconciliation frameworks, testing strategy, control implementation, implementation coordination	Licensing submissions, legal decisions, accounting-policy determinations, report attestation, regulator-facing commitments
Phase three	Post-acquisition integration with concurrent capital-planning modernization, data-governance remediation and supervisory remediation	Integration strategy, capital-planning modernization, target-state architecture, source-system rationalization, standardization of reporting definitions, data-lineage improvement, reconciliation and validation design, control enhancement, remediation planning and governance	Final regulatory interpretation, filing ownership, management attestation, accounting policy, formal supervisory communication

That boundary is stated explicitly because it is load-bearing for the rest of the paper. The methodology described here is a technology and delivery methodology. Regulatory interpretation, attestation and supervisory commitment sit with authorized Legal, Compliance, Finance, Regulatory Reporting and executive stakeholders, and the framework is designed to produce what those functions need rather than to substitute for them.

Section 2 positions this against the literature. Section 3 sets out the obligation set. Sections 4 through 7 cover data migration, control re-establishment, entity remapping and supervisory communication. Section 8 consolidates the framework, Section 9 states limitations, which are substantial.

All institution names, internal project names, platform identifiers, transaction details and non-public reporting information are omitted or generalised throughout. Regulatory report names are retained because they are public instruments.

2. BACKGROUND AND RELATED WORK

Three literatures touch this problem. None addresses it.

Regulatory and supervisory instruments define the required outcome. BCBS 239 [9] sets the principles, the ECB guide [10] sets the current articulation including attribute-level lineage, and the Basel core principles [15] frame supervisory expectations more broadly. These are normative documents. They say what must be true, not how to keep it true while the institution is being rebuilt.

Technical data-management research solves the component problems in isolation and in other domains. Provenance has a mature survey literature: Herschel et al. [3] organise it around why provenance is captured, in what form and from what source, and Bose and Frew [4] survey lineage retrieval for scientific data processing. Both are grounded in computational reproducibility rather than regulatory attestation, and the difference is not cosmetic. Scientific reproducibility asks whether a result can be recomputed. Regulatory reproducibility asks whether a filed figure can be explained, reconciled and defended years later to a supervisor, by people who were not there.

Record linkage and deduplication [5] supplies exactly the machinery needed to detect one exposure arriving through two feeds, but that literature addresses entity resolution over customer or bibliographic records, not consolidation perimeters with effective-dated legal-entity mappings. Schema and ontology matching [6] frames the crosswalk problem generally and under-treats the effective-dated case, where the correct mapping depends on the reporting date. Data quality assessment [2] and data governance [1] are framework contributions that assume an organisation whose ownership structure holds still.

Organisational research on post-merger integration is the thinnest strand. Chang et al. [7] is among the few peer-reviewed treatments of information systems integration following acquisition, and the recurring observation in this area is that the practically useful knowledge stays inside institutions. RegTech scholarship, surveyed by Firmansyah and Arman [8], has concentrated on compliance automation, transaction monitoring and machine-learning applications rather than on reporting continuity through structural change.

The gap is therefore specific. Regulation states the required end state. Technical research solves adjacent problems for other purposes. Organisational research documents that integration is hard without reaching the reporting layer. The execution method connecting them exists largely as undocumented institutional knowledge, held by the relatively small number of practitioners who have carried a reporting estate through a structural transition without missing a filing.

This paper is built on adjacent literature because direct predecessors do not appear to exist. That is stated here rather than concealed through selective citation.

3. What Does Not Pause

The reporting obligations that continue through a transition vary by legal entity, jurisdiction, asset size, organisational structure and supervisory category. Across the transitions described here they included monthly exposure-level submissions, quarterly capital-planning schedules, quarterly holding-company reports, quarterly Call Report and related FFIEC schedules, capital, liquidity, credit-risk and prudential returns, local legal-entity and consolidated reports, jurisdiction-specific financial and statistical returns, annual and event-driven capital-planning submissions, and the public financial and regulatory disclosures supported by the same underlying data [11][12][13].

Figure 2 plots the cadence classes across a twelve-month transition window.

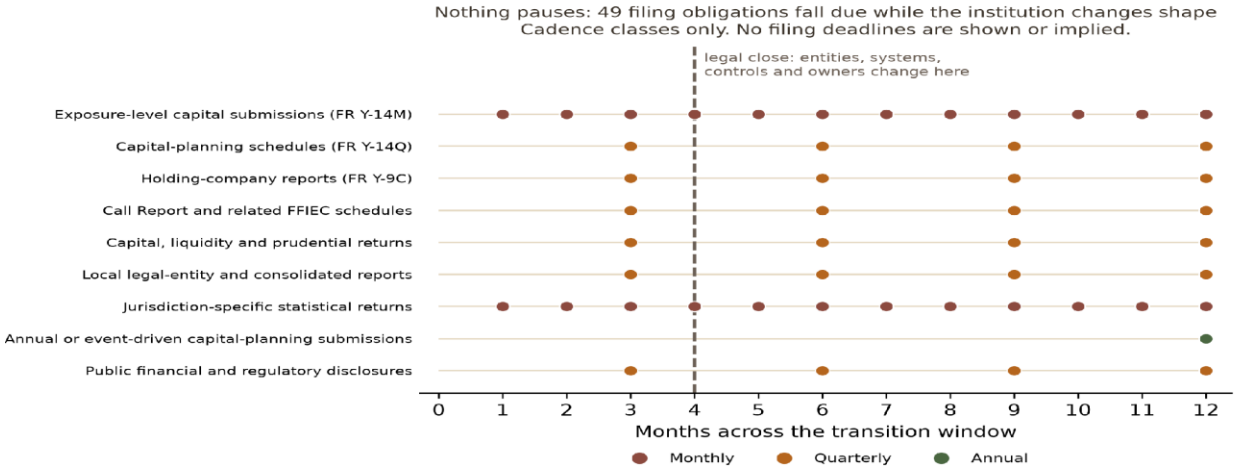


Figure 2. Filing obligations falling due across a transition window, with the legal close in month four. Forty-nine filing events occur in twelve months on cadence classes alone. Cadence classes only: no filing deadlines are shown or implied, since exact deadlines vary by report, jurisdiction and period and must be confirmed against the applicable instructions.

The figure is included for the arithmetic rather than for effect. Each of those forty-nine events requires complete data, working controls, reconciled balances and an authorized approver on the day it falls due, and on this cadence roughly two thirds of them land after the structural change has already started.

The consequences of failure escalate faster than is generally appreciated outside the function. Table 2 sets out the escalation observed across these transitions.

Table 2. Escalation path following a missed, incomplete or materially inaccurate filing.

Stage	**Consequence**
Immediate	Late-filing notification
Correction	Amended or resubmitted reports
Supervisory enquiry	Respondent queries and additional validation requests
Formal finding	Matters Requiring Attention, then Matters Requiring Immediate Attention [14]
Programme	Mandated remediation programme with tracked commitments
Governance	Control and governance findings; challenges to management attestation
Supervisory posture	Enhanced scrutiny; restrictions or conditions imposed
Financial	Civil monetary penalties
Where reporting data supports public disclosure	Incorrect public disclosure, restatement, investor concern, litigation exposure, reputational damage

The last row is the one that changes the risk calculus. Regulatory reporting data frequently feeds public financial disclosure, so a reporting defect during an integration is not contained within the supervisory relationship. In severe,

repeated or systemic cases the combined supervisory, financial, operational and reputational consequences are significant.

Figure 1 states the structural problem the rest of the paper addresses.

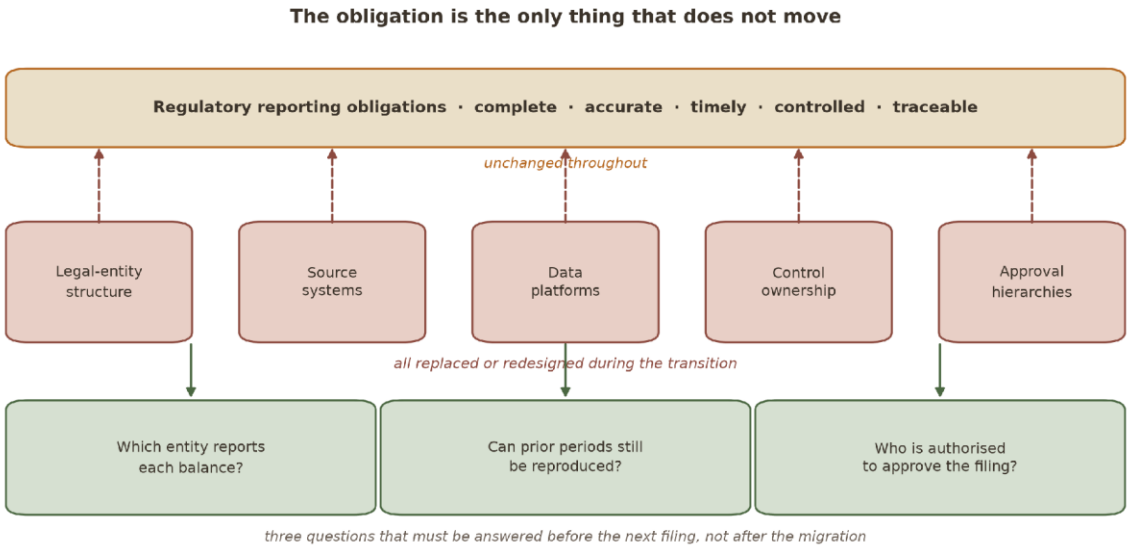


Figure 1. The obligation is the only fixed element. Legal-entity structure, source systems, data platforms, control ownership and approval hierarchies are all replaced or redesigned during the transition, and three questions must be answered before the next filing rather than after the migration.

4. Data Migration and Historical Reproducibility

4.1 What moves

The transitions required migration, integration or remapping of general-ledger balances, loan and deposit information, credit-facility data, customer and counterparty information, product and account attributes, collateral information, credit-risk and capital data, legal-entity reference data, organisational and consolidation hierarchies, intercompany positions, regulatory classifications, manual adjustments, historical reporting attributes, data-quality exceptions, report-specific calculations, and reporting population and exclusion rules.

Data moved from multiple legacy operational, lending, deposit, finance, risk and reporting systems into enterprise finance-data environments, consolidated data platforms and target regulatory reporting architectures.

4.2 The part that is not a data transfer

Treating this as a migration of balances is the most common error and the most expensive one to correct later. A regulatory reported figure is a balance that has been classified, mapped, adjusted, included or excluded, and then aggregated, according to rules that themselves have a history. If the logic does not move with the balance, the target environment can generate a number without being able to explain how it arrived at it.

The elements most at risk were consistently not the raw balances, which live in governed systems and move reliably. They were the contextual artifacts surrounding them. Table 3 sets these out.

Table 3. Fragile contextual artifacts, why they are exposed, and how they were preserved.

Artifact at risk	Why it is exposed	Preservation mechanism

Historical product and account mappings	Superseded by target mappings; prior-period meaning lost	Effective-dated mapping tables
Legal-entity assignments	Redrawn at close; prior perimeter no longer inferable	Legacy-to-target crosswalks with effective dates
Manual reporting adjustments	Held outside core systems, often in spreadsheets	Manual-adjustment logs retained with rationale
Report-specific exclusions	Encoded in reporting layer, not in source	Documented business-rule inventory
Undocumented transformation logic	Exists only in code or in practice	Version-controlled transformation rules
Spreadsheet-based calculations	Not covered by change management	Archived with the period-end extract
Historical exception-handling decisions	Precedent, not policy	Documented data-quality exception records
Locally maintained reporting definitions	Team-level, not enterprise-level	Standardised definitions in the target
Knowledge held by individual subject-matter experts	Leaves with the person	Lineage documentation and business-rule capture
Explanations kept in email or informal procedure	Not a record	Retained control evidence

The common property is that all of these sat outside centrally governed systems. That is precisely why they were vulnerable, and why a migration programme measured on balance-level completeness will pass its own tests and still destroy the institution's ability to reproduce a prior period.

Historical reproducibility was maintained through period-end source extracts, archived regulatory reporting datasets, effective-dated mapping tables, version-controlled transformation rules, lineage documentation, source-to-target mappings, reconciliation evidence, manual-adjustment logs, historical report versions, filing and submission archives, documented data-quality exceptions, business-rule documentation, retained control evidence and legacy-to-target crosswalks.

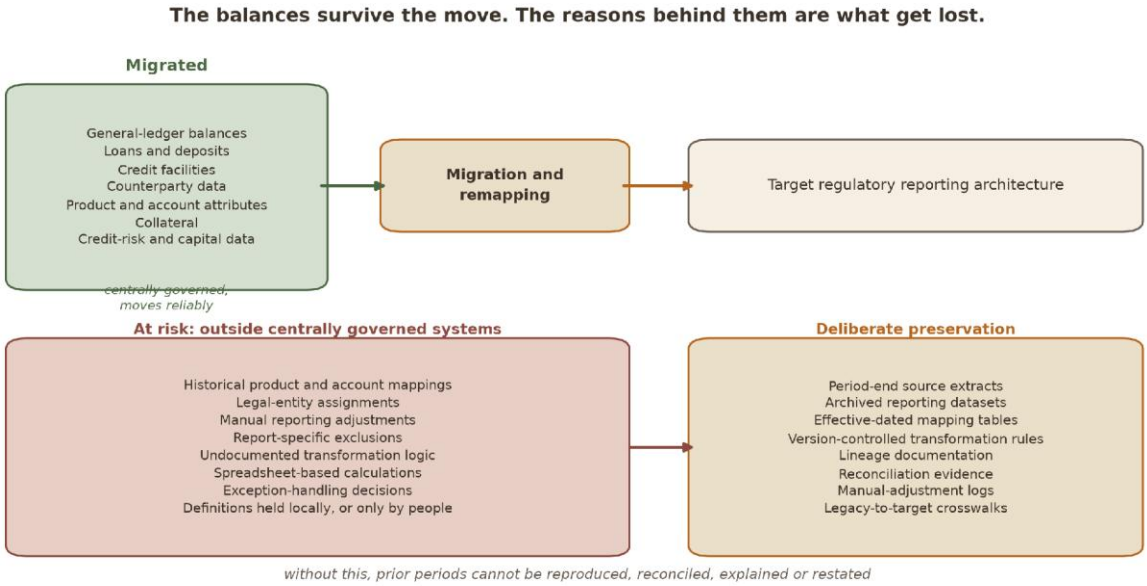


Figure 3. Balances migrate reliably because they are centrally governed. The artifacts that explain them are the ones at risk, and they survive only through deliberate preservation.

The test worth applying is not whether the target can produce the current period, which it generally can well before anyone is comfortable with it. It is whether the target can reproduce, reconcile, explain and if necessary restate a period that closed before the migration began. That test is considerably harder to pass, and it tends to be the one a supervisor gets around to asking about.

5. Control Re-Establishment

5.1 Why controls lapse

A control depends on four things at once: its design, the population it covers, an accountable owner, and an evidence trail. A structural transition can invalidate any one of them while leaving the control apparently in place, which is what makes this failure mode hard to spot. Nothing has been decommissioned, the control still runs on schedule, and it has stopped working.

The observed causes were that the source system was retired or replaced, data began arriving from a different platform, the reporting population expanded to include acquired entities, an existing control was not extended to newly integrated data, the original owner changed roles, the legal-entity hierarchy changed, the approval chain no longer represented the combined organisation, or different entities applied different validation rules and materiality thresholds.

Figure 4 maps eighteen control classes against those eight failure modes.

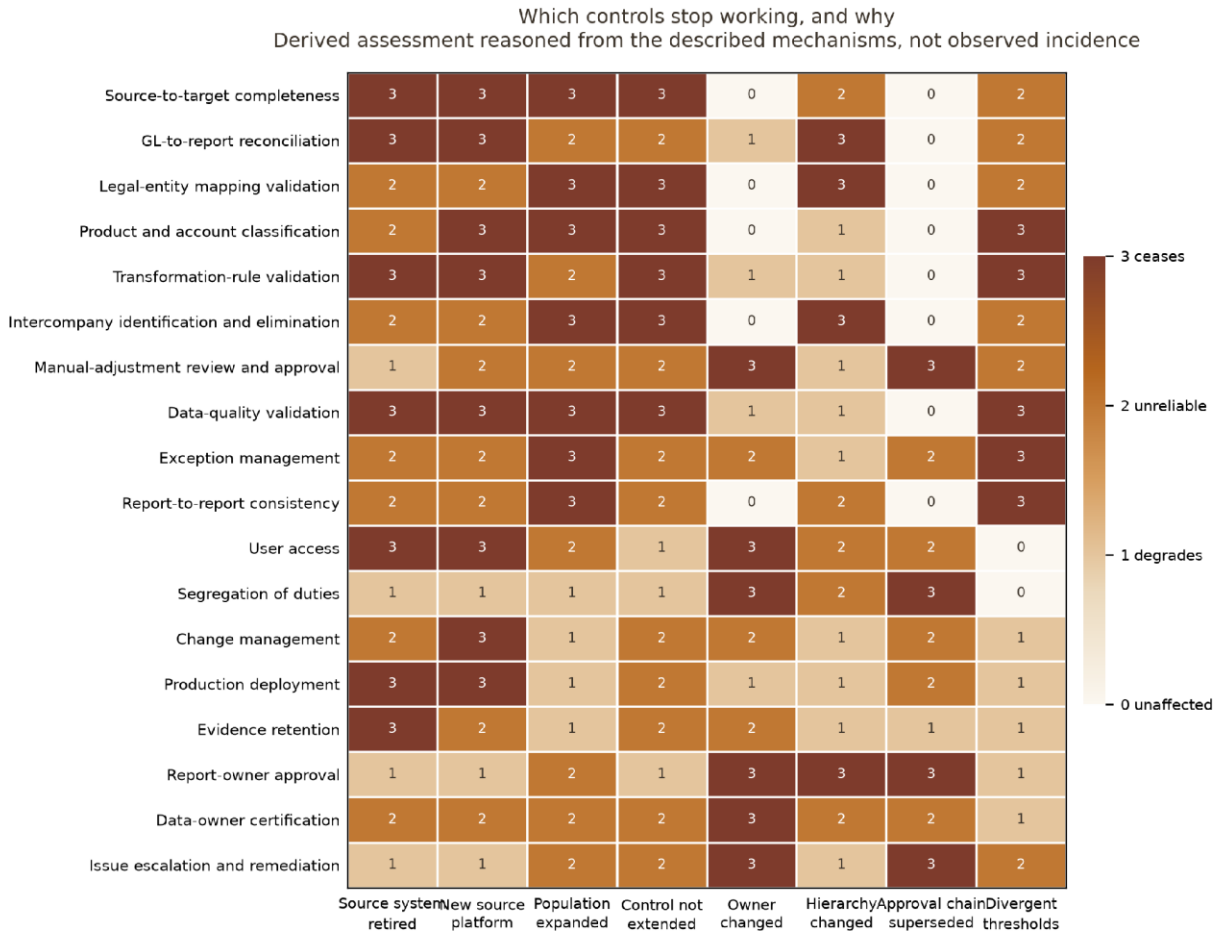


Figure 4. Control-class exposure to transition-driven failure modes. Ratings are a derived assessment reasoned from the described mechanisms, not an observed incidence rate.

Three readings follow from it. First, data-quality validation and exception management carry the highest aggregate exposure, because they are sensitive to almost every mode: they depend on the source, the platform, the population, the coverage and the thresholds simultaneously. Second, the single most destructive mode is data arriving from a new platform, which alone compromises source-to-target completeness, general-ledger reconciliation, product classification, transformation-rule validation, data-quality validation, user access, change management and production deployment. Third, the ownership and approval modes attack a different and often neglected set: segregation of duties, report-owner approval, data-owner certification and issue escalation are largely immune to platform change and highly exposed to personnel and hierarchy change. Programmes that scope control remediation around systems will systematically miss that second cluster.

5.2 The re-establishment sequence

Controls were re-established through a fourteen-step sequence: inventory legacy controls; identify design, ownership, coverage and execution gaps; map controls to the target reporting process; revise control descriptions and procedures; assign accountable owners; align controls to authoritative data sources; update populations, thresholds and validation rules; introduce compensating controls where target automation was incomplete; perform parallel reporting cycles; test control design and execution; resolve exceptions and defects; retain execution evidence; obtain approval or re-certification; and monitor the control through formal governance

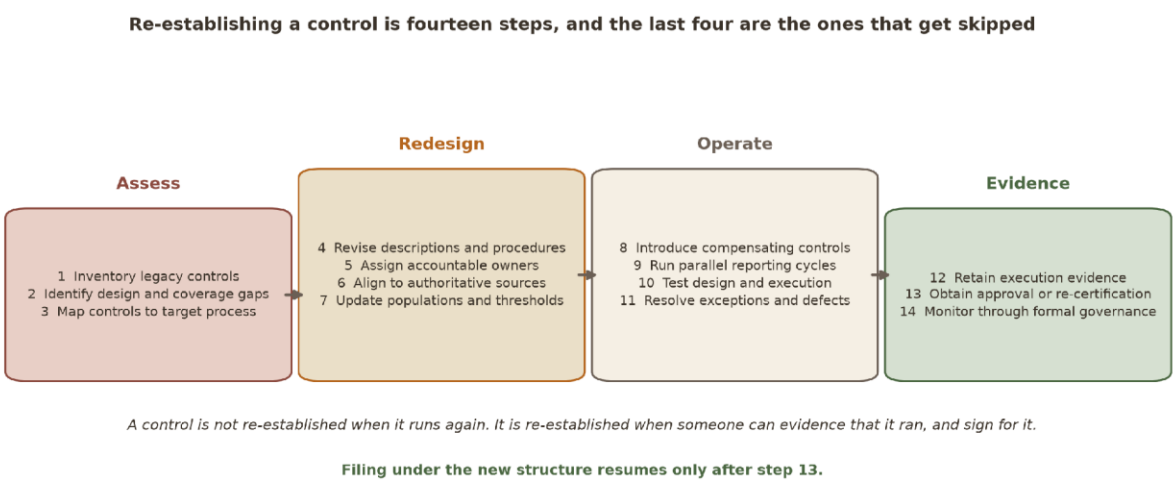


Figure 5. The control re-establishment sequence, grouped into assess, redesign, operate and evidence stages.

The last three steps are the ones under schedule pressure that get deferred, and they are the ones that matter to a supervisor. A control that runs but whose execution cannot be evidenced is, from a supervisory standpoint, a control that has not been demonstrated to run. Filing under the new structure resumes at step thirteen, not at step ten.

Approval or attestation was required from the applicable data owners and stewards, business process owners, regulatory report owners, Finance or Regulatory Reporting controllers, technology control owners and authorized filing officers. Independent Risk, Compliance, Internal Audit or Quality Assurance functions could separately evaluate the control environment before the target process was treated as sustainable.

Governance frameworks in the literature [1] assume identifiable and stable ownership. A transition suspends that assumption, and the re-establishment sequence exists precisely to rebuild it.

6. Legal-Entity Remapping and the Prevention of Double Counting

6.1 The remapping process

Legal-entity restructuring is not the substitution of one identifier for another. Accounts, facilities, customers, products, counterparties, general-ledger balances, risk positions and reporting attributes must each be assigned to the correct legal entity and reporting perimeter as of the applicable effective date.

The process was: establish an authoritative crosswalk between legacy and target legal-entity identifiers; apply effective dates to each mapping; position each entity within the new legal and consolidation hierarchy; map customers, accounts, products and facilities to the appropriate reporting entity; identify intercompany counterparties and balances; apply elimination rules at the appropriate consolidation level; separate pre-transition and post-transition populations; distinguish transferred balances from newly originated activity; reconcile stand-alone entity totals to consolidated totals; detect duplication across legacy and target feeds; preserve lineage to the originating entity and system; and confirm that every reportable balance is included once, in the correct perimeter.

The effective-dating requirement is worth isolating. The correct mapping for a given record is a function of the reporting date, so a crosswalk without effective dates cannot reproduce a prior period even when it is correct for the current one. The schema-matching literature [6] treats mapping largely as a static correspondence problem and under-treats this temporal case.

6.2 The hardest case

The recurring difficulty was one exposure appearing in more than one dataset during the migration period. An exposure could remain active in a legacy entity-level feed while also being migrated into an enterprise finance environment and subsequently included in a consolidated regulatory reporting dataset.

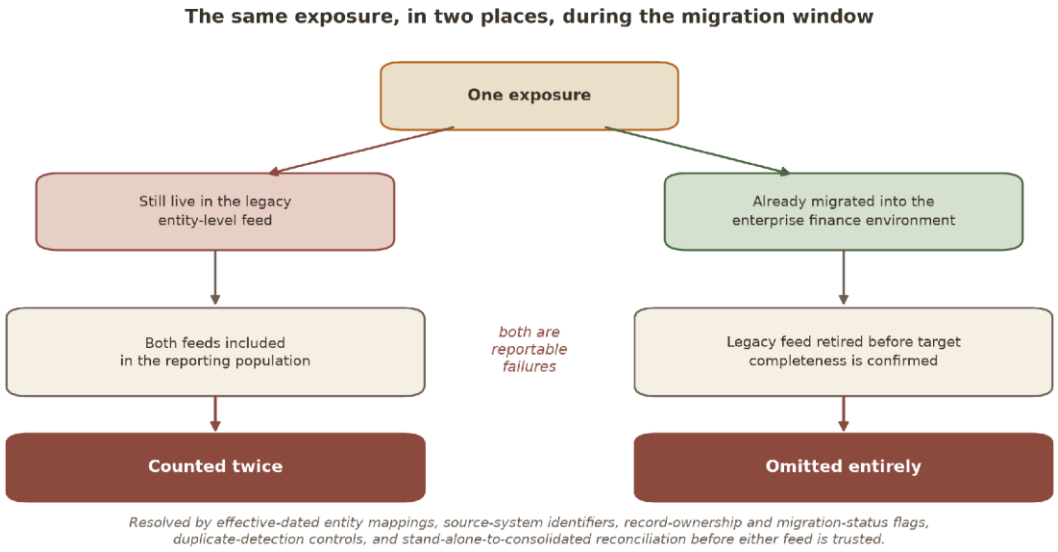


Figure 6. The dual-feed problem and its two failure paths. Including both feeds counts the exposure twice. Retiring the legacy feed before target completeness is confirmed omits it. Both are reportable failures.

The two failures are symmetric and they pull in opposite directions, which is what makes the timing decision genuinely difficult rather than merely careful. Resolution required effective-dated legal-entity mappings, source-system identifiers, record-ownership indicators, migration-status flags, intercompany classifications, consolidation rules, elimination indicators, duplicate-detection controls, stand-alone-to-consolidated reconciliations, and traceability from the final reported value back to the originating record.

Listing 1 shows the classification at the centre of this.

Listing 1. Dual-feed outcome classification for one reporting date. An exposure appears in the legacy feed until that feed is retired, and in the target feed once it has migrated.

```
    def classify_at_reporting_date(population, reporting_day, legacy_retirement_day):
  df = population.copy()
  df["in_legacy_feed"] = reporting_day < legacy_retirement_day
  df["in_target_feed"] = df["migration_day"] <= reporting_day

  df["report_count"] = df["in_legacy_feed"].astype(int) + df["in_target_feed"].astype(int)
  df["outcome"] = np.select(
     [df.report_count == 0, df.report_count == 1, df.report_count >= 2],
     ["omitted", "reported_once", "double_counted"],
     default="unclassified",
  )
  df["elimination_required"] = df["intercompany_flag"] & (df.report_count > 0)
  return df
```

6.3 An analytical model of the overlap decision

The practical question is when to retire the legacy feed. To examine the structure of that decision, the classification above was applied across a synthetic exposure population over the successive filing periods falling inside a transition window. The model contains no measured data. It is a reasoning artifact derived from the failure logic described above, and every row it emits is flagged as synthetic.

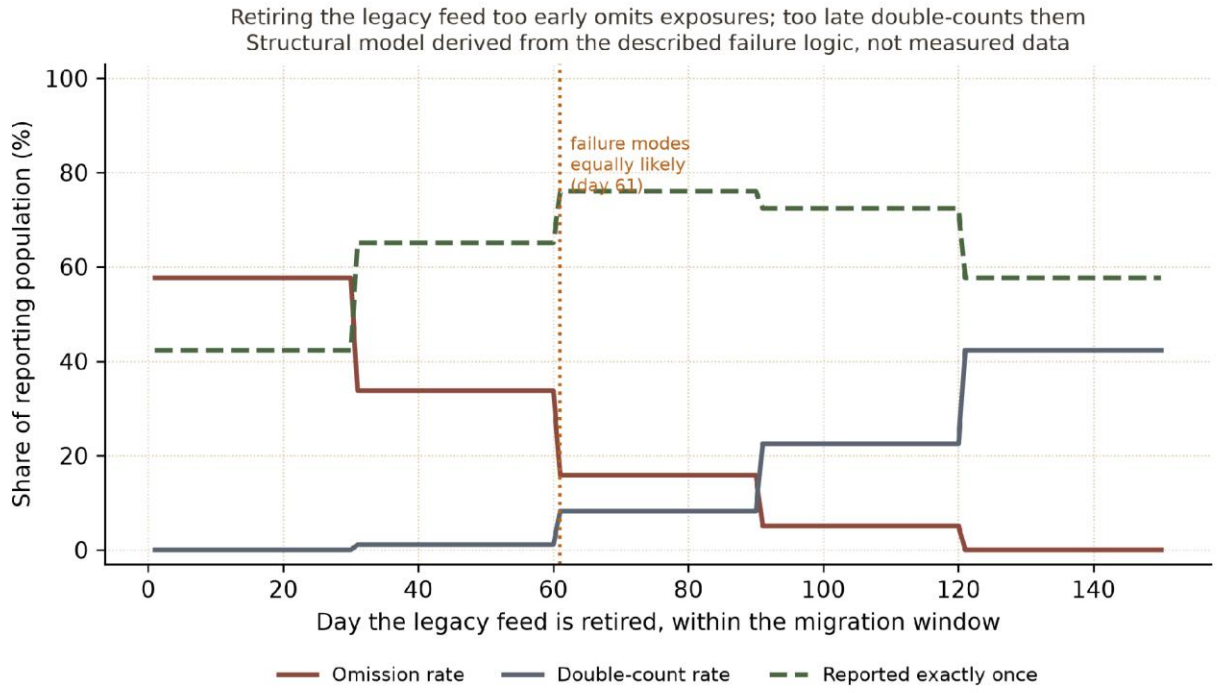


Figure 7. Omission and double-count exposure as a function of legacy-feed retirement timing, accumulated across the filings falling inside the window. Structural model, not measured data.

Table 4. Modelled dual-feed outcomes by retirement timing, accumulated across four filing periods. Synthetic population, derived model.

Scenario	Reported once	Omitted	Double counted
Legacy retired at first period end	42.4%	57.6%	0.0%
Legacy retired mid-window	76.0%	15.8%	8.2%
Legacy retired at window close	72.4%	5.1%	22.5%
Legacy retained one period beyond	57.6%	0.0%	42.4%

Three observations follow, and they are properties of the structure rather than of any institution.

The failure rate is not monotonic in retirement timing. Both the earliest and the latest options are worse than an interior choice, so neither the instinct to decommission promptly nor the instinct to retain the legacy feed for safety is correct on its own.

Retaining the legacy feed is not the conservative option. It merely exchanges omission risk for double-count risk. In the modelled window, holding the legacy feed one filing period beyond the migration produced a higher total failure rate than retiring it at window close.

Because the failure modes cross, the correct retirement point cannot be selected from the calendar. It depends on the completeness of the target population at the moment of decision, which is knowable only from reconciliation evidence. This is the practical argument for gating legacy-feed retirement on stand-alone-to-consolidated reconciliation rather than on a migration plan date.

A related structural point: duplicate detection here is a record-linkage problem, and the algorithmic literature [5] has addressed it thoroughly for customer and bibliographic records without ever being applied to consolidation perimeters. The blocking and indexing techniques surveyed there transfer directly once the linkage key is understood to be the combination of source-system identifier, effective-dated entity mapping and migration-status flag.

7. Supervisory Communication

Supervisory attention during these transitions concentrated on whether the institution could demonstrate that its reporting remained complete, accurate, timely, controlled and traceable throughout the organisational change [15].

Supervisory stakeholders sought assurance that required reports would continue without interruption; that the institution understood the reporting impact of the acquisition or restructuring; that reporting responsibilities were clearly assigned; that transferred populations and opening balances were complete; that data definitions were consistent across combined entities; that authoritative data sources had been established; that the same attribute was not sourced inconsistently from multiple systems; that data lineage remained traceable from report to originating source; that stand-alone and consolidated balances reconciled; that intercompany balances were identified and eliminated appropriately; and that historical reporting periods remained reproducible.

The structural observation is that every item on that list is an evidence-production requirement rather than a reporting requirement. None of it is satisfied by filing correctly. Each is satisfied by being able to show, on request and after the fact, how the correct filing was produced and who was accountable for each step. That is why the artifacts described in Sections 4 and 5, the effective-dated mappings, the reconciliation evidence, the retained control evidence, the manual-adjustment logs, are not administrative overhead. They are the substance of the supervisory relationship during a transition.

Where findings were raised, the escalation ladder in Table 2 applied, and formal findings carry tracked remediation commitments with evidence obligations of their own [14].

Scope note. This section is deliberately shorter than the others and is grounded in published supervisory frameworks rather than in the author's account of specific engagements. Supervisory communications are non-public, and the confidentiality position adopted throughout this paper excludes them. Section 9 records the resulting limitation.

8. The Framework, Consolidated

The four components combine into a sequence that can be applied to a transition of any of the three classes in Table 1.

Table 5. Consolidated execution framework.

Stage	Question it answers	Primary artifacts produced
Obligation inventory	What must keep filing, on what cadence, under which entity	Report inventory by entity, jurisdiction and cadence
Impact assessment	Which obligations does this transition touch, and how	Reporting impact assessment per report family

Entity crosswalk	Which entity reports each balance, as of when	Effective-dated legacy-to-target entity mappings
Data-sourcing strategy	Which system is authoritative for each attribute	Source-to-target mappings, authoritative-source register
Reproducibility preservation	Can prior periods still be explained and restated	Period-end extracts, versioned rules, adjustment logs, lineage documentation
Control re-establishment	Which controls stopped working, and who signs for the new ones	Control inventory, gap assessment, re-certification evidence
Dual-feed governance	Is every balance reported once, and once only	Duplicate-detection controls, stand-alone-to-consolidated reconciliation
Evidence retention	Can we show a supervisor how this was produced	Retained control and reconciliation evidence

Two properties of the sequence are worth naming. It is gated rather than phased, in that the entity crosswalk cannot complete before the obligation inventory and dual-feed governance cannot conclude before reconciliation evidence exists. And what it ultimately produces is evidence rather than systems, which is an uncomfortable thing to tell a programme sponsor who is measuring progress in platform milestones. A transition that ends with a working target platform and no evidence trail still has work left in it.

9. LIMITATIONS

Single-practitioner account, no measurements. This paper reports no quantitative results because none were available to report. There are no counts of entities, reports, controls, records or defects, no durations, no error rates and no cycle counts. The contribution is a methodology and a set of structural observations. It is not empirical evidence, and nothing here supports a claim about effectiveness relative to any alternative approach.

Derived rather than observed figures. Figure 4 is an analytical mapping reasoned from described mechanisms, not an incidence rate. Figures 7 and Table 4 come from a synthetic population exercising the described failure logic. Both are labelled as such in the artifacts, which flag every emitted row. They illustrate structure. They are not evidence about any institution.

Supervisory communication is under-treated. One of the three components named in the paper's own title is the thinnest section in it, for two reasons. Supervisory communications are non-public and excluded by the confidentiality position adopted here. Separately, the source material available for this section was incomplete. Section 7 is consequently built on published supervisory frameworks rather than on direct account, and a fuller treatment would require material this paper does not have.

Confidentiality constrains specificity. Institution names, project names, platform identifiers, transaction details and jurisdiction-specific arrangements are omitted throughout. This is necessary and it has a cost: some claims that would be more persuasive with a concrete instance are stated at the level of mechanism instead.

No filing deadlines are stated. Exact deadlines vary by report, jurisdiction, asset size and reporting period, and are subject to change. Nothing here should be relied on for deadline purposes; the applicable regulatory instructions govern.

Generalisability is asserted, not demonstrated. The framework derives from transitions within a particular regulatory environment. The failure modes appear structural rather than institution-specific, which is the basis for offering the framework at all, but that is an argument rather than a finding.

10. CONCLUSION

The obligation does not pause. That single fact determines the shape of every decision in a reporting-estate transition, and it is why a migration programme that would be sound in another context can be unsound here.

Three conclusions carry beyond the specific transitions described. Data migration in a regulatory context is the migration of reporting logic, not of balances, and the artifacts most likely to be lost are the ones that were never in a governed system to begin with: the manual adjustments, the local definitions, the exception precedents, the explanation that lived in somebody's inbox. A control is not re-established when it runs again but when its execution can be evidenced and someone with authority will sign for it, which places the last three steps of the re-establishment sequence, the ones most often compressed under schedule pressure, at the centre rather than the periphery. And the legacy-feed retirement decision is not a scheduling matter, because omission and double-count risk move in opposite directions and cross, so the decision has to be gated on reconciliation evidence rather than on a plan date.

Future work should test whether the control exposure mapping in Section 5 holds across institutions, which would require incident data no single practitioner has. The dual-feed model in Section 6 would benefit from calibration against actual migration telemetry. Most valuable would be work connecting the record-linkage literature to regulatory consolidation perimeters, where the algorithms already exist and the application does not.

REFERENCES

1. R. Abraham, J. Schneider and J. vom Brocke, "Data governance: A conceptual framework, structured review, and research agenda," International Journal of Information Management, vol. 49, pp. 424-438, 2019. doi: 10.1016/j.ijinfomgt.2019.07.008
2. C. Batini, C. Cappiello, C. Francalanci and A. Maurino, "Methodologies for data quality assessment and improvement," ACM Computing Surveys, vol. 41, no. 3, 2009. doi: 10.1145/1541880.1541883
3. M. Herschel, R. Diestelkämper and H. Ben Lahmar, "A survey on provenance: What for? What form? What from?," The VLDB Journal, vol. 26, pp. 881-906, 2017. doi: 10.1007/s00778-017-0486-1
4. R. Bose and J. Frew, "Lineage retrieval for scientific data processing: a survey," ACM Computing Surveys, vol. 37, no. 1, pp. 1-28, 2005. doi: 10.1145/1057977.1057978
5. P. Christen, "A survey of indexing techniques for scalable record linkage and deduplication," IEEE Transactions on Knowledge and Data Engineering, vol. 24, no. 9, pp. 1537-1555, 2012. doi: 10.1109/TKDE.2011.127
6. P. Shvaiko and J. Euzenat, "Ontology matching: State of the art and future challenges," IEEE Transactions on Knowledge and Data Engineering, vol. 25, no. 1, pp. 158-176, 2013. doi: 10.1109/TKDE.2011.253
7. S.-I. Chang, I.-C. Chang and T. Wang, "Information systems integration after merger and acquisition," Industrial Management & Data Systems, vol. 114, no. 1, pp. 37-52, 2014. doi: 10.1108/imds-03-2013-0157
8. B. Firmansyah and A. A. Arman, "A systematic literature review of RegTech: Technologies, characteristics, and architectures," in Proc. International Conference on Information Technology Systems and Innovation (ICITSI), 2022. doi: 10.1109/icitsi56531.2022.9970830
9. Basel Committee on Banking Supervision, Principles for effective risk data aggregation and risk reporting, BCBS 239. Bank for International Settlements, January 2013.
10. European Central Bank, Guide on effective risk data aggregation and risk reporting. ECB Banking Supervision, May 2024.
11. Board of Governors of the Federal Reserve System, Capital Assessments and Stress Testing (FR Y-14A, FR Y-14Q, FR Y-14M), reporting forms and instructions.
12. Board of Governors of the Federal Reserve System, Consolidated Financial Statements for Holding Companies (FR Y-9C), reporting form and instructions.
13. Federal Financial Institutions Examination Council, Consolidated Reports of Condition and Income (Call Report), FFIEC 031, 041 and 051, instructions.
14. Board of Governors of the Federal Reserve System, supervisory guidance concerning supervisory findings, Matters Requiring Attention and Matters Requiring Immediate Attention.
15. Basel Committee on Banking Supervision, Core Principles for Effective Banking Supervision. Bank for International Settlements, revised April 2024.
16. Deloitte, BCBS 239 Benchmark Survey, 2024. Appendix A: Accompanying Artifacts Every row emitted by dual_feed_reconciliation.py carries is_synthetic=True. Every row emitted by the control exposure mapping carries is_derived=True. Neither contains measured data, and the paper does not present either as evidence about any institution.