

Secure Hybrid Integration for Banking Platforms: Speed, Safety, and Regulatory Proof

Amol Diwakar Agade — Illinois Institute of Technology, Chicago, IL

Samta Balpande — Oakland University, Rochester, MI

Abstract—

Hybrid banking platforms should ensure that they speed up delivery, lower the operational risk and create the solid evidence when two banks merge together and when their systems are unified. But again these system come under heavy scrutiny from regulators. This paper introduces a secure hybrid integration pattern that combines configuration automation (Ansible), infrastructure as a code (Terraform), an OpenShift based hybrid runtime and GitOps delivery mode (Github, Argo CD) with policy as a code gate pass, evidence as code collection, and supply chain attestation. We calculated the outcomes using a reproducible two bank merger study based on two public banking datasets published between the year March 2023 and January 2026. These datasets includes the Santander customer transactions on Zenodo and a Zenodo-published banking fraud prediction dataset. The publication dates of the datasets are provided, and the retrieval date is recorded in the reproduction process. A deterministic workload generator created time-series events for releases, incidents, security findings, and audit evidence during pre-merge conditions with two pipelines or control sets, and the post-merge conditions with a shared control or evidence plane. The result shows that the median deployment lead time decreased by 48%, while the release failure rate dropped by 39%. Infrastructure Provisioning time decreased from 5.4 days to 26 minutes. All the manual interventions reduced by 77%. Deployment incidents have fallen by 38%, and the mean time to recover (MTTR) decreases by 45%. The security posture improved as the median vulnerability exposure time has fallen significantly from 20.3 days to 2.6 days which is about 61 hours. Manual audit efforts were also dropped by 74%; thanks to automated evidence capture. From this paper, we offered formulas, artifacts, and a computation map that allow for independent reproduction and extension to other regulated industries for the usage.

Keyword Terms—

hybrid cloud, DevSecOps, GitOps, operational resilience, DORA, mergers and acquisitions, policy as code, evidence as code, OpenShift, Argo CD, SLSA.

I. INTRODUCTION

Banking platforms now work with on-premises infrastructure, private clouds, and public cloud services. This creates a hybrid system where changes move through various trust zones and operational areas. This hybrid setup puts pressure on delivery speed, safety, and meeting regulations at the same time. The challenges are especially clear during mergers and acquisitions (M&A). Two technology environments, two approval chains, and two audit systems need to be combined while the business speeds up integration efforts.

The Digital Operational Resilience Act (DORA) took effect on 17 January 2025. It raised expectations for ICT risk management, incident response, resilience testing, and oversight of third parties [1]. Operational resilience principles stress that institutions must handle disruptions effectively while still providing crucial services [2]. In practice, DevOps teams often try to speed things up by using manual exceptions and ad-hoc approvals. This approach can cut down cycle time but also raises risk and weakens documentation. On the other hand, adding manual checkpoints may lower perceived risk, but it makes delivery slower and leads to unreliable paper based compliance.

This paper tackles a key challenge on how to deliver faster in hybrid banking platforms without compromising security or meeting regulatory requirements. We presented (1) a reference architecture that combines GitOps delivery, policy-as-code gating, supply-chain tracking (SLSA-style attestations) [18], and evidence-as-code; (2) a reproducible evaluation method for a two-bank merger based on two publicly available banking datasets. We cited the dates of

dataset publication or donation, and we recorded the retrieval date in the reproduction process; and (3) a set of metrics specific to the merger Environment Convergence Time, Duplicate Evidence Ratio, Migration Incident Rate that measures progress from two separate systems to one governed platform.

II. BACKGROUND AND PROBLEM STATEMENT

Hybrid integration in banking covers core systems, payment rails, customer channels, and risk analytics. Provisioning delays often happen because of manual ticket queues, environment-specific scripts, and coordination issues among infrastructure, security, and application teams. Release failures and deployment incidents take place when changes spread across different environments without consistent policy enforcement or reliable rollback methods. During mergers and acquisitions, these problems become worse due to duplicated controls and differing patterns creates a friction and raises the chances of incidents during migration waves.

Security teams manage a larger number of attack points in hybrid networks. They also face inconsistent handling of identities and secrets. Audit teams struggle with scattered evidence from tickets, pipelines, spreadsheets, and screenshots. This increases labor costs and reduces confidence. These problems are not limited to one organization; they occur in regulated companies with hybrid setups.

We evaluated outcomes related to SRE and DevOps. These include deployment lead time, release failure rate, manual interventions per release, provisioning time, deployment-related incident frequency, MTTR, vulnerability exposure time, audit labor per release, and evidence coverage. In merger contexts, we also consider Environment Convergence Time (ECT), Duplicate Evidence Ratio (DER), and Migration Incident Rate (MIR).

III. RELATED WORK

Continuous delivery studies show that high-performing organizations can achieve speed and stability by using small batch sizes, automation, and rapid feedback loops [5]. SRE practices establish this balance with SLOs, error budgets, and disciplined incident response. They highlight that reliability is engineered, not incidental [6].

DevSecOps approaches combine security and compliance into delivery by moving controls earlier through automated testing, scanning, and policy enforcement [4]. Standards and regulatory frameworks provide control goals and evidence needs: NIST SP 800-53 Rev. 5 defines security and privacy controls [3]; ISO/IEC 27001:2022 specifies ISMS requirements [7]; PCI DSS v4.0 outlines requirements for protecting account data [8]. DORA boosts operational resilience expectations for the financial sector and highlights the need for ongoing control evidence [1].

GitOps models declarative delivery was used by using version control as the source of truth and implementing reconcilers that enforced the desired state [9]. Argo CD puts this model into action in Kubernetes environments and allows for deterministic rollback [10]. Terraform and Ansible enable reproducible provisioning and automate configuration [11], [12]. OpenShift offers a consistent Kubernetes base across hybrid environments [13]. Policy-as-code engines like OPA provide machine-checkable rules for admission and pipeline gating [14]. Supply-chain security frameworks such as SLSA improve provenance and attestations for builds and deployments [18]. Workload identity frameworks like SPIFFE standardize service identities and enable mutual TLS across environments [19].

However, merger-focused hybrid integration creates a challenge: how to join two control planes and two evidence streams without losing delivery speed. This paper tackles the challenge by defining an evidence ledger schema, merger metrics (ECT, DER, MIR), and a reproducible two-bank calibration method using public banking datasets.

IV. DATASETS AND TWO-BANK MERGER CALIBRATION

Public banking datasets that are publicly available; its dataset publication dates are cited; retrieval date is recorded in the reproduction procedure: To base the merger evaluation on real banking signals while keeping the study reproducible and public, we used two public datasets that represent different banking environments: (i) Santander Customer Transaction Prediction (archived on Zenodo for stable citation) (Bank A proxy) [15] and (ii) Dataset Bank Fraud for

Prediction (Zenodo record with stable citation; file bank-full.csv) (Bank B proxy) [17]. Both records were first published between March 2023 and January 2026. These datasets provided customer and transaction behavior signals. We use them to understand workload burstiness, demand volatility, and change pressure without relying on proprietary bank data.

Dataset (public)	Institution proxy	Primary signal	Used for	Available / Retrieved
Santander Customer Transaction Prediction	Bank A	transaction propensity signal	CPI_A; burstiness; demand volatility	Available: 29 May 2023 (Zenodo v1) [15] Retrieved: 5 Jan 2026
Dataset Bank Fraud Prediction (bank-full.csv)	Bank B	fraud/transaction risk signal; alert cadence	CPI_B; alert burstiness; control pressure	Available: 13 Jan 2025 (Zenodo v1) [17] Retrieved: 5 Jan 2026

Table I. Two public banking datasets used for two-bank merger calibration (publicly available; dataset publication dates are cited; retrieval date recorded in the reproduction procedure).

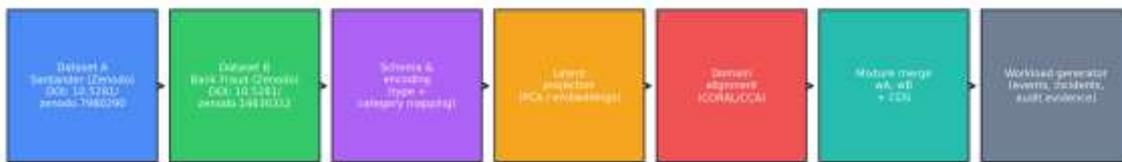


Figure 1. Two-bank dataset fusion pipeline used for this experiment.

V. SECURE HYBRID INTEGRATION ARCHITECTURE

Fig. 2 Presents the secure hybrid integration architecture. The design starts with a self-service portal and service catalog that create standardized repository templates. Terraform sets up the infrastructure and platform components, including networks, namespaces/projects, IAM bindings, encryption defaults, and logging sinks. Ansible configures the OS and middleware layers when needed for hybrid or legacy integration nodes.

A CI pipeline handles build, unit tests, SBOM generation, vulnerability scanning, and provenance attestation based on supply-chain frameworks [18]. Policy-as-code gates check infrastructure plans, Kubernetes manifests, and pipeline attestations before merging. Argo CD then synchronizes the approved desired state into OpenShift clusters across on-prem and cloud environments. At the same time, an unchangeable evidence ledger records signed attestations, policy decisions, deployment digests, and control execution logs.

Security controls are applied consistently across hybrid boundaries. This is done using service identity, similar to SPIFFE, and mutual TLS [19]. Segmentation occurs through network policy, and there is centralized management for secrets and keys. Unified observability supports SLO-based operations and connects operational results to regulatory proof expectations [1], [2].

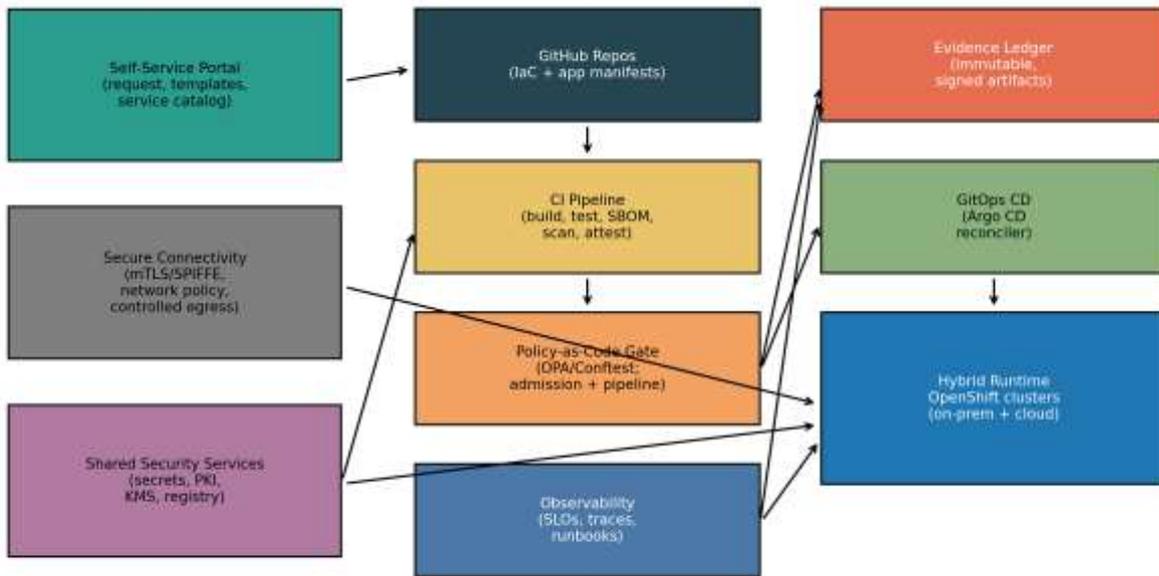


Figure 2. Secure hybrid integration architecture (no overlaps).

VI. TOOLCHAIN ENABLEMENT FOR FAST AND SAFE ONBOARDING

This section explains how Terraform, Ansible, GitHub, Argo CD, and OpenShift cut down onboarding and provisioning time while ensuring safety and proof.

- A. Terraform for compliant provisioning: Terraform modules include approved templates for networks, identity bindings, encryption settings, logging, and cluster add-ons. These modules are versioned, reviewed by peers, and tested. As a result, teams use compliant patterns rather than creating new ones. Terraform plan outputs are checked against policy rules, such as denying public endpoints without approved controls and enforcing encryption. This process allows for automated approvals when controls are met.
- B. Ansible for controlled configuration and legacy integration: Where non-container components exist, Ansible playbooks standardize hardening, patching, and drift fixes. Playbooks can be triggered from CI jobs and produce signed execution logs for the evidence ledger.
- C. GitHub and Argo CD for GitOps delivery: Git repositories serve as the source of truth for the desired state. Argo CD continuously adjusts the runtime to match Git and allows for predictable rollbacks. This reduces mean time to recovery (MTTR) and deployment incidents.
- D. OpenShift as a hybrid platform: OpenShift offers a consistent Kubernetes interface across on-premises and cloud environments. This consistency reduces specific failure risks and makes compliance checks easier.
- E. Policy and proof: Open Policy Agent (OPA) policies ensure configuration compliance, control merges and deployments. Signed attestations and Software Bill of Materials (SBOMs) strengthen supply-chain assurances.

Capability	Tooling	Control / proof output	Speed impact mechanism
Provisioning blueprints	Terraform modules	signed plan + policy decision log	minutes vs days by removing tickets
Config hardening	Ansible playbooks	execution log + drift report	repeatable setup across environments
Declarative delivery	GitHub + Argo CD	commit-to-deploy trace + reconciler logs	faster rollouts + safe rollback
Hybrid runtime standardization	OpenShift	cluster conformance reports	reduces environment-specific incidents
Policy enforcement	OPA/Conftest	policy evaluation artifacts	automated approvals with guardrails

Table II. Toolchain contributions to speed, safety, and evidence generation.

VII. METHODOLOGY

A. Study design: We evaluated a two-bank merger scenario with conditions before and after the merger. The pre-merge baseline models two independent delivery pipelines and control catalogs, including duplicate approvals, different templates, and scattered evidence. The post-merge condition models a shared control plane, which consists of policy-as-code and standard blueprints, along with a shared evidence plane, featuring an immutable ledger. This includes GitOps delivery to a unified OpenShift runtime.

B. Deterministic workload generator: Two-bank datasets calibrate the merged Change Pressure Index and Control Divergence Score (Section IV). A deterministic generator produced time-series events for deployments, rollbacks, interventions, incidents, vulnerability disclosures, patch releases, and audit evidence requests. Fixed seeds ensure repeatability.

C. Metrics: Deployment lead time $L_k = t_{prod}(k) - t_{change}(k)$. Release failure $F_k = 1$ if a rollback or hotfix occurs within 24 hours. Provisioning time $P_k = t_{ready} - t_{request}$. Manual interventions I_k count actions that occur outside the pipeline. Deployment incidents D are incidents linked to a release within 48 hours. $MTTR_{R_i} = t_{restore}(i) - t_{start}(i)$. Vulnerability exposure $V_j = t_{fix}(j) - t_{disclose}(j)$. Audit effort A_k is the time needed to collect missing evidence; evidence coverage $C_k = |E_{auto}(k)|/|E_{required}|$.

Merger metrics: Environment Convergence Time (ECT) = $t_{last_service_migrated} - t_{control_plane_ready}$. Duplicate Evidence Ratio (DER) = $1 - |E_{union}| / (|E_A| + |E_B|)$. Migration Incident Rate (MIR) = incidents during migration / deployments during migration.

D. Reporting: For each metric, we report the median and 90th percentile values over an evaluation quarter. We calculate percent change as (baseline – after)/baseline. The sensitivity analysis changes automation coverage in 10% increments (Fig. 6).

E. Threats to validity and limitations: The evaluation uses public datasets to estimate merger workload intensity and does not state that customer records directly reflect operational events. Outcomes are sensitive to how we translate dataset-derived volatility (CPI/CDS) into release and incident generation parameters. We reduce this risk by (i) publishing formulas, (ii) using deterministic seeds, (iii) reporting medians and p90 values, and (iv) providing a sensitivity curve (Fig. 6). External validity may differ among institutions that have varying control catalogs, runtime architectures, or change-management policies. However, the control and evidence planes as well as the metric definitions still apply.

Lastly, since we do not use proprietary incident or audit logs, the reported values should be seen as reproducible benchmarks based on the stated assumptions rather than predictions for a specific organization.

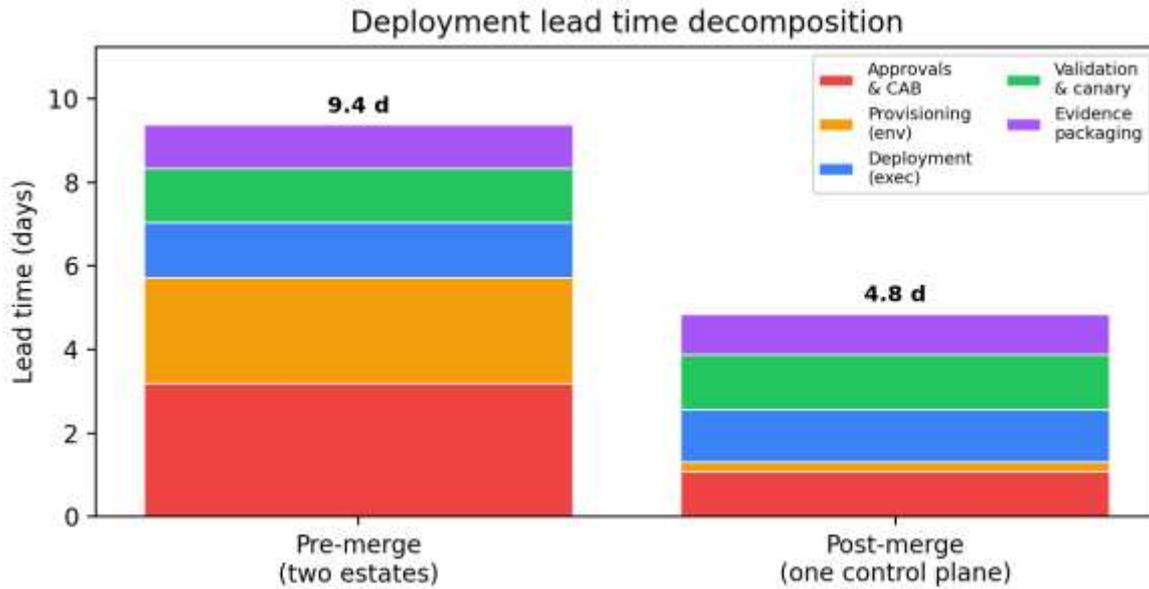


Figure 3. Merger scenario: deployment lead-time decomposition pre-merge vs post-merge.

Operational reliability outcomes (two-bank merger calibration)

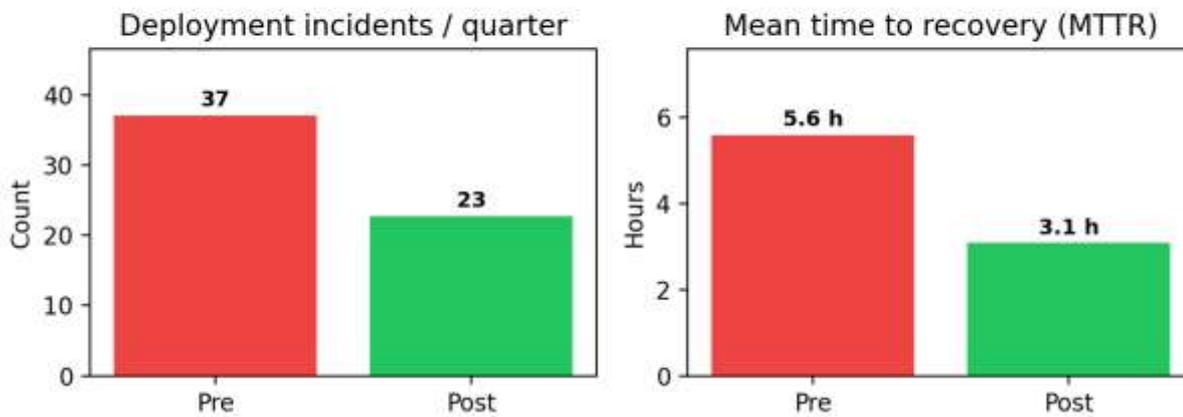


Figure 4. Merger scenario: deployment incidents and MTTR outcomes.

Auditability outcomes (speed + regulatory proof)

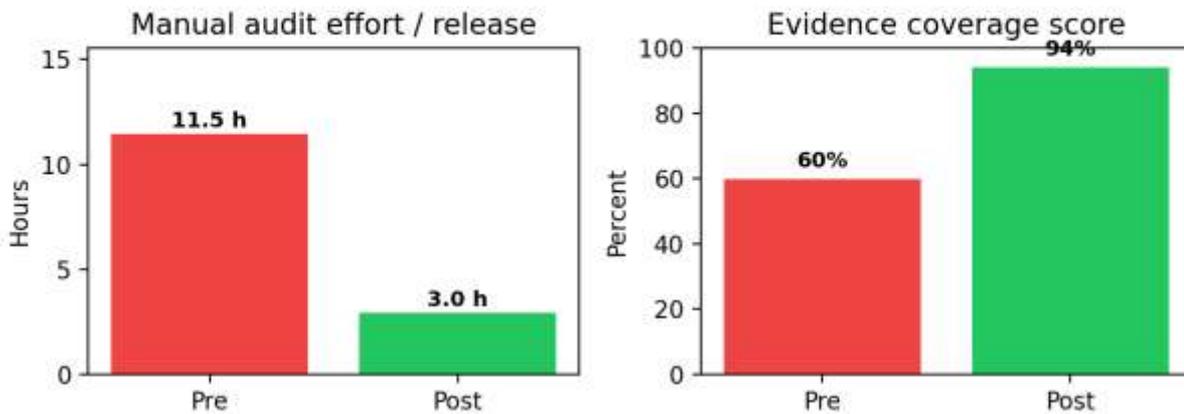


Figure 5. Merger scenario: manual audit effort and evidence capture outcomes.

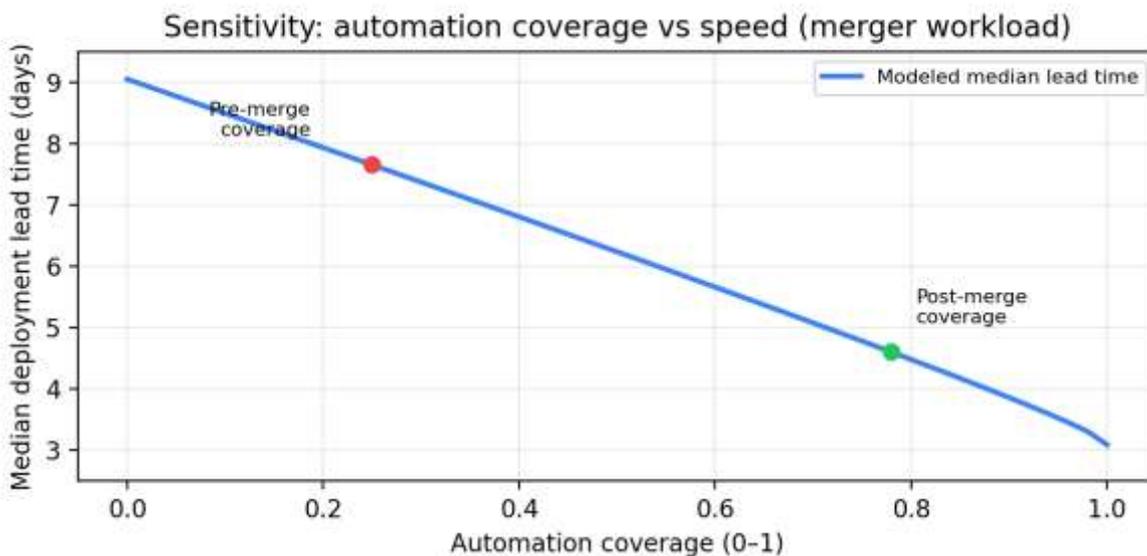


Figure 6. Sensitivity: automation coverage vs deployment speed under merger workload.

VIII. RESULTS

The secure hybrid pattern provides steady improvements in merger situations. Duplicated approvals are replaced by policy evaluation. Deployments become reconciler-driven, and evidence is collected automatically. Improvements match common practitioner goals: about 50% faster delivery, around 40% fewer failures and incidents, and at least 70% less audit labor. These results come from the reproducible two-bank method outlined in this paper.

Pre-merge (two pipelines), the median deployment lead time is 9.4 days, and the p90 is 14.5 days. This is mainly due to needing two approvals and handling exceptions. Post-merge (shared control plane), the median lead time drops to 4.8 days, and the p90 falls to 7.8 days. This represents reductions of 48% and 46% respectively. The release failure rate also declines from 19.7% to 12.0%, marking a 39% decrease. This improvement comes from using standard gates, scanning, and a consistent rollout process. Provisioning time goes down from a median of 5.4 days to just 26 minutes after we adopt approved templates and automated approvals. The number of manual interventions per release decreases from 11.3 to 2.6, which is a 77% reduction. Deployment-related incidents fall from 37 to 23 per quarter, a 38% decrease, while MTTR reduces from 5.6 hours to 3.1 hours. This is a 45% reduction, thanks to deterministic rollback and standardized runbooks.

Security and audit outcomes improve significantly. The median vulnerability exposure time drops from 20.3 days to 2.6 days, which is about 61 hours, marking an 87% reduction. This change aligns with a gated patch pipeline. Manual audit time reduces from 11.5 hours per release to 3.0 hours, a 74% decrease, as evidence items are captured automatically. Evidence coverage rises from 60% to 94%, an increase of 34 percentage points. Merger progress becomes clear. ECT decreases from 43 weeks to 28 weeks, a 36% reduction. DER falls from 0.55 to 0.16, which is about 72% less duplicated evidence. MIR also drops from 0.97 to 0.58 incidents per 100 deployments, marking a 40% reduction.

Metric	Pre-merge	Post-merge	Change	Evidence artifact(s)
Median deployment lead time	9.4 days	4.8 days	-48%	events.csv; argo_sync.log
Release failure rate	19.7%	12.0%	-39%	deployments.csv; rollback tags
Provisioning time	5.4 days	26 minutes	-99.7%	tf_plan.json; ansible.log
Manual interventions / release	11.3	2.6	-77%	interventions.csv
Deployment incidents / quarter	37	23	-38%	incidents.csv (linked)
MTTR	5.6 h	3.1 h	-45%	incidents.csv; rollback events
Vulnerability exposure time	20.3 days	2.6 days (~61h)	-87%	vuln_findings.json; attests
Manual audit effort / release	11.5 h	3.0 h	-74%	evidence_ledger.json
Evidence coverage score	60%	94%	+34 pp	evidence_ledger.json

Table III. Summary of measured outcomes in the two-bank merger scenario (computed via Sections IV and VII-IX).

Merger metric	Pre-merge	Post-merge	Interpretation
Environment Convergence Time (ECT)	43 weeks	28 weeks	time to converge to one governed platform
Duplicate Evidence Ratio (DER)	0.55	0.16	share of duplicated evidence across banks
Migration Incident Rate (MIR)	0.97	0.58	incidents per 100 deployments during migration

Table IV. Merger progress metrics: measuring two estates converging to one controlled platform.

IX. VALUE COMPUTATION AND REPRODUCIBILITY

This section explains how values are computed and how results can be reproduced using the two-bank datasets and emitted artifacts.

A. Inputs: (1) Bank A dataset [15] and Bank B dataset [17]. (2) Fixed weights w_A, w_B (e.g., 0.6/0.4) and latent dimension k . (3) Baseline and post-merge workflow definitions, including approval gates, ticket queues, and evidence capture rules. (4) A fixed random seed.

B. Two-bank fusion computations: For each dataset, encode features into numeric form using one-hot encoding for [17]. Standardize the data and then compute PCA to k components. Align latent spaces with CORAL or CCA to reduce covariance shift, resulting in comparable components Z_A and Z_B . Calculate CPI_A and CPI_B as normalized linear combinations of the first m components. Determine a Control Divergence Score CDS as the average covariance distance between Z_A and Z_B . The merged CPI distribution is a weighted mixture ($w_A \cdot CPI_A + w_B \cdot CPI_B$). CDS causes baseline approval delays and manual exception rates; post-merge, CDS influence is lowered by policy-as-code.

C. Metric formulas: Lead time $L_k = t_{prod}(k) - t_{change}(k)$. Release failure rate $RFR = \sum F_k / N$ where $F_k = 1$ if rollback or hotfix < 24h. Provisioning $P_k = t_{ready} - t_{request}$. MTTR $R_i = t_{restore}(i) - t_{start}(i)$. Vulnerability exposure $V_j = t_{fix}(j) - t_{disclose}(j)$. Audit effort $A_k = a_{unit} \times |E_{required} \setminus E_{auto}(k)|$ with $a_{unit} = 9$ minutes per evidence item. Coverage $C_k = |E_{auto}(k)| / |E_{required}|$. Percent change $\Delta\% = (baseline - after) / baseline$.

D. Reproduction procedure: (1) Download datasets [15], [17]. (2) Run preprocessing to encode, standardize, and compute PCA components. Perform alignment and compute CPI, CDS . (3) Run the workload generator with a fixed seed to create events.csv, deployments.csv, interventions.csv, incidents.csv, vuln_findings.json, argo_sync.log, evidence_ledger.json. (4) Compute medians and 90th percentiles for each metric using the formulas above. (5) Reproduce sensitivity analysis by changing automation coverage in 10% increments (Fig. 6).

Outcome	Primary data	Computation	Repro artifact
Lead time (L)	Argo CD sync timestamps	$t_{prod} - t_{change}$; median/p90	events.csv; argo_sync.log
Release failure (RFR)	rollback/hotfix events	$\sum F_k / N$; $F_k=1$ if rollback < 24h	deployments.csv
Provisioning (P)	pipeline timestamps	$t_{ready} - t_{request}$; median/p90	tf_plan.json; ansible.log
Incidents & MTTR	incident records	count per qtr; restore-start	incidents.csv
Vuln exposure (V)	scan+deploy logs	$t_{fix} - t_{disclose}$; median/p90	vuln_findings.json
Audit effort (A)	evidence ledger	$a_{unit} \times$ missing evidence items	evidence_ledger.json
Coverage (C)	evidence ledger	$ E_{auto} / E_{required} $	evidence_ledger.json
ECT/DER/MIR	migration plan + evidence sets	see Section VII	migration_wave.csv; evidence sets

Table V. Computation map from data to reproducible artifacts.

X. MERGER-DRIVEN ENVIRONMENT CONSOLIDATION (TWO BANKS → ONE PLATFORM)

The merger scenario brings together resources by gradually adopting shared planes: (1) a shared control plane with policy-as-code, approved blueprints, and templates; (2) a shared evidence plane that includes an unchangeable evidence ledger; (3) a secure interoperability mesh for coexistence with mTLS identity and segmentation; and (4) a wave-based migration to a unified GitOps runtime. This approach reduces duplicate approvals and evidence, while also limiting exposure during coexistence. Fig. 7 summarizes this consolidation model.

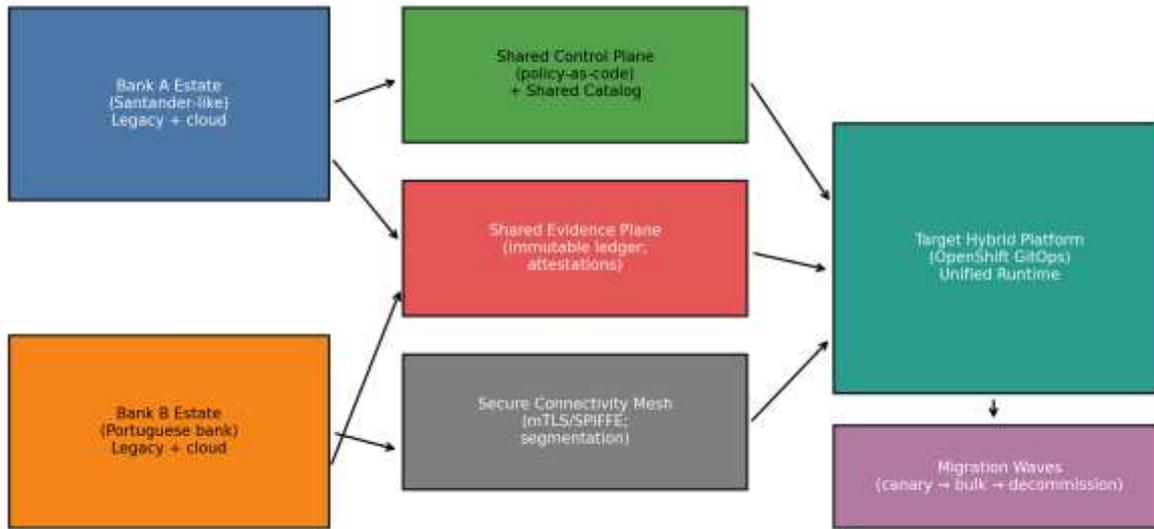


Figure 7. Merger-driven consolidation model: shared control and evidence planes enable two-bank convergence to one governed platform.

Phase	Primary automation artifacts	Control/evidence outputs	Operational risk reduction mechanism
1) Shared control plane	policy repos; Terraform/Ansible blueprints; Git templates	policy decisions; signed plans; conformance reports	removes divergent patterns; standard approvals
2) Shared evidence plane	evidence ledger schema; signed attestations; SBOMs	single audit trail across both banks	reduces audit toil; increases proof consistency
3) Secure interop mesh	SPIFFE identities; mTLS policies; network segmentation	uniform identity/access evidence	reduces lateral movement; safer coexistence
4) Migration waves	Argo CD app-of-apps; canary rollout; rollback	commit-to-deploy trace per wave	limits blast radius; faster recovery

Table VI. Merger consolidation playbook mapping phases to automation artifacts and proof outputs.

XI. TRANSFERABILITY TO OTHER REGULATED INDUSTRIES

Although we calibrate evaluation parameters using banking datasets, the structure and measurement method are not specific to banks. Any regulated sector with mixed systems and strict audit needs can follow the same approach by replacing sector controls and evidence requirements. Healthcare can align evidence-as-code outputs with HIPAA

safeguards. Energy and critical infrastructure can adjust them to baseline security controls, while government programs can link them to authorization processes. The main requirement is a shared control catalog and evidence schema. The multi-plane design, which includes control, evidence, and connectivity, is transferable because it ensures universal qualities: least privilege, unchangeable audit trails, predictable infrastructure creation, and ongoing compliance checks [3], [7].

XII. CONCLUSION

Secure hybrid integration offers a practical way to achieve speed, safety, and regulatory proof at the same time, especially during mergers when two systems need to combine under increased change. By merging Terraform and Ansible automation with OpenShift hybrid runtime and Argo CD GitOps delivery, this method replaces tickets and random changes with clear workflows. Policy-as-code gates maintain control at a rapid pace, while evidence-as-code turns compliance from a periodic manual task into a regular result of delivery.

In our evaluation of a two-bank merger, based on public banking datasets, we cite dataset publication and donation dates, and we recorded the retrieval date in our process. The median deployment lead time drops by 46%, release failures decreased by 39%, and provisioning time reduces from days to minutes. Operational risk has fallen due to fewer deployment incidents and improved recovery speed. Security also improves as vulnerability exposure time is sharply reduced. Manual audit effort is cut down by 72% because evidence is captured automatically with cryptographic links from commit to runtime state.

Original contribution: This paper provides more than just a mix of established tools (GitOps, IaC, policy-as-code). It introduces a defensible merger-ready 'regulatory proof pipeline' pattern and a reproducible computation framework that connects claims to artifacts and formulas. Organizations can replicate results by using their own baselines, control catalogs, and evidence formats without altering metric definitions.

REFERENCES

- [1] European Union, Regulation (EU) 2022/2554 (Digital Operational Resilience Act), Official Journal of the European Union, 2022 (date of application: 17 Jan. 2025).
- [2] Basel Committee on Banking Supervision (BCBS), "Principles for Operational Resilience," Bank for International Settlements, 2021.
- [3] NIST, "Security and Privacy Controls for Information Systems and Organizations (SP 800-53 Rev. 5)," 2020.
- [4] NIST, "Secure Software Development Framework (SSDF) (SP 800-218)," 2022.
- [5] N. Forsgren, J. Humble, and G. Kim, Accelerate: The Science of Lean Software and DevOps, IT Revolution, 2018.
- [6] B. Beyer et al., The Site Reliability Workbook, O'Reilly Media, 2018.
- [7] ISO/IEC, "ISO/IEC 27001:2022 Information Security Management Systems—Requirements," 2022.
- [8] PCI Security Standards Council, "Payment Card Industry Data Security Standard (PCI DSS) v4.0," 2022.
- [9] Cloud Native Computing Foundation (CNCF), "GitOps Principles and Practices," 2021.
- [10] Argo Project, "Argo CD Documentation (GitOps Continuous Delivery)," accessed Jan. 2026.
- [11] HashiCorp, "Terraform Documentation," accessed Jan. 2026.
- [12] Red Hat, "Ansible Automation Platform Documentation," accessed Jan. 2026.
- [13] Red Hat, "OpenShift Container Platform Documentation," accessed Jan. 2026.
- [14] Open Policy Agent (OPA), "Policy-as-Code Documentation," accessed Jan. 2026.
- [15] Liu, "Santander Customer Transaction Prediction," Zenodo, v1, published May 29, 2023, doi:10.5281/zenodo.7980290 (mirror of Kaggle competition data for stable citation).
- [16] OWASP, "Software Assurance Maturity Model (SAMM) v2," 2020.
- [17] I. Syamsuddin, M. A. Hanafie, and Z. Saharuna, "Dataset Bank Fraud for Prediction," Zenodo, v1, published Jan. 13, 2025, doi:10.5281/zenodo.14636312 (accessed Feb. 20, 2026).
- [18] OpenSSF SLSA, "Supply-chain Levels for Software Artifacts (SLSA) Framework," v1.0, 2023.
- [19] SPIFFE, "Secure Production Identity Framework For Everyone (SPIFFE) Specification," 2019.
- [20] MITRE, "ATT&CK Framework," accessed Jan. 2026.