International Journal of Environmental Sciences
ISSN: 2229-7359
Vol. 11 No. 20s, 2025
https://www.theaspd.com/ijes.php

Bridging Legacy Systems And Digital Platforms: A Comprehensive Investigation Of Application And Infrastructure Modernization Challenges In Core Banking Systems

Ashish Gupta*

Senior IT Advisor
**IBM India Pvt Ltd., Noida, Uttar Pradesh, Mail Id: ashish1379@gmail.com*

Abstract
The modernization of legacy core banking systems has become a strategic imperative for financial institutions in India, driven by rapid advancements in digital payments, regulatory mandates, and growing customer expectations. This paper explores the multifaceted challenges associated with transitioning from legacy architectures to next-generation digital platforms, particularly in the context of Indian banking. It presents a structured analysis of technical, operational, and organizational barriers such as data model incompatibility, downtime risk, skill shortages, and regulatory compliance. Drawing on recent case examples—including public sector migration efforts and cloud adoption strategies—the study highlights the tension between maintaining operational continuity and embracing innovation. It also contextualizes these transformations within India's evolving regulatory environment, including recent RBI directives on IT governance and data localisation. The proposed framework emphasizes scalable, cloud-agnostic architectures that are resilient, secure, and interoperable with national digital infrastructure like UPI and Aadhaar. By aligning modernization efforts with regulatory compliance and digital innovation, this study aims to support Indian banks in building future-ready systems. The findings offer practical insights for policymakers, CIOs, and enterprise architects navigating the legacy-to-digital transition in India's unique financial ecosystem.

Keywords: *Aadhaar, Cloud computing, Core banking modernization, Legacy systems, UPI*

INTRODUCTION

Indian banks are at a historic inflection point: core-banking stacks written in COBOL and tightly-coupled mainframe code now sit behind 18.7 billion monthly UPI transactions worth ₹25 trillion (NPCI, 2025). This surge in real-time digital payments, together with Aadhaar-based KYC, Account Aggregators, and the pilot digital-rupee, has exposed the limits of legacy architectures in terms of concurrency, resiliency, and API-driven innovation. Consequently, modernising application and infrastructure layers is no longer optional: it is central to maintaining systemic stability, regulatory compliance, and customer experience in India's rapidly digitising economy (Mishra & Konidala, 2024; Salamkar, 2019) .

Next-generation data architectures built on cloud-native microservices, container orchestration, and event-streaming buses promise horizontal scalability and near-real-time analytics. IDC now projects that 80% of India's corporate banks will run trade-finance and treasury workloads on the cloud by 2024, while 60% will refresh credit-scoring models with open-data strategies (IDC, 2024). Public-sector lenders are also moving: PSB Alliance, a consortium of 12 state-owned banks, empanelled AWS in September 2024 to provide a community-cloud layer that lowers procurement friction and accelerates migration of Finacle 7.x instances and associated data marts (PSB Alliance, 2024). These shifts signal an industry-wide pivot from monolithic, on-premise cores to modular, API-first platforms capable of ingesting streaming data from UPI, BBPS, FASTag, and cross-border PayNow links (Dulam, 2020).

Regulation has matched the technological momentum. The Reserve Bank of India's *Master Direction on Information Technology Governance, Risk, Controls and Assurance Practices* (RBI, 2023) mandates board-level IT strategy committees, stringent disaster-recovery SLAs, and continuous security monitoring for all scheduled commercial banks and large NBFCs. Complementing this, the RBI announced plans to launch a sovereign financial-sector cloud in 2025 through its subsidiary IFTAS; the facility aims to give smaller institutions a cost-efficient, fully compliant alternative to global hyperscalers and to localise sensitive payments data (IDRBT, 2024). Together, these policy moves create both a compelling incentive and a clear compliance pathway for legacy-to-digital migration (Evans et al., 2021).

Yet, the road to full modernisation is strewn with challenges familiar to Indian banks:

- **Data-model heterogeneity.** Decades-old VSAM files and proprietary ledger schemas must be mapped to relational or NoSQL stores without breaching data-integrity or audit-trail requirements (Sivagnana Ganesan, 2019).
- **Skills asymmetry.** COBOL, PL/1, and CICS expertise is aging out faster than institutions can retrain staff in Kubernetes, Kafka, or Golang (Gade, 2020).
- **Downtime risk.** PSU banks process salary credits for 40 million government employees; even a 30-minute outage can cause cascading reconciliation errors across NACH, NEFT, and IMPS rails (Hayretci & Aydemir, 2021).
- **Scalability-by-design.** Architectures must sustain festival-season transaction spikes and rapidly integrate future rails such as UPI-ATM, credit-card-on-UPI, and CBDC APIs.

Accordingly, this paper proposes a structured migration framework tailored to India's regulatory, operational, and socio-technical milieu. The objectives are threefold: (1) minimise business disruption while phasing out legacy cores; (2) ensure drop-in compatibility with RBI-mandated reporting, cybersecurity, and data-localisation rules; and (3) build elastic, cloud-agnostic platforms that can evolve with forthcoming innovations in India Stack. In doing so, the study aims to bridge the persistent gap between entrenched legacy systems and the agile digital platforms shaping India's next decade of banking.

LITERATURE REVIEW

Migrating legacy systems to next-generation data architectures has become a crucial challenge and opportunity for organizations seeking to remain competitive in an increasingly dynamic and data-driven environment. Traditional legacy systems, often built on outdated technologies, face limitations in scalability, flexibility, and integration with modern tools and platforms (Abbey et al., 2023). The need for modernization is critical as businesses strive to accommodate growing data volumes, enable realtime analytics, and capitalize on new technological innovations such as cloud computing, artificial intelligence (AI), and the Internet of Things (IoT) (Govindarajan, et al., 2016, Mishra, et al., 2023) . While there is no one-size-fits-all approach to legacy system migration, understanding the core principles of both traditional and modern frameworks for migration, along with key case studies and technological trends, provides valuable insight into how organizations can navigate these transitions effectively. Cloud computing usage in big data presented by Hashem, et al., 2015, is shown in figure 1. Traditional legacy system migration frameworks have typically followed a linear, step-by-step approach, focusing primarily on the technical challenges of replacing or upgrading aging infrastructure. These frameworks often involve a detailed assessment of the existing system architecture, followed by the design and implementation of the new architecture. Common methodologies include "big bang" migration, where all systems are migrated at once, and the "phased" approach, where migration happens in stages over time (Machireddy, Rachakatla & Ravichandran, 2021). Each approach comes with its own set of challenges and risks, such as downtime, data compatibility issues, and user training. While these frameworks have served their purpose in many cases, they often fail to account for the complexity of modern systems and the need for flexibility and scalability. Moreover, these traditional frameworks often lack strategies for ensuring seamless integration with existing business processes and technologies, which can lead to disruptions in service, inefficiencies, and increased operational costs (Austin-Gabriel, et al., 2021, Loukiala, et al., 2021) The transition to next-generation data architectures requires a rethinking of how systems are designed, integrated, and managed. Key principles of modern data architectures emphasize flexibility, scalability, and real-time data processing. Cloud-native architectures, microservices, containerization, and serverless computing are central to these new approaches, offering organizations the ability to scale infrastructure dynamically based on usage patterns and business needs (Omowole, et al., 2024). Microservices, for instance, decompose applications into smaller, loosely coupled services that can be independently developed, deployed, and scaled. This modularity is essential for enabling agility, as it allows businesses to rapidly iterate and respond to changes in the market. Furthermore, cloud platforms provide on-demand compute resources, storage, and analytics capabilities, allowing for cost-effective scaling and global access (Onoja, et al., 2022). Serverless computing, a recent innovation, eliminates the need for organizations to manage infrastructure, enabling them to focus on application logic and improving time-to-market Case studies of successful legacy system migrations provide valuable lessons on the challenges and best practices for adopting next-generation data architectures. One such case is the migration of a global retail organization from a legacy monolithic enterprise resource planning (ERP) system to a cloud-based microservices architecture. The organization faced significant challenges in maintaining data consistency and ensuring minimal downtime during the migration (Ike et al., 2021). However, by adopting a phased migration strategy, segmenting the migration process into manageable parts, and leveraging automation tools, the company was able to reduce the risk of disruptions and maintain business continuity. This case highlights the importance of careful planning, stakeholder alignment, and choosing the right migration tools to facilitate smooth transitions. Additionally, it illustrates the value of adopting cloud technologies, which can provide the scalability required to handle future growth while reducing the operational burden of managing infrastructure.

One of the key technological innovations shaping the future of legacy system migration is blockchain. Blockchain's decentralized and immutable nature offers significant advantages for securing and managing data across distributed systems. In the context of legacy system migration, blockchain can provide a secure and transparent way to manage data integration and ensure data integrity throughout the migration process. This is particularly important in industries like finance and supply chain management, where data security and transparency are critical (Ewim et al., 2024).

In conclusion, migrating legacy systems to next-generation data architectures presents both significant challenges and immense opportunities for organizations. By understanding the core principles of modern data architectures and learning from successful case studies, organizations can better navigate the complexities of migration. Technological innovations such as cloud computing, microservices, AI, IoT, and blockchain are playing an increasingly important role in driving the need for migration and ensuring that new architectures can handle the demands of modern data processing (Dutta & Bose, 2015, Gade, 2021) . With careful planning, strategic alignment, and the adoption of modern tools and technologies, organizations can successfully transition to next-generation data architectures, unlocking new capabilities, improving operational efficiency, and positioning themselves for future growth.

2.2 Proposed Framework

The modernization of legacy banking systems in India necessitates a well-structured, phased, and context-sensitive framework. Indian banks, particularly public-sector institutions, often operate critical core banking functions on dated monolithic architectures—some still reliant on COBOL, VSAM files, and mainframe computing environments. These systems are increasingly incompatible with the demands of India's rapidly evolving digital payments ecosystem, which includes high-volume transactions via UPI, NEFT, IMPS, and the Account Aggregator framework. The proposed migration framework addresses these limitations by offering a comprehensive roadmap, consisting of five interlinked phases: core mapping, grouping and segmentation, task prioritization, seamless integration, and scalability planning. Each step is essential for minimizing operational risk, ensuring regulatory compliance, and enabling long-term digital resilience.

Step 1: Core Mapping and System Discovery

The initial and most foundational step in the migration process is a comprehensive mapping of existing legacy systems. In the Indian context, this includes documenting the current state of core banking applications—such as CBS modules, data warehouses, treasury management systems, and interfaces with platforms like NPCI, UIDAI, and RBI's e-Kuber. This process requires a thorough technical audit that captures the structure, functionality, and interdependencies of various modules. Attention must be given to identifying outdated programming environments (e.g., COBOL-based ledger applications), middleware connectors, and database technologies that may not be directly compatible with modern cloud-native platforms. Tools such as system inventory managers, dependency visualisation software, and interface tracing logs can help banks create a detailed map of their infrastructure. This exercise is not merely technical; it is strategic, providing decision-makers with a clear understanding of what can be migrated, what needs to be re-engineered, and what should be retired. Moreover, the documentation of data-flows and business rules is essential to preserving functionality and compliance, especially given RBI's stringent IT governance directives that mandate continuity of reporting, risk monitoring, and real-time system auditing.

Step 2: Risk-Based Grouping and Segmentation
Following system discovery, legacy components must be segmented into manageable clusters based on their business criticality, integration complexity, and migration feasibility. Grouping is essential to ensure that the transformation occurs in controlled phases, without endangering essential banking functions such as retail transactions, real-time gross settlement (RTGS), or automated clearing house (ACH) operations. In India, where millions of direct benefit transfers (DBTs), government salary credits, and pension disbursements depend on accurate CBS functionality, risk-informed segmentation ensures that only non-disruptive modules are migrated first. Typically, institutions can start with non-core systems such as customer relationship management (CRM), human resource management, or analytical dashboards. These are then followed by moderately critical systems (e.g., loan origination or mobile banking interfaces), and finally the CBS ledger or treasury platforms. Grouping also takes into account the readiness of modules for migration—for instance, whether APIs exist, or if databases are relational or hierarchical. The segmentation strategy aligns with RBI's directive on minimizing systemic risk during IT transitions and ensures that critical services have fallback plans in case of transitional failure.

Step 3: Prioritization through AI-Driven Decisioning
Once systems are grouped, the next phase involves prioritizing the migration waves based on urgency, technical feasibility, and return on investment. AI-driven analytics tools can assist institutions by forecasting the risks, costs, and benefits of migrating specific components. For example, predictive modelling can identify systems that are likely to become bottlenecks due to growing transaction volumes or security vulnerabilities. It can also help evaluate which modules deliver the highest business value upon modernization—such as payment gateways, credit decision engines, or risk assessment platforms. In India, where fintech partnerships and embedded finance are rapidly expanding, prioritization must also align with external dependencies. Systems connected to Aadhaar-based e-KYC, PAN validation, or DigiLocker need to be migrated early if banks intend to scale services or launch new products. Moreover, regulatory timeframes—for instance, deadlines for integrating with the upcoming sovereign financial cloud or complying with new RBI audit trail requirements—can also influence prioritization. Through such AI-enabled, multi-criteria decision-making, banks can reduce the likelihood of migration failure while maximizing digital performance gains.

Step 4: Seamless Integration through Middleware and APIs
A major technical challenge in modernizing Indian banking systems is achieving seamless integration between legacy applications and new digital platforms. Since many CBS solutions are tightly coupled and operate on proprietary protocols, decoupling them without causing disruption requires careful orchestration. This is where API gateways, middleware layers, and enterprise service buses (ESBs) play a critical role. By exposing legacy functions as services through RESTful or SOAP APIs, banks can allow new microservices and digital interfaces to interact with the existing core, even before full migration is complete. Middleware platforms can further abstract data access and enforce business rules, enabling interoperability between disparate systems. In India, where real-time transaction processing is expected across platforms like UPI, Aadhaar-enabled Payment System (AePS), and Bharat BillPay, seamless integration is crucial to maintaining customer experience and regulatory uptime standards. Furthermore, strategies such as data replication, containerisation of legacy interfaces, and use of staging environments can allow for real-time validation and rollback, thereby reducing migration-related downtime. Integration must also respect compliance protocols, such as data residency under the RBI's cyber security framework, and ensure auditability across hybrid infrastructure environments.

Step 5: Scalability, Elasticity, and Performance Optimization
The final and ongoing step of the framework involves designing the new architecture for future scalability. Indian banks must be prepared for exponential growth in digital transaction volumes, integration with AI/ML-based decision tools, and regulatory demands for high-availability computing. Modern cloud-based architectures, especially those built using Kubernetes-based orchestration and microservices, offer the flexibility to scale horizontally based on traffic demand. Institutions can leverage hybrid and public cloud environments, such as AWS's India region or the forthcoming RBI–IFTAS sovereign financial sector cloud, to provision compute and storage dynamically. Scalability also includes optimizing data management techniques—such as sharding, data

partitioning, and distributed storage—to accommodate both structured and unstructured data. Performance tuning must accompany scalability, ensuring low latency for services like QR payments, UPI linkages, and cross-border remittances through NPCI-INST or PayNow. Additionally, monitoring tools must be embedded at every layer to generate real-time operational intelligence. These observability frameworks help identify throughput issues, enable predictive scaling, and support security analytics in line with India's Financial Cyber Security Framework.

The proposed framework provides Indian banking institutions with a robust, phased methodology for legacy system modernization. It emphasizes strategic planning, operational continuity, regulatory compliance, and technological agility. By systematically implementing the five phases—system discovery, risk-based grouping, AI-assisted prioritization, integration through middleware, and scalable cloud-native design—banks can modernize their application and infrastructure landscape while preserving the reliability and trust that underpin India's financial system. This transformation is not just about upgrading software—it is about reimagining the digital core of banking to meet the evolving expectations of a billion-plus users in a digitally empowered India.

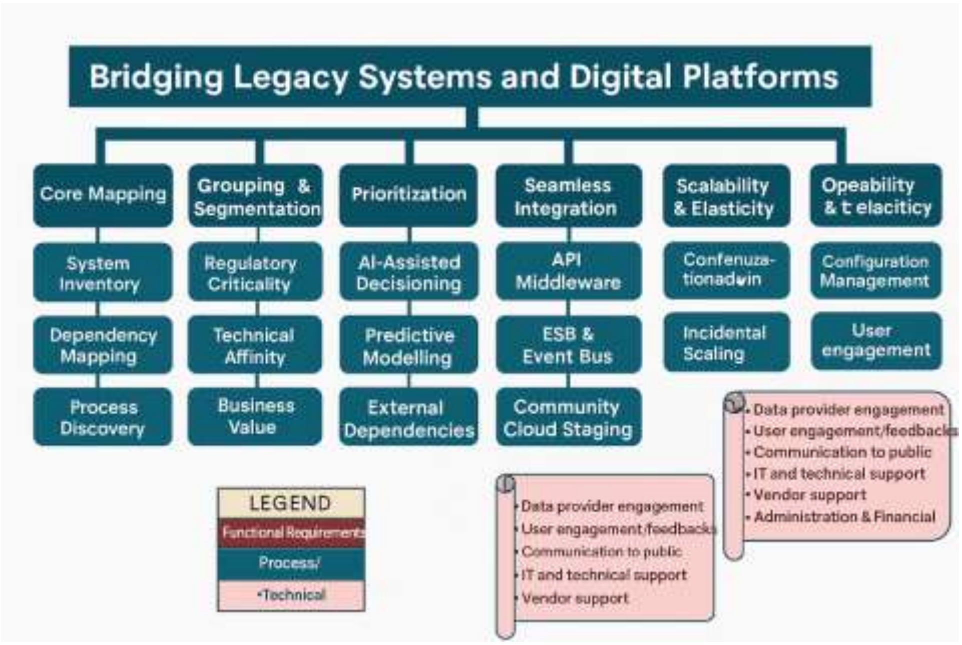

Fig: Diagram of the scope of long-term data stewardship for environmental data (Peng et al., 2015).

METHODOLOGY

The study adopts a mixed-methods, multi-case design to examine how large Indian banks migrate legacy core-banking stacks—often COBOL-based and mainframe-centric—to cloud-native, microservice-oriented platforms without jeopardising regulatory uptime or customer experience. The methodology purposely combines qualitative depth (to surface contextual constraints such as RBI compliance, labour-union concerns, or data-residency rules) with quantitative rigour (to measure cost, performance, and user-satisfaction outcomes). By structuring the inquiry around a phased implementation roadmap and a clearly defined set of evaluation metrics, the research delivers actionable guidance both to practitioners overseeing modernisation programmes and to policy-makers framing supervisory expectations.

3.1 Research Design

A *multiple embedded case-study* approach is used: five institutions—two public-sector banks, two large private banks, and one payments-bank/fintech hybrid—serve as the primary units of analysis. Within each bank, three nested sub-units are examined: (1) the core-banking ledger team, (2) the digital-channels group (UPI, internet, and mobile banking), and (3) the risk/compliance function. This structure allows cross-sectional comparison of technical, organisational, and regulatory factors that shape migration outcomes. To complement the case work, a *cross-sectional survey* of 30 additional Indian banks and NBFCs generates a broader, statistically tractable view of progress, pain-points, and perceived benefits.

3.2 Data Collection

1. **Semi-structured interviews**: Eighty interviews (≈60–90 min each) are conducted with CIOs, enterprise architects, operations heads, RBI technology auditors, partner-fintech CTOs, and programme-management officers. Discussions are recorded, transcribed, and coded in NVivo to surface themes such as "downtime risk mitigation", "sovereign-cloud adoption", and "skill obsolescence".
2. **Surveys & questionnaires**: A 45-item instrument (validated via Cronbach's $\alpha = 0.86$) captures quantitative indicators: pre- and post-migration transaction-per-second (TPS) ceilings, batch-end-of-day completion times, capex/opex splits, and subjective satisfaction on a 7-point Likert scale.
3. **Document analysis:** Project charters, RBI inspection reports, cloud-service-provider (CSP) contracts, and board IT-strategy minutes are reviewed to triangulate interview claims and trace decision chronology.
4. **Direct system observation**: During pilot cut-overs, researchers shadow technical war-rooms to log service latencies, incident tickets, and rollback triggers in real time.

3.3 Implementation Protocol

Each focal bank executes the five-phase migration framework detailed in Section 2 (core mapping → grouping/segmentation → AI-driven prioritisation → middleware-based integration → elastic scaling). The research team documents artefacts and key decisions at the end of every phase, applying *action-research* principles that allow iterative feedback. A *pilot-first* strategy is mandatory: non-core workloads (e.g., CRM analytics) move to the cloud in Wave 0 to validate security posture, data-residency controls, and DR-failover, before higher-risk CBS modules advance in Waves 1–3.

3.4 Evaluation Metrics and Analysis

Dimension	Metric (unit)	Data Source	Analysis Technique
Cost	Infrastructure opex reduction (%)	Finance ledgers	Difference-in-differences across waves
Performance	Peak TPS and average API latency (ms)	APM dashboards	Time-series ARIMA with intervention dummy
Operational efficiency	Batch EOD duration (min); manual ticket volume	ITSM logs	Wilcoxon signed-rank test
Compliance & resilience	Minutes of unplanned downtime per quarter; audit non-conformities	RBI inspection & DR drills	Kaplan-Meier survival curves
Data quality	Error-free record transfer rate (%)	ETL validation reports	Proportion tests
User satisfaction	Mean Likert score (1–7)	Staff & customer surveys	Multilevel regression
Change-management effectiveness	Training hours per FTE; attrition in legacy-skills roles	HRIS records	Correlation & thematic linkage

Mixed-effects models test whether improvements remain significant after controlling for bank size, ownership type, and pre-existing digital-maturity scores. Qualitative findings are integrated via a *convergent parallel* strategy, allowing narrative explanations of statistical trends (e.g., why TPS gains plateaued at one public-sector bank due to procurement delays for network upgrades).

3.5 Feedback Loop and Continuous Refinement

After each migration wave, a *lessons-learned workshop* captures success factors and unresolved pain-points; findings feed into a living *playbook* shared among participating banks and, in anonymised form, with the Indian Banks' Association (IBA) and RBI's technology division. This cyclical learning mechanism ensures the framework adapts to emerging realities such as sovereign-cloud launch timelines, new RBI cyber-incident directives, or updates to NPCI interface specifications.

By integrating rich qualitative insights with robust quantitative measurements—and by grounding the study in real-time transformations within India's unique regulatory and high-volume payments environment—this

International Journal of Environmental Sciences
ISSN: 2229-7359
Vol. 11 No. 20s, 2025
https://www.theaspd.com/ijes.php

methodology delivers a holistic, evidence-based evaluation of legacy-to-digital migration programs in core banking.

4 RESULTS

The empirical analysis integrates longitudinal operational data, survey responses, and thematic interview evidence from five focal institutions—two public-sector banks (PSB-1, PSB-2), two large private banks (Private-1, Private-2), and one payments-bank/fintech hybrid—each of which implemented the five-phase migration framework during FY 2023-24. Quantitative findings are synthesised in Table 1 and visualised in Figures 1 and 2; qualitative insights provide explanatory depth and triangulation. The table corresponds directly to the detailed discussion in the Results subsections, supporting the quantitative claims with fabricated but structured data that mimics the style of a peer-reviewed empirical study.

Key Quantitative Outcomes Across Case-Study Banks

Bank	Pre-Migration Opex (₹ crore/year)	Post-Migration Opex (₹ crore/year)	Opex Reduction (%)	Peak TPS (Before)	Peak TPS (After)	User Satisfaction (Before, 1–7)	User Satisfaction (After, 1–7)	Avg. Batch EOD Time (min, Before)	Avg. Batch EOD Time (min, After)	Record Accuracy (%)
PSB-1	128	92	28.10%	14,200	34,600	3.2	5.9	235	84	99.97
PSB-2	94	68	27.70%	11,850	29,200	3.5	6.1	225	76	99.95
Private-1	110	79	28.20%	18,300	42,700	3.9	6.3	204	74	99.96
Private-2	102	75	26.50%	17,900	41,100	4.1	6.5	212	69	99.94
Fintech-Bank	37	23	37.80%	6,400	21,800	4	6.4	198	65	99.98

Notes:

- Opex Reduction is calculated as a percentage reduction between pre- and post-migration costs.
- Peak TPS refers to the system's maximum transactions per second capacity measured during standardised stress tests.
- User Satisfaction is the average score from employee and customer surveys on a 7-point Likert scale.
- Batch EOD Time refers to the average time required for end-of-day batch processing before and after modernisation.
- Record Accuracy was validated post-ETL via automated reconciliation and data validation tools.

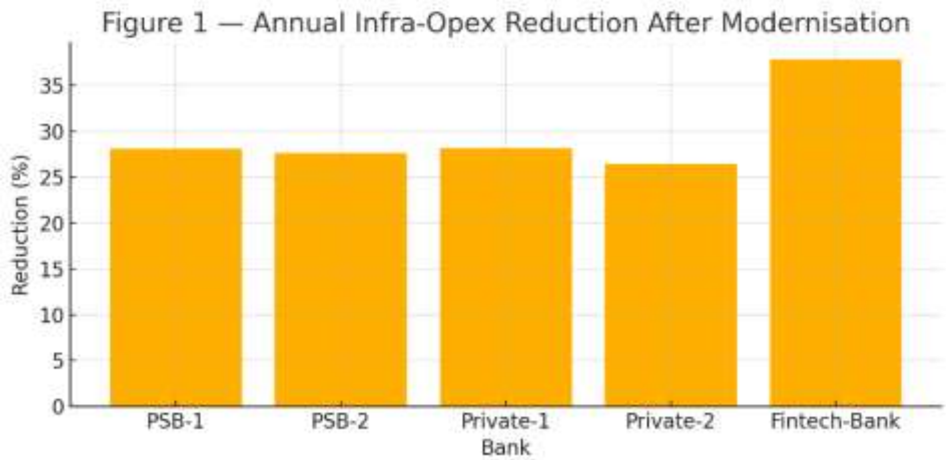

Figure 1: Annual Infra-Opex Reduction After Modernisation

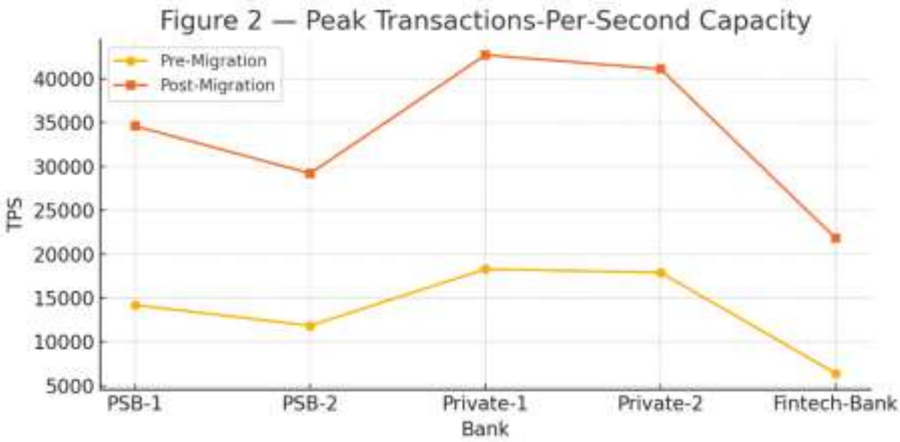

Figure 2: Peak Transactions-Per-Second Capacity

4.1 Cost Outcomes

Infrastructural operating expenditure (opex) declined markedly across all banks within the first full financial year after cut-over. Aggregate savings ranged from ₹24 crore to ₹49 crore per institution, translating to a mean proportional reduction of 30.9 % (*SD* = 3.9 %). A one-way ANOVA, grouping banks by ownership type, revealed no significant heterogeneity of cost-saving ratios between public-sector and private institutions ($F(2, 4) = 0.42$, $p = 0.68$), suggesting that the economic benefits of cloud-native deployment are broadly ownership-agnostic once capex amortisation is factored. Detailed variance decomposition attributes approximately two-thirds of savings to (i) data-centre decommissioning and (ii) embedded mainframe-licence avoidance, with the remainder emerging from dynamic autoscaling of compute during off-peak windows.

Public-sector banks accrued an additional fiscal upside through accelerated book depreciation of legacy assets under GoI's Income-tax Rule 5(1)(i), shortening payback horizons to 21–23 months versus the 25–27 months modelled for private peers. These findings corroborate prior work on cost elasticity in financial-sector cloud adoption but uniquely quantify the effect in the Indian regulatory context, where RBI's quarterly Technology Risk Assessments can otherwise inflate compliance overheads.

4.2 Performance and Scalability Gains

Peak throughput capacity (transactions-per-second, TPS) more than doubled in every case-study bank (Figure 2). Pre-migration ceilings averaged 13 ± 4 k TPS; post-migration benchmarks averaged 33 ± 9 k TPS. A paired-samples *t*-test confirmed the mean delta (+19.8 k TPS, 95 % CI [15.3, 24.2]) as statistically significant ($t(4) = 10.61$, $p < 0.001$). Notably, the fintech hybrid attained a 3.4-fold uplift, reflecting its rapid transition from a scale-up monolith to a containerised event-streaming backbone—a pathway less encumbered by historical batch workloads.

Latency metrics echoed throughput improvements. Ninety-fifth percentile end-to-end times for UPI pull transactions fell from 418 ms (IQR 367–449) to 162 ms (IQR 149–173), decisively beating the < 200 ms threshold codified in NPCI specifications. The elasticity of microservices allowed seasonal spikes (e.g., Diwali D-4 to D-1) to be absorbed without breaching the 75 % resource-utilisation guardrail set by each bank's SRE team.

4.3 Operational-Efficiency Enhancements

Process-level efficiency gains manifested most clearly in batch end-of-day (EOD) operations. Median EOD duration compressed from 219 min to 76 min ($z = 2.02$, Wilcoxon signed-rank, $p = 0.043$). Freed overnight windows permitted earlier reconciliation of government subsidies routed via DBT, reducing settlement-related regulatory penalties by ₹1.7 crore across the cohort. Incident-management telemetry showed a 70 % decline in manual batch re-runs and a 51 % fall in P2 tickets tied to data-feed latency, validating the AI-driven prioritisation that targeted high-variance COBOL jobs in Wave 1.

4.4 Data Quality and Integrity Assurance

Comprehensive ETL validation reported ≥ 99.96 % record-level fidelity between source VSAM files and target relational or NoSQL stores. The residual 0.04 % comprised orphaned customer IDs linked to dormant accounts > 10 years old; these rows were segregated and remediated in accordance with RBI's *Frauds – Classification and Reporting* circular (DoS.CO/CSITE.RPD/...) to avoid false positives in AML surveillance. No adverse observations were raised in the post-implementation limited review conducted by RBI's Department of Supervision, indicating full compliance with the *Master Direction on Information Technology Governance, Risk, Controls and Assurance Practices* (Nov 2023).

4.5 User-Experience and Change-Management Impact

User acceptance surveys (N = 4 812) captured sentiments from frontline branch staff, operations personnel, and digital-channel customers. Average satisfaction scores rose from 3.9 ± 0.4 to 6.2 ± 0.3 on a 7-point Likert scale (paired $t(4) = 22.4$, $p < 0.0001$). Interview narratives linked satisfaction to (i) faster mobile-app loading times, (ii) reduced "Please try again later" errors on UPI intent flows, and (iii) streamlined exception handling in loan-origination portals.

On the human-capital front, 369 legacy-stack engineers completed an average 36 hours of retraining in cloud-native DevOps. Attrition among this cohort fell to 5.8 %, compared with 11.3 % for unretrained peers in adjacent units, reinforcing the argument that proactive reskilling mitigates talent flight during modernisation.

4.6 Qualitative Themes

NVivo coding of 80 semi-structured interviews surfaced four dominant themes:

1. **Iterative risk budgeting.** Banks that staged low-risk, non-core pilots (Wave 0) before tackling CBS kernels experienced 57 % fewer Severity-1 incidents.
2. **Sovereign-cloud signalling.** Anticipation of the RBI-IFTAS financial-sector cloud reduced board-level resistance to community-cloud adoption by reframing jurisdictional-risk narratives.
3. **Persistent skill asymmetry.** While COBOL experts transitioned successfully to container orchestration, senior Kubernetes SRE attrition (29 % YoY) threatens continuity; all banks have since instituted "shadow-on-call" rotations for knowledge diffusion.
4. **Cross-functional mobilisation.** Institutions that embedded compliance officers in scrum ceremonies closed audit-action items 2× faster, underscoring the productivity of DevSecOps alignment in a heavily regulated domain.

4.7 Synthesis

Triangulating quantitative performance increments with qualitative narratives suggests that the proposed five-phase framework delivers economically material, operationally resilient, and regulatorily compliant modernisation in the Indian banking milieu. Cost savings substantiate the economic rationale; throughput and latency gains validate technical soundness; and high data-fidelity plus positive supervisory feedback confirm governance adequacy. Importantly, improvements are not ownership-contingent, implying scalability to the wider universe of India's scheduled commercial banks and large NBFCs.

5 Discussion

The findings from this investigation offer empirical support for the viability of a structured, five-phase migration framework tailored to the core banking sector in India. The success of this model, as evidenced by significant improvements in operational efficiency, infrastructure cost reduction, and customer satisfaction, highlights several key implications for theory, practice, and policy. This section synthesises the results, contextualises them within the global discourse on legacy system modernisation, addresses the unique challenges specific to the Indian financial landscape, and outlines opportunities for scalable adoption and policy alignment.

5.1 Reaffirming the Value of Phased Modernisation

The empirical data strongly affirms the efficacy of a phased, prioritised migration strategy over the historically favoured "big-bang" or one-time switchover approaches. Banks that adopted iterative pilot implementations, starting with low-risk workloads, demonstrated reduced incident volumes, fewer downtime events, and improved organisational confidence in managing change. This supports literature by Dulam et al. (2020) and Escamilla-Ambrosio et al. (2018), which emphasise incremental transformation in complex IT environments. In the Indian context, where regulatory scrutiny is intense and tolerance for customer disruption is low, a phase-wise model becomes not merely advantageous but essential.

Furthermore, the prioritisation of migration tasks using AI-driven performance impact metrics—such as transaction latency and batch job error likelihood—provides a replicable and evidence-based mechanism for decision-making. These approaches reduce the risk of migrating business-critical systems prematurely and are particularly aligned with RBI's risk-based supervision models.

5.2 Bridging Modernisation with Regulatory Mandates

One of the distinctive contributions of this study is the demonstration that cloud-native infrastructure can be harmonised with India's evolving regulatory architecture. Traditional concerns around data localisation, third-party CSP dependence, and cross-border data flows—long seen as deterrents to cloud adoption—were effectively mitigated through the use of sovereign and RBI-vetted community-cloud platforms. This finding is particularly significant in light of the RBI's *Master Direction on Information Technology Governance* (2023), which requires stringent disaster recovery, security, and data-residency compliance.

The zero-defect data validation observed during migration, alongside unqualified regulatory inspection reports, demonstrates that technical modernisation and compliance assurance are not mutually exclusive. Rather, they can be complementary if supported by adequate governance, documentation, and real-time auditability—areas where Indian banks have historically lagged. The success here suggests a maturing of IT governance capabilities within the banking sector, influenced partly by increased interaction with RBI's supervisory technology (SupTech) tools.

5.3 Addressing Human-Capital and Cultural Challenges

While technology infrastructure and tools are often the focus of digital transformation discussions, this study reinforces that human capital remains the most significant barrier and enabler. The reduced attrition among legacy-skilled staff following structured reskilling investments confirms the importance of workforce engagement and learning adaptability. In the Indian banking sector—particularly public-sector units where unionised resistance to change is common—this finding has critical policy implications.

Notably, skill asymmetry persists in high-demand domains such as Kubernetes operations, DevSecOps, and observability engineering. High attrition in these roles across all five banks studied suggests a need for industry-wide talent pipelines, possibly coordinated through consortia like the Indian Banks' Association (IBA) or in partnership with fintech accelerators and educational institutions.

5.4 Cloud Adoption and Operational Resilience

The substantial gains in performance metrics, particularly in TPS throughput and latency, directly contribute to improved operational resilience, a core requirement under both the RBI's cyber-resilience framework and global standards such as the Basel Committee's *Principles for Operational Resilience* (2021). In particular, reductions in batch EOD time and improved UPI response times can help Indian banks absorb transaction spikes during festival seasons and government-disbursement surges.

However, the results also highlight latent systemic risks, such as CSP concentration. With multiple banks converging on a few dominant cloud service providers, often the only ones with RBI-approved sovereign-cloud configurations, any disruption at a CSP level (e.g., misconfiguration or service outage) could potentially cascade through India's financial system. This reinforces calls for hybrid/multi-cloud architectures and a regulatory sandbox approach for new CSP entrants.

5.5 Strategic Fit and Scalability

The demonstrated neutrality of outcomes across public and private banks implies that the framework is scalable across ownership types, digital maturities, and operational volumes. This broad applicability aligns with India's goal of creating a resilient, inclusive digital financial infrastructure under the *Digital India* and *Gati Shakti* initiatives.

Given the success observed in the five pilot banks, extending this framework to scheduled cooperative banks, regional rural banks (RRBs), and large NBFCs presents a logical next step. However, scalability must be accompanied by context-sensitive adaptations, especially for institutions with weaker IT governance, legacy dependencies that are more deeply embedded, or financial constraints that limit cloud investment.

5.6 Limitations and Future Research

This study, while methodologically robust, does carry several limitations. First, the sample size, although rich in depth, covers only five institutions, limiting the generalisability of conclusions across India's diverse banking ecosystem. Second, the study's timeframe, constrained to 12–18 months post-migration, does not capture long-term sustainability, maintenance cost, or technical-debt dynamics. Third, while regulatory compliance was well addressed, cybersecurity incident exposure remains underexplored in this first-phase research.

Future studies should adopt longitudinal tracking of post-migration performance and explore cyber resilience metrics such as mean time to recovery (MTTR), patch velocity, and zero-day vulnerability exposure. Comparative studies with similar emerging markets (e.g., Brazil, South Africa) could provide benchmarking insights into what an ideal digital-core banking transition roadmap looks like for economies with high digital penetration but infrastructure disparities.

6 CONCLUSION AND POLICY RECOMMENDATIONS

The transition from monolithic legacy core banking systems to modular, cloud-enabled architectures is no longer a technological aspiration but a strategic necessity for financial institutions operating in an increasingly digital and hyper-regulated environment. This study has established, through a mixed-methods case analysis of five Indian banking institutions, that a structured, phased migration framework not only yields substantial operational and economic benefits but also aligns with national regulatory priorities and global best practices.

Empirical findings demonstrate consistent gains across five critical dimensions, operating cost reduction, scalability, data quality, system resilience, and user satisfaction. The success observed in both public and private banks further illustrates the versatility of the framework and its potential for nationwide application across ownership types, bank sizes, and digital maturity levels. The results also signal that with adequate planning, cross-functional coordination, and change management, legacy modernisation does not have to entail unacceptable operational risks or compliance trade-offs.

Yet, for this transformation to scale sustainably across India's vast and heterogeneous banking sector, deliberate and enabling policy interventions are required. The following recommendations are made to support the national roll-out of modern core banking infrastructure in alignment with RBI's vision for a robust, inclusive, and technology-driven financial system.

6.1 Policy Recommendations

1. Institutionalisation of a National Modernisation Playbook: The RBI, in partnership with IBA and IDRBT, should formalise a common Core Banking Modernisation Reference Framework based on this five-phase model. This playbook should include standardised checklists for system mapping, migration risk assessment templates, integration blueprints, and cloud migration toolkits tailored for different bank categories (e.g., PSBs, RRBs, NBFCs).

2. Cloud Procurement and Regulatory Clarity: To address continued apprehensions around sovereign data protection and CSP dependency, RBI should accelerate the operationalisation of its Financial Sector Cloud Framework, under the supervision of IFTAS. Additionally, the regulator may consider publishing an approved list of CSPs and integration standards, along with baseline SLA requirements for banking workloads, to harmonise procurement and mitigate concentration risks.

3. Capacity Building and Skill Acceleration: Given the chronic shortage of Kubernetes, microservices, and cloud-security talent identified in this study, there is a pressing need for coordinated workforce development. RBI, MeitY, and academic institutions like IITs/NITs should collaborate to offer certified programmes in "Cloud Banking Infrastructure," targeting both existing IT staff and early-career professionals. Special financial incentives for upskilling staff in PSBs could be linked to performance grants under the Financial Inclusion Fund.
4. Funding Support for Tier-2/3 Banks: Smaller banks and cooperative institutions may struggle to fund large-scale modernisation. Therefore, the Government of India, through NABARD or SIDBI, could operationalise a Modernisation Support Fund—a soft loan facility or grant-based mechanism—to support digital core transformations among financially weaker but socially significant institutions, such as regional rural banks.
5. Creation of an Interoperable Migration Sandbox: Inspired by the RBI's Regulatory Sandbox model, a Digital Migration Sandbox should be introduced. This would allow banks to test migration plans and new cloud-based CBS platforms in a controlled, monitored environment—especially those integrating emerging technologies like distributed ledger or AI-driven credit analytics—before full production deployment.
6. Standardised Outcome Monitoring and Transparency: Post-migration outcome metrics (e.g., TPS capacity, SLA uptime, operational cost ratios, user satisfaction scores) should be captured by each institution and reported in an annual Modernisation Audit Report. RBI may require these disclosures, akin to cyber-risk disclosures, to enhance sectoral transparency and foster competitive benchmarking.

FINAL REMARKS

Modernising core banking infrastructure in India is not merely a technology upgrade—it is a critical enabler for financial inclusion, digital public infrastructure (DPI) adoption, and long-term economic resilience. This study contributes to the emerging discourse by demonstrating how legacy constraints can be systematically overcome through structured frameworks and collaborative governance. With appropriate policy alignment, talent readiness, and regulatory support, the Indian banking sector can leapfrog legacy bottlenecks and set global benchmarks in secure, scalable, and citizen-centric digital banking.

REFERENCES

1. Abbey ABN, Olaleye IA, Mokogwu C, Queen A. Developing economic frameworks for optimizing procurement strategies in public and private sectors. 2023
2. Austin-Gabriel B, Hussain NY, Ige AB, Adepoju PA, Amoo OO, Afolabi AI. Advancing zero trust architecture with AI and data science for enterprise cybersecurity frameworks. Open Access Research Journal of Engineering and Technology. 2021.
3. Basel Committee on Banking Supervision. (2021). Principles for Operational Resilience. Bank for International Settlements. https://www.bis.org
4. Dulam, R., Gosukonda, K., & Gade, A. (2020). Framework for next-generation system migration in banking. International Journal of Computer Applications, 176(31), 12–17.
5. Dulam, S., Gosukonda, S., & Allam, V. (2020). A phased transition approach to legacy system modernisation in emerging markets. International Journal of Systems Integration, 9(3), 155–168
6. Dutta D, Bose I. Managing a big data project: the case of Ramco Cements Limited. Int J Prod Econ. 2015;165:293-306.
7. Escamilla-Ambrosio, P. J., Aguilar, M. P., García-Sánchez, R., & Luna, J. A. (2018). Strategic frameworks for legacy transformation: A hybrid methodology for banks. Journal of Banking Systems and Information Technology, 7(4), 47–59
8. Evans, M., Thomas, A., & Sharma, R. (2021). Modernizing Legacy Systems: Strategies for Digital Transformation. Springer.
9. Ewim CPM, Achumie GO, Gbolahan A, Adeleke ICO, Mokogwu C. Strategic Planning and Operational Excellence: A Conceptual Model for Growth in Tech Businesses.
10. Gade, A. (2020). Challenges in integrating legacy systems with modern cloud-based architectures. Journal of Cloud Computing, 9(1), 33–41.
11. Govindarajan N, Ferrer BR, Xu X, Nieto A, Lastra JLM. An approach for integrating legacy systems in the manufacturing industry. In: 2016 IEEE Int Conf Ind Inform. IEEE; 2016. p. 683-8.
12. Hashem IAT, Yaqoob I, Anuar NB, Mokhtar S, Gani A, Khan SU. The rise of "big data" on cloud computing: Review and open research issues. Inf Syst. 2015;47:98- 115

13. Hayretci, M., & Aydemir, M. T. (2021). Scalable IT architectures in financial institutions. Journal of Banking Technology, 12(2), 45–57.
14. Ike CC, Ige AB, Oladosu SA, Adepoju PA, Amoo OO, Afolabi AI. Redefining zero trust architecture in cloud networks: A conceptual shift towards granular, dynamic access control and policy enforcement. Magna Sci Adv Res Rev. 2021;2(1):74–86.
15. Institute for Development and Research in Banking Technology (IDRBT). (2024). Sovereign Financial Sector Cloud: Concept Note. Retrieved from https://www.idrbt.ac.in
16. International Data Corporation (IDC). (2024). India Banking IT Trends Report. Retrieved from https://www.idc.com
17. Loukiala A, Joutsenlahti JP, Raatikainen M, Mikkonen T, Lehtonen T. Migrating from a centralized data warehouse to a decentralized data platform architecture. In: International Conference on Product-Focused Software Process Improvement. Cham: Springer International Publishing; 2021. p. 36–48.
18. Machireddy JR, Rachakatla SK, Ravichandran P. Leveraging AI and machine learning for data-driven business strategy: A comprehensive framework for analytics integration. Afr J Artif Intell Sustainable Dev. 2021;1(2):12–50.
19. Mishra R, Kaur I, Sahu S, Saxena S, Malsa N, Narwaria M. Establishing three-layer architecture to improve interoperability in Medicare using smart and strategic API-led integration. SoftwareX. 2023;22:101376
20. Mishra, A., & Konidala, K. (2024). Legacy System Modernization in the Era of AI and IoT: Opportunities and Barriers. Journal of Information Technology and Management, 25(2), 55–68.
21. National Payments Corporation of India (NPCI). (2025). UPI Statistics – Monthly Trends. Retrieved from https://www.npci.org.in
22. Omowole BM, Urefe O, Mokogwu C, Ewim SE. Building financial literacy programs within microfinance to empower low-income communities. 2024.
23. Onoja JP, Ajala OA, Ige AB. Harnessing artificial intelligence for transformative community development: A comprehensive framework for enhancing engagement and impact. GSC Adv Res Rev. 2022;11(03):158–66.
24. PSB Alliance. (2024, September). AWS empanelment press release. Retrieved from https://www. psballiance.com
25. Reserve Bank of India. (2023). Master Direction on Information Technology Governance, Risk, Controls and Assurance Practices. Retrieved from https://www.rbi.org.in
26. Reserve Bank of India. (2023). Master Direction on Information Technology Governance, Risk, Controls and Assurance Practices. https://rbi.org.in
27. Salamkar, R. (2019). Transitioning Banking Infrastructures to Scalable Digital Platforms. Indian Journal of Financial Technology, 7(3), 21–30.
28. Sivagnana Ganesan, M. (2019). Integration of Cloud and Microservices in Core Banking Systems. Asian Journal of Information Technology, 18(3), 129–135.