International Journal of Computational and Experimental Science and ENgineering (IJCESEN)

Vol. 12-No.1 (2026) pp. 1414-1420
http://www.ijcesen.com

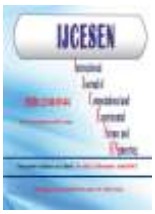

ISSN: 2149-9144

Research Article

Mainframe Zero Strategy: Eliminating Legacy Dependencies through Cloud-Native Data Platforms

Naga Malleswara Babu Velpuri*

Independent Researcher, USA
* **Corresponding Author Email:** naga.velpuri9@gmail.com - **ORCID:** 0000-0002-5007-2277

Article Info:

DOI: 10.22399/ijcesen.5047
Received : 11 January 2026
Revised : 01 March 2026
Accepted : 08 March 2026

Keywords

Mainframe Modernization,
Cloud-Native Architecture,
Legacy System Transformation,
Change Data Capture,
Strangler Fig Pattern

Abstract:

Critical business processes in financial services, retail, healthcare, and government sectors have decades of experience and processing of significant commercial volumes of transactions on enterprise mainframe systems and have been entrusted with the custody of sensitive organizational information. Nevertheless, these old platforms create major limitations to organizational agility, integration capacities, and responsibilities to changing market requirements despite their transferable reliability and security features. Targeting both the technical change and the organizational change, governance restructuring, and the development of workforce capabilities, the Mainframe Zero (MFZ) strategy introduces a systematic approach to the systematic removal of mainframe dependencies by migrating to cloud-native data platforms. Frameworks include discovery processes to catalog the workloads and business dependencies, risk-based phasing to order activities during migration by business value and technical complexity, data migration with both bulk historical backfill and change data capture frameworks, and application refactoring with the use of strangler fig patterns to break down the project incrementally. Target-state architectures use managed cloud databases, such as BigQuery to support analytics workloads, Cloud Spanner to support transactions where the world is one, and event-driven messaging infrastructure to replace batch-based processing models. Security and compliance integration integrates identity management, encryption, network controls, and monitoring features in line with the NIST Cybersecurity Framework across the execution of transformation. Organizational enablement focuses on domain-related team architecture, platform engineering competencies, federated governance frameworks, and communication mechanisms that maintain stakeholder trust in the course of long-term move programs. An example of a case study has shown quantifiable results such as cost reduction of infrastructures by eradicating mainframe licensing, better recovery targets by utilization of multi-region deployment structures, and higher velocity of feature delivery by contemporary development methods. Decision frameworks help technology leaders to choose the suitable migration patterns in terms of workload characteristics, strategic goals, and risk tolerance, and prioritize the transformation domains and control the risk accordingly during cloud-native modernization programs.

1. Introduction

Enterprise mainframe systems are still the foundation of the most vital business processes in financial services, retail, healthcare, and government sectors in the global arena. These systems have proved to be highly persistent, and they handle most of the business deals in the world, and they hold sensitive organizational information that has taken decades to produce. A study of mainframe infrastructure points out that the systems have tremendous transaction processing power, data integrity controls, and data security functionality that have enabled them to become invaluable to mission-critical processes that need unquestionable reliability [1]. Nevertheless, the same features that make it stable also constrain the ability of the organization to be agile and integrate with new cloud services and react quickly to changing market demand.

The Mainframe Zero (MFZ) approach is an overall approach to the systematic dissolution of the

mainframe reliance by transferring it to the cloud-native information infrastructures. The strategy considers that mainframe modernization does not just imply technical change but also organizational change, governance reorganization, as well as development in workforce capability. Empirical studies of the modernization of legacy systems have also pointed out that effective transformation efforts must be organized in ways that can deal with both technical complexity and human and organisational issues, which often determine a project's success or failure [2]. There is a much larger rate of failures in companies that engage in unorganized modernization activities than in companies that use disciplined methodologies, with well-defined phases, gates, and validation conditions.

This paper defines MFZ as a repeatable model including discovery, risk-based phasing, data migration, application refactoring, security integration, and organizational enablement. Later paragraphs provide practical recommendations to enterprise architects and technology executives who go through mainframe retirement efforts, offer target-state architecture, decision models, and summaries of success stories with real results.

2. Discovery and Risk-Based Phasing Methodology

The first step in the effective mainframe retirement process is an extensive process of discovery, which catalogs the current workloads, charts the business process interdependencies, and forms a baseline knowledge of the current operational environment in need of transformation. The discovery initiatives should deal with the high complexity that is in mainframe environments that have grown across decades and experienced accumulating layers of functionality, integration points, and undocumented dependencies. Empirical studies of the modernization of COBOL-based systems have shown that in most cases, enterprise mainframe systems have very large codebases, built over long durations, and big chunks of the system have since been lost in organizations due to retirement and attrition [3]. Such a knowledge gap poses a significant danger to modernization efforts that are undertaken without proper discovery.

The discovery stage uses both automated scanning systems and manual analysis methods to list application components, database schemas, data flows, and integration interfaces. Compiler tools read the source code and extract program structure, data references, and inter-program communications, building dependency graphs that present dependencies between applications in the portfolio. Manual analysis is used to supplement automated discovery by including business stakeholders knowledgeable of functional requirements, operational procedures, and other exceptional handling procedures that might not be clear in a code study. Important deliverables are detailed application inventories that record technical attributes and business environment, such as ownership, criticality level, regulatory provisions, and business alignment.

Risk-based phasing contextualizes discovery knowledge into migration roadmaps, which ranked workload transformation in a balanced approach to business value, technical complexity, and operational risk. Good phasing strategies acknowledge that any transformation initiative should keep the business running, but gradually phase out of the old systems, and this change should be well orchestrated and must not disrupt the interdependent systems. Studies that have studied management of organizational change in technological shifts focus on gradual processes that enable demonstration of success at an early stage, develop confidence in stakeholders and organizational momentum that may support long-term undertakings [4]. Companies that strive to achieve global simultaneous change often run out of resources, exhaust the patience of stakeholders, and add to technical issues that set back progress.The first stages of migration are normally directed towards the analytics and reporting loads that are not as complex as the transactional systems and have lower operational risk. Business intelligence systems, enterprise data warehouses, and decision support systems tend to use batch update cycles and allow a limited number of brief inconsistency that makes migration implementation easier. These workloads create tangible business value by increasing the level of analytics, and transformation benefits could be seen in these workloads, which would justify future investment in the next step. After analytics migration, organisations will move to operational systems such as inventory management, supply chain planning, and customer relationship platforms that are moderately complex with well-established validation processes.

Last migration stages work with transactional core systems, where consistency requirements, latency sensitivity, and business criticality create extremely rigid requirements of success. Order management, financial ledgers, and payment process systems demand broad parallel-run validation, thorough rollback facilities, and detailed cutover planning. Risk-based phasing makes sure that risky migrations are preceded by the team that has gained experience, perfected their playbooks, and

developed organizational confidence through the successful previous phases.

3. Data Migration Methodology and Quality Assurance

Data migration is the workstream basis that allows mainframe retirement to take place, and it is the workstream that has to be carefully planned and executed to maintain data integrity between the heterogeneous source and target environments. Typically, enterprise mainframe systems are systems of record that have decades of transactional history, customer relationships, financial records, and operational data that need to be transferred correctly to cloud-native systems. The literature on enterprise information management highlights that the quality of data has a direct influence on organizational decision-making, operational efficiency, and regulatory adherence, and therefore, the accuracy of migration is critical to success in the transformation process [5]. Before starting the migration process, organizations should have elaborate data governance frameworks, quality standards, validation processes, and remediation processes.

The MFZ methodology involves a single migration process that is based on bulk historical backfill and change data capture mechanisms that synchronize migration processes across long periods of migration windows. Bulk backfill takes historical data out of mainframe databases such as DB2, IMS hierarchical structure, and VSAM file systems, and stages the extracted data in the cloud storage to be transformed and loaded. The extraction processes should be able to handle mainframe-specific data formats such as packed decimal numeric formats, EBCDIC character encoding, and record structures defined in COBOL copybooks. These legacy representations are transformed into cloud-native formats by transformation logic and with data quality rules identifying anomalies, enforcing business constraints, and finding records that need manual remediation.

Change data capture provides the consistency of mainframe sources and cloud targets continuously throughout the migration cycles that can span several months in a large-scale activity. Changes in the database are automatically replicated to subsequent systems with little latency through CDC mechanisms, which intercept changes as they happen to the database. This real-time synchronization allows parallel-run environments where the mainframe and cloud environments can run workloads at the same time, allowing a comparative validation to ensure that the functionality is equivalent to that of the production cutover. Studies about real-time data integration architectures point out that CDC-based models have a cutover risk that is significantly smaller than in point-in-time migration models that demand longer periods of system unavailability [6].

Reconciliation frameworks offer operational release of the migration progress and data congruence in the source and target environments. The process of automated reconciliation runs continuously as part of migration, comparing the number of records in an environment, calculating the hash-based integrity, and testing business rules. Mismatches are the initial events provoking investigation processes that identify the root causes of the problem, such as logic errors in transformations, time variations, and the quality of source information. Pipeline stages have quality gates that block propagation of data that has been corrupted to downstream systems, which would have led to corrupt data being sent to the target environments where it could be needed in the production process.

Data quality frameworks incorporate validation rules applying business semantics across migration pipelines. These rules ensure that referential integrity exists between related entities, that domain constraints on single fields are valid, and that business rule constraints are met upon calculated values and derived attributes. During the process of migration, organizations often find out that old systems have data quality problems that have been hidden behind application logic, or have been treated as working fine. Migration projects offer chances to address these problems, enhancing the quality of the data in the target systems in comparison to source systems.

4. Application Refactoring and Cloud-Native Target Architecture

Application refactoring refines mainframe applications to cloud-native services based on incremental decomposition patterns that address risk and gradually eliminate old-fashioned components. Rewrites that aim at simultaneous modification of entire application portfolios have high failure rates because of the complexity of scope, long schedule, and compounded technical issues. The pattern of the strangler fig offers an established solution that allows the gradual migration process through intercepting the requests at the system boundaries and redirecting the traffic to the modernized implementations in gradual stages. The AWS prescriptive advice on the strangler fig pattern underlines that this method can minimize transformation risk by allowing for the migration in small steps, constant validation, and easy rollback once problems occur [7]. New cloud-

native services become operable gradually until the original mainframe units are put into an idle state and shut down.

The initiation of the strangler fig pattern is through laying down facade layers that intercept requests that are directed to mainframe applications. These facades only forward the traffic to the existing systems without making changes to them, and they provide infrastructure on which the routing changes will be done without affecting the existing operations. With the availability of modernized services, the facade configurations reroute certain types of requests to the cloud implementation and leave the rest of the traffic to be routed to the mainframe components. This progressive redirection can be used to granularly validate the functionality that has been modernized, and rollback only requires a straightforward configuration change in routing, as opposed to system restoration.

Target-state architectures take advantage of containerized cloud services with a zero-overhead infrastructure management orchestration with enterprise-grade reliability, scalability, and security levels. The analytics of mainframe-based reporting and decision support workloads are migrated to the analytics basis of the BigQuery platform. According to Google Cloud documentation, BigQuery is a fully managed enterprise data warehouse that allows running super-fast SQL queries based on Google infrastructure processing power, and serverless architecture does not need capacity planning and infrastructure management [8]. By adopting managed services, organizations that have migrated out of batch based mainframes reporting have realized a significant performance enhancement as well as simplification of operations.

Cloud Spanner also supports the globally uniform transactional demands typical of fundamental business applications such as order processing, inventory management, and financial processing. The distributed architecture of Spanner offers high levels of consistency at geographical boundaries as well as providing horizontal scalability that is not limited to the capacity of an architecture to support an increase in transaction volume. AlloyDB is a PostgreSQL-compatible relational database that is optimized to serve transactional workloads that need low-latency reads and high-throughput writes.

Event-based architectures substitute the mainframe-based models of batch processes. Pub/Sub messaging infrastructure circulates business events throughout distributed services to support real time processing to remove delays inherent in the batch accumulation cycles. The adjustment of inventory elicits immediate downstream changes instead of waiting until time windows at the end of the day, making the end-to-end latency run-time (in hours) hours-to-seconds, and enhancing the resiliency of the system by having loosely coupled parts.

5. Security, Compliance, and Organizational Change Management

The security and compliance requirements incorporated during MFZ implementation can also be used to guarantee that cloud-native target conditions meet regulatory requirements and improve the overall state of security. Cloud platforms offer holistic security functions that are often beyond the control of mainframe environments when correctly implemented and managed. Studies conducted on cloud security models indicate that organizations should adopt defense-in-depth approaches involving identity management, network controls, encryption, monitoring, and incident response features in order to meet the security goals [9]. To achieve security requirements, migration planning should consider security requirements based on initiation and controls that need to be incorporated all through transformation and not try to have a retrofit after deployment.

Introduction to identity and access management. Identity and access management applies the principles of least privilege using the granular policies that limit access to the resources according to the job function, project assignment, and the classification of sensitivity of the data. Role-based access control systems specify the set of permissions adapted to the organizational tasks, making it easier to administer and at the same time guaranteeing a uniform enforcement. The need for multi-factor authentication ensures the security of paths of privileged access, minimizing the risk of credential theft. Service accounts that allow automated processes have narrow permissions that are restricted to a set of required resources and actions.

Encryption secures the privacy of data both in storage and transmission settings. Cryptographic data control: Customer-managed keys allow organizations to have key custody but use cloud-native storage services. Transport Layer Security offers security for moving data along the network communications. The major management processes deal with rotation and access controls, and recovery processes that guarantee cryptographic continuity.

The framework developed by NIST Cybersecurity gives systematic advice to organizations that deal with security risk in terms of identification, protection, detection, response, and recovery services [10]. Asset inventory and risk assessment

are the starting points of the framework implementation that ensure initial awareness of the security situation. The protection controls relate to access control, awareness training, data security, and protective technology implementation. Detection capabilities allow security events to be detected in a timely manner by using constant monitoring and anomaly detection. Response procedures: Response procedures are needed to make sure that the incident is handled effectively, minimizes its impact, and enables recovery. Recovery planning is concerned with the restoration of capabilities and services after security attacks.

Organizational change management treats human factors that predetermine a successful migration over and above technical implementation. The domain-aligned teams take ownership of modernized services and include business knowledge and technical capacity. Platform engineering teams create standardized templates, deployment pipelines, and observability settings that make teams onboard faster and maintain architectural integrity. Enablement programs provide training on cloud platform capabilities, current development practices, and tooling specific to migration.

Federated governance strikes a balance between central control and domain control and has guardrails that impose security, conformity, and architectural controls but allow flexibility in implementation. The communication plans keep stakeholders informed about the progress made in the project by making frequent updates and reporting risks clearly.

6. Case Study: Measurable Outcomes and Decision Framework

An example of a case study inspired by MFZ programs at one of the multinational retail companies illustrates quantifiable results that can be achieved through methodological execution. The organization had mainframe systems that supported point-of-sale processing, inventory control, merchandise planning, and financial reconciliation in its operations all over the world. Mainframe infrastructures were devouring huge shares of technology budgets in terms of hardware maintenance contracts, software authorization charges, and specialized human resources wages. The discovery analysis showed a large size of the COBOL codebase, thousands of batch job definitions, and many real-time transaction programs that need to be modernized.Migration was done over a series of years in stages, starting with analytics workloads such as enterprise reporting, demand forecasting, and customer analytics. The next stages dealt with inventory, warehouse, and merchandise planning systems. The last stages dealt with point-of-sale transaction processing and financial systems that were most complex and operationally critical. The data migration involved large amounts of data being moved between phases, with the CDC synchronization maintaining consistency at the points of running in parallel. High accuracy rates were confirmed at initial cutover in reconciliation processes, with any remaining discrepancies managed by the already developed data stewardship processes.

Post-migration evaluations ensured that there have been significant gains in the areas of operation, financial, and agility. The cost of infrastructure was reduced by removing the mainframe license and placing it on the cloud services on a consumption basis. Multi-region deployment architectures with automatic failover properties were able to achieve significant improvements in recovery time objectives. The feature delivery velocity was brought up significantly, as measured by how often the deployments were done, and teams delivered updates continuously, no longer had to wait till quarterly release windows to make their updates.

Decision matrices help organisations to choose suitable migration patterns in relation to the nature of workloads and organisational strategic goals. Lift-and-shift strategies are well-suited to workloads with high-speed migration and low modification needs, and rehosting applications on cloud technologies with very few changes in architecture. Re-platforming is a moderate change using cloud-native services without changing the underlying application logic, suitable for the workload that has a database dependency that can be migrated to a managed service. The optimal long-term value is achieved with architectural change which comes with more investment and higher execution risk, which is suitable in strategic systems where long-term modernization commitment is worth the risk.

Pattern selection is done based on such factors as business criticality, technical complexity, integration dependencies, and timeline available. Workloads with a high strategic value and moderate complexity are good candidates of refactoring. The end-of-life systems with minimal strategic differentiation can be carefully considered to be replaced by commercial packages instead of the alternative, which is a custom modernization investment.

Figure 1: *MFZ Risk-Based Phasing Framework [3, 4]*

Table 1: *Data Migration Mechanisms [5, 6]*

Mechanism	Function	Use Case
Bulk Backfill	Extract historical data from DB2/IMS/VSAM	Initial data load
Change Data Capture	Real-time synchronization during migration	Parallel-run periods
Reconciliation	Compare source/target consistency	Validation checkpoints
Quality Gates	Enforce business rules, detect anomalies	Pipeline stages
Hash Verification	Integrity checks across environments	Cutover validation

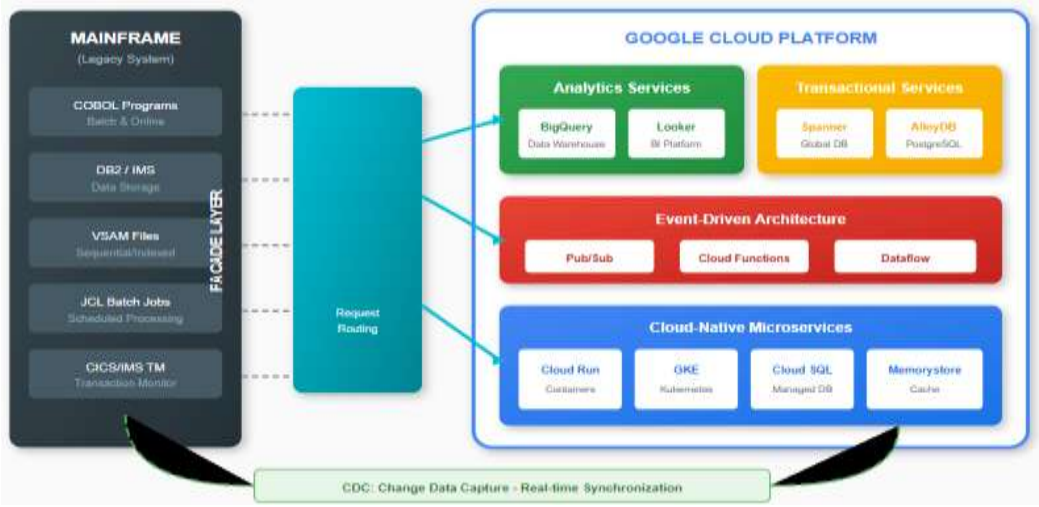

Figure 2*: Cloud-Native Target Architecture [7, 8]*

Table 2*: NIST CSF Security Controls [9, 10]*

NIST Function	Control Area	Implementation
Identify	Asset inventory	Resource cataloging, risk assessment
Protect	Access management	IAM policies, RBAC, MFA, CMEK
Detect	Monitoring	Audit logging, anomaly detection
Respond	Incident handling	Response procedures, containment
Recover	Service restoration	Backup recovery, multi-region failover

4. Conclusions

The Mainframe Zero plan provides a rigorous system of eliminating the dependence of legacy mainframes with cloud-native changeover to allow organizations to remove technical debt and achieve significant operational and financial upgrades in the infrastructural cost, increase system resilience, and speed of innovation. Effective implementation requires architectural rigor with a focus on achieving target-state results despite tactical complications in large-scale migration, and organizational change management aligning teams to modernization goals and developing cloud-native

capabilities using formal enablement programs. Risk-based phasing based on discovery processes that develop comprehensive workload inventories and dependency mappings, and so on, is the basis of sequencing transformation activities to develop organizational confidence by demonstrating early success before taking on mission-critical transactional systems. Bulk extraction with change data capture data migration provides continuity and accuracy in longer parallel-run periods, whereas incremental decomposition with application refactoring to strangler fig patterns provides continuity and stability in controlling risk by continual validation and easy rollback. An identity management based on security integration that includes encryption and monitoring is in line with the accepted cybersecurity frameworks to assure regulatory compliance during transformation and to not hinder delivery progress. The analytics, transaction processing, and event-driven messaging capabilities enabled by cloud-native target architecture using managed services are no longer limited by the capabilities of mainframe infrastructure, such as real-time processing of data, scalability in response to changes in demand, and the ability to implement features quickly due to modern development practices. Organizations that undergo MFZ transformations place themselves in a position of sustained innovation, no longer limited by legacy requirements, having up-to-date platforms that adapt as business needs evolve, eventually enabling technology leaders to achieve the full potential of cloud-native architectures to offer sustained business differentiation and competitive advantage in high-speed market settings.

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper
- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.
- **Use of AI Tools:** The author(s) declare that no generative AI or AI-assisted technologies were used in the writing process of this manuscript.

References

[1] Peter Rutten and Ashish Nadkarni, "IBM LinuxONE: A Secure Data-Serving and Hybrid Cloud Infrastructure," IDC, 2019. [Online]. Available: https://bcxpartners.co.za/wp-content/uploads/2020/09/IDC-Whitepaper-IBM-LinuxONE-A-Secure-Data-Serving-and-Hybrid-Cloud-Infrastructure-.pdf

[2] Sunil Khemka and Arunava Majumdar, "Legacy Modernization with AI - Mainframe modernization," arxiv. [Online]. Available: https://arxiv.org/pdf/2512.05375

[3] Ashish Upadhaya, "Understanding Legacy Software: The Current Relevance of COBOL," VU, 2023. [Online]. Available: https://ictinstitute.nl/wp-content/uploads/2023/12/COBOL_Thesis_Dec4_Ashish.pdf

[4] Karen Yeung and Lee A. Bygrave, "Demystifying the modernized European data protection regime: Cross-disciplinary insights from legal and regulatory governance scholarship," John Wiley & Sons, 2022. [Online]. Available: https://onlinelibrary.wiley.com/doi/pdfdirect/10.1111/rego.12401

[5] Szymon Dziubak, "Review of Cloud Database Benefits and Challenges," Czasopisma, 2023. [Online]. Available: https://czasopisma.prz.edu.pl/mmr/article/view/868

[6] Greeshma Suryadevara, "Real-Time Data Integration With Change Data Capture (Cdc): Techniques, Challenges, And Applications In Modern Data Architectures," International Research Journal of Modernization in Engineering, Technology and Science, 2025. [Online]. Available: https://www.irjmets.com/uploadedfiles/paper//issue_1_january_2025/66092/final/fin_irjmets1736666365.pdf

[7] AWS, "Strangler fig pattern." [Online]. Available: https://docs.aws.amazon.com/prescriptive-guidance/latest/cloud-design-patterns/strangler-fig.html

[8] Google Cloud, "From data warehouse to autonomous data and AI platform." [Online]. Available: https://cloud.google.com/bigquery

[9] Tulasiram Yadavalli, "Addressing Security and Compliance Challenges in Google Cloud Storage for Regulated Industries," SSRN, 2025. [Online]. Available: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5225342

[10] NIST, "The NIST Cybersecurity Framework (CSF) 2.0," 2024. [Online]. Available: https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.29.pdf