Integrating Rehosted IBM i (AS/400) RPG Applications with the Mambu Cloud Banking Platform on Google Cloud

A Reference Architecture for Malaysian Financial Institutions

Dr. Bruce Acacio

June 2026

Abstract

Malaysian financial institutions operate substantial portfolios of business-critical applications written in RPG and COBOL for the IBM i (AS/400) platform. Rehosting these applications to commodity cloud infrastructure using the Infinite i environment preserves decades of embedded business logic while eliminating dependence on proprietary hardware. In parallel, many institutions are adopting Mambu, an API-first software-as-a-service core banking platform that operates natively on Google Cloud. The two estates must coexist, often for years, as products migrate progressively to the new core. This paper presents a reference architecture for bidirectional integration between rehosted RPG applications and Mambu when both workloads operate on Google Cloud Platform (GCP). The architecture combines Mambu's event distribution mechanisms (webhooks and the Streaming API) with an asynchronous staging pattern that preserves legacy transactional validation, mediated by a stateless integration tier. Regulatory considerations specific to Malaysia, including Bank Negara Malaysia's Risk Management in Technology (RMiT) policy and data residency, are analysed, and a phased implementation roadmap is proposed.

Keywords: *IBM i, AS/400, RPG, application rehosting, Infinite i, Mambu, core banking, composable banking, Google Cloud Platform, RMiT, legacy coexistence*

1. Introduction

The IBM i platform, successor to the AS/400, remains widely deployed in banking across Southeast Asia, commonly hosting deposit processing, lending, general ledger, payments, and regulatory reporting functions written in RPG and COBOL. These systems are reliable and deeply encoded with institutional business rules, yet they present growing challenges: a shrinking pool of platform specialists, proprietary hardware refresh cycles, and architectural isolation from the API-driven ecosystems in which modern banks operate.

Two modernization strategies are frequently pursued in parallel. The first is rehosting, in which RPG and COBOL applications are recompiled and executed on x86 cloud infrastructure without functional change, using an emulation and runtime environment such as Infinite i from Infinite Corporation [1]. The second is progressive core replacement, in which new products are launched on a cloud-native core banking platform—here, Mambu—while existing books of business remain on the legacy estate until migrated [4].

These strategies create a coexistence problem. For an extended transition period, the rehosted RPG applications and the Mambu core each hold authoritative data for different products and must exchange transactions, balances, and customer records continuously and reliably. Rehosted RPG applications retain their original batch- and database-oriented interfaces; they do not natively consume the REST APIs and event streams that Mambu exposes. This paper proposes a reference architecture that bridges this gap when both environments are deployed on Google Cloud Platform (GCP) and examines the regulatory obligations that apply to a Malaysian financial institution supervised by Bank Negara Malaysia (BNM).

2. Background

2.1 The Infinite i Rehosting Environment

Infinite i is a suite of compilers, utilities, and operating system services that permits applications developed in RPG, COBOL, CL, and DDS for the AS/400 environment to be migrated at source-code level, recompiled, and executed on x86 Linux virtual machines on public clouds including AWS, Azure, and GCP [1]. Significantly for integration purposes, Infinite i includes a complete functional replication of the DB2/400 database, allowing applications to execute against familiar physical and logical file structures without rewriting data access code [2]. Because this replicated database resides on conventional infrastructure, it can be reached through industry-standard SQL interfaces (ODBC/JDBC), providing a stable integration surface that does not require modification of the migrated applications.

2.2 The Mambu Cloud Banking Platform

Mambu is an API-first, software-as-a-service core banking platform supporting deposits, lending, and payments through a composable architecture. Programmatic access is provided by RESTful v2 APIs covering clients, accounts, and transactions [4]. For outbound event distribution, Mambu offers two mechanisms: webhooks, which push configurable HTTP callbacks to a registered endpoint when specified events occur (for example, a loan account entering arrears), and the Streaming API, an enterprise feature providing pull-based, configurable event feeds suited to high volumes and multiple consuming applications [5, 6]. Mambu operates end to end on Google Cloud—its banking engines run on Google Kubernetes Engine—and is generally available across AWS, Google Cloud, and Azure, with regional deployment options that support local data residency requirements [7, 8].

2.3 Google Cloud in Malaysia

Google announced a US$2 billion investment in Malaysia in May 2024, comprising its first Malaysian data centre and a dedicated Google Cloud region in Greater Kuala Lumpur, with completion targeted for 2026 [11]. An in-country region is material for latency and for data residency posture under Malaysian regulatory expectations (Section 4). Google Cloud also publishes specific guidance on aligning its services with BNM requirements [10]. Institutions should verify the operational status of the Malaysian region, and Mambu's deployment options within it, at the time of implementation.

3. Proposed Integration Architecture

The architecture comprises the rehosted legacy estate, a stateless integration tier, and the Mambu SaaS platform, all operating on GCP. Two guiding principles apply. First, each system remains the undisturbed system of record for the products it owns; integration synchronizes, it does not duplicate authority. Second, all interaction with the legacy estate occurs at the database boundary that Infinite i already exposes, so no RPG or COBOL source code requires modification.

3.1 The Rehosted Estate

Infinite i virtual machines are deployed on Compute Engine within a private VPC subnet with no external IP addresses. The replicated DB2/400 database holds the legacy ledgers and customer files. Read access for integration is obtained through the environment's SQL interfaces over ODBC/JDBC [2, 12], confined to network paths within the VPC.

3.2 The Integration Tier

A stateless mediation service, deployed on Cloud Run and connected to the private subnet through Serverless VPC Access [13], performs protocol translation in both directions: REST and event semantics on the Mambu side, relational semantics on the legacy side. The tier holds no persistent business state—durable buffering is delegated to Pub/Sub—which simplifies scaling, certification, and disaster recovery.

3.3 Eventflow from Mambu to the Legacy Estate

Where legacy applications must reflect activity originating in Mambu (for example, the general ledger consolidating postings from products already migrated), the integration tier consumes Mambu events and writes them to staging tables in the replicated DB2/400 database, which existing RPG batch

programs process through their normal validation and posting routines. For low-volume, point-to-point notifications, webhooks suffice [6]; where multiple consumers require the same events or volumes are high, the pull-based Streaming API is preferred [5]. Events are landed first on Pub/Sub to decouple Mambu's delivery from legacy batch windows and to provide replay.

3.4 Transaction Flow from the Legacy Estate to Mambu

Where legacy applications originate transactions affecting Mambu-held accounts (for example, a branch application capturing a deposit to a migrated product), the pattern is reversed: RPG programs write outbound instructions to dedicated request tables—analogous to the native data queue idiom of IBM i—which the integration tier polls, transforms, and posts to Mambu's v2 APIs using idempotency keys to guarantee exactly-once effect under retry [4]. Responses and errors are written back to result tables for legacy consumption. This asynchronous staging preserves the legacy estate's batch integrity while meeting Mambu's API contract.

3.5 Reconciliation

Coexistence demands independent verification. A scheduled reconciliation job compares balances and transaction counts across both estates and publishes exceptions to an operational dashboard; it is the primary control evidencing integration integrity to auditors and to BNM.

4. Security and Regulatory Considerations for Malaysian Financial Institutions

4.1 Bank Negara Malaysia's RMiT Policy

BNM's Risk Management in Technology (RMiT) policy document governs technology risk for licensed financial institutions in Malaysia. The June 2023 update introduced enhanced requirements for the management of cloud technology risk, effective 1 June 2023 for licensed digital banks and 1 June 2024 for other institutions [9]. A core banking platform is unambiguously a critical system: institutions adopting public cloud for such systems are expected to perform a comprehensive risk assessment, conduct due diligence on cloud service providers, incorporate minimum contractual safeguards, and address the key risks and control measures enumerated in Appendix 10 of the policy [9]. Institutions must retain ownership, control, and management of customer data and cryptographic keys. In the proposed architecture this maps to customer-managed encryption keys through Cloud KMS for institution-controlled projects, contractual clarity with Mambu on data ownership and exit, and early engagement with BNM given the criticality of the workload.

4.2 Data Residency

Mambu's multi-region deployment model, demonstrated by deployments such as Bank Jago on the Jakarta region [8], combined with the forthcoming Malaysian cloud region [11], permits both the rehosted estate and the Mambu tenant to operate in-country. Cross-border processing that remains should be assessed under the Personal Data Protection Act 2010 and the institution's RMiT obligations [9, 14]. Confining integration traffic to private networking within a single national region materially simplifies this analysis.

4.3 Technical Security Controls

The control set follows defense in depth. The Infinite i estate is reachable only from the integration tier across private networking; calls to Mambu authenticate with API keys or OAuth held in Secret Manager and rotated; webhook endpoints validate signatures and source restrictions; Cloud Audit Logs and VPC Flow Logs provide the evidentiary trail expected under RMiT's audit provisions; and all data is encrypted in transit (TLS 1.2+) and at rest. These controls align with the guidance Google Cloud publishes for BNM-regulated workloads [10].

5. Implementation Roadmap

A phased approach mitigates delivery risk. Phase 1 (discovery, four to six weeks) confirms the SQL connectivity profile of the specific Infinite i version deployed, the Mambu event and API entitlements held, and completes the RMiT risk assessment and engagement with BNM. Phase 2 (read-only pilot, six to eight weeks) delivers one event flow end to end—typically Mambu transaction events feeding the legacy general ledger staging tables—with reconciliation from day one. Phase 3 (transactional integration) introduces legacy-originated postings to Mambu using the staging pattern with idempotent delivery, verified by parallel run. Phase 4 (industrialization) extends coverage to the remaining data domains, hardens monitoring and disaster recovery, and establishes the migration runbook by which products move from the legacy estate to Mambu, progressively retiring integration flows as they become redundant.

6. Conclusion

Rehosting with Infinite i and adopting Mambu are independently rational decisions for a Malaysian bank; the architecture presented here demonstrates that they are also mutually compatible without modification of legacy RPG source code. By treating the replicated DB2/400 database as the legacy integration surface, consuming Mambu's webhooks and Streaming API through a thin stateless tier, and preserving legacy validation through asynchronous staging with continuous reconciliation, an institution can run old and new cores side by side with integrity throughout a multi-year migration. The regulatory pathway under RMiT is well defined, and an in-country Google Cloud region strengthens the residency posture of the design. The recommended next step is a discovery engagement culminating in a code analysis of the IBM i-based applications pilot with a subsequent read-only event-synchronization pilot.

References

[1] Infinite Corporation, "Infinite i — IBM i (AS/400) Application Migration and Modernization." *https://www.infinitecorporation.com/*

[2] Infinite Corporation, "Database Migration & Integration Technology for DB2/400." *https://www.infinitecorporation.com/data-migration*

[3] AWS Marketplace, "Infinite i" product listing. *https://aws.amazon.com/marketplace/pp/prodview-pxlvwrdp5k6pq*

[4] Mambu, "Mambu APIs," Mambu Documentation Hub. *https://docs.mambu.com/docs/mambu-apis/*

[5] Mambu, "Streaming API," Mambu Documentation Hub. *https://docs.mambu.com/docs/streaming-api/*

[6] Mambu, "Webhooks Overview," Mambu Documentation Hub. *https://docs.mambu.com/docs/webhooks-overview/*

[7] Mambu, "Mambu Forms Global Partnership with Google Cloud," press release. *https://mambu.com/en/insights-hub/press/mambu-forms-global-partnership-with-google-cloud*

[8] Google Cloud, "Mambu Case Study." *https://cloud.google.com/customers/mambu*

[9] Bank Negara Malaysia, "Risk Management in Technology (RMiT)" Policy Document, as updated June 2023 (including Appendix 10, Management of Cloud Technology Risks). *https://www.bnm.gov.my/*

[10] Google Cloud, "Bank Negara Malaysia (BNM) Compliance." *https://cloud.google.com/security/compliance/bank-negara-malaysia*

[11] Google, "Advancing Malaysia Together: Google Announces US$2 Billion Investment in Malaysia, Including First Google Data Center and Google Cloud Region," press release, 30 May 2024. *https://www.googlecloudpresscorner.com/*

[12] IBM Corporation, "Accessing IBM i Db2 (AS/400) Databases via JDBC," IBM Support documentation. *https://www.ibm.com/support/pages/accessing-iseries-db2-as400-database-jdbc*

[13] Google Cloud, "Serverless VPC Access" and "Cloud Run" product documentation. *https://cloud.google.com/vpc/docs/serverless-vpc-access*

[14] Government of Malaysia, Personal Data Protection Act 2010 (Act 709).

All online sources accessed June 2026. This paper is provided for discussion purposes; institutions should validate vendor capabilities and regulatory interpretations with Infinite Corporation, Mambu, Google Cloud, and their own legal and compliance advisers.