ALDEN ADVISORY

RESEARCH PAPER

The Future of Core Banking

Why continuity, not technology, is the real migration risk

Sarah Chen — Head of Digital Transformation

Research Paper — Alden Advisory

Alden Advisory — London — aldenadvisory.example (concept site)

Executive Summary

Core banking replacement is consistently framed, in board papers and vendor pitches alike, as a technology decision: which platform, which vendor, which architecture. This paper argues that framing is backwards. The technology choice is rarely what determines whether a migration succeeds. Continuity is.

A bank undergoing core system replacement has to keep taking payments, processing loans, and staying open for the two or three years the migration typically takes. Most implementation plans treat that continuity requirement as a constraint to be managed around the technology plan, rather than the central design problem the technology plan exists to solve. This paper draws on patterns observed across Alden Advisory's core banking engagements to argue for continuity-first migration design.

Why Big-Bang Cutovers Underperform

Big-bang cutover — replacing the entire core system in a single weekend event — remains attractive to boards because it promises a clean end date. In practice, it concentrates all migration risk into a single, high-stakes window with limited ability to roll back cleanly if a critical process fails to behave as tested.

Phased migration, sequencing by product line from lowest to highest operational risk, consistently produces better outcomes across the engagements we have observed — not because it eliminates risk, but because it isolates it. A failure in phase two of five is a contained, recoverable event. A failure in a single big-bang cutover is an institutional crisis.

The Governance Question Nobody Resolves in Advance

The most consistent failure pattern we observe is not technical. It's that no single person has been given clear authority to accept operational risk during a migration window before the migration begins. Every business unit retains informal veto power over changes that might affect its numbers during the transition, and migration plans get designed to avoid ever triggering that veto — which produces plans that are technically elegant and organisationally undeliverable.

Resolving this requires naming, before migration planning begins, a single executive with explicit authority to approve each phase's go-live independent of business-unit sign-off. This is an uncomfortable governance conversation to have early, and a far more comfortable one than having it during a live cutover.

Measuring Success Beyond Go-Live

A migration that goes live is not the same as a migration that succeeded. We recommend measuring against the operational baseline established before migration began — batch processing time, customer-facing incident rate, and staff time spent on manual workarounds — rather than against the narrower question of whether the new system is technically live.

Institutions that build this measurement discipline into the migration plan from the outset are better positioned to identify genuine adoption problems early, rather than discovering them months after the programme has officially closed.