RE-VIVE

Customer Onboarding Process Intelligence

Large Multinational Bank — from raw case records to executive insights in two weeks

8.28M	953K	130+	2 wks
events analysed	service requests	process maps generated	elapsed, ~2 hrs client time

www.re-vive.com

The problem — a 130-product onboarding journey no one could see end to end

The bank set out to cut onboarding cycle times by 25% across more than 130 products and services — account opening, wire transfer, online business banking, remote deposit — for existing clients opening new accounts. Before targeting anything, it needed a factual picture of how that journey actually ran.

LONG, VARIABLE CYCLES

Some requests closed in days, others ran for months

INVISIBLE REWORK

Repeat steps and loops buried inside case records

NO CROSS-PRODUCT BASELINE

130+ products with no common way to compare them

FRAGMENTED HANDOFFS

Document, KYC and AML steps spanning teams and systems

The question was not whether onboarding worked.
It was: where were the days actually going, and how many of them were rework?

© RE-ViVE | Confidential

The engagement at a glance

Customer onboarding for existing clients opening new accounts · 130+ products and services

8.28M
events
from 54 source tables

953K
service requests
across product history

130+
process maps
one per product type

2 wks
to findings
~2 hrs client time

What was in scope

THE JOURNEY

Request through documentation, verification and account setup — captured in ServiceNow and OpenText, which run the journey in two different regions

Account opening, wire transfer, online business banking, remote deposit and more

THE WINDOW

Nine months of 2023 activity analysed in depth — 70,840 requests opened in that period — against process history reaching back to 2011

Deep-dive product: Account Opening, 13,685 requests

THE CONSTRAINTS

Extracts handed over as-is with personally identifiable information scrubbed to bank policy, plus roughly two hours of subject-matter support

No data warehouse build, no process workshops

Objective: reduce onboarding cycle times across the product set by 25% without adding headcount.

© RE-ViVE | Confidential

Step 1 — source evidence to execution foundation

The same workflow ran in two systems across two regions — read-only extracts, PII scrubbed, mapped into one Execution Data Model.

SOURCE SYSTEMS

ServiceNow

The full onboarding workflow, end to end

One region of the organisation

OpenText

The same workflow, separately instrumented

The other region

WHAT EACH ONE HELD

- Service requests and statuses
- Activities and timestamps
- Assignees and product types

54 Oracle tables in total

EXECUTION RECONSTRUCTION

1. **Extract**
   Both systems pulled as-is, no formatting by the bank
2. **Scrub**
   All personally identifiable information removed to bank policy
3. **Stitch**
   Both systems mapped into a common Execution Data Model
4. **Normalise**
   Timestamps, product types and rework flags standardised

RECONSTRUCTED EXECUTION

Request	Activity	Timestamp
SR-40218	Request Created	03 Apr 09:12
SR-40218	Collect Documents	05 Apr 14:30
SR-40218	Documents Received	11 Apr 10:08
SR-40218	Perform KYC Tasks	17 Apr 08:55
SR-40218	Send for E-Sign	24 Apr 16:40
SR-40218	AML Approval	02 May 11:20

Every request also carries product type, client, assignee, status, SLA target and rework flags — so any insight can be sliced by any dimension.

953K requests · 130+ products

Two systems, two regions, one Execution Data Model — no process remodeling and no process workshops.

© RE-ViVE | Confidential

Step 2 — from execution data to Execution Intelligence

RE-ViVE algorithms reconstruct actual execution and generate Execution Intelligence automatically.

1. **Process reconstruction**
   Rebuilds the true end-to-end onboarding journey per product from raw case records

2. **Variant & path analysis**
   Separates real execution paths from unique activity flows, ranked by volume and time

3. **Cycle time & delay isolation**
   Quantifies where days are consumed activity by activity, with SLA deviation views

4. **Rework & self-loop detection**
   Finds repeat executions and duplicated information requests inside the same case

5. **Productivity & dimension analysis**
   Benchmarks by assignee, client, product and period, and drills to the individual request

WHAT CAME OUT

- 130+ process maps, drillable to the individual service request
- Flow comparison across products, periods and clients
- Automatic flagging of ideal versus problematic workflows
- Views by volume, time, SLA deviation and rework
- Recommendation engine ranking automation opportunities

2 weeks

from data handover to findings review, against roughly two hours of subject-matter support.

© RE-ViVE | Confidential

Finding 1 — 60% of the biggest product's cycle was rework

Account Opening: 13,685 requests, one fifth of everything opened in the analysis window.

20%

of all requests

13,685 of the 70,840 service requests opened in the period

16.62

days on average

From request creation to completed account opening

60%

of that time was rework

10.14 rework days on an average request

4 of 29

activities drove it

Roughly 70% of all rework days concentrated in four steps

WHERE THE 16.62 DAYS WENT



THE FOUR REWORK HOTSPOTS

- Received Documents
- Perform KYC Tasks
- Documents Received / AML Approval
- Send Documents for E-Sign

96,675 of 138,769 total rework days sat in these four steps alone.

A separate view found three sequential steps carrying a further 14% of rework.

© RE-ViVE | Confidential

Finding 2 — a small tail of requests ran for months

Most Account Opening requests closed inside ten days. A stubborn minority took more than fifty.

RE-VIVE

THE DELAY TAIL

602 requests over 50 days

Against a majority that completed within ten. The same product, the same system, the same designed process — but a tail running to 200 days.

WHAT THE PROCESS MAP SHOWED

Three activities each averaged more than ten days on their own inside the delayed population:

- **Submit to KYC Approval**
- **Send Documents for E-Sign**
- **Documents Received / AML Approval**

WHY IT MATTERS

Averages hide it

A 16.62-day mean says nothing about a client waiting seven weeks for an account. The tail is where the complaints and attrition live.

It is concentrated

The delay does not spread evenly across 29 activities. Three verification and document steps carry it, and all three are already known rework hotspots.

It is addressable

Because every delayed request is traceable to its activity sequence, the bank can work the 602 as a named list rather than a statistic.

© RE-ViVE | Confidential

Finding 3 — 13,685 requests took roughly 3,000 different paths

Strip out the rework and only 46 genuine activity flows remain. The sprawl was rework, not product complexity.

PATHS THROUGH ONE PRODUCT

~3,000

distinct paths actually walked, once rework is counted

46

unique activity flows when rework is stripped out

~23

target flows after consolidation — a 50% reduction

5 rework steps

could be removed from an average request, based on analysis of the top ten variants. The bank sets the target number of flows; the data shows roughly half are redundant.

ONE PRODUCT, TWO PATHS — ONLINE BUSINESS BANKING

30 variations exist. The dominant one runs 57% of volume through just six activities.

V1 **2.82** days

Straight through, no rework

V2 **6.87** days

Adds a return step and rework loops

Same product, same system — 4 extra days bought nothing but rework.

© RE-ViVE | Confidential

Finding 4 — the same steps ran again and again

Cancelled requests exposed duplication that completed requests had hidden.

WHERE CANCELLED REWORK SAT

30% in ten products

- Analysis of cancelled requests showed rework concentrated in the top ten products by volume
- The concentration made a short, targeted remediation list possible
- Cancelled cases proved a useful lens — they expose loops that closed cases mask

WIRE TRANSFER SERVICE

6.13 rework days

- Just 11 activities end to end, yet 6.13 days of rework per request
- That is roughly 53% of the time spent on a typical request
- A short process is not automatically an efficient one

SELF-LOOPS

90% of activities repeated

- Ten of the eleven activities were executed multiple times within the same request
- Consistent with duplicated effort, or the same information being asked for more than once
- Dimension analysis traced the loops to specific assignees and workflow items

WHAT IT POINTS TO

A design problem

- Repetition at this rate is rarely individual error — it reflects how the workflow is structured
- Business validation was the recommended next step before any automation build
- Every loop is traceable to the request and the person who executed it

Nine tenths of a short, well-understood process was being repeated inside itself — invisible until execution was reconstructed.

© RE-ViVE | Confidential

The opportunities we identified

Five prioritised plays. Figures are opportunities quantified from the bank's own execution data — projected, not yet realised.

01

Automate the document and verification chain

Target the activities carrying the bulk of rework, starting with document transfer steps that need no human judgement.

Measured opportunity: ~30% faster across ~6,000 requests

02

Eliminate rework in the four hotspot activities

Received Documents, KYC, AML approval and e-signature together hold roughly 70% of all rework days.

Measured opportunity: move average cycle from 10 days toward 6

03

Consolidate the variant landscape

Reduce 46 unique flows toward a target set of about 23, and hold new requests to the proven path.

Measured opportunity: remove ~5 repeated steps per request

04

Retrain analysts onto the fastest proven path

The largest variant of each product runs around 50% above expected completion time; the fast path already exists.

Measured opportunity: ~25% improvement on the largest variant

05

Instrument SLA, rework and productivity continuously

Keep the process maps live so delay, rework and duplication surface as they happen rather than in the next review.

Proposed action: keep execution views current

The opportunity was already visible in the execution data. RE-ViVE made it measurable and actionable.