EXECUTIVE SUMMARY

Business Intelligence Portfolio

Sales Dashboard Findings and Banking Warehouse Design

Reporting Period	Project Window: 2024 - 2025
Document Version	1.0
Date Prepared	April 2026
Classification	Portfolio - Internal Reference
Prepared By	Data & Analytics Team

DISTRIBUTION

Hiring Managers - Recruitment Partners - Portfolio Reviewers

PREPARED BY	REVIEWED BY	APPROVED BY
_____ *Signature / Date*	_____ *Signature / Date*	_____ *Signature / Date*

1. Portfolio Overview

This portfolio contains two distinct data analytics projects, packaged as a single bundle[1]. The first project converts raw transactional data from a gourmet food distributor (the Distributor) into an executive Tableau dashboard. The second designs a dimensional warehouse for a retail bank (the Bank), starting from a flat transaction file and producing a production-ready star schema with full ETL specification.

Both projects share a common discipline: take messy operational data, define the analytical questions precisely, design the structure that answers them, and document every decision so a downstream engineer or analyst can pick up the work without context loss.

Project A - Sales Performance Dashboard

Attribute	Detail
Source organization (genericized)	The Distributor - a gourmet food wholesaler operating across 21 countries.
Dataset	7 CSV files: orders, order_details, products, categories, customers, employees, shippers. 2,155 line items across 830 orders.
Period	July 2013 to April 2015 (22 months).
Deliverable	Three-page Tableau dashboard: Sales & Revenue Overview, Product & Category Intelligence, Regional/Operational/People Performance.
Audience	Executive leadership and sales operations management.
Tool	Tableau Public.

Table 1. Project A scope summary.

Project B - Banking Transaction Warehouse

Attribute	Detail
Source organization (genericized)	The Bank - a retail bank with branch, ATM, mobile, POS, and internet banking channels.
Source data	Flat transaction file with 15 columns covering customer, branch, product, channel, transaction type, amount, and balance.
Deliverable	PostgreSQL star schema (8 tables), ETL specification, data quality framework, monthly aggregate table.

[1]Identifying details have been genericized for portfolio submission. Customer names, employee names, shipper names, and the source organization name are not the original values from the underlying datasets.

Attribute	Detail
Grain	One row per banking transaction.[3]
Audience	Data engineers, data quality, internal audit, analytics leadership.
Tool	PostgreSQL 18, pgAdmin 4.

Table 2. Project B scope summary.

2. Project A - Headline Findings

Topline metrics

Metric	Value	Reading
Total revenue	**$1.27M**	Across 830 delivered orders.
Average order value	$1,525	Stable across the period.
Active customers	89	Concentrated customer base.
Average fulfillment days	8.5	From order date to shipped date.
Freight as % of revenue	5.1%	Reasonable shipping overhead.
Sales reps	9	Top performer carries 18% of revenue.

Table 3. Topline metrics from the dashboard.

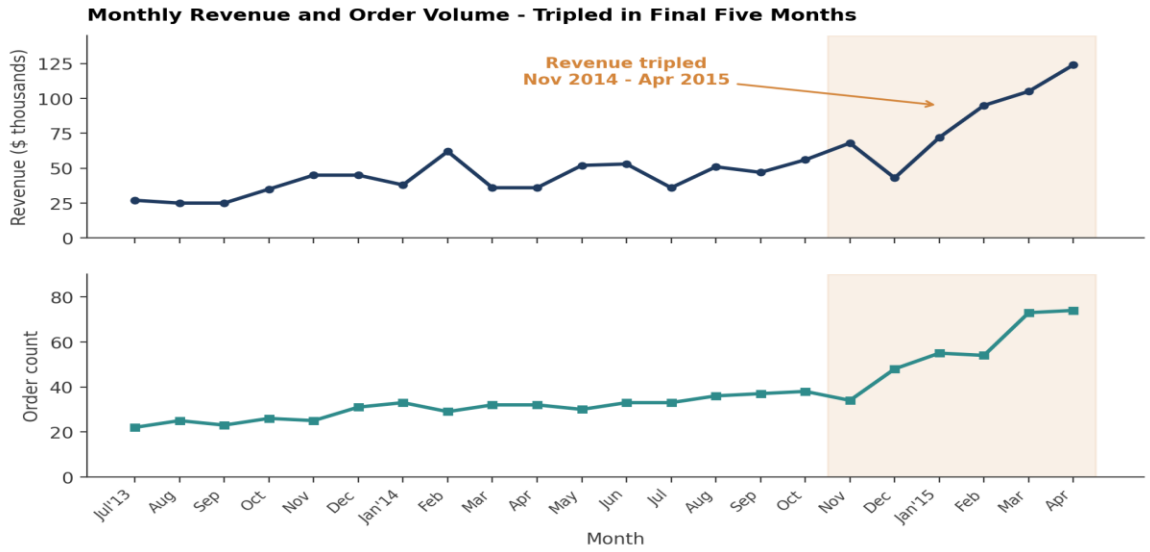

Figure 1. Monthly revenue and order volume. Both metrics inflected sharply from November 2014 onward.[object Object]

THE SINGLE MOST IMPORTANT FINDING

Revenue tripled from November 2014 to April 2015. Order volume doubled in the same window. Both metrics inflected together, which rules out price as the driver - this was real demand growth.

What caused it is not in the data. The dashboard surfaces the pattern; the recommendation is to investigate the cause (sales process change, new acquisition channel, market expansion) before the next planning cycle.

Five additional findings worth attention

- Beverages and Dairy Products drive 40% of all revenue. The remaining 6 categories sit between $95K and $167K each. Concentration risk is real.
- 75% of revenue comes from undiscounted orders. Heavy discounts (over 20%) still command an average order value of $1,374 - comparable to the 1-10% band. Discounts target large orders, not broad demand.
- Germany ($189K) and USA ($172K) deliver 31% of revenue. The next 19 countries contribute to the rest. Either market disappearing is a material hit.
- Carrier B is fastest (7.5 days) at $80 per order. Carrier C is slowest (9.2 days) at $87 per order. Carrier C underperforms both metrics; the contract is a candidate for renegotiation or replacement.
- Rep A generated $233K, more than the bottom three reps combined. The variance is wide enough to be a coaching and territory-design opportunity, not a pure individual-performance story.

3. Project B - Design Highlights

The Bank's source file is a flat list of transactions with 15 columns. To answer the business questions - customer behaviour, branch performance, churn signals, product mix, time-based reporting - it needs to be reshaped into a dimensional warehouse. **The decisions taken during that reshape are the actual deliverable**; the SQL file is just the implementation.

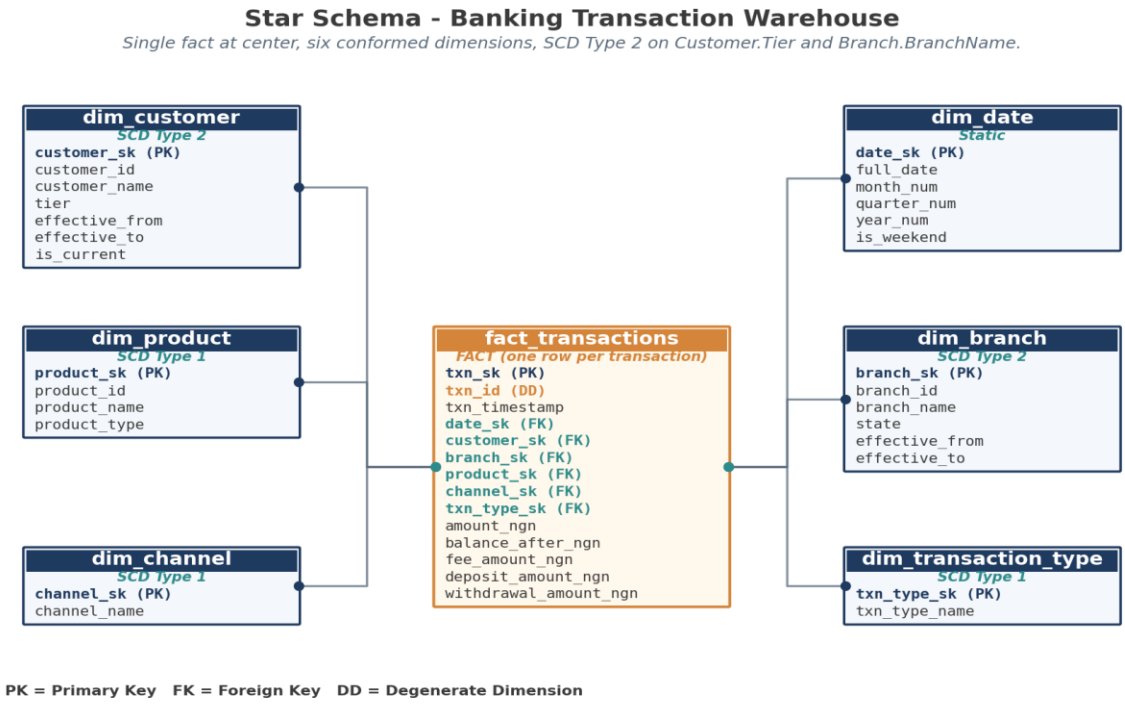

Figure 2. The star schema. One fact at center, six dimensions around it.

Design decisions worth highlighting

Decision	Rationale
Grain: one row per transaction	The most important modeling decision. COUNT(*) means transaction count; SUM(amount) means transaction value. Every measure is anchored to this grain.
Separate dim_channel and dim_transaction_type	Channel (how it happened) and transaction type (what happened) are independent dimensions. Combining them into one would force every analytical question to pick a side.
SCD Type 2 on Customer.Tier and Branch.BranchName	Customer tier change is central to churn analysis - reports must know the tier at transaction time, not just the current tier. Branch rebrands need historical name preservation for old reports.
Surrogate keys on every dimension	Source keys can change, be reused, or have gaps. Surrogate keys also enable SCD Type 2 by allowing multiple rows for the same natural key.

Decision	Rationale
Embedded fee_amount_ngn as nullable	The source file lacks a fee column. The model includes the column nullable so fee data can be added later without redesigning the fact table.
Monthly aggregate table	Most banking reports are time-based. Pre-aggregating to month/branch/product_type cuts repeated full-table scans.

Table 4. Six design decisions and their rationale.

4. Strategic Recommendations Across Both Projects

From Project A

Recommendation 1 - Investigate the late-2014 revenue acceleration

Element	Detail
Action	Identify the trigger for the November 2014 inflection. Interview the sales team, audit campaign records, check for product launches or partnership starts. Document the cause so it can be sustained or repeated.
Owner	Head of Sales, with support from Marketing and Operations.
Expected outcome	Either a sustained-growth playbook (if the cause is repeatable) or a risk register entry (if the cause is external and won't repeat).
Risk if ignored	The trend may continue or it may reverse. Without a known cause, leadership cannot plan capacity, inventory, or hiring.

Table 5. Project A Recommendation 1.

Recommendation 2 - Renegotiate or replace Carrier C

Element	Detail
Action	Move volume from Carrier C to Carrier A (faster, similar cost) or Carrier B (slightly slower, 25% cheaper). Test for 60 days; measure delivery time and customer complaint rate.
Owner	Head of Operations.
Expected outcome	Either improved fulfillment time at flat cost (Carrier A route) or reduced freight spend at marginally higher delivery time (Carrier B route).
Risk if ignored	Customer experience remains the worst-of-both on speed and cost.

Table 6. Project A Recommendation 2.

From Project B

Recommendation 3 - Implement the warehouse in three phases

Phase	Scope
Phase 1	Deploy the staging table and dimensional structure. Run initial historical load. Validate row counts and referential integrity. Phase target: 2 weeks.
Phase 2	Wire incremental load against the transaction source. Implement SCD Type 2 close-and-replace logic. Phase target: 2 weeks.

Phase	Scope
Phase 3	Build the agg_monthly_branch_activity table and downstream reporting layer. Roll out partitioning by month on fact_transactions. Phase target: 1 week plus rollout.

Table 7. Suggested rollout phasing.

5. What These Two Projects Demonstrate Together

Project A is a build for an executive audience - storytelling, headline metrics, design choices that prioritize readability over completeness. Project B is a build for an engineering audience - schema design, key strategies, ETL ordering, quality controls.

Both projects start with the same step: define the questions, then design the structure that answers them. The difference is what gets built next. Project A produces a visualization layer; Project B produces a data layer.

> **THE SHARED DISCIPLINE**
>
> Define the grain or the unit of analysis before doing anything else.
>
> Anchor every measure to the grain.
>
> Document the limitations openly (cost data missing in Project A; fee data missing in Project B).
>
> Make the design defensible without the author present - because the author won't always be there.

Document Control

Version	Date	Change Summary	Author
0.1	2026-02-10	Initial draft of both project summaries	Data & Analytics Team
0.5	2026-02-22	Recommendations and topline metrics added	Data & Analytics Team
0.9	2026-03-08	Internal review; cross-project comparison section added	Data & Analytics Team
1.0	2026-03-15	Approved as portfolio submission	Data & Analytics Team

Table 8. Revision history.