METHODOLOGY REPORT

Sales Performance Dashboard

Data Sources, Calculations, and Design Decisions

Reporting Period	Data Window: July 2013 - April 2015
Document Version	1.0
Date Prepared	April 2026
Classification	Internal - Analytics Reference
Prepared By	Data & Analytics Team

DISTRIBUTION

Analytics Team - BI Developers - Hiring Managers

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

1. Purpose

This report documents how the three-page Tableau dashboard was built. It is the companion to the Dashboard Findings Report (Document 3) and is intended for an audience that needs to reproduce, extend, or audit the work.

The dashboard answers six executive questions:

- How are revenue and order volume trending over time?
- Which product categories drive revenue?

- Which countries and cities contribute most to sales?
- Which shippers perform best on speed and cost?
- How do discounts affect revenue and average order value?
- Which employees generate the most sales?

Each question maps to one or more charts on the dashboard. The mapping is documented in Section 5.

2. Data Sources

The dashboard runs on seven CSV files from the Distributor's operational database. Together, they contain transactional, product, customer, and operational reference data.

File	Rows	Role
orders.csv	830	Order header. One row per order. Holds order_date, shipped_date, customer_id, employee_id, shipper_id, freight, ship-to country.
order_details.csv	2,155	Order line items. One row per product per order. Holds quantity, unit_price, discount. This is the fact-level table.
products.csv	77	Product catalogue. Holds product_name, category_id, unit_price, discontinued flag.
categories.csv	8	Category lookup. Holds category_name.
customers.csv	91	Customer master. Holds company_name, country, city.
employees.csv	9	Sales employee master. Holds first_name, last_name, hire_date.
shippers.csv	3	Carrier master. Holds shipper company name.

Table 1. Source CSV files used by the dashboard.

3. Analytical Grain

The grain of the dashboard is one row per order line. That is, the level at which every measure is calculated. Revenue is computed at the line level (price multiplied by quantity, less discount) and then aggregated up to the chart's grouping.

This choice has downstream consequences:

- order_details is the fact-level table; everything else provides descriptive context.
- orders, customers, products, categories, employees, and shippers are dimensions.

- COUNTD(order_id) gives distinct order count (used for AOV). COUNT(*) on order_details would give line count, not order count.

4. Data Model in Tableau

Tableau's relationship model (not the classic join model) is used to preserve grain[2]:

From	To	Relationship on
order_details	orders	order_id
order_details	products	product_id
products	categories	category_id
orders	customers	customer_id
orders	employees	employee_id
orders	shippers	shipper_id

Table 2. Relationship structure. order_details sits at the center; orders is the bridge to customer, employee, and shipper context.

> **WHY RELATIONSHIPS, NOT JOINS**
> Classic joins in Tableau force a single grain at the data-source level. A customer-level metric and a line-level metric cannot both be computed correctly from the same data source if it was built with joins. Relationships defer the join logic until query time, so each chart picks its own grain.

5. Key Calculations

Every metric on the dashboard is a calculated field in Tableau. The definitions are listed below in the order they appear on the dashboard.

5.1 Revenue

```
// Tableau calculated field: Revenue
[unit_price] * [quantity] * (1 - [discount])
```

Listing 1. Revenue calculation.[object Object]

This is the primary measure across all three dashboard pages. It captures the price actually paid by the customer, after any line-level discount. The product cost is not available in the source data, so this is revenue, not margin or contribution.

5.2 Average Order Value

```
// Tableau calculated field: Average Order Value
SUM([Revenue]) / COUNTD([order_id])
```

Listing 2. AOV calculation.

AOV is computed by dividing total revenue by distinct order count. Note the use of COUNTD: a single order with 5 line items must count as one order, not five. This is the most common AOV error in Tableau and was caught during dashboard QA.

5.3 Fulfillment Days

```
// Tableau calculated field: Fulfillment Days
DATEDIFF('day', [order_date], [shipped_date])
```

Listing 3. Fulfillment days calculation.

Used in the operational chart. shipped_date is the date the order left the warehouse, not the customer-received date - so this is a fulfillment metric, not a delivery metric. Orders where shipped_date is NULL (about 21 records, mostly very recent orders) are excluded from the average.

5.4 Freight as % of Revenue

```
// Tableau calculated field: Freight % of Revenue
SUM([freight]) / SUM([Revenue]) * 100
```

Listing 4. Freight percentage calculation.

Used in the KPI strip on page 1. Note that freight is at the order level (orders.freight), not the line level - and revenue is at the line level (Revenue calculation above). SUM aggregates both up to the same level (the chart's grouping) before dividing.

5.5 Discount Band

```
// Tableau calculated field: Discount Band
IF [discount] = 0 THEN "0%"
ELSEIF [discount] <= 0.10 THEN "1-10%"
ELSEIF [discount] <= 0.20 THEN "11-20%"
ELSE ">20%"
END
```

Listing 5. Discount band classifier.

Bucketing decision: bands chosen to give meaningful distribution. With 1 percentage point bands, the distribution would be too granular to read; with two bands (zero/nonzero), the dashboard would miss the *heavy discounts on largest orders* finding. Four bands balance readability against detail.

5.6 Product Status

```
// Tableau calculated field: Product Status
IF [discontinued] = 1 THEN "Discontinued"
ELSE "Active"
END
```

Listing 6. Active/discontinued classifier.

Surfaces the finding that discontinued products still carry $185K of revenue. Without this calculation, that revenue is invisible in the category breakdown.

6. Dashboard Design Decisions

6.1 Three pages, not one

The decision to split into three pages instead of one mega-dashboard was deliberate. Each page answers a coherent set of questions and can be read in 60 seconds. A single 12-chart page would force the reader to scan rather than read.

Page	Focus	Questions answered
Sales & Revenue Overview	Topline performance	Q1 - revenue/volume trend. Plus KPI strip at top.
Product & Category Intelligence	What drives revenue	Q2 - category revenue. Q5 - discount impact on revenue and AOV. Active vs. discontinued split.
Regional, Operational & People Performance	Where and who	Q3 - country revenue. Q4 - shipper speed and cost. Q6 - employee revenue.

Table 3. Page structure and question mapping.

6.2 Titles state the finding

Standard Tableau chart titles are descriptive: "Monthly Revenue Trend." Those titles force the viewer to interpret the chart from scratch. A title that states the finding does the cognitive work upfront.

Standard title (rejected)	Finding-led title (used)
Monthly Revenue Trend	Revenue tripled Nov 2014 to Apr 2015. Order volume doubled.
Revenue by Category	Beverages and Dairy drive 40% of revenue.
Discount Analysis	75% of revenue comes from undiscounted orders, but heavy discounts cluster on the largest orders.
Country Revenue	Germany and USA dominate - 31% of global revenue.
Shipper Performance	Carrier A leads on speed. Carrier C underperforms on both cost and speed.
Employee Sales	Rep A dominates the salesforce.

Table 4. Title-rewriting examples. Each rewrite shifts the reader from interpretation to confirmation.

6.3 Restrained color palette

The dashboard uses a near-monochrome palette (navy and slate, with amber for emphasis only). This is deliberate. Color in dashboards should carry information; using many colors purely for visual interest dilutes that signal.

- Navy: primary data series, top performers, dominant values.
- Slate: secondary/mid-tier categories and comparisons.
- Amber: warnings, the inflection point in the trend chart, underperformance flags.
- Sage / Teal: used sparingly for cost (sage = good, teal = neutral) on the operational page.

6.4 Filters apply across all three pages

Three filters (Year, Country, Category) apply as cross-dashboard filters. This means a viewer who filters to Germany on page 1 sees Germany-only data on pages 2 and 3 as well. It is the difference between a dashboard and three loosely-related charts.

7. Boundary Decisions

7.1 May 2015 excluded

The dataset ends May 6, 2015, so May 2015 is a partial month. Including it would create a sharp downward dip on the trend chart that misrepresents the actual business state[1]. The dashboard filters to order_date through April 30, 2015. April 2015 (the visible peak) is a complete month and accurately represents the trend.

7.2 Profit margin not calculated

The source data does not include product cost or COGS. Margin cannot be calculated without that input. Rather than fabricate a margin figure, the dashboard reframes the original brief's "discount impact on profit margin" question as "discount impact on revenue and AOV". This is documented openly on the dashboard, not hidden.

> **DATA LIMITATION - DOCUMENTED, NOT INVENTED**
> If product cost data becomes available later, the existing dashboard can incorporate margin without redesign - the calculation would be Revenue minus (unit_cost multiplied by quantity), grouped by the same dimensions. Until then, revenue is the largest measure that can be defended.

7.3 Customers, not consumers

The Distributor sells to gourmet retailers and restaurants, not end consumers. "Customers" in this dashboard refers to commercial accounts, not individual shoppers. The 89 active customers figure should be read with that context.

8. Quality Assurance

Three QA passes were run on the dashboard before sign-off.

Pass 1 - Number reconciliation

Metric	Dashboard value	Cross-check method
Total revenue	$1,265,793	SQL: SUM(quantity * unit_price * (1 - discount)) on order_details.
Order count	830	SQL: COUNT(DISTINCT order_id) on orders.

[1]May 2015 contains only 6 days of data (the dataset ends May 6). Including it as a full month would create a sharp downward dip on the trend line that misrepresents the actual business state.

Metric	Dashboard value	Cross-check method
Active customers	89	SQL: COUNT(DISTINCT customer_id) on orders with at least one delivered order.
Average fulfillment days	8.5	SQL: AVG(DATEDIFF) excluding orders where shipped_date IS NULL.

Table 5. Numbers reconciled against direct SQL queries on the source CSVs.

Pass 2 - Filter interaction

Every cross-page filter was tested in three states: all-selected (default), single-country, and single-category. The expected behavior is that all charts on all three pages update consistently. No filter combination produced empty charts or mismatched totals.

Pass 3 - Edge cases

- Months with zero revenue (none in this dataset - confirmed by query).
- Orders with NULL shipped_date (excluded from fulfillment metric, included in revenue).
- Customers with zero orders (excluded from active count).
- Categories with no recent sales (still visible on the category chart).

9. Known Limitations

Limitation	Impact	Mitigation
No product cost data	Margin and contribution analysis impossible.	Reframed to revenue/AOV. Documented openly.
No customer cost-of-acquisition data	Cannot compute CAC, LTV, or payback period.	Out of scope for this dashboard.
Dataset window only 22 months	Cannot establish multi-year seasonality patterns.	Year-over-year comparison shown where possible (2013 vs 2014 vs 2015 partial).
No session/browse data	Cannot explain why discounts work or why certain countries underperform.	Recommendation 1 calls out the need to instrument behavioral tracking.
No product return data	Revenue is gross; returns and refunds are not netted out.	If material, this would be reflected in customer complaint volumes (also out of scope).

Table 6. Known limitations and how they are handled.

10. Reproducibility

To reproduce the dashboard from scratch:

- Open Tableau. Connect to the seven CSV files.
- Build relationships per Section 4.
- Create the six calculated fields per Section 5.
- Build each chart with the title and grouping shown on the published dashboard.
- Add the three cross-dashboard filters.
- Apply the color palette from Section 6.3.
- Validate using the three QA passes in Section 8.

Estimated time from clean Tableau to deployed dashboard: 6-8 hours, including QA.

Document Control

Version	Date	Change Summary	Author
0.1	2026-02-12	Initial draft - data sources and grain documented	Data & Analytics Team
0.5	2026-02-25	All six calculations listed; QA section added	Data & Analytics Team
0.9	2026-03-09	Design decisions section added	Data & Analytics Team
1.0	2026-03-15	Approved as reference document	Data & Analytics Team

Table 7. Revision history.