XFlowData | Case Study: Data Quality Validation Pipeline (Business Overview)

Data leaders know that poor data quality has a direct impact on decision-making, compliance, and customer experience.

Traditional controls often rely on static rules that quickly become outdated, requiring manual effort to maintain and adjust.

Our solution introduces an automated, adaptive, and policy-driven approach to data quality control.

Case Overview: E-commerce Order Validation

We'll examine how the system validates e-commerce order data, using actual sample data that contains both valid records and quality violations to showcase the platform's capabilities.

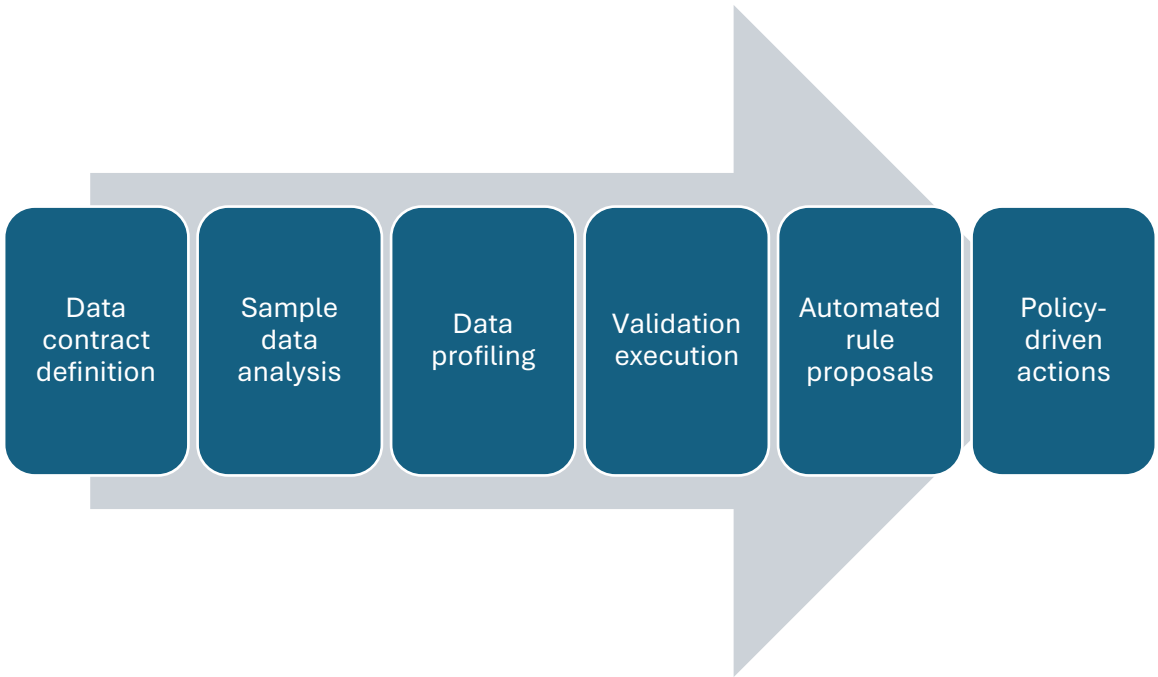

Step 1. Defining the "Data Contract"

A *data contract* is a formal agreement that specifies what "good data" looks like. For e-commerce orders, this includes rules such as:

- **Order ID** must be unique.
- **Order date** must be recent (e.g., within the last 1 year).
- **Amounts** must be positive and below €100,000.
- **Currency** must be limited to EUR, USD, or GBP.

This sets clear expectations between data producers and consumers, creating a common language around quality.

Step 2. Detecting Quality Issues

When new order data arrives, the system automatically checks it against the contract. In our example dataset, it detects:

- **Negative amount** (-7 EUR) violating the positive range expectation.
- **Unsupported currency** (MAD) not in the allowed set
- **New column** (extra_note) not defined in the contract.
- **Data type mismatch** (dates stored as strings instead of datetime objects)

These issues are flagged instantly, before the data is propagated into downstream systems.

order_sample.csv

order_id	order_date	amount	currency	**extra_note**
1	2024-10-01T10:00:00	12.50	EUR	ok
5	2025-09-01T08:00:00	**-7**	EUR	bad_amount
7	2025-09-04T00:00:00	42000	**MAD**	future_date

Step 3. Profiling the Data

Beyond simple rule checking, the system profiles the dataset to understand its real characteristics:

- Confirms the range of values **(–7 to 42,000 EUR**).
- Identifies all currencies used, including unexpected ones.
- Detects schema drift (**new or missing columns**).

This gives data leaders a statistical overview of where the dataset diverges from business expectations.

Step 4. Validating and Reporting

The validator generates a clear quality report. It categorises issues into:

- **Hard errors** → violations that require immediate action (e.g., rejecting records with invalid currency).
- **Soft warnings** → deviations that may be acceptable but should be monitored (e.g., slightly unusual amounts).
- **Schema changes** → alerts when new columns appear.

This structured reporting provides transparency and traceability.

Step 5. Proposing New Rules

The system doesn't just highlight problems — it learns from the data. It can automatically suggest adjustments, such as:

- Expanding the accepted amount range if the business agrees that negative refunds should be valid.
- Adding a new currency to the list if the company has expanded into new markets.
- Creating rules for unexpected fields that prove useful.

This adaptive capability reduces manual maintenance and accelerates alignment between business reality and data governance.

Step 6. Policy-Driven Actions

Every organisation has different tolerance levels for data issues. Policies allow leaders to define responses:

- **Quarantine** data with structural mismatches.
- **Reject** records that break hard rules.
- **Warn** when a soft expectation fails.
- **Propose** new rules for unexpected fields.

This ensures that controls are not only automated but also aligned with governance strategy.

Why It Matters for CIOs & CDOs

This system delivers concrete business benefits:

- **Automated Quality Gates**: Bad data is stopped at the source.
- **Adaptive Rule Management**: Rules evolve with the business, reducing costly manual intervention.
- **Transparent Reporting**: Quality insights are accessible to both technical and business stakeholders.
- **Policy Flexibility**: Different data domains can have different governance strategies.
- **Real-Time Monitoring**: With streaming support, data quality can be monitored continuously, not just in batch.

Bottom Line

XFlowData' approach transforms data quality management from a reactive, manual burden into a proactive, automated capability.

It gives CIOs and CDOs the confidence that their data is accurate, reliable, and aligned with business needs — even as the data landscape evolves.

Contents