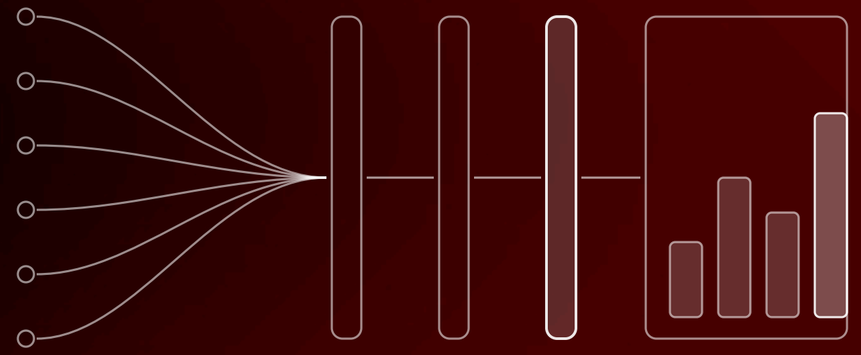
CASE STUDY | DATA & ANALYTICS | LOGISTICS

When Revenue Reporting Slows, Look Below the Dashboard

How ICS Compute modernized the data foundation behind revenue analytics for a major Indonesian logistics company.



Data architecture | Business logic | Governance | Validation | Performance

INTRODUCTION

When a revenue dashboard slows down, the cause often sits several layers below it. For a major Indonesian logistics company, those layers were the data architecture and the business logic that had built up around its reports.

As the company grew, more teams asked for more views of revenue and operations, and each request added another layer of logic to the reporting environment. Definitions, calculations and exceptions collected inside it because that is where people first needed them. As data volumes and report counts rose, the environment became harder to optimize and maintain.

ICS Compute modernized the data foundation behind the company's revenue and operational analytics, starting below the dashboard layer. The work redesigned the architecture and reorganized data processing while preserving critical business logic and validating reporting consistency.

Executive summary

SITUATION

A major Indonesian logistics company used a reporting environment that had grown together with the business. Revenue and operational analytics depended on it.

CHALLENGE

Rising reporting demand and data complexity increased the load on the environment. Important business logic had accumulated inside it, which made analytical workloads harder to optimize and maintain.

APPROACH

ICS Compute redesigned the data architecture into clear layers, reorganized data processing, moved reporting logic into a structured analytical model, and strengthened governance. Results were validated against the business definitions the company already trusted.

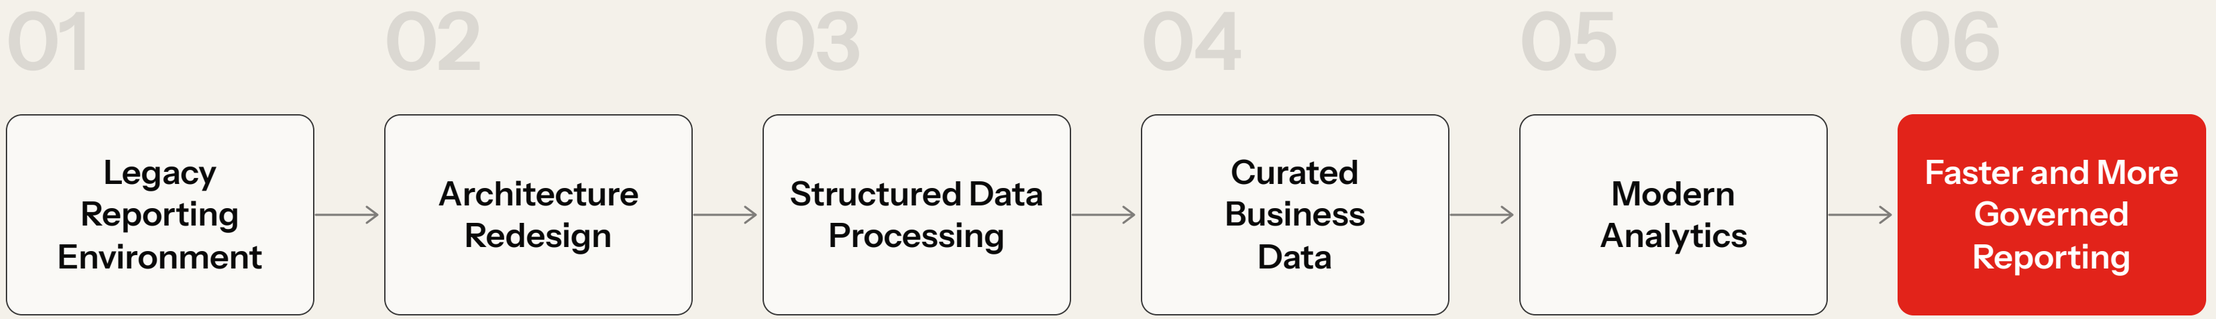
OUTCOME

Analytical workloads run significantly faster. Reporting is more consistent and easier to maintain, governance is stronger, and the company has a foundation for future Data & AI initiatives.

At a Glance

INDUSTRY	Logistics
BUSINESS AREA	Revenue and operational analytics
CORE CHALLENGE	Growing reporting complexity, business logic accumulated in the reporting environment, and rising pressure on analytical performance
MODERNIZATION FOCUS	Data architecture, data engineering, analytics modeling, performance and governance
ICS COMPUTE ROLE	Architecture, engineering, business logic refactoring, validation and handover
BUSINESS OUTCOME	Faster analytical workloads, more consistent reporting, stronger governance and a base for future Data & AI work

The Modernization Path



The architecture changes first, and reporting improves because the data underneath it is structured.

The Business Challenge

Enterprise analytics environments grow in layers. A finance team may ask for revenue by customer segment. An operations team may ask for the same revenue by service type. A regional leader may need a variation for a monthly review. Each request is reasonable, and each is usually answered where the data is already available, which is often the reporting layer itself.

After several years of this pattern, the reporting environment holds a significant amount of business logic. Calculations, definitions and exceptions sit inside reports and the queries behind them. The logic works and the business relies on it, so it has to be handled with care.

Growth also changes the load. More data and more reports mean the same logic runs more often over more records. Workflows that were easy to maintain at an earlier scale take more effort to change, and performance tuning becomes harder because the logic and the reporting layer depend on each other.

Governance needs rise at the same time. A larger organization needs to know who can see which data, where a number comes from, and whether two reports that show revenue use the same definition. Those questions are difficult to answer when business logic is spread across the reporting layer.

This pattern is a normal result of business growth. Reporting that served the company at one stage meets new demands at the next, and at some point the environment needs a data foundation designed for the current scale and the questions ahead.

Why This Was More Than a Migration

Moving data to a new platform is a well-understood task. Moving it while keeping the numbers people rely on intact is harder, and it was the central requirement of this project.

A reporting environment records decisions as well as data: how revenue is attributed, which transactions count toward a figure, how exceptions are handled, how periods are compared. Those decisions live in calculations. Over the years people across the company learned to trust the resulting numbers, and they use them in reviews, planning and performance discussions.

If the data moved to a new platform and the calculations changed even slightly, the same report could show a different figure the next morning. Finance, operations and leadership would then need to settle which number is right before they could discuss anything else. A faster dashboard that causes that question has reduced trust in reporting.

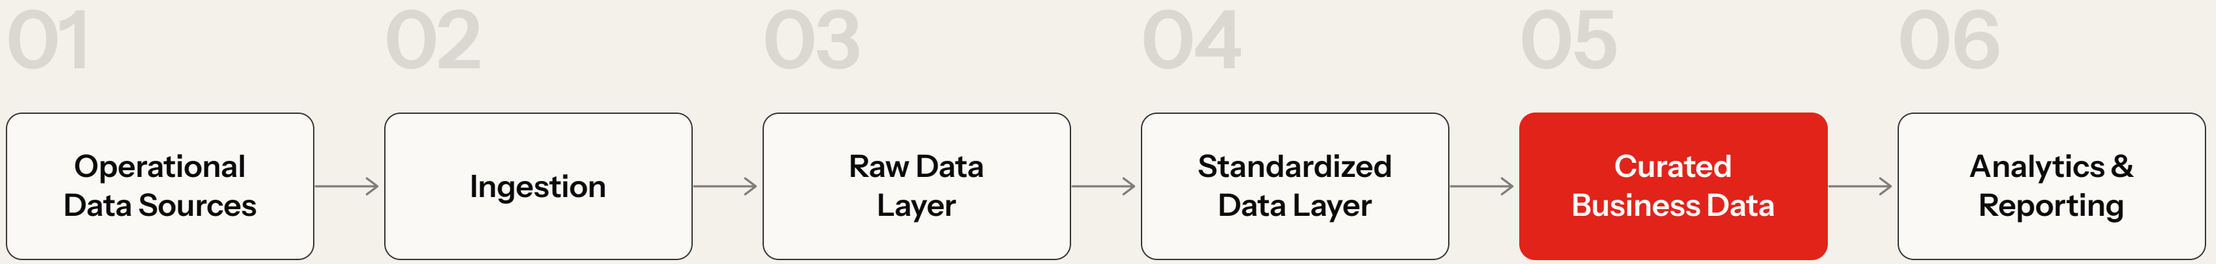
Key idea

Modernization succeeds when the business meaning of the data survives the move.

For that reason the team treated the existing calculations as requirements for the new design. Each business definition had to be carried into the new environment, and KPI consistency became a test condition. Validation was a defined part of the delivery, so the business could compare outputs from the new environment with the figures it already trusted.

High-Level Modernization Approach

ICS Compute redesigned the data environment as a set of layers with defined purposes. Data moves through the layers in one direction, and business logic is placed in specific layers. The design follows a lakehouse-style architecture that keeps raw, standardized and curated data separate.



Data moves in one direction through the layers, and reports read from the curated layer where business definitions are applied once.

Each layer can be tested and changed on its own, and a problem in a report can be traced back through curated and standardized data to the raw record.

RAW DATA

Every downstream result can be traced back to the record as received

STANDARDIZED DATA

Consistent structures and quality controls before business logic is applied

CURATED BUSINESS DATA

Business definitions and calculations applied once and reused across reports

ANALYTICS AND REPORTING

Dashboards that read prepared data and have less work to do when opened

Business Logic Refactoring

In the legacy reporting environment, much of the business logic sat close to the reporting applications. Calculations were applied inside reports and the queries behind them, and similar transformations could appear in more than one place. This is a common pattern in environments that grow with the business, because logic is added where it is first needed.

Refactoring moved that logic into structured data processing. The team identified the calculations that reports depended on, separated those that belong in data preparation from those that belong in presentation, and rebuilt them as reusable processing steps that produce curated analytical datasets. Reports then read from those datasets.

The engineering principle is to implement each business definition once, in a governed place, and let everything else consume it. When a definition changes, the change is made once and tested once.

	LEGACY ENVIRONMENT	MODERNIZED DATA FOUNDATION
Business logic	Closer to reporting applications	**Structured transformation layers**
Transformations	Repeated in several places	**Reusable curated datasets**
Traceability	Harder to follow	**Clearer lineage**
Processing and reporting	Greater dependency between them	**Separated**

Reporting Modernization

Reporting modernization in this project involved the logic and structure behind each report as well as its appearance. Dashboards are where business users meet the data, so their figures have to match the numbers people already use. The work covered four areas.

01 Business definitions were preserved, so reports continued to calculate figures according to the definitions the company already used.

02 Outputs were validated, with results from the modernized reports checked against trusted reference results.

03 Analytical structures were improved, so reports read from datasets designed for the questions people ask.

04 Processing close to the dashboard was reduced, because calculations that reports once performed at query time are now prepared in the data layers.

Performance Optimization

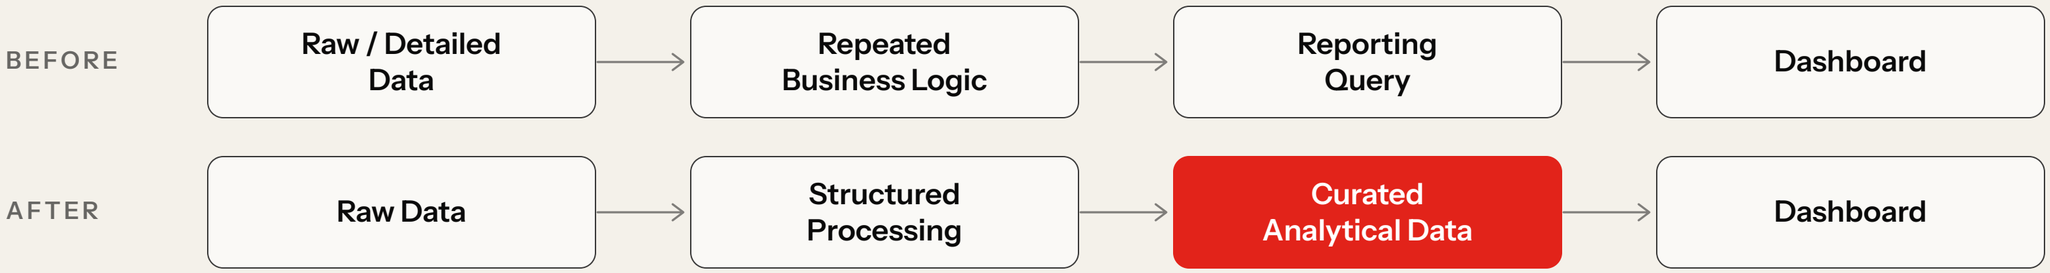
Important analytical workloads became significantly faster, and business users now have faster access to business-critical reporting.

Key idea

A slow dashboard is a visible symptom, and its cause often sits several layers below it.

Tuning a dashboard in isolation has limits when the delay comes from the data underneath it. A report that applies business logic to detailed records every time it runs spends much of its effort on that work, and changes to the visualization layer cannot remove it. Performance improved here because the work moved into the data processing layers, where logic is applied in advance and stored in analytical structures built for reporting queries.

Layering also helps with future tuning. Each step can be measured separately, so the step that limits performance is easier to find.



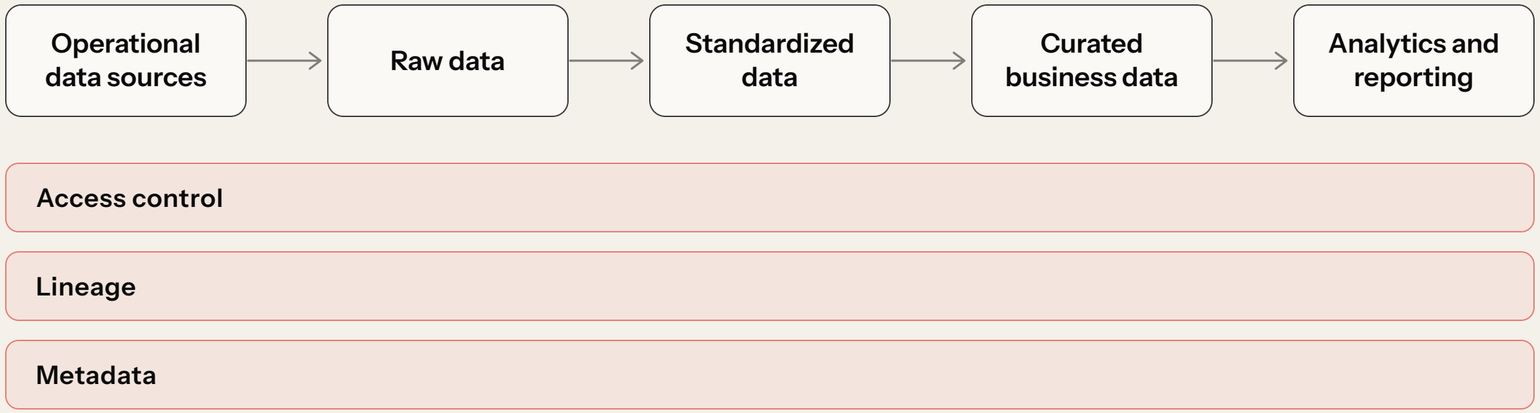
Move repeated processing upstream so reporting consumes business-ready data.

Data Governance

Governance is part of the architecture of the modernized environment. Access control determines who can see each layer and each dataset, so sensitive data reaches only the people who need it. Lineage records how data moves from source through each layer to the final report. Metadata describes what each dataset contains and who is responsible for it.

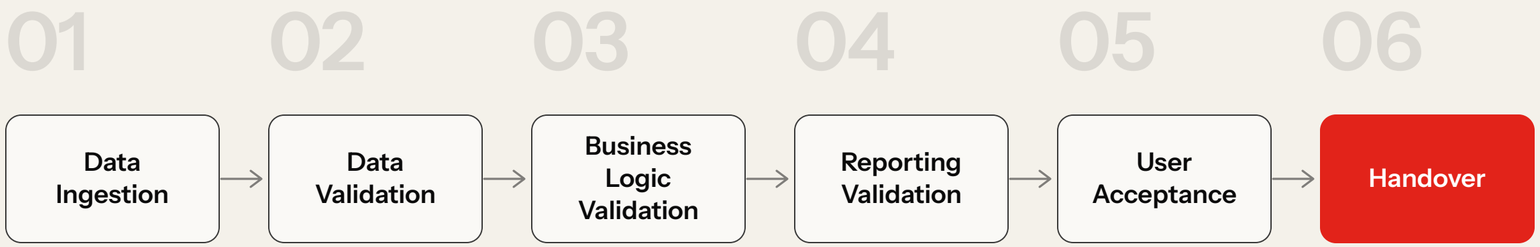
With these controls, a figure in a report can be traced back to its origin, and a reviewer can follow that path without asking the engineers who built it. Data that passes through the standardized and curated layers also carries stronger preparation and quality controls, which improves reporting consistency.

Consistent business definitions come from applying each definition once in curated data, so reports that use it show the same result.



Validation and Handover

For this engagement, modernization extended beyond technical implementation. The business needed to accept the reporting outputs, and the internal team needed to be ready to operate the environment. Validation followed the path of the data, from ingestion to the final report, and ended with acceptance by business users.



Each validation step builds on the one before it, and the business accepts the results before handover.

- **Data integrity validation** was used to verify that data moves through the layers completely and accurately.
- **Business logic validation** compared calculated results with the figures the business already trusted.
- **Functional validation** was used to verify that pipelines and reports behave as designed.
- **Performance testing** measured how analytical workloads run in the new environment.
- **Operational testing** was used to verify that the environment runs reliably in day-to-day conditions.
- **User acceptance** gave business users the chance to confirm that reports show the figures they expect.
- **Documentation and knowledge transfer** prepared the company's team to operate, support and extend the environment.

Business Impact

The modernization improved six areas of the company's analytics environment.

Analytics Performance	Analytical workloads run significantly faster, giving business users quicker access to business-critical reporting
Reporting Consistency	Business definitions are applied once in curated data, so reports that use them show consistent results
Maintainability	Logic sits in structured, reusable layers, which makes changes easier to implement and test
Governance	Access control, lineage and metadata give clearer control over data and where it comes from
Scalability	A layered architecture gives the environment room to support more data sources, reports and use cases
Future Data & AI Readiness	Governed, consistent data provides the foundation that later analytics and AI initiatives depend on

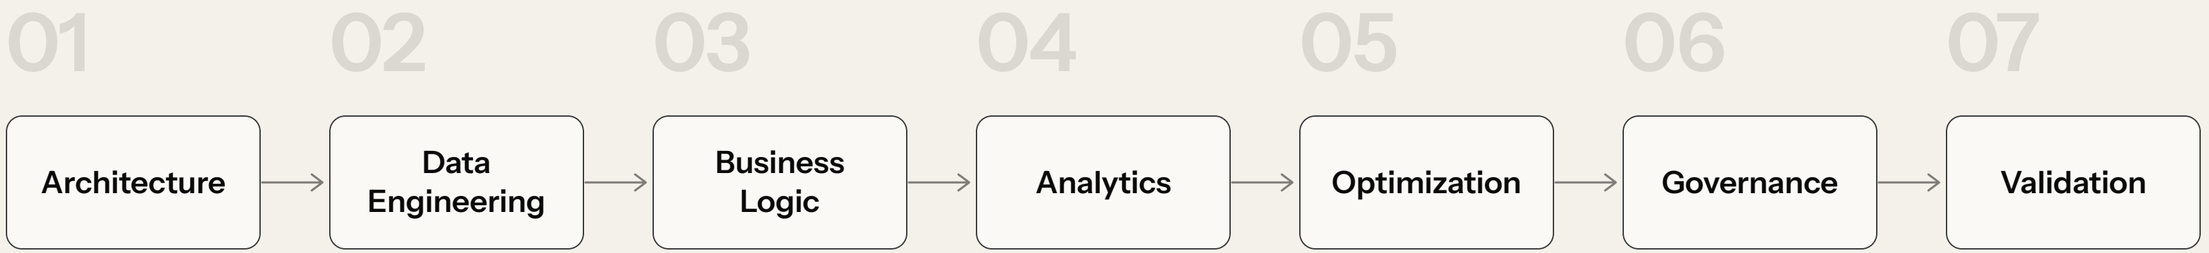
Why This Matters for AI

AI cannot compensate for inconsistent data definitions. A conversational analytics tool that answers questions about revenue can give different answers on different days when the data behind it defines revenue in more than one way. Intelligent assistants, predictive systems and AI workflows depend on the same thing: enterprise information that is consistent, governed and trusted.

The layered architecture, curated business data, consistent definitions, lineage and access control created in this project are the foundation those AI use cases require. This engagement focused on building the trusted data foundation required for future AI workloads.

ICS Compute's Role

ICS Compute delivered the modernization end to end, from architecture to handover. The work covered data architecture, ingestion design, data engineering, data modeling, business logic refactoring, analytics modernization, performance optimization, governance, validation and handover.



One engagement covered each stage, so architecture, business logic and reporting decisions were made together.

The same capabilities apply across ICS Compute's data practice, which builds data foundations on AWS and Databricks and connects them to analytics, governance and production operations.

Key Lessons

01

A slow dashboard points to the layers below it.

When reports are slow, the useful question is where the logic and the processing sit. Tuning the visualization layer alone leaves the cause in place.

02

A migration needs to preserve business logic.

Moving data without the calculations behind it creates risk, because the same report can show a different number on the new platform.

03

Governance belongs in the architecture.

Access control, lineage and metadata work best when the layers are designed with them from the beginning.

04

A data foundation is a practical first step toward AI readiness.

AI initiatives depend on consistent, governed data, and that work can start with the reporting environment a company already has.

Conclusion

The project modernized the foundation underneath the company's dashboards while preserving critical business logic and validating reporting consistency. Faster analytics followed from a layered architecture, reusable business logic, clear governance and careful validation. The company now has a reporting environment it can maintain and extend, and a governed data foundation future AI initiatives can build on.

Is your reporting environment carrying more business logic than it should?

ICS Compute helps enterprises modernize the data foundation behind their reporting and analytics, from architecture and data engineering to governance, validation and operations, so that future Data & AI initiatives start from data the business trusts.

DATA FOUNDATION
Curated and governed data

→

ANALYTICS
Consistent, faster reporting

→

AI
Future AI workloads

Discuss Your Data & AI Roadmap →

ICS Compute
PT. Innovation Cloud Services · 2026

www.icscompute.com sales@icscompute.com