

Feltex

CLIENT

Industry &
Infrastructure

INDUSTRY

Design · Develop

DELIVERED AS

• Microsoft Fabric

• OneLake

• Delta Lake

• PySpark

POWER BUSINESS DECISIONS WITH CLOUD SCALE ANALYTICS

A governed data foundation for Feltex on Microsoft Fabric

A metadata-driven ingestion platform on Microsoft Fabric that lands data from fourteen of Feltex's systems in one governed Bronze layer, ready for the analytics that come next.

01

THE CHALLENGE

Feltex is one of South Africa's leading automotive component manufacturers, a KAP business supplying acoustic, trim and foam components to the country's vehicle makers and the aftermarket. Its operational data lived in many places: ERP and finance platforms, manufacturing execution and production monitoring, and HR and safety records kept in SharePoint, each reached a different way. Feltex wanted one controlled place for all of it to land, loaded the same way every time, with clear sight of what ran, what failed and what changed. It also needed the access control and South African data residency an enterprise platform has to have before anyone builds reports on it.

02

THE APPROACH

We designed and built the Bronze layer of Feltex's data platform on Microsoft Fabric: the raw, source-aligned foundation every later layer reads from. Rather than write a pipeline per table, we built four generic, parameterised ingestion engines, plus a PySpark notebook for REST API data, and drive them all from a central control register. Each source has a thin parent pipeline that only knows how to connect. What to load, and how, lives in the register as data. Loads are incremental, keyed on watermarks that move forward only when a table lands successfully, so a failed load keeps its recovery point, and the largest systems run in sequenced load waves. Every record carries its source, pipeline, run and landing time, and every run is logged per stream and per table, with row counts for reconciliation. Access runs through Entra ID groups, scheduled work runs under a workspace identity, secrets live in Azure Key Vault, and settings for development and production sit in variable libraries so changes promote cleanly.



Feltex

CLIENT

Industry & Infrastructure

INDUSTRY

Design · Develop

DELIVERED AS

• Microsoft Fabric

• OneLake

• Delta Lake

• PySpark

03

THE OUTCOME

Feltex formally accepted the platform on 15 September 2026. It now has a single, controlled entry point for enterprise data: fourteen sources and 830 Delta tables landing in dedicated Bronze lakehouses in OneLake, each with its own SQL analytics endpoint. Onboarding a new table is mostly a matter of adding a row to the register rather than writing extraction code, and every load can be traced to its run, checked by row count and recovered from its last good watermark. We documented the platform and handed it over to Feltex's own team, and it is now the governed base for the Silver and Gold layers, semantic models and reporting that Feltex builds next.

04

RESULTS

- ✓ 14 — source systems on one ingestion pattern
- ✓ 830 — Bronze Delta tables, landed and logged
- ✓ 4 — generic engines in place of per-table pipelines
- ✓ 1 — central register driving every load

05

SUMMARY & BENEFITS

- ✓ One metadata-driven framework for every source: four generic engines driven by a central control register, with no per-table extraction code.
- ✓ Incremental loads on watermarks that move only on success, so a failed load never loses its place.
- ✓ Traceable end to end: lineage metadata on every record, run and table logs, and row-count reconciliation.
- ✓ Governed from the first run: Entra ID role groups, a workspace identity, Key Vault secrets and data hosted in South Africa.
- ✓ Ready for Silver, Gold and reporting, through a SQL analytics endpoint on every Bronze lakehouse.

Want a result like this?

Talk to us about your own challenge — info@firsttech.digital or +27 10 501 0800.

Let's Talk →