

IBERGRID 2026 · SANTIAGO DE COMPOSTELA · 27–29 OCTOBER 2026

A Data Lake for unified data exploitation at the Gran Telescopio Canarias

Josué Toledo Castro, Manuel Huertas López, Romano Corradi, Luis Alberto Rodríguez García, Sebastián Fernández García, Henry Brito Martín, Ismael Álvarez Candelario

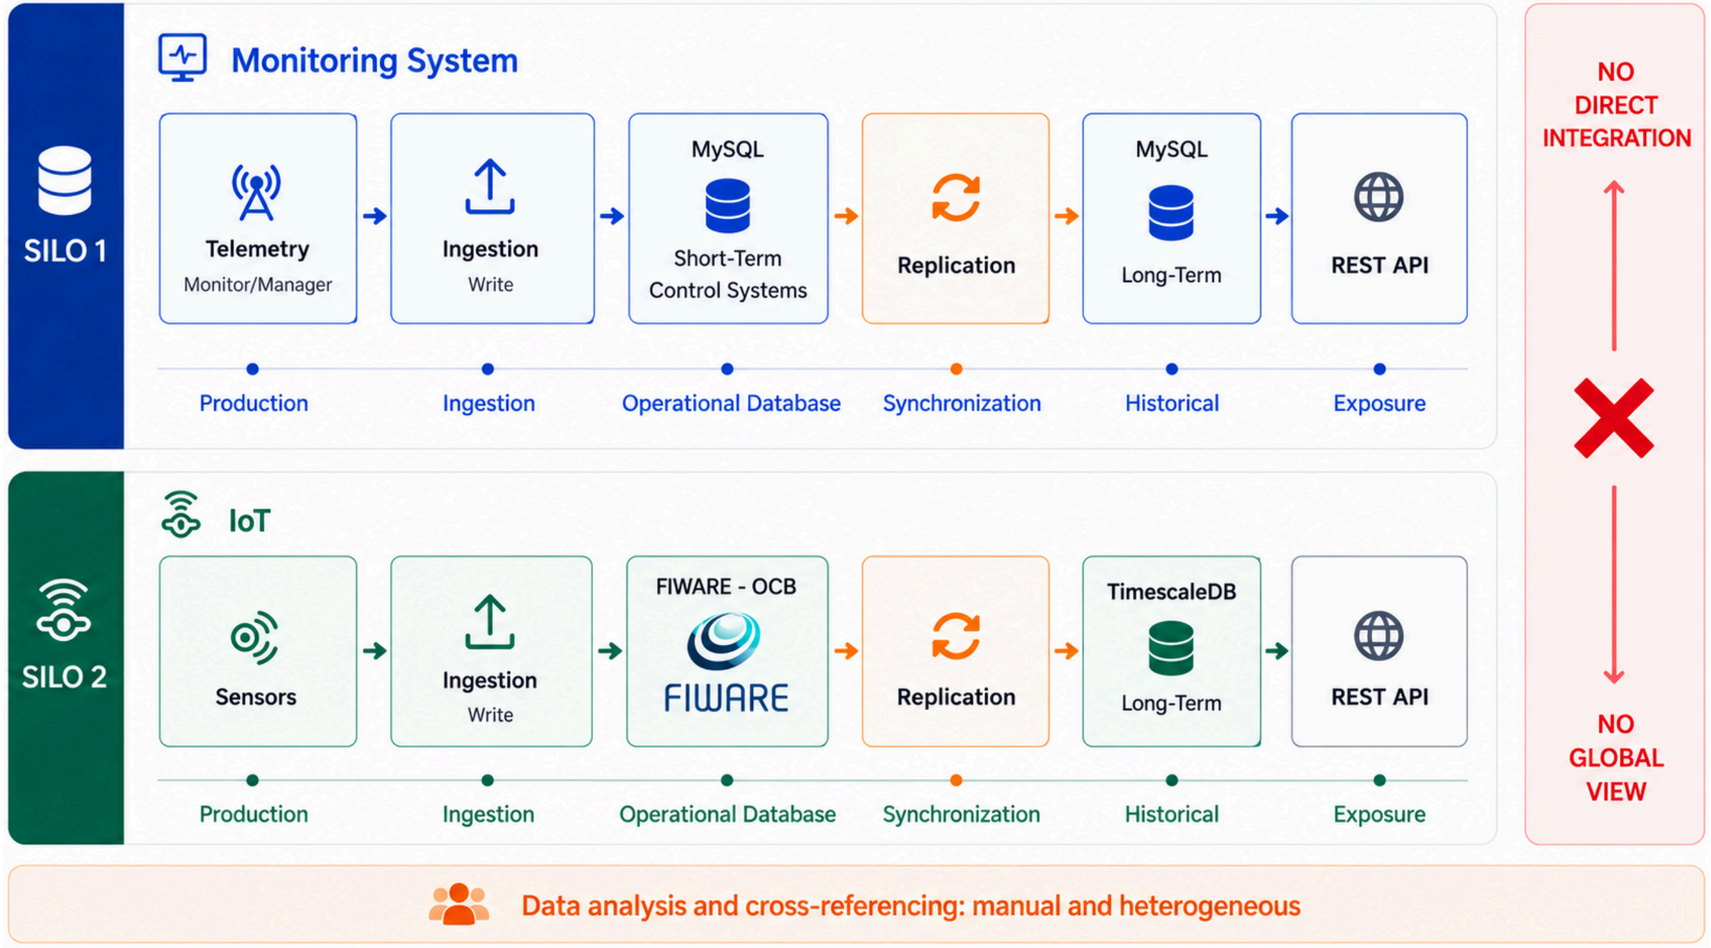
Gran Telescopio Canarias (GRANTECAN) | Engineering Operations Group (GOI) | josue.toledo@gtc.iac.es

Image: Daniel López / IAC

Starting point I

A heterogeneous, fragmented data ecosystem built on data silos.

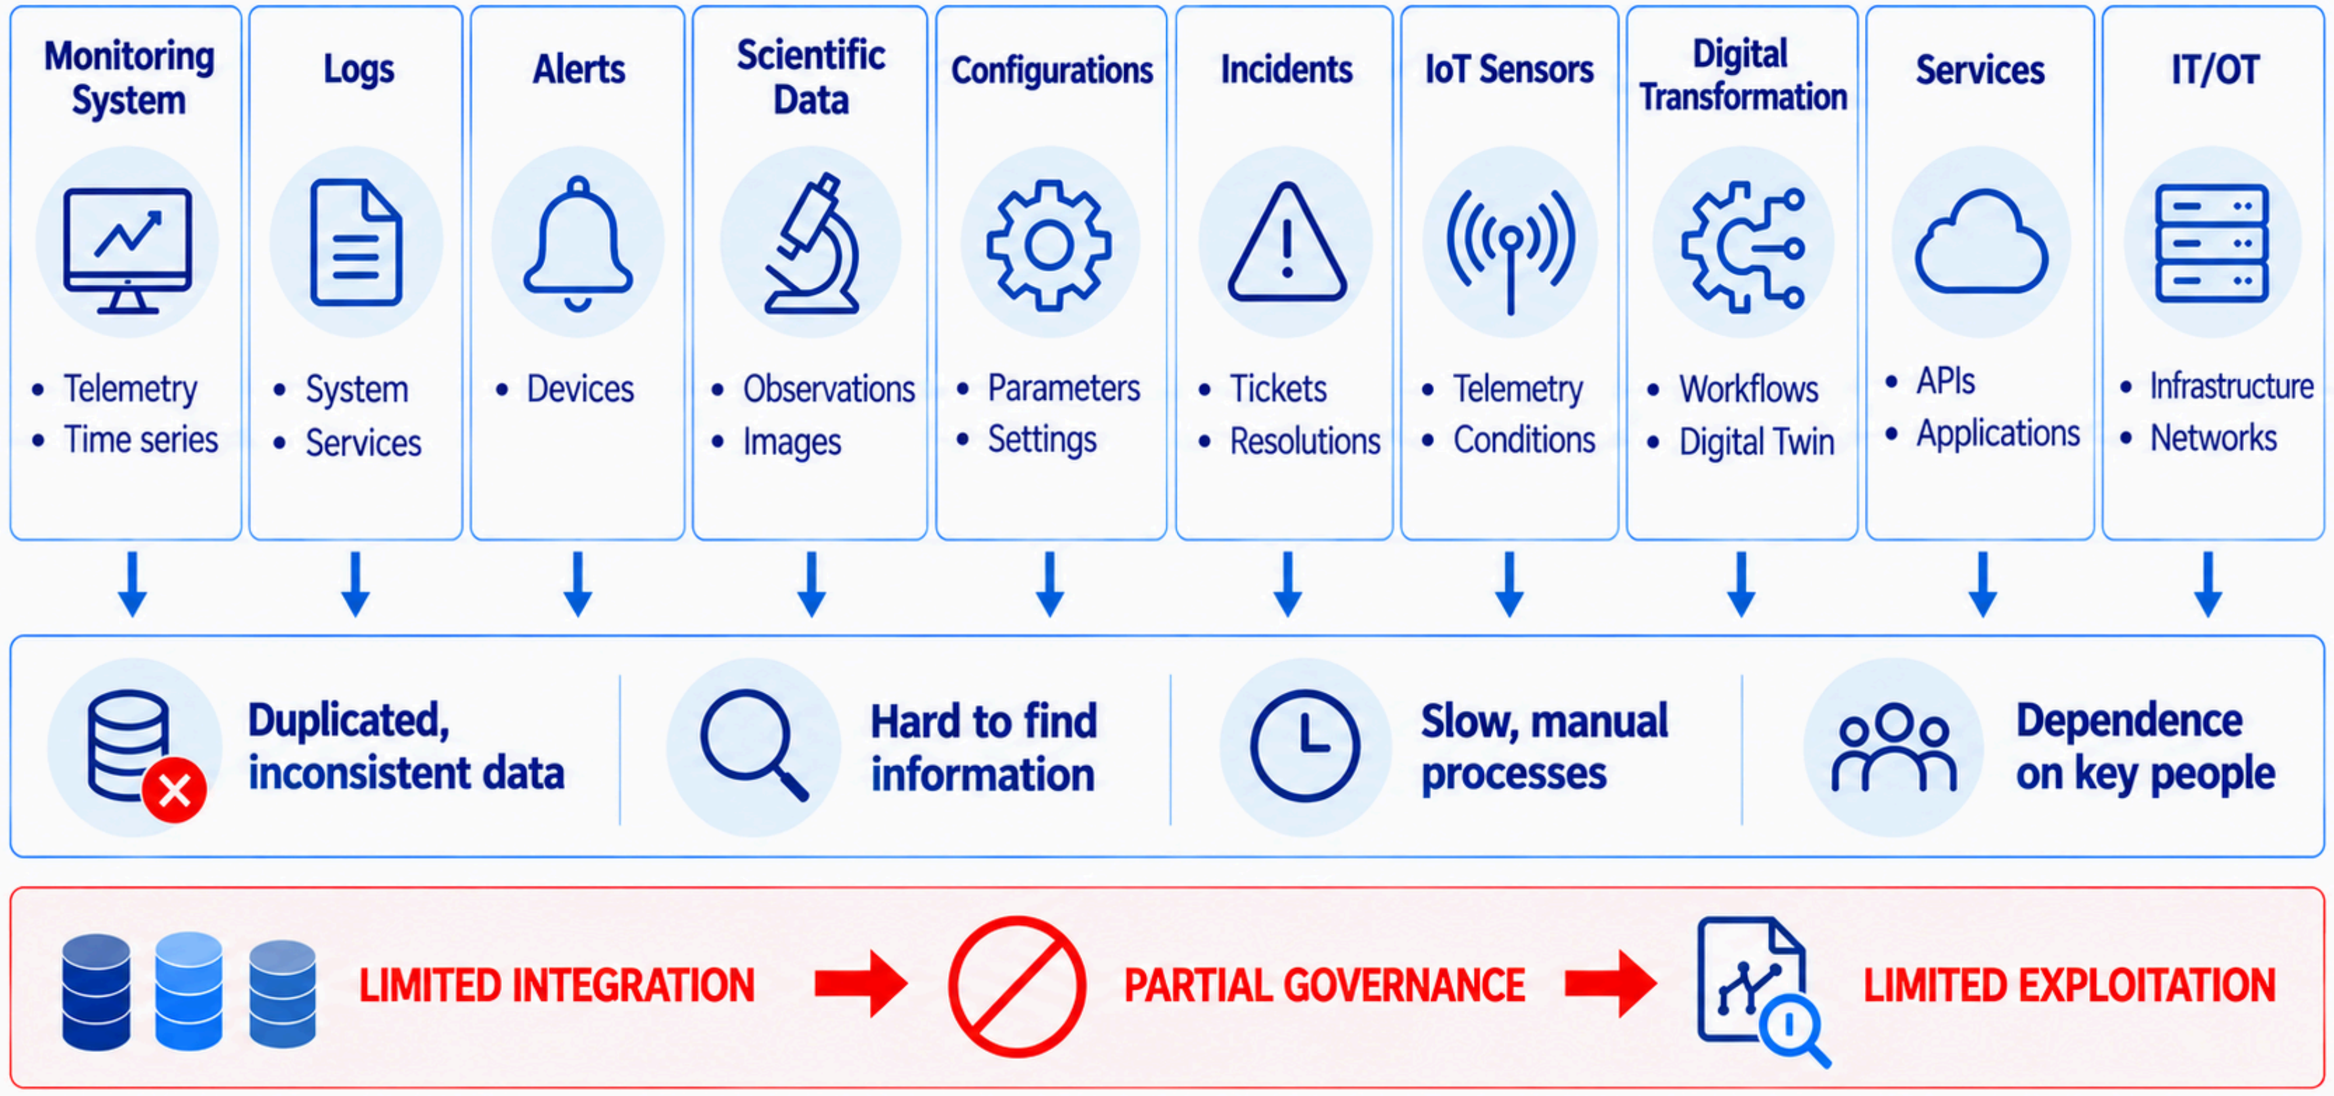
- **Disconnected** data silos
- **Heterogeneous** formats
- **No unified** access
- **Complex** cross-analysis
- **Limited** scalability



Technical data (3 Vs) | Volume · Velocity · Variety

Starting point II

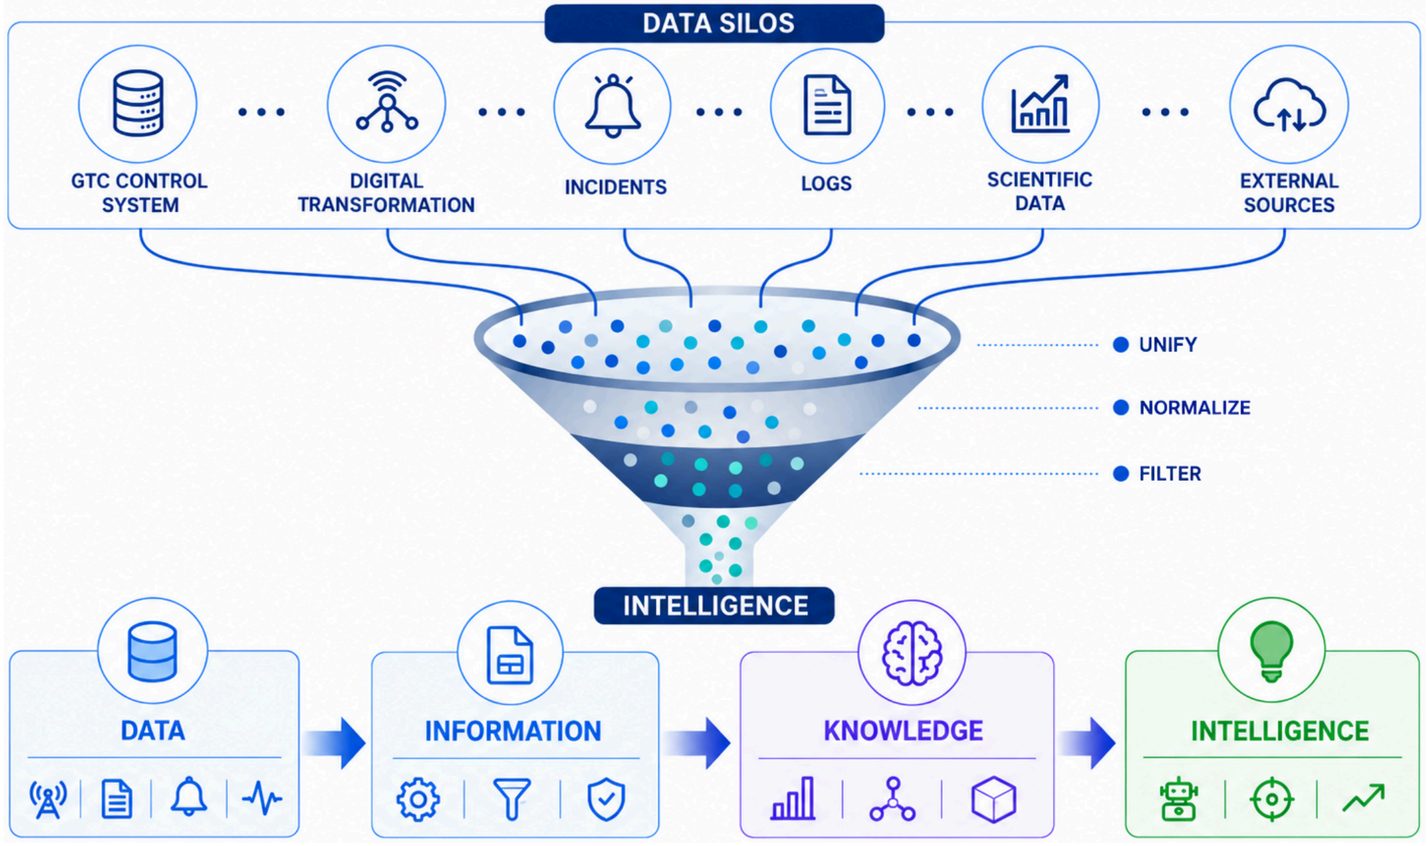
Many sources, each with its own storage and tools.



Objectives

From isolated data to connected, governed and useful knowledge.

1. **Connect** data now spread across silos
2. **Normalise** formats, structures and meaning
3. **Unify** access, querying and exploitation
4. **Build in** governance, quality and security
5. **Enable** advanced analytics, AI and collaboration
6. **Democratise** the use of data

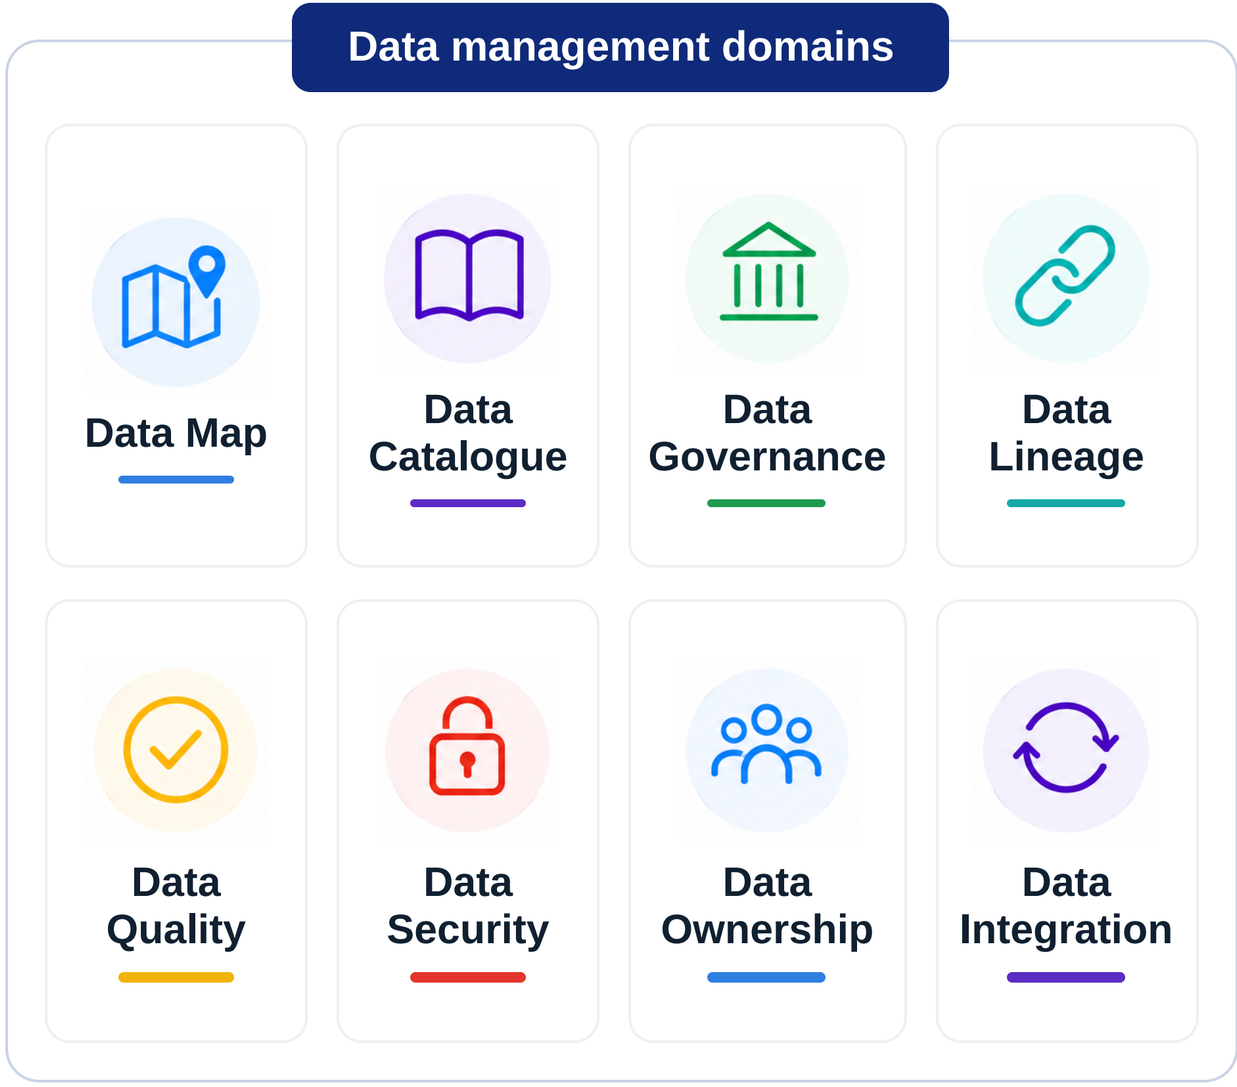


Data Management Plan + **Governed Data Lake**

Data Management Plan

Eight practical domains to manage data across its whole lifecycle.

Data management domains



INVENTORY

Map · Catalogue

- **Map** every source and dataset
- **Catalogue** schema, owner, meaning

GOVERN

Governance · Ownership

- **Owners** per domain
- **Rules** of use and approvals

SECURE

Quality · Security · Lineage

- **Quality** checks and lineage
- **Access** control and audit

INTEGRATE

Integration

- **One platform**, open formats
- **Homogeneous access**, no new silos

FAIR principles

Data that can be found, accessed, understood and reused.

F Findable

- Identifiers
- Metadata
- Catalogue

A Accessible

- Access protocols
- Role-based permissions

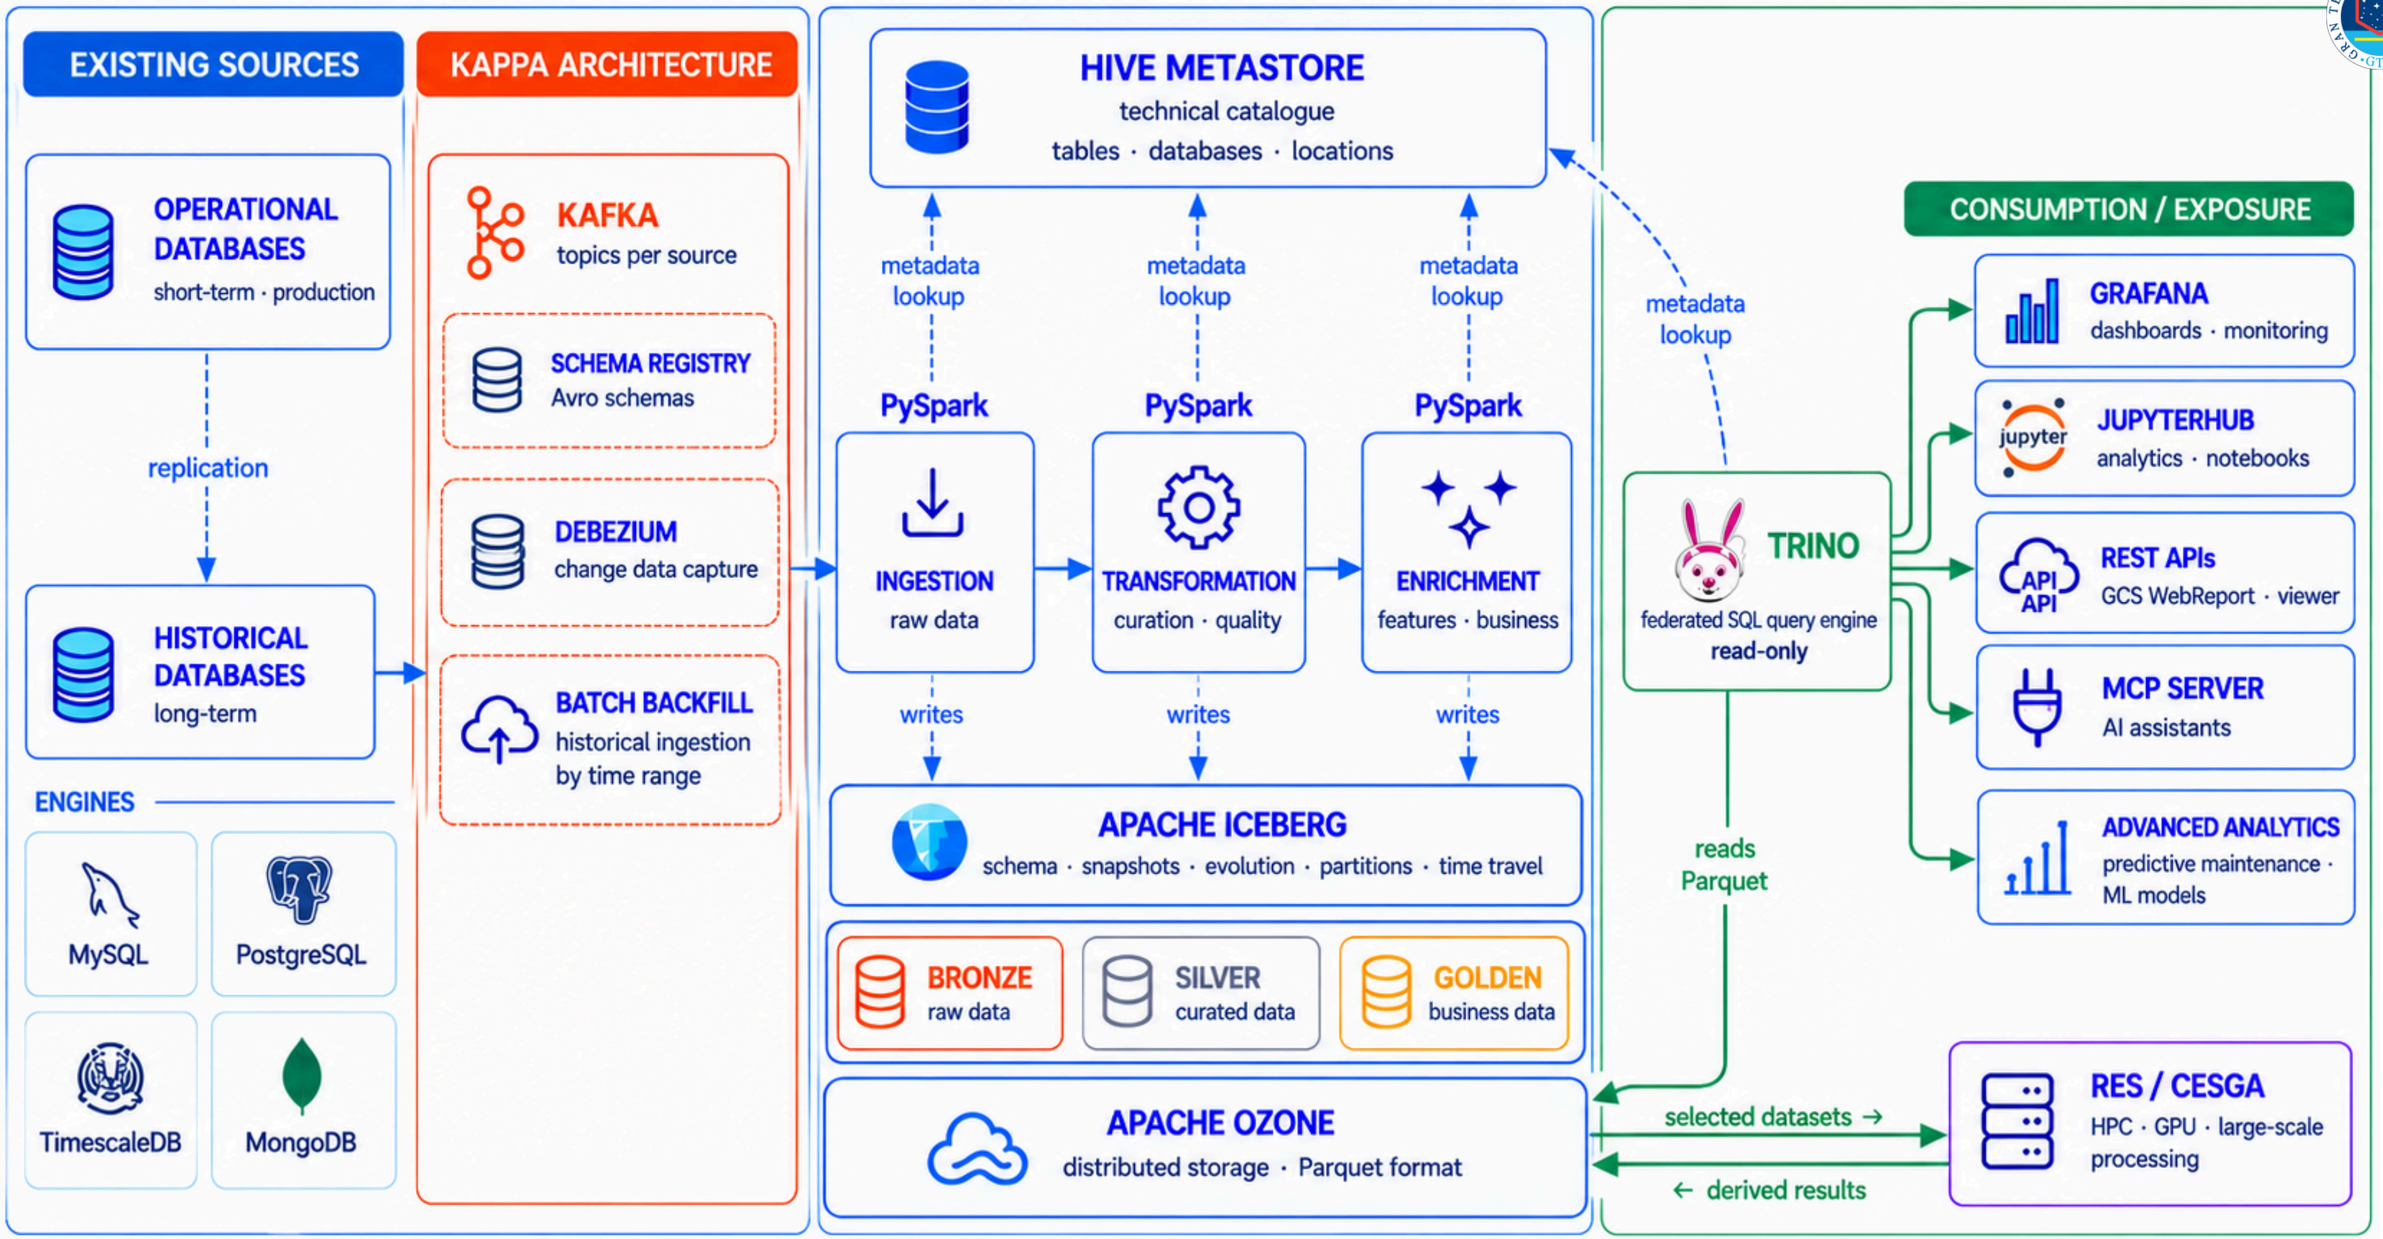
I Interoperable

- Open formats
- Common vocabularies

R Reusable

- Licences
- Context
- Provenance

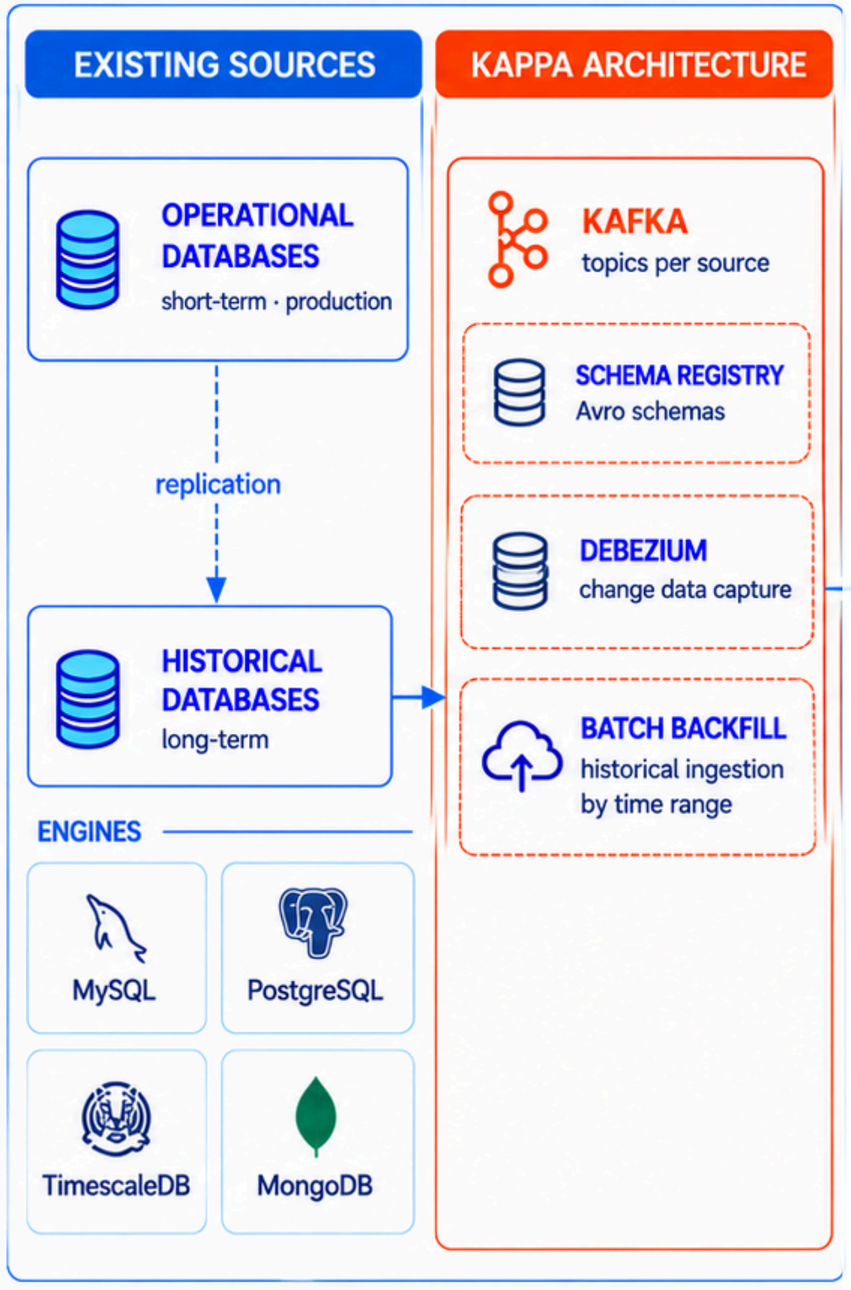
Data contracts | One agreement per source: **what** is delivered, **what it means**, **who owns it** and **who can use it**



1 · Ingestion layer

2 · Processing & storage layer

3 · Exposure layer



1 · Ingestion

Sources

- MySQL: telemetry, logs, alarms, scientific data
- PostgreSQL / TimescaleDB: IoT and monitoring
- MongoDB: configurations and documents

Change Data Capture

- Debezium reads each database's change log
- One Kafka topic per table
- Avro schemas in Schema Registry

Historical backfill

- Spark batch jobs by time range
- Progress tracked in ingestion_progress
- Idempotent re-runs

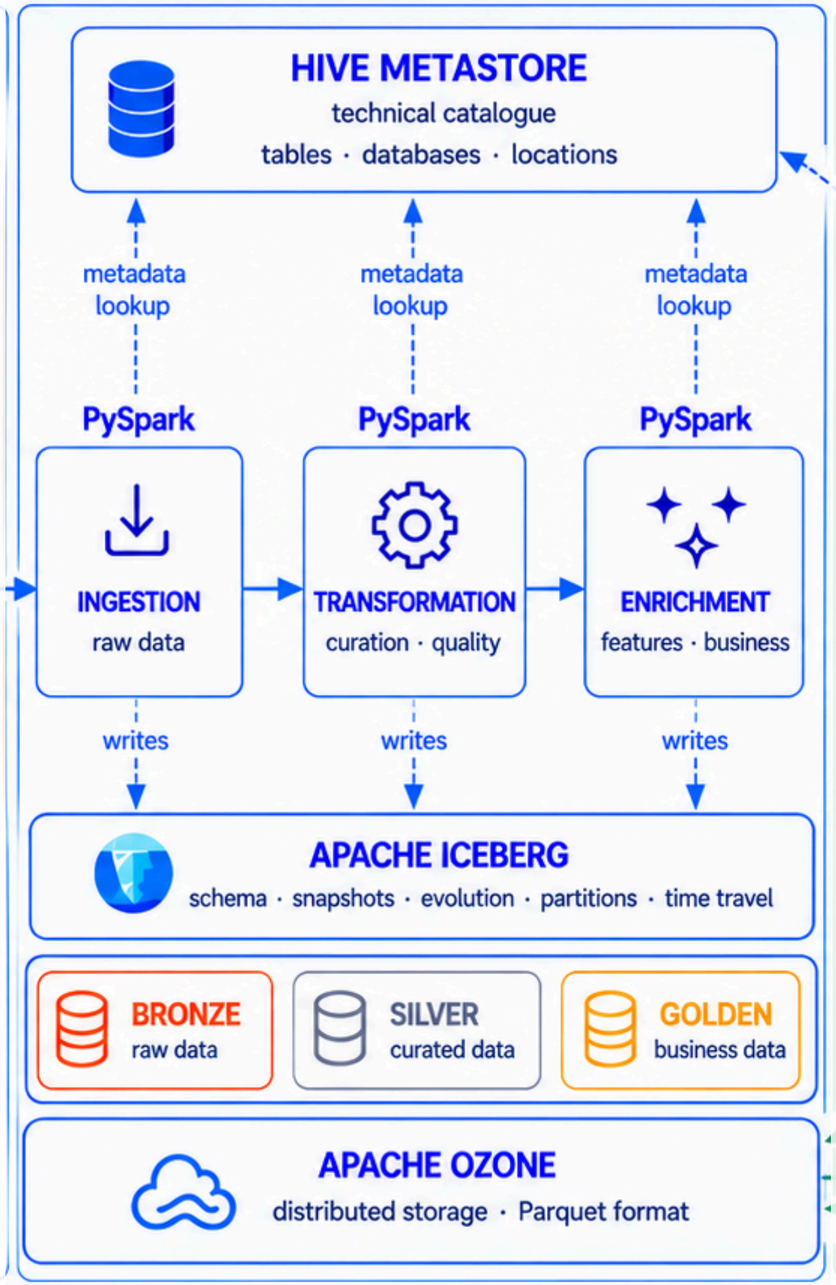
2 · Processing & storage

Spark on Kubernetes

- Streaming apps
- Batch apps
- Generic apps, configured per source

Medallion layers

- Bronze: raw, as in source
- Silver: curation and transformation
- Golden: use-case datasets



Lakehouse storage

- Iceberg tables
- Parquet files on Ozone
- Hive Metastore catalogue
- Daily compaction

Quality & observability

- Completeness checks
- Kafka Connect alerts
- Prometheus metrics
- Grafana dashboards

3 · Exposure

Trino

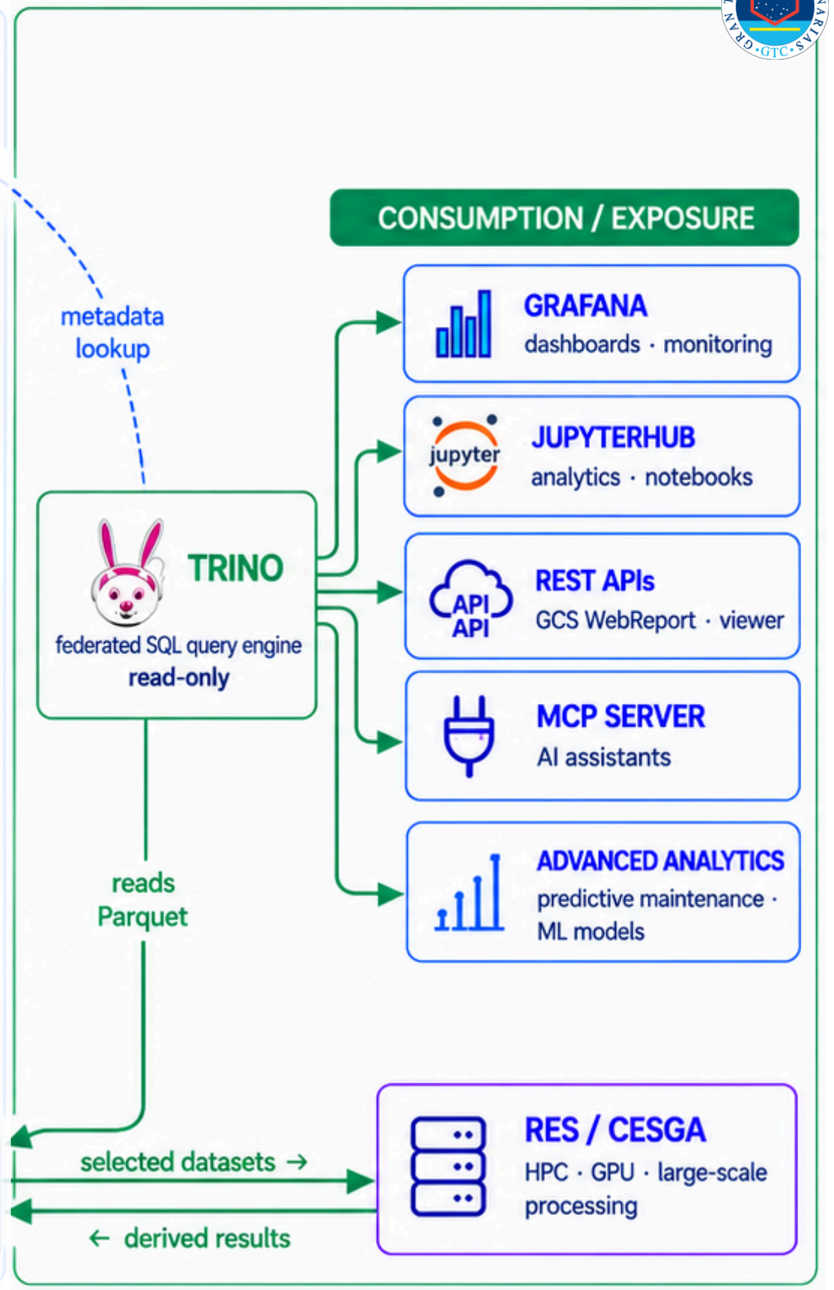
- Single SQL entry point
- Read-only by default
- Entra ID tokens over HTTPS

Consumers

- JupyterHub notebooks
- Grafana dashboards
- REST APIs
- AI assistants through MCP

Catalogue

- Hive Metastore: tables and locations
- Schemas and tables visible in SQL
- Same catalogue for Spark and Trino



The life of a telemetry sample

From a row in the source database to a committed Iceberg snapshot in Bronze.

1 Source row

MySQL · Monitor Manager

```
id_monitor 4711
time_sample …456
value      12.83
```

→

2 Change event

Debezium · binlog

```
Insert event
Avro · schema 87
```

→

3 Kafka message

One topic per table

```
Partition 3
Offset 918204
```

→

4 Micro-batch

Spark · every 5 min

```
Avro → columns
Append to Bronze
```

→

5 Iceberg commit

Ozone + Hive Metastore

```
Parquet written
New snapshot
```

What lands in Ozone

```
deltalake-bronze-monitor-manager/monitor_sample_double/
├── data/
│   └── time_sample_day=2026-10-21/        Partition
│       └── 00042-…-00001.parquet          Rows + statistics
└── metadata/
    ├── 01234-….metadata.json              Table state
    ├── snap-…-1-….avro                    Manifest list
    └── …-m0.avro                          Manifest
```

What we gain

- **Hidden partitioning** by day
- **Pruning** with min/max statistics
- **Atomic commits**
- **Time travel** and rollback
- **Compaction** to 128 MB files

Setting up a new data source

A repeatable procedure from the producer to SQL, the same for every source.

1 Source ready

- Change log enabled
- Replication user

`MySQL · PostgreSQL`

→

2 CDC connector

- Debezium per source
- One topic per table

`Kafka Connect`

→

3 Live to Bronze

- Raw, as in source
- Streaming

`Spark · Iceberg`

→

4 Backfill

- History by time range
- Resumable and idempotent

`Spark batch`

↓

5 Audit

- Coverage checks
- Counts against the source

`Iceberg metadata`

←

6 Silver

- Curated data
- One table per device

`Spark`

←

7 Golden

- Use-case datasets
- Anomaly scores

`Spark`

←

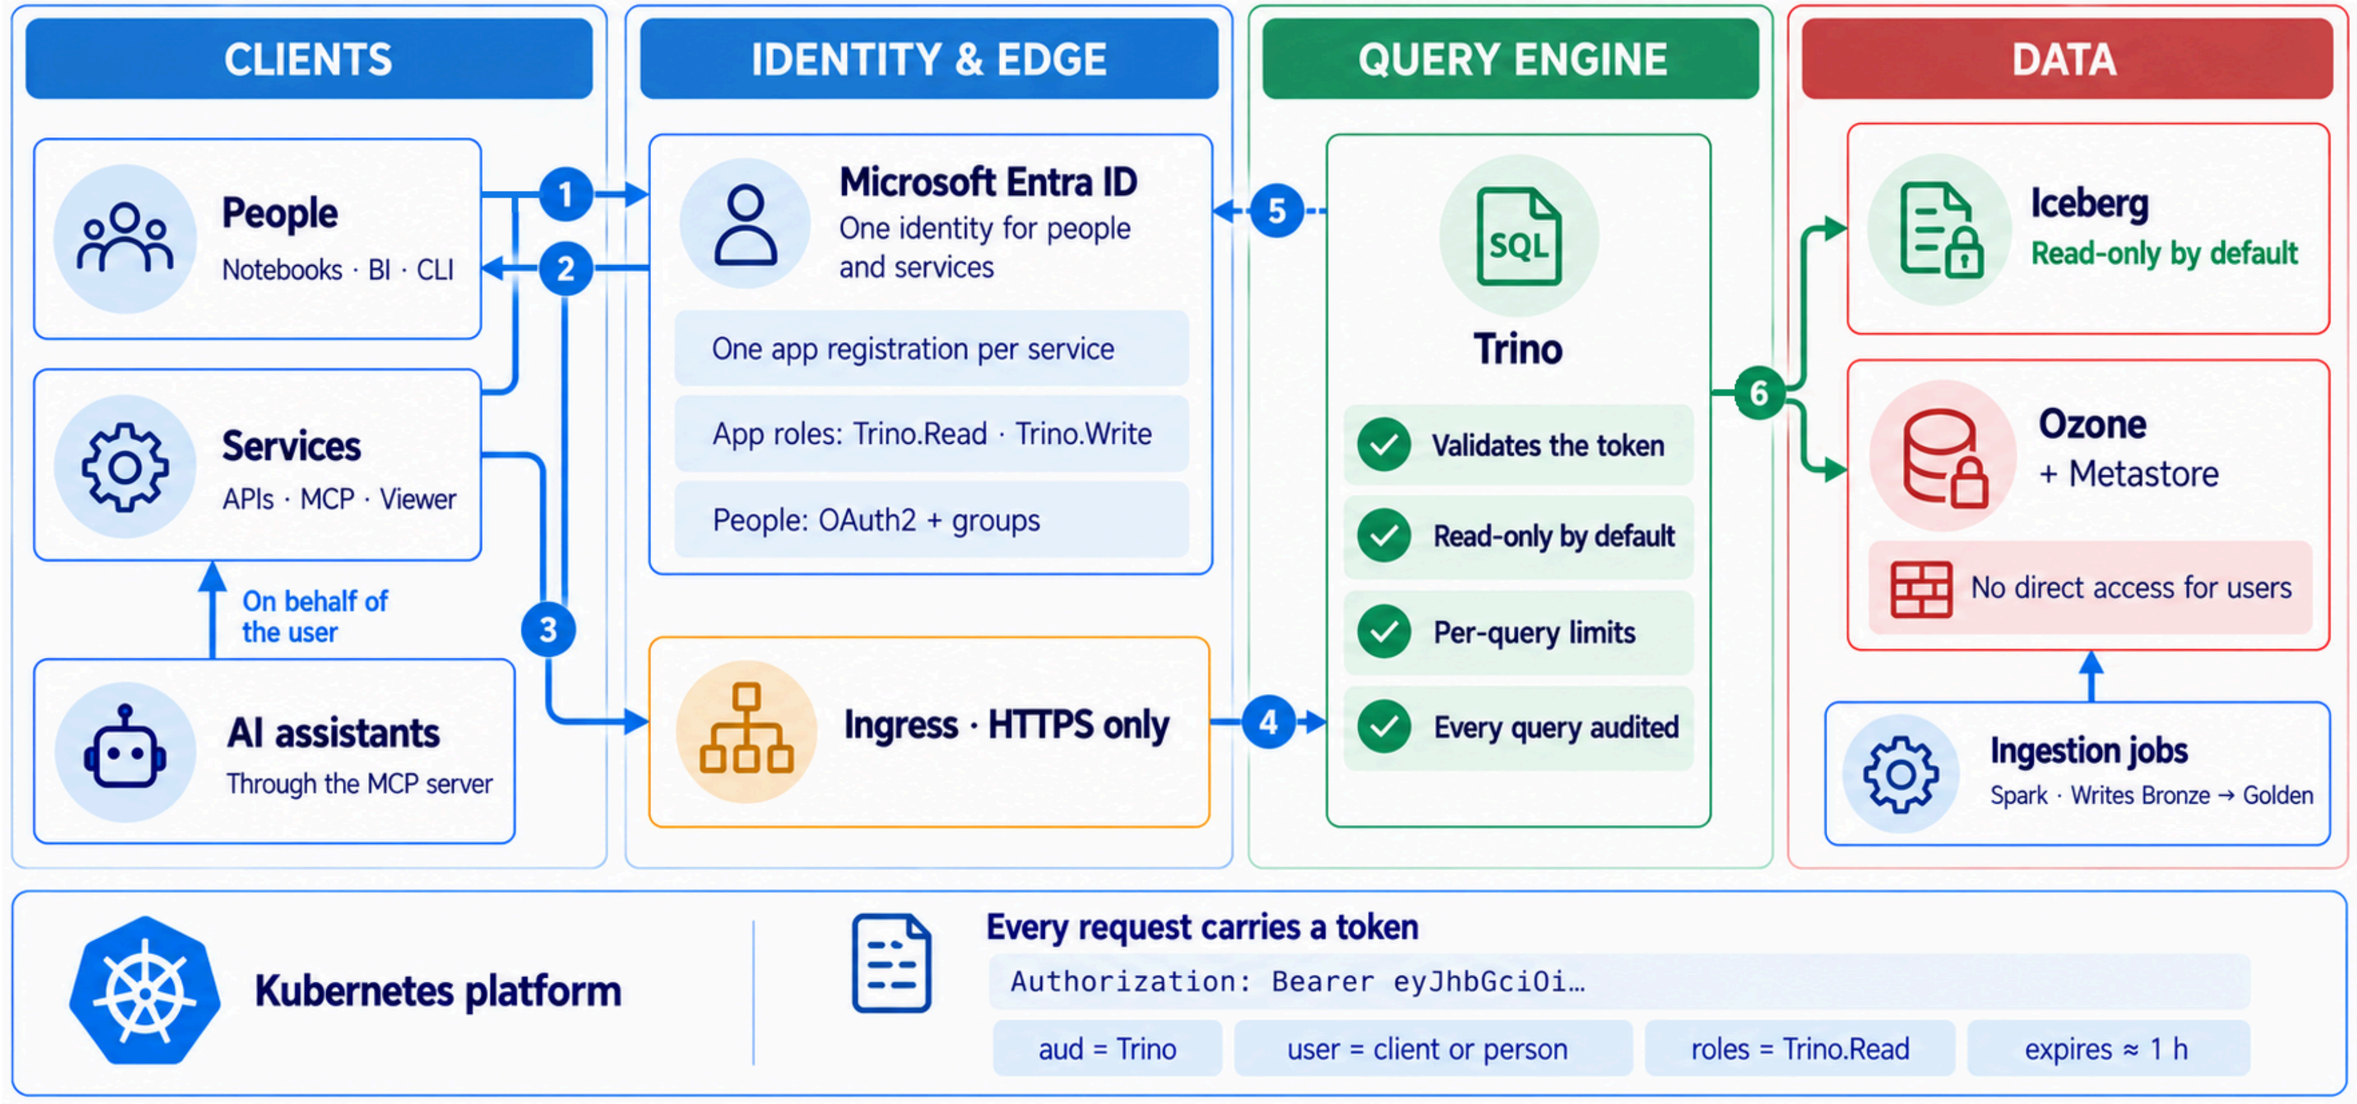
8 Publish

- SQL through Trino
- Notebooks · Dashboards
- APIs · AI assistants

`Trino`

Security by design

Target design, being rolled out in phases: one identity, one HTTPS entry point, read-only by default.



Scale today

The lake as it stands today, measured on 2026-10-23.

101 tables

In 9 buckets, three layers

- **Bronze**: 75 tables
- **Silver**: 8 tables
- **Golden**: 18 tables

300+ M / day

New rows via CDC

- Streaming from 5 source systems
- One Kafka topic per table
- Same tables for live and history

Since 2009

All available history

- Telemetry, logs, alarms and scientific data
- History loading in progress
- Live data on top, every day

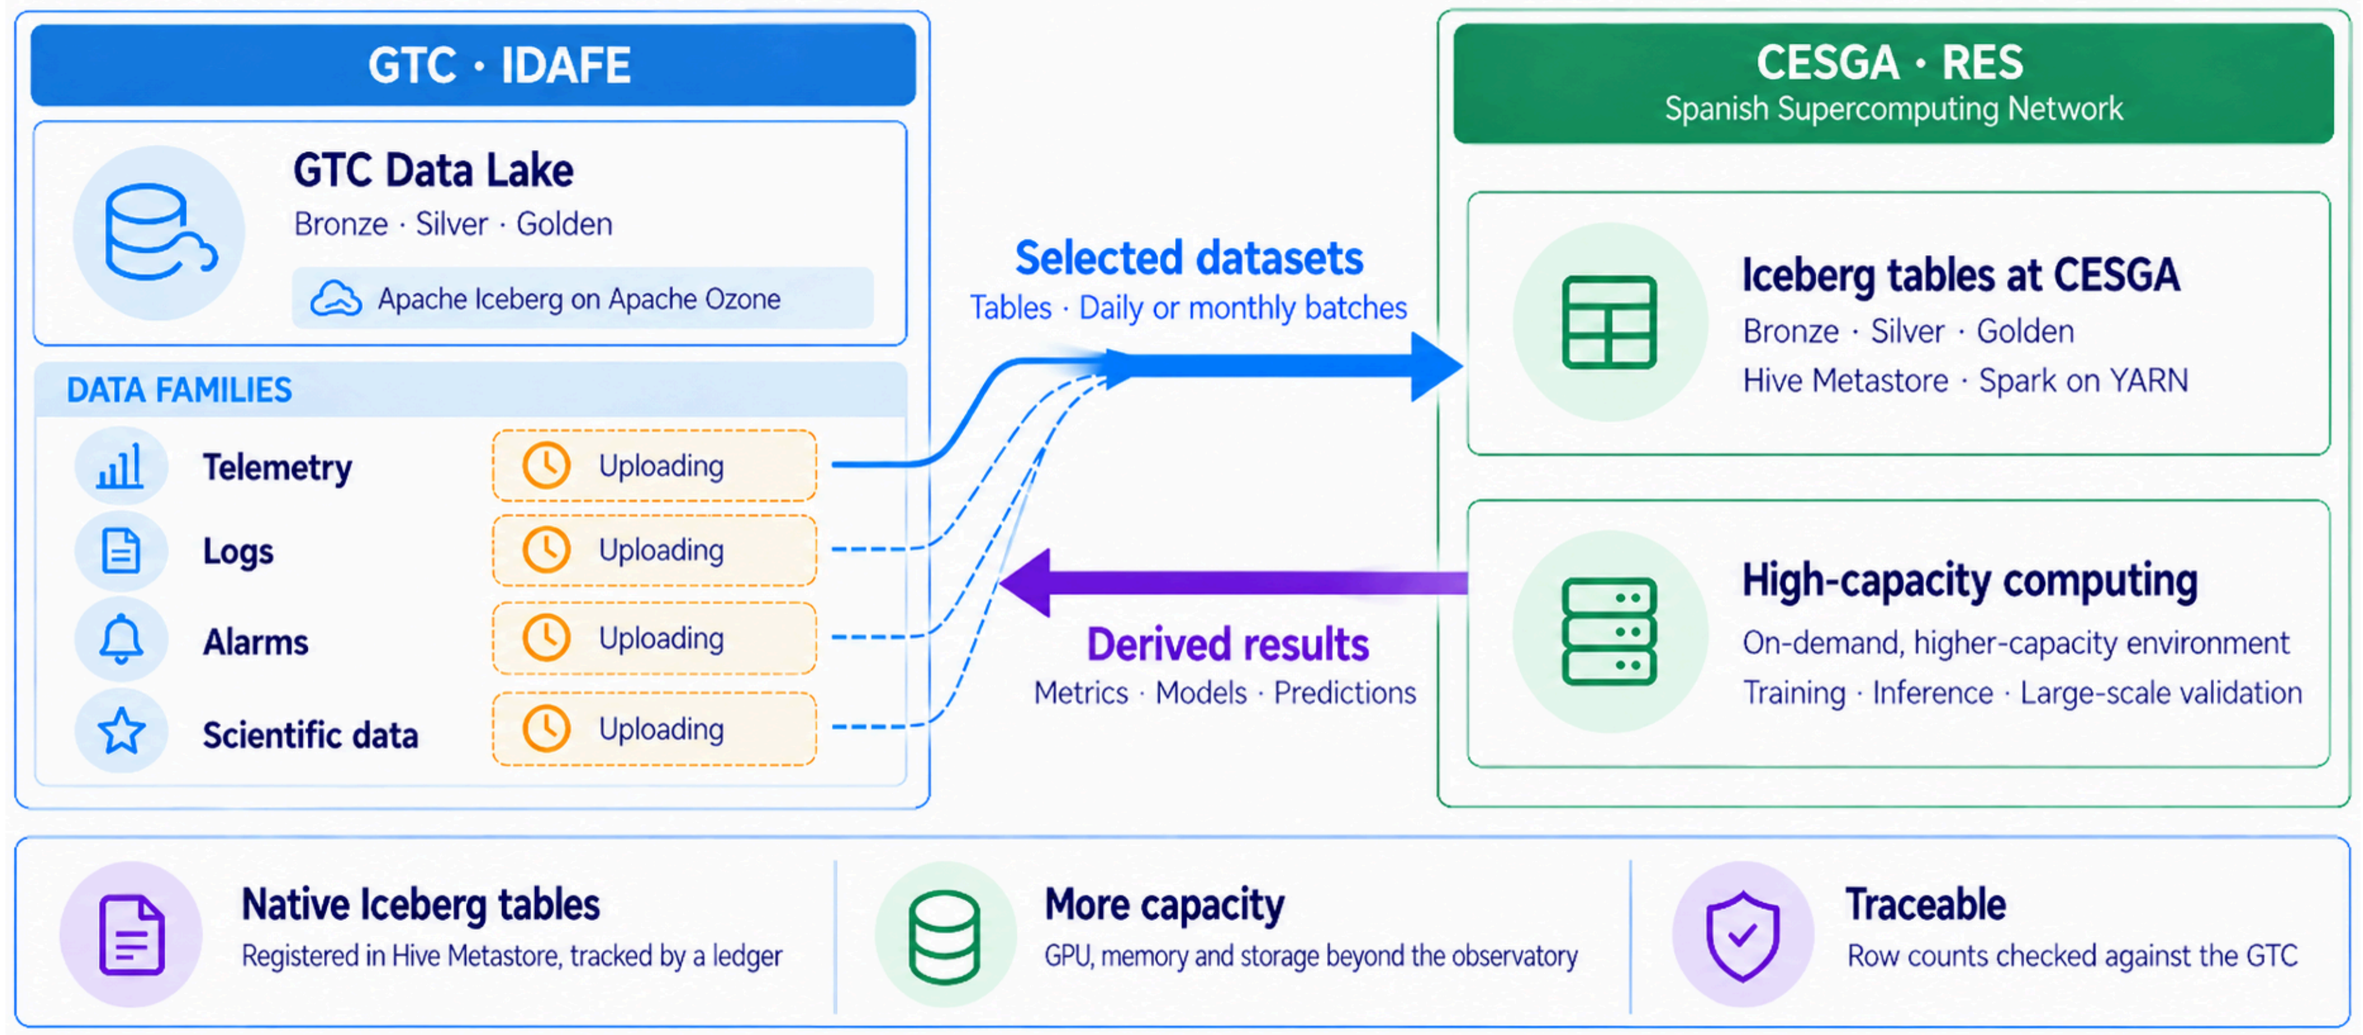
Next

More sources, same procedure

- MongoDB · Zabbix
- Weather data
- Digital transformation products
- Incident reports
- And more to come

Collaboration GTC – RES/CESGA

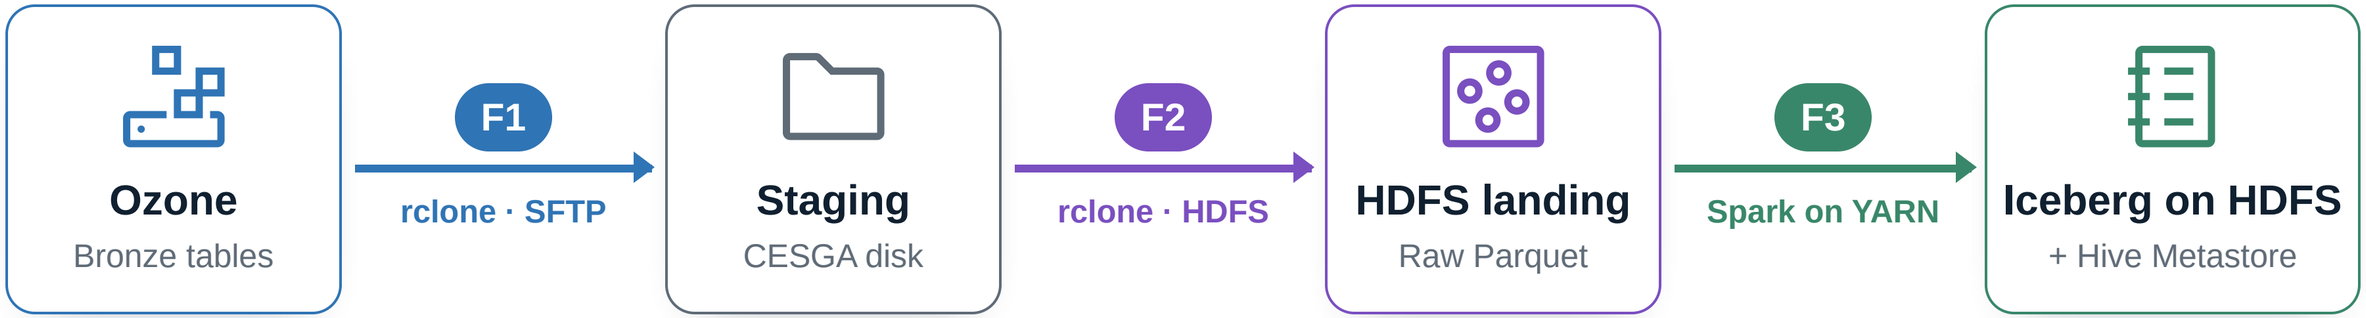
Selected datasets replicated to CESGA as native Iceberg tables, for HPC-scale analysis.



Ingestion with CESGA — the three phases

Bronze replicated to CESGA as native Iceberg tables: decoupled, idempotent and resumable.

GTC · IDAFE | CESGA · RES



F1 · Transfer

- Ozone → **CESGA disk** over SFTP
- **SFTP**: the way in from the GTC
- **Add-only**: never modifies or deletes

F2 · Landing

- Disk → **HDFS**, same paths
- Where **Spark on YARN** reads
- **Temporary**, emptied every year

F3 · Iceberg tables

- Stored on **HDFS**
- Registered in **Hive Metastore**
- **Schemas and tables** created in advance
- **Daily or monthly** jobs, **progress table**

Coverage

- **2021 → 2025**: complete
- **2026**: history + daily sync
- **2010 → 2020**: in progress

Sanity

- **Row counts** at the three stages
- CESGA **must equal** IDAFE
- Mismatch → **reload** the partition

Team

Main profiles and collaboration behind the Data Management Plan and the Data Lake.

Data platform architecture

Technical vision and evolution

Manuel Huertas

Vision, evolution and design criteria.

Data platform engineering

Building and running the platform

Manuel Huertas · Josué Toledo · Ismael Álvarez · Sebastián Fernández · Francisco Rodríguez · Jorge de La Rosa

Ingestion, platform, automation and operations.

Data engineering and AI

Analytics, models and artificial intelligence

AI group led by Luis Rodríguez

Advanced analytics, models and use cases.

Data Management Plan

Data management domains

Josué Toledo Castro

Data management principles, data map and catalogue.

TOWARDS A DATA CULTURE AT THE GTC

The data was always there.
Now it has one place, one language and one door.



Image: Daniel López / IAC

Thank you

josue.toledo@gtc.iac.es