



# Learn how to reduce your Snowflake cost by 50%+ with a lakehouse

Using Dremio's Unified Lakehouse Platform for  
Self-Service Analytics

 **dremio**


WHITEPAPER

# Data Warehousing at Less than Half the Cost of Snowflake

Dremio's easy and open Iceberg data lakehouse provides  
analytics at less than half the cost of Snowflake



dremio.com






## Using Dremio to Reduce your Snowflake Data Warehouse Costs



DREMIO BLOG



# What Snowflake Customers Tell Us



## \$\$ Data Lock-In

Can't access data efficiently unless it's in Snowflake's proprietary format. Migrating data can be a long process and is expensive.

## \$\$ Not Built for BI & Self-Service

Due to performance, Snowflake customers often copy data into BI extracts and cubes for self-service. This cost is not often talked about with TCO analysis.

## \$\$ Expensive to Maintain

Data teams spend a lot of time and resources maintaining expensive queries and optimizing materialized views.

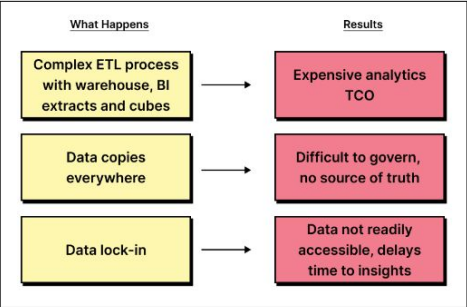
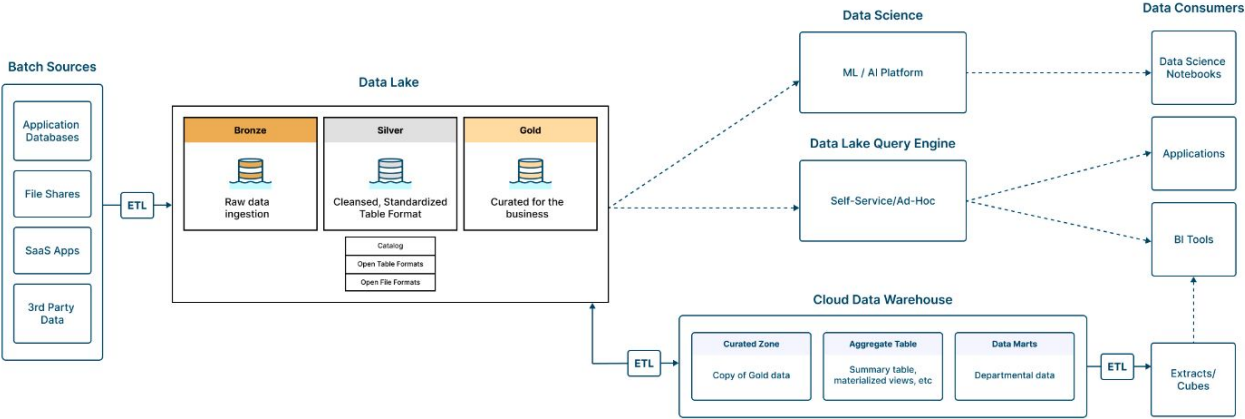
## \$\$ Expensive ETL

Data ingestion in and out of Snowflake is another layer of compute cost.

*"I spend 30% of my time managing and optimizing expensive SF queries and materialized views. We have to ingest data from materialized views into BI extracts to get more snappy performance [with this architecture]."*

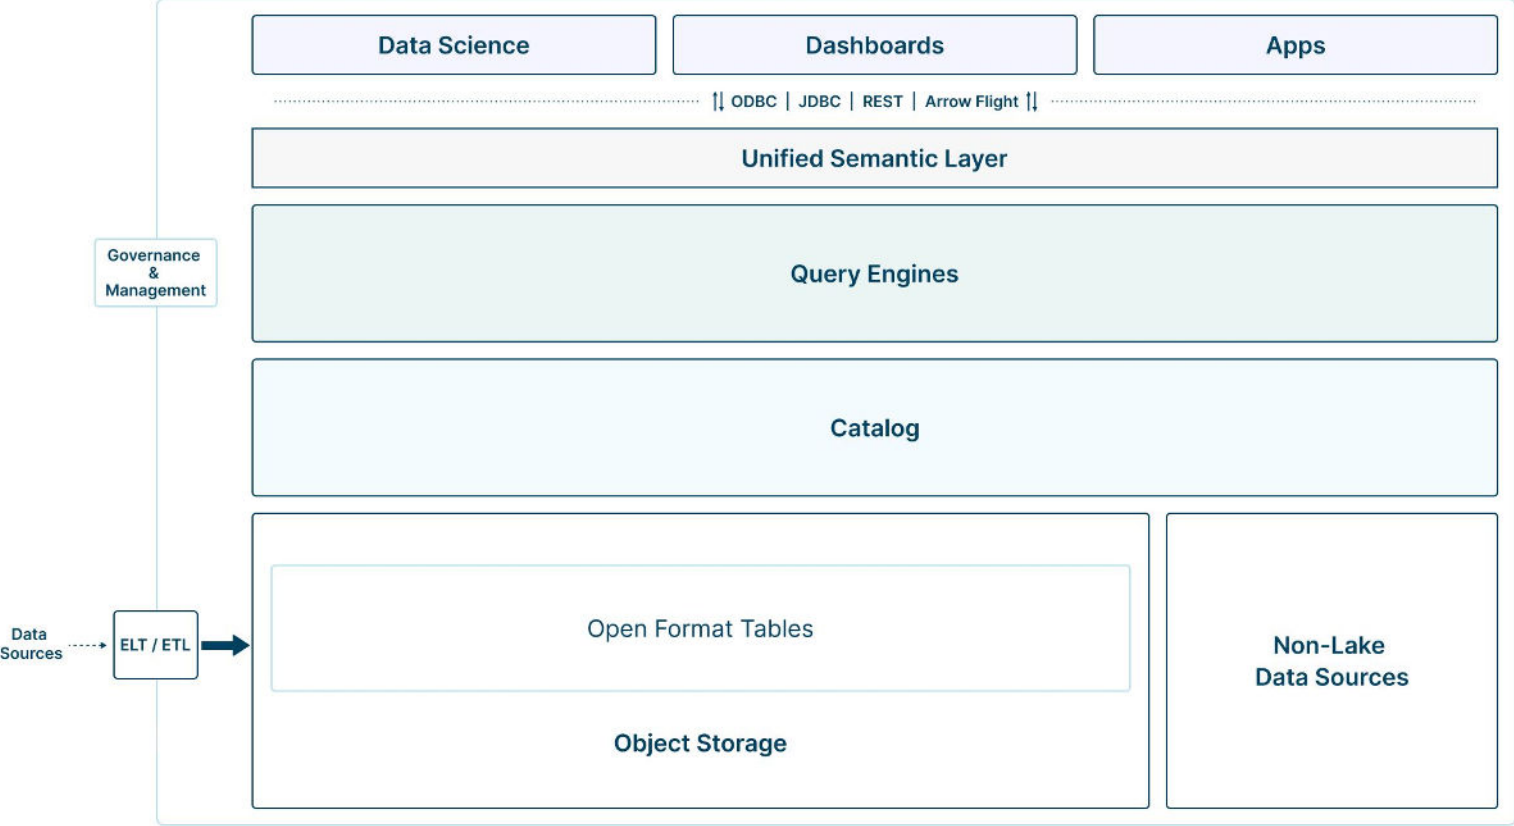
Data Engineer at Highspot

# Current Approaches to Data Management

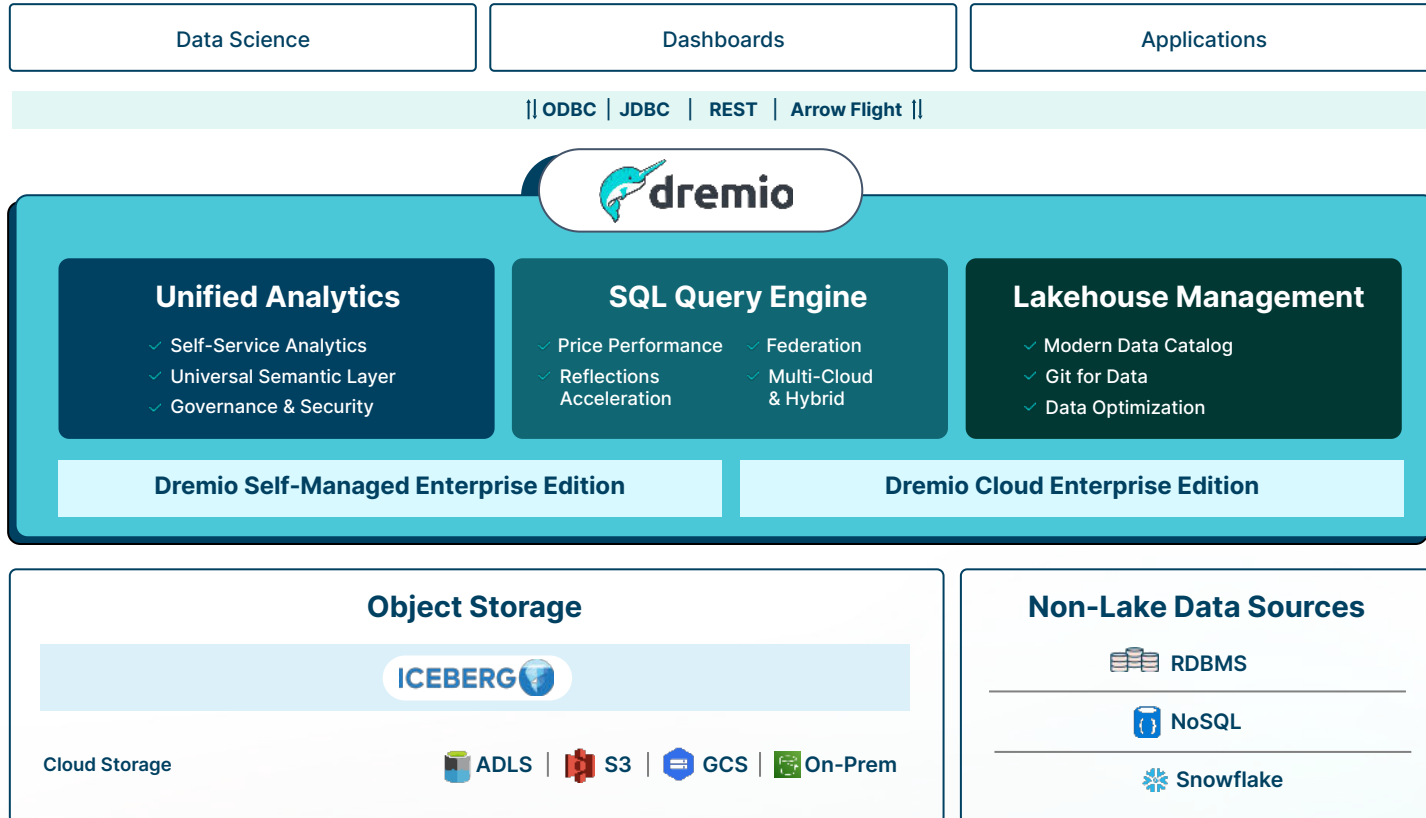


**Common two-tiered architecture based on cloud data warehouse**

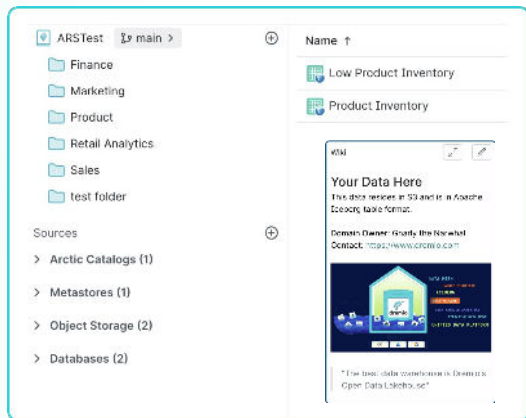
# Ideal Enterprise-Grade Lakehouse



# The Unified Lakehouse Platform for Self-Service Analytics

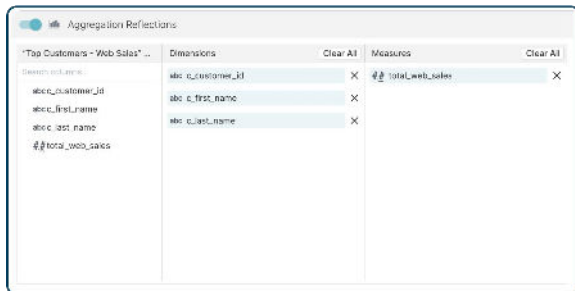


# How Dremio Enables Enterprise-Grade Lakehouse



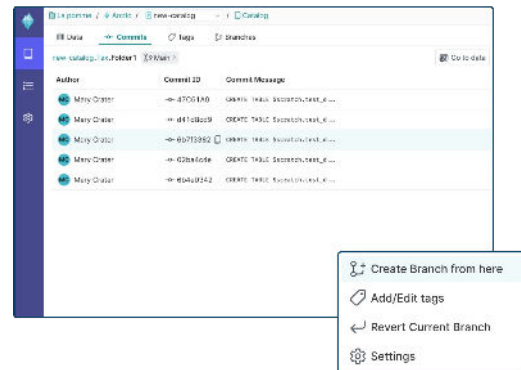
## Unified Access Layer

- Deliver meaningful data products to end users while preserving business context and logic



## SQL Query Engine

- Accelerate analytical workloads without copying data into warehouses, BI extracts, or offline spreadsheets.




## Lakehouse Management


- Git-inspired versioning capabilities such as branches, tags, and commits to simplify data lakehouse management. Give data consumers a consistent view of the data at all times.

# The Dremio Difference


**Best-in-Class TCO**




**Fastest Time to  
Insight**



**Ease of Use  
Through  
Self-Service Data**



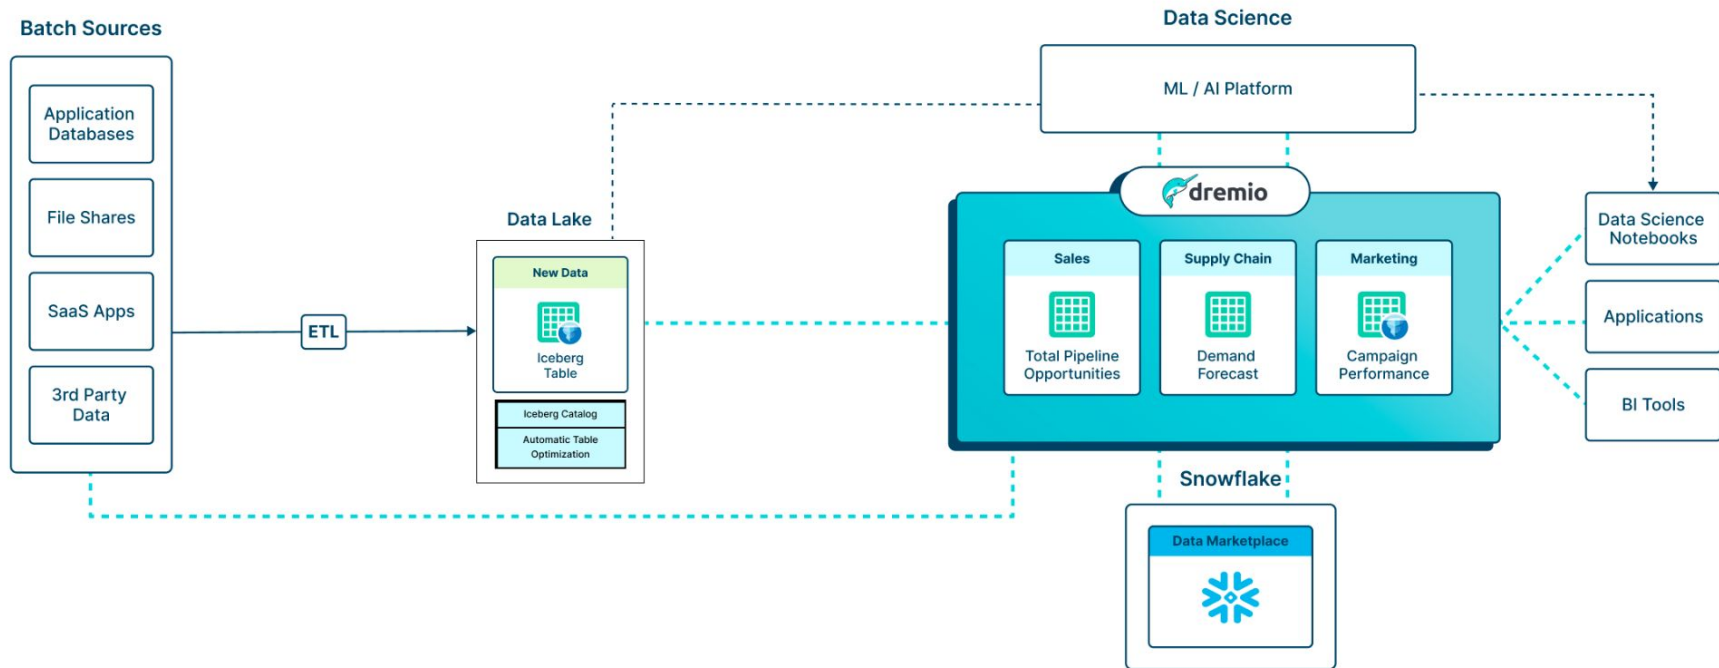
**Flexible & Open  
Architecture**





# Warehouse and Lakehouse - Better Together

Cost Optimize | Deliver Unified and Self-Service Analytics | Open Architecture



# Analytics on Dremio is Less Than Half the Cost of Snowflake

## 3 Year Analytics of Total Cost of Ownership

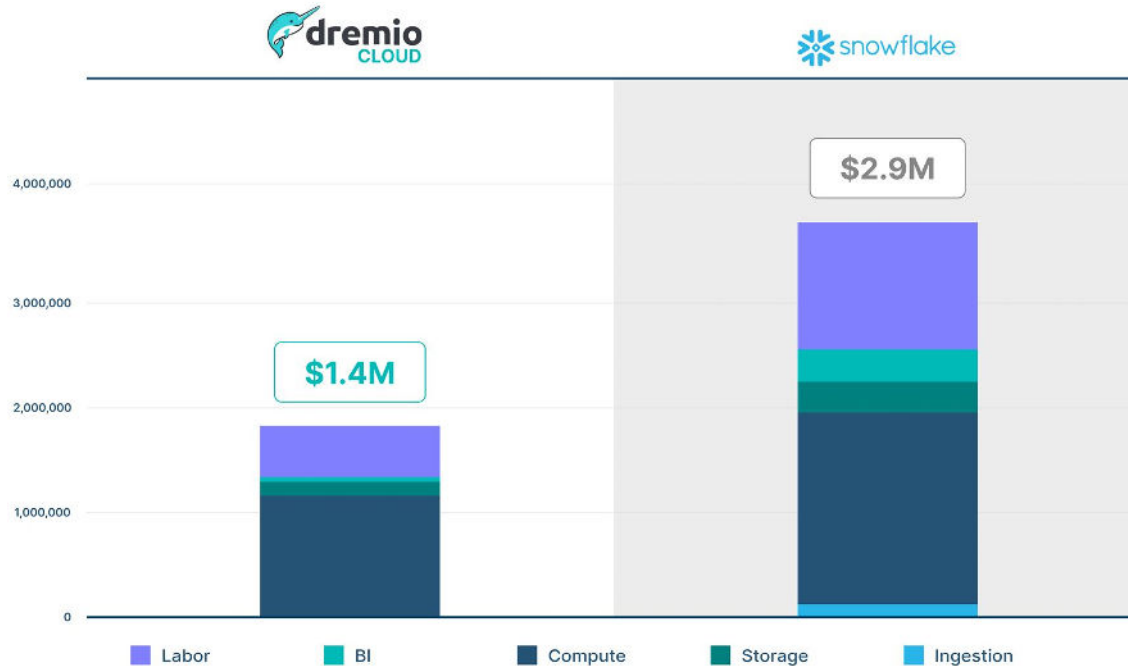
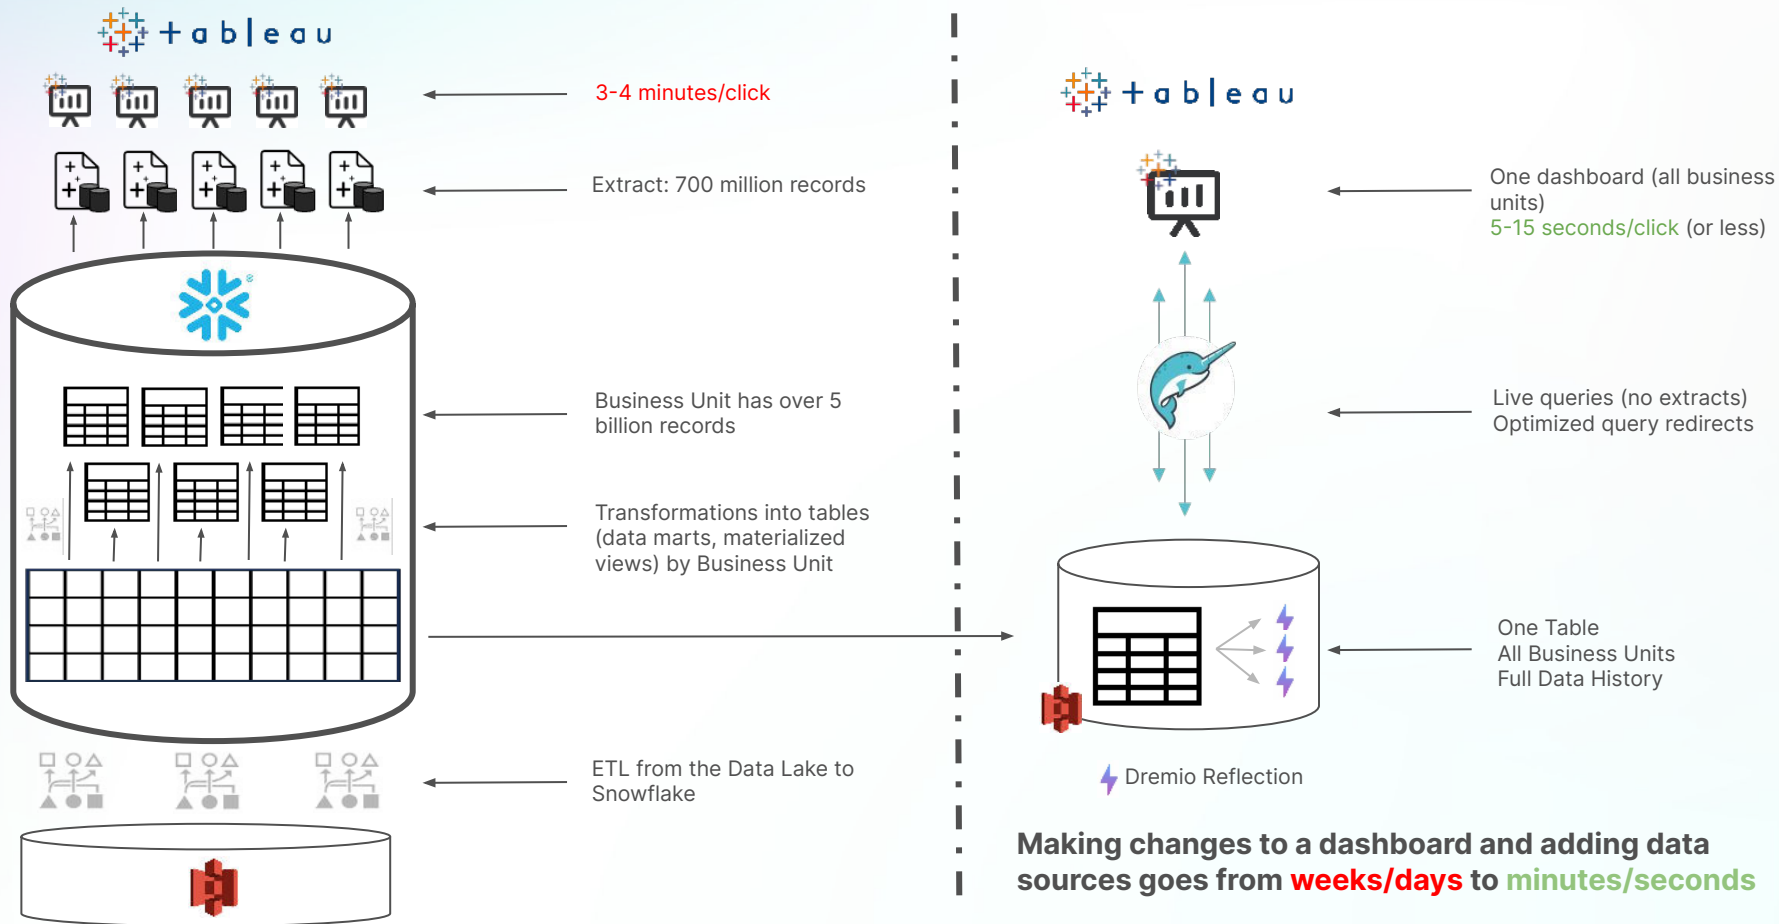


Figure 1: Three-year Analytics TCO comparison between Dremio Cloud and Snowflake

- End-to-End Analytics TCO
- 3 TB Daily Data Ingestion
- Dremio Large (8 nodes) vs Snowflake Large Warehouse

# Fortune 10 Customer - 75% TCO Savings, \$3M Savings In Just One Dept




# 91% Lower TCO with Dremio Compared to Snowflake

*Global Leader in manufacturer of commercial vehicles*

- Self-Service BI Workloads on **Azure** and **AWS**
- **Dremio**
  - Data already in S3 and ADLS
    - No movement
    - No copy
  - 5 nodes of m5d.2xlarge
  - \$4.00 per hour for compute
- **Snowflake**
  - Medium Warehouse
  - **\$47.20** to run the same workload on Snowflake
    - **\$24.40** per hour for compute
    - **\$22.80** per hour ingesting data into Snowflake

Total Cost Per Hour for Sales Analytics




# Dremio is over 3x faster than Snowflake

*Largest pharmaceutical and biotech company in the world*

- BI Workloads on Azure
- **Dremio**
  - Parquet files on ADLS
  - 10 nodes, 128 GB 16 CPU on Azure with ARM installation
- **Snowflake**
  - Clustered Table
  - X-Large Warehouse

Customer BI Performance on Azure (seconds)



# Dremio vs Snowflake

|  | Dremio  | Snowflake  |
|--|---|--|
| Data Engineering                       |   |  |
| <b>Data Ingest</b>                     | Not required; no data movement  | Requires ETL into Snowflake, incurs ingestion costs  |
| <b>Data Transformation - last mile</b> | Virtual; Self-Service   | Physical; IT assistance required   |
| <b>Data Transformation - long haul</b> | Not core Dremio capability yet -- Spark recommended                                       | Typical use case but expensive   |
| User Experience                        |   |  |
| <b>Semantic Layer</b>                  | Core functionality  | Not available; 3rd-Party Options with BI cubes/extracts  |
| <b>Acceleration</b>                    | C3 Caching; Reflections*<br>(*any number of tables)                                       | Materialized aggregates*<br>(*single table)  |
| <b>Data Curation/ Federation</b>       | Virtual Joins; External reflections   | ETL; Sandbox environment   |
| <b>Query redirect</b>                  | Auto rewrite for any query*<br>(*User does not need to know about reflection to leverage) | Rewrites only for Materialized Views*<br>without joins<br>(*User must direct query to materialized view or base table) |

 **dremio**


WHITEPAPER

# Data Warehousing at Less than Half the Cost of Snowflake

Dremio's easy and open Iceberg data lakehouse provides  
analytics at less than half the cost of Snowflake



[dremio.com](https://dremio.com)





## Using Dremio to Reduce your Snowflake Data Warehouse Costs



DREMIO BLOG

