

Streamlining Legacy Migrations: A Comparative Analysis of Teradata to Snowflake Transformation

Santosh Vinnakota

Senior Technical Consultant

California, USA

Santosh2eee@gmail.com

Abstract

The increasing need for scalable, cost-effective, and cloud-native data solutions has propelled organizations to migrate from legacy systems like Teradata to modern platforms such as Snowflake. This paper delves into comprehensive strategies, explores technical challenges, and presents robust solutions for migrating Teradata workloads to Snowflake. It emphasizes advanced techniques for schema conversion, SQL translation, query performance tuning, and cost optimization. Detailed workflows, in-depth technical comparisons, and actionable insights are accompanied by diagrams, flowcharts, and examples to guide data engineers and architects through successful migrations.

Keywords: Legacy Systems, Teradata, Snowflake, Cloud Migration, Schema Conversion, SQL Translation, Query Performance Tuning, Data Optimization

1. INTRODUCTION

Legacy data platforms like Teradata have long supported enterprise analytics and reporting. However, with growing business demands for real-time data processing, scalability, and reduced operational costs, transitioning to modern platforms like Snowflake has become a strategic imperative. Snowflake's unique architecture, which decouples storage and compute, and features such as micro-partitioning, provide significant advantages over traditional systems.

This paper outlines a technical roadmap for transitioning Teradata workloads to Snowflake, addressing challenges and implementing best practices for a seamless and optimized migration process.

2. COMPARATIVE ANALYSIS: TERADATA VS. SNOWFLAKE

2.1 Teradata Architecture:

A Teradata deployment operates as a cluster of Access Module Processors (AMPs), which serve as the core engines driving its data processing capabilities. Each AMP resides on a dedicated server instance, handling a distributed share of both computational and storage workloads across the cluster. This architecture inherently couples storage and compute, ensuring that each AMP stores and processes its designated data segment locally, optimizing data locality and processing efficiency.




Fig. 1. Teradata Architecture

2.2 Snowflake Architecture

Snowflake's cloud-native, multi-cluster, shared data architecture offers a distinct advantage by separating compute and storage functions. This separation enables flexible scaling and optimized resource utilization.

With the decoupling of storage and compute layers, each can be scaled independently and on demand. For instance, you can dynamically adjust compute resources by spinning up additional virtual warehouses during peak usage, such as when data analysts refresh dashboards in the morning, and then scale them down overnight when data demand is low.

This flexibility ensures efficient resource allocation, aligning seamlessly with fluctuating business requirements for data processing. Additionally, Snowflake simplifies operations with intuitive controls accessible via its API or user interface, making it easy to manage without requiring extensive technical expertise.




Fig. 2. Snowflake Architecture

2.3 Architectural Differences

Feature	Teradata	Snowflake
Deployment	On-Premises	Cloud-Native
Storage	Row-based	Columnar
Compute Scaling	Limited	Elastic, On-Demand
Performance Tuning	Manual	Automated

Technical Insight: The decoupling of compute and storage in Snowflake enables independent scaling, reducing costs and improving performance for dynamic workloads.

3. MIGRATION STRATEGIES

3.1 Assessment and Planning

3.1.1 Legacy System Analysis(Teradata):

- *Workload Profiling:* Analyze the existing Teradata workloads, including query patterns, data volumes, and peak usage.
- *Schema Analysis:* Review the Teradata schema design (tables, indexes, and views) to identify dependencies and complexities.
- *ETL Jobs and Pipelines:* Inventory ETL jobs, stored procedures, and scripts used in the data workflows.
- *Resource Utilization:* Evaluate system performance metrics, including CPU, I/O, and memory usage.

3.1.2 Target Platform Analysis (Snowflake)

- *Feature Compatibility:* Map Teradata features (e.g., stored procedures, indexing, and partitioning) to Snowflake equivalents.
- *Cost Modeling:* Estimate Snowflakes compute and storage costs based on anticipated usage patterns.
- *Performance Expectations:* Define performance benchmarks for queries, ETL jobs, and workloads on Snowflake.

3.1.3 Gap and Impact Analysis

- *Functionality Gaps:* Identify gaps where Teradata-specific features do not have direct Snowflake counterparts (e.g., BTEQ scripts or proprietary functions).
- *Data Format and Precision:* Assess potential data format mismatches or precision issues during the migration.
- *Downstream System Impact:* Evaluate the impact of migration on downstream systems consuming data from Teradata.

Key Tip: Use automated tools like Teradata Profiler to assess current workloads and generate detailed performance baselines.

3.2 Schema Conversion

- Extract schemas using tools or scripts tailored to Teradata.
- Normalize schemas and translate data types to Snowflake equivalents. For example:
 - Teradata's INTEGER maps to Snowflake's NUMBER.
 - Use Snowflake's clustering keys to replace primary index optimizations in Teradata.




Fig. 3. Migration ETL

Advanced Technique: Leverage Snowflake's semi-structured support (e.g., JSON, Parquet) for complex data models previously flattened in Teradata.

4. CHALLENGES AND SOLUTIONS

4.1 Schema Complexity

Challenge: Complex joins, denormalized tables, and indexes in Teradata may not translate directly to Snowflake.

Solution: Simplify schemas by normalizing tables, consolidating redundant datasets, and leveraging Snowflake's micro-partitioning for efficient query execution.

4.2 SQL Translation

Challenge: Teradata-specific SQL constructs may require significant rework.

Solution:

- Develop a library of translated SQL functions.
- Automate conversion using tools like SQLMorph and validate converted queries against sample datasets.

```

WITH recent_orders AS (
  SELECT *
  FROM orders
  WHERE order_date >= ADD_MONTHS(CURRENT_DATE, -6)
)
SELECT customer_id,
       SUM(order_amount) AS total_order_amount,
       RANK() OVER (ORDER BY SUM(order_amount) DESC) AS customer_rank
FROM recent_orders
GROUP BY customer_id
QUALIFY customer_rank <= 10
ORDER BY customer_rank;
  
```

Fig. 4. Teradata Syntax

```
WITH recent_orders AS (  
  SELECT *  
  FROM orders  
  WHERE order_date >= DATEADD(MONTH, -6, CURRENT_DATE)  
)  
customer_totals AS (  
  SELECT customer_id,  
         SUM(order_amount) AS total_order_amount,  
         RANK() OVER (ORDER BY SUM(order_amount) DESC) AS customer_rank  
  FROM recent_orders  
  GROUP BY customer_id  
)  
SELECT customer_id, total_order_amount, customer_rank  
FROM customer_totals  
WHERE customer_rank <= 10  
ORDER BY customer_rank;
```

Fig. 5. Snowflake Syntax

4.3 Data Validation and Integrity

Challenge: Ensuring data accuracy and consistency during migration.

Solution: Implement robust validation frameworks:

- Row counts and checksum validations.
- Column-level sampling and anomaly detection using Snowflake's query profiling tools.

5. OPTIMIZATION TECHNIQUES

5.1 Query Performance Tuning

- Optimize queries with Snowflake's clustering keys for high-cardinality columns.
- Rewrite queries to leverage Snowflake's in-memory compute optimizations.
- Use **Query Profile** for performance diagnostics and improvement.
- Cluster large fact tables by frequently queried dimensions to improve query latency.

5.2 Data Loading Best Practices

- Compress source data using formats like Parquet for faster ingestion.
- Automate real-time data ingestion pipelines with Snowpipe and monitor using Snowflake's Task Monitoring.

5.3 Cost Optimization

- Schedule compute-intensive workloads during off-peak hours.
- Use Snowflake's resource monitors to set cost thresholds and alerts.




Fig. 6. Cost Optimization Workflow

6. TECHNICAL DIAGRAMS AND FLOWCHARTS

Phase 1	Assess Current Environment and Inventory Assets	Existing Teradata Env to understand data structures, schemas, dependencies, overall complexity	Identify the data volumes, types and patterns to determine the scope of migration
Phase 2	Define Objects and Requirements	Document specific requirements for Snowflake, including data storage, processing, security and compliance	Goals and Objectives to improve performance, reducing costs and enabling new capabilities
Phase 3	Data Profiling and Cleansing	Data profiling to identify inconsistencies, anomalies, data quality issues in Teradata	
Phase 4	Migration Approach	ETL, ELT, CDC etc.	
Phase 5	Data Mapping and Transformation	Map the source data schema from Teradata to the Target Snowflake schema	Define data transformation rules to convert and load the data accurately in Snowflake
Phase 6	Implement Data Migration Tools	Utilize data migration tools or platforms that support Teradata and Snowflake migration	
Phase 7	Establish Monitoring and Validation Mechanisms		

Fig. 7. Migration Phases Overview




Fig. 8. Snowflake Architecture Post-Migration

Components: Ingestion pipelines, scalable compute clusters, secure storage layers, and real-time analytics dashboards.

7. CONCLUSION

Migrating from Teradata to Snowflake is a transformative journey that enables businesses to harness the power of modern cloud-native data platforms. By addressing schema conversion, SQL translation, and query performance optimization with the techniques outlined in this paper, organizations can achieve seamless transitions while maximizing cost savings and performance gains.

REFERENCES

- [1] T. White, *Hadoop: The Definitive Guide*, 4th ed., O'Reilly Media, 2015.
- [2] J. Gorman, *Cloud Data Warehousing for Dummies*, Wiley, 2020.
- [3] J. Van Der Lans, *Data Virtualization for Business Intelligence Systems*, Morgan Kaufmann, 2012.
- [4] D. Loshin, *Enterprise Knowledge Management: The Data Quality Approach*, Morgan Kaufmann, 2011.
- [5] Snowflake Documentation: <https://docs.snowflake.com/>
- [6] Teradata Migration Guide: <https://www.teradata.com/>
- [7] SQLMorph for SQL Translation: <https://sqlmorph.com/>