



# SNOWFLAKE PERFORMANCE AUTOMATION AND TUNING

THREE-DAY COURSE

24C21



UNIVERSITY

DATASHEET

## OVERVIEW

This three-day accelerated course presents key performance capabilities, Snowflake best practices, and tuning techniques to help participants use the Snowflake Data Cloud to develop diverse, high-performance workloads. The course illustrates performance design best practices and pitfalls so participants can apply the methodology and features to their diverse workloads.

The course is delivered on the latest version of the Snowflake Data Cloud with new performance features and some emerging features in public preview.

This course consists of lectures with many practical examples in instructor-led demos, followed by experiential exercises.

## ACQUIRED SKILLS

- Work effectively and efficiently with the key capabilities and best practices behind Snowflake’s Data Cloud four-tier architecture designed for performance and scale.
- Apply the appropriate Snowflake-provided tools for performance assessment and optimization.
- Use Snowflake’s tuning methodology on established performance features, including best practices and avoiding pitfalls to work efficiently with the platform.
- Explore emerging performance features with suitable use cases and adoption of best practices.
- Formulate diverse, high-performance, and efficient workloads, including data transformations, analytic applications, and data sharing.
- Achieve cost optimization by working effectively and monitoring the Snowflake Data Cloud.

## WHO SHOULD ATTEND

- Data Application Developers
- Data Architects
- Data Engineers
- Snowflake Administrators
- Technical Team Leads

## PREREQUISITES

This course is designed for participants with extensive Snowflake experience or participants who have completed other relevant Snowflake courses such as “Snowflake Fundamentals” or “Snowflake Advanced”.

## DELIVERY FORMAT

Instructor-led Public or Private classes are available.

## TOPICS COVERED

### Anatomy of a Query

- Query Compilation
- Query Execution

### Constructing Performant Queries

- Filtering Data
- Joining Data
- Aggregating, Ordering, and Grouping Data
- Subqueries and CTEs
- Estimating and Sampling
- Accessing External Data

### Virtual Warehouse Optimization

- Virtual Warehouse Types
- Virtual Warehouse Settings
- Monitor Virtual Warehouse Efficiency

### Automatic Clustering Service

- Overview
- Evaluate Ordering (Clustering)
- Implement and Test Cluster Keys
- Create a Cluster Key
- Monitor Clustering Cost
- Clustering: Common Misconceptions

### Materialized Views

- Why Materialized Views?
- Monitor Materialized Views

## Dynamic Tables

- Monitor Dynamic Tables
- Dynamic Tables vs. Materialized Views

## Search Optimization Service

- Overview
- How it Works
- Add a Search Optimization
- Monitor Search Optimization Cost

## Query Acceleration Service

- Overview
- Identify Eligible Queries
- Configure Query Acceleration
- Monitor Query Acceleration

## Looking For Trouble (Issues and Inefficiencies)

- Finding Trouble in the Query\_History
- Finding Problematic Query Operators
- Finding Patterns and Trends in Queries
- Finding Inefficiencies