

SNOWFLAKE PERFORMANCE AUTOMATION AND TUNING

THREE-DAY COURSE

25F16



OVERVIEW

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

The course is delivered on the latest version of the Snowflake AI Data Cloud, including new performance features and 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 AI Data Cloud architecture which is 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 recommended best practices and avoiding pitfalls to work efficiently with the platform.
- Explore emerging performance features with suitable use cases and adoption of recommended 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 AI Data Cloud.

WHO SHOULD ATTEND

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

PREREQUISITES

- Recommended completion of the "Snowflake Multi-Factor Authentication (MFA) Essentials" free on-demand course.

- Completion of “Snowflake Fundamentals,” “Snowflake Advanced,” or equivalent Snowflake knowledge.

DELIVERY FORMAT

Instructor-led Public or Private classes are available.

TOPICS COVERED

Anatomy of a Query

- Query Optimizations
- Query Execution

Constructing Performant Queries

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

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

- Create 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

Other Performance Considerations

- Hybrid Tables Overview and Performance
- Accessing External Data
- Apache Iceberg™ Tables in Snowflake Overview and Performance
- Memoizable Functions Overview and Performance

Looking For Trouble (Issues and Inefficiencies)

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