



SNOWFLAKE PERFORMANCE AUTOMATION AND TUNING

3-DAY COURSE

22F12

OVERVIEW

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

The course is delivered on the latest version of the Snowflake Cloud Data Platform 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 hands-on exercises.

ACQUIRED SKILLS

- Work effectively and efficiently with the key capabilities and best practices behind Snowflake's Cloud Data Platform three-tiered architecture designed for performance and scale.
- Apply the right set of 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 Cloud Data Platform.

WHO SHOULD ATTEND

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

PREREQUISITES

This course is designed for participants with extensive hands-on experience with Snowflake or participants that have taken other relevant Snowflake courses such as *Snowflake Fundamentals (FUND)* or *Snowflake Advanced (ADV)*.

TOPICS COVERED

Performance Built from the Ground Up in Snowflake Architecture

- Automatic micro-partitioning
- Natural data clustering
- Workload segmentation
- Vectors of compute scaling
- What the execution engine can do for you
- What the optimizer can do for you
- Lightning-fast metadata service
- Caching features

Snowflake Performance Diagnostic and Tuning Toolset

- Query profile and explain plan
- Query profile and query scenarios
- Explain plan and query scenarios
- Understanding micro-partitioning and statistics
- Assessing clustering quality
- Using INFORMATION_SCHEMA functions and account usage share database
- Common investigation scenarios

Best Practices Across the Snowflake Cloud Data Platform

- Writing effective queries
- Scaling up, right-sizing a warehouse, and concurrency
- Metadata operations and optimizations

Best Practices for Raw Data Zone

- Snowflake as a data lake
- Batch data ingestion and DML
- Streaming ingestion

Best Practices for Conformed Data Zone

- Incremental data processing and CDC
- Evaluating data formats and semi-structured data
- Data engineering pipelines

- Data clustering and tuning

Best Practices for Modeled Data Zone

- Materialized views and use cases
- High-performance query features
- Data unloading

Emerging Performance Features

- Highlight preview features and appropriate use cases
- Explore best practices for new features

Monitoring and Optimizing Costs

- Virtual warehouse compute
- Serverless compute
- Cloud services metadata operations
- Storage
- Replication
- Data transfer