KOENIG
step forward

Snowflake Performance Automation and Tuning

Course outline

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 e□ective 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