



SNOWFLAKE ARCHITECT

THREE-DAY COURSE

24H01

OVERVIEW

This three-day course for Architects and Technical Leaders provides the skills, knowledge, and Snowflake best practices to deploy and operate Snowflake, insights and recommendations based upon real-world customer experiences, and the confidence to get the very best out of Snowflake's technology.

ACQUIRED SKILLS

- Examine the tradeoffs associated with the available environment and Snowflake Account deployment options.
- Use the Snowflake data security framework to balance the often-conflicting needs of protecting sensitive data while democratizing access and facilitating sharing.
- Apply Snowflake best practices to maximize performance and efficient use of resources during data ingestion, transformation, and end user queries.
- Analyze Snowflake metadata to identify performance and cost issues and recommend remedial action.

WHO SHOULD ATTEND

- Solution Architects
- Data Architects
- Database Architects
- Enterprise Data Architects
- Senior Data Engineers
- Technical Team Leads

PREREQUISITES

- SQL skills, a background in database management, and, ideally, experience in designing and deploying analytic-based systems are required.
- Completion of "Snowflake Fundamentals" or equivalent Snowflake knowledge.

DELIVERY FORMAT

Instructor-led Public or Private classes are available.

TOPICS COVERED

Architecture Overview

- Snowflake Architecture

- Snowflake's Layered Architecture

Deployment Considerations

- Organization
- Geographic Account Considerations
- Snowflake Security Domains
- Environment Deployment Options
- Cross Environment Data Transfer
- Options
- Environment Separation
- Logical Data Architecture (Layers)
- Physical Architecture Options
- Database Considerations
- Database Reference Options
- Summary and Recommendations

Data Architecture

- Overall Data Flow
- Reference Data Architecture
- Handling Raw History
- Integration with Data Lake
- Create External Table
- Query External Table
- Partitioned External Tables
- Snowpipe Streaming Overview
- External Network Access
- Native Apps
- Change Data Capture, Creating and Managing Streams
- Dynamic Tables
- Hybrid Tables
- Iceberg Tables in Snowflake

Virtual Warehouse Management

- Objectives
- Workload Challenges
- Scale Up for Large Workloads
- Key Concepts: Scaling Up

- Key Concepts: Diminishing Elapsed Time Improvements
- Scale Out for Multiple Concurrent Users
- Speed vs. Throughput
- Right-Sizing Virtual Warehouses
- Virtual Warehouse Deployment Approach
- Measuring Workloads
- Summary
- Case Study

Data Security Framework

- Data Security Framework
- Data Classification
- Identify Data Sensitivity

Role-Based Access Control (RBAC)

- Overview
- RBAC Requirements
- RBAC Hierarchy Design
- RBAC Role Design
- Naming Standards
- RBAC Script Building
- Snowflake RBAC Best Practices

Data Governance Features

- Data Masking and Row Access Policies
- Dynamic Data Masking
- Row Access Policies
- Summary

Sharing Architecture

- Use Case
- Direct Share
- With Replication
- Data Mesh
- Relevant Snowflake Capabilities for a Data Mesh
- Data Mesh Architecture Options with Snowflake

- Data Products in Snowflake
- Auto-fulfillment
- How Auto-fulfillment Works

Snowflake Data Storage

- Data Storage Methods
- Snowflake Data Storage
- Data Storage Implications
- Time Travel
- Key Point
- Time Travel and Storage
- Clones
- Data Recovery
- Agile Data Management
- Development
- System Testing
- Deployment

Table Clustering

- What Is Table Clustering?
- Partition Pruning (Elimination)
- Overlapping Values
- Evaluating Clustering
- Implement and Test Clustering Keys

Performance Optimization

- Search Optimization
- How Does It Work?
- Materialized Views
- Materialized View Use Cases
- Query Acceleration Service (QAS)

Management and Observability

- Observability on Snowflake
- Outbound Notifications
- Snowflake Alerts

- Observability Within Snowsight
- Budgets