

SNOWFLAKE ARCHITECT

Prerequisites: Basic knowledge of SQL and snowflake is required.

Duration: 3 Days (8 Hrs./Day)

Course Objective: 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.

Lab Requirement: Koenig DC/Linux.

Module 1 - Architecture Overview

Snowflake Architecture

Snowflake's Layered Architecture

Module 2 - 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

Module 3 - 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

Module 4 - 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

Module 5 - Data Security Framework

Data Security Framework

Data Classification

Identify Data Sensitivity

Module 6 - Role-Based Access Control (RBAC)

Overview

RBAC Requirements

RBAC Hierarchy Design

RBAC Role Design

Naming Standards

RBAC Script Building

Snowflake RBAC Best Practices

Module 7 - Data Governance Features

Data Masking and Row Access Policies

Dynamic Data Masking



Row Access Policies

Summary

Module 8 - 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-fulfilment

How Auto-fulfilment Works

Module 9 - 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

Module 10 - Table Clustering

What Is Table Clustering?

Partition Pruning (Elimination)

Overlapping Values

Evaluating Clustering

Implement and Test Clustering Keys

Module 11 - Performance Optimization

Search Optimization

How Does It Work?

Materialized Views

Materialized View Use Cases

Query Acceleration Service (QAS)

Module 12 - Management and Observability

Observability on Snowflake

Outbound Notifications

Snowflake Alerts

Observability Within Snowsight

Budgets