

Snowflake - Data Cloud

Session Number	Session Title	Objective	Topics	Details
Session-1				
1	Introduction to Cloud Data Warehousing	Data Warehouse and Cloud concepts Need of the hour - Cloud Data Warehouse How Snowflake provides the solution for current demand	An overview of Data Warehousing	OLTP vs OLAP Advantages of DW ETL and Staging Dimensional and Tabular Modelling
			Introduction to Cloud Computing	What is Cloud? Advantages of Cloud Computing Types of Services
			Challenges of traditional Data Warehouse Systems	Increasing Volume, Variety and Complexity of data Beyond the limit expectations from Analysts Limited Elasticity and Significant Cost
			Cloud Data Warehousing - A solution	Flexibility of Growth Easy Implementation Less worries on Management Efficient Recovery and Accessibility
Session-2				
2	Key Concepts of Snowflake	Innovative Architecture of Snowflake Why Snowflake scores higher than its competitors	History and Core Design	Snowflake Founders Economical Growth of the Organization Current Management Portfolio Market Leader in Cloud DW Services of Public Cloud Providers
			The Pillars of Snowflake	Shared Disk and Shared-Nothing Architectures Multi-cluster, Shared Data Architecture Three 'linked' independent layers
			Storage Layer	Blob Storage Services Micro-partitions Table Clustering High Availability Security
			Compute Layer	Virtual Compute Services Flexible Scale up & down Multi-Clustering Pay only for the actual usage
			Services Layer	Authentication and Authorization Query Optimization Smart Cache Management Metadata Management Secured Data Sharing

Session-3

3	Snowflake - Fundamentals	Understanding Snowflake Accounts Multiple way to connect with Snowflake Base objects of Snowflake How Standard and Extended SQL Support used to manage data	Account Information	Supported Cloud Partners Snowflake Editions Geographical Locations
			Connecting to Snowflake	Web User Interface Understanding Sessions Command Line Interface Third party Applications
			Snowflake Objects	Virtual Warehouse Databases Schemas, Tables and Views Container Hierarchy Stages File Formats
			SQL and other reference	Parameters Data Types Collation support SQL Variables Object Identifiers Constraints Table Literals

Session-4

4	Loading / Unloading Data	Loading Data from Cloud Storages Loading Data from Local Systems Unloading Data from Snowflake into Cloud and Local Storages	Internal and External Stages	Type of internal stages External stage locations Accessing stage objects
			Bulk Loading	Preparing Data for loading Syntax of COPY command Putting data into stages from local systems Load data from stages Managing regular data load
			Data Unloading	Using COPY command to unload data into stages Options for unloading data Getting data into local system from stages Unloading semi-structured data

Session-5

5	External Data Management	How to Query data in external stages Apply Transformations while loading data from stages Concepts behind External Tables	Querying Stages	Query syntax and Parameters Supported Functions Usage of File Formats
			Transforming data while loading	Supported Stages and File formats CSV data transformations Transforming Semi-structured data Supported Functions
			External Tables	Schema on Read Virtual Columns Partitioned External Tables Refreshing external table metadata External Table DDL

Murali Kannan Thinakaran

Session-6

6	Semi-structured Data	Loading Semi-structured data into standard tables Transforming and Querying Semi-structured data Working with Command Line Interface	Working with Semi-structured data	Semi-structured Data Types Data Size Limitations
			Transforming and Querying Semi-structured data	Casting Key-Values Loading semi-structured data into relational type tables Semi-structured Data Functions
			SnowSQL	Download and Install SnowSQL PUT and GET files Running batch files Configuration file settings SnowSQL Commands

Session-7

7	Continuous Data Protection & Streaming Data Ingestion	Overview of Security Options Understanding the Role Based Access Control Snowflake way of Historical Data Maintenance How Automatic Disaster Recovery works Configuration of external cloud services for automatic ingestion How to enable Snowpipe for continuous data ingestion	Security Options	Network Policies Multi-Factor Authentication Federated Authentication & SSO Data Encryption
			Access Control	Access Control Framework System Roles Custom Roles Privileges Enforcement Model
			Time Travel	How Time Travel works Data Retention period SQL Extensions for Time Travel
			Fail Safe	What is Fail-safe Fail-safe vs Backup Viewing Fail-safe storage
			Configure Snowpipe	Using Cloud Messaging Services (AWS SNS, Azure Event Grid, GCS Pub/Sub) Creating a Storage Integration File Size, Order of loading and data duplication Snowpipe DDL

Session-8

8	Performance Considerations	Understanding Cache management in Snowflake Design Considerations of Table & Warehouse Controlling the order of data stored in micro-partitions	Cache Management	Metadata Cache Result Cache Local Disk Cache Centralized Storage Intelligent Query Optimization
			Warehouse & Table Designing	Suspending the Warehouse Scale out options settings Data Type for columns Cluster Key management Column length specification VARIANT vs FLATTEN
			Data Clustering	Impact of Micro-partitions Clustering Keys Clustering Information Clustering Depth Re-clustering
			Query Profile	Query History Getting information from Information Schema Query Pruning Disk Spilling

Session-9

9	Special Features of Snowflake	The uses of temporary objects Concepts and Consideration of Object Cloning The usage and advantages of Materialized Views	Temporary Stages & Tables	Using a temporary stage Temporary table usages Transient tables Comparison of Table types
			Zero Copy Cloning	How cloning is different from copying Access Control privileges for cloned objects Clonable Objects Impact of DDL and DML Cloning of Historical Objects
			Materialized Views	Working with Materialized Views Advantages of Materialized View Limitations on Materialized View Joining Materialized View Clustering a Materialized View

Session-10

10	Consume Snowflake data & Account Management	Connect Snowflake with other applications How to use Snowflake drivers Sharing data securely with other Snowflake account users and non-account users Understanding Snowflake Credit & Storage How to monitor credits on different levels	Integration and BI Tools	ODBC & JDBC Driver configurations Consider performance with ETL tools Consume data using BI tools Integrating with JAVA and Python
			Secured Data Sharing	Data Providers Sharing data with other account users Data Consumers Managing Reader Accounts Secure Objects and Data Access Control
			Snowflake Credits & Storage Usage	Snowflake Database Account Usage Schema System Views for Usage calculation Account level, Database level, User level statistics Working with Resource monitors

Murali Kannan Thinakaran