Snowflake DEA-C01

SNOWFLAKE SNOWPRO ADVANCED - DATA ENGINEER CERTIFICATION QUESTIONS & ANSWERS

Exam Summary – Syllabus – Questions

DEA-C01

Snowflake Certified SnowPro Advanced - Data Engineer Certification

65 Questions Exam – 75% Cut Score – Duration of 115 minutes

www.VMExam.com

Table of Contents

Know Your DEA-C01 Certification Well:

The DEA-C01 is best suitable for candidates who want to gain knowledge in the Snowflake Advance. Before you start your DEA-C01 preparation you may struggle to get all the crucial SnowPro Advanced - Data Engineer materials like DEA-C01 syllabus, sample questions, study guide.

But don't worry the DEA-C01 PDF is here to help you prepare in a stress free manner.

The PDF is a combination of all your queries like-

- What is in the DEA-C01 syllabus?
- How many questions are there in the DEA-C01 exam?
- Which Practice test would help me to pass the DEA-C01 exam at the first attempt?

Passing the DEA-C01 exam makes you Snowflake Certified SnowPro Advanced - Data Engineer Certification. Having the SnowPro Advanced - Data Engineer certification opens multiple opportunities for you. You can grab a new job, get a higher salary or simply get recognition within your current organization.

Snowflake DEA-C01 SnowPro Advanced - Data Engineer Certification Details:

Exam Name	Snowflake SnowPro Advanced - Data Engineer
Exam Code	DEA-C01
Exam Price	$375 USD
Duration	115 minutes
Number of Questions	65
Passing Score	750 + Scaled Scoring from 0 - 1000
Recommended Training / Books	**Snowflake Data Engineering Training** **SnowPro Advanced: Data Engineer Study Guide**
Schedule Exam	**PEARSON VUE**
Sample Questions	**Snowflake DEA-C01 Sample Questions**
Recommended Practice	**Snowflake Certified SnowPro Advanced - Data Engineer Certification Practice Test**

DEA-C01 Syllabus:

Section	Objectives	Weight
Data Movement	- Given a data set, load data into Snowflake. • Outline considerations for data loading • Define data loading features and potential impact - Ingest data of various formats through the mechanics of Snowflake. • Required data formats • Outline Stages - Troubleshoot data ingestion. - Design, build, and troubleshoot continuous data pipelines. • Design a data pipeline that forces uniqueness but is not unique • Stages • Tasks • Streams • Snowpipe • Auto ingest as compared to Rest API - Analyze and differentiate types of data pipelines. • Understand Snowpark architecture (client vs server) • Create and deploy UDFs and Stored Procedures using Snowpark • Design and use the Snowflake SLQ API - Install, configure, and use connectors to connect to Snowflake. - Design and build data sharing solutions. • Implement a data share • Create a secure view • Implement row level filtering - Outline when to use External Tables and define how they work. • Partitioning external tables • Materialized views • Partitioned data unloading	**25-30%**
Performance Optimization	- Troubleshoot underperforming queries.	**20-25%**

Section	Objectives	Weight
	• Identify underperforming queries • Outline telemetry around the operation • Increase efficiency • Identify the root cause - Given a scenario, configure a solution for the best performance. • Scale out as compared to scale in • Clustering as compared to increasing warehouse size • Query complexity • Micro partitions and the impact of clustering • Materialized views • Search optimization - Outline and use caching features. - Monitor continuous data pipelines. • Snowpipe • Stages • Tasks • Streams	
Storage and Data Protection	- Implement data recovery features in Snowflake. • Time Travel • Fail-safe - Outline the impact of Streams on Time Travel. - Use System Functions to analyze Micro-partitions. • Clustering depth • Cluster keys - Use Time Travel and Cloning to create new development environments. • Backup databases • Test changes before deployment • Rollback	**10-15%**
Security	- Outline Snowflake security principles. • Authentication methods (Single Sign On (SSO), Key Authentication, Username/Password, Multi-factor Authentication (MFA)) • Role Based Access Control (RBAC)	**10-15%**

Section	Objectives	Weight
	• Column Level Security and how data masking works with RBAC to secure sensitive data - Outline the system defined roles and when they should be applied. • The purpose of each of the system defined roles including best practices usage in each case • The primary differences between SECURITYADMIN and USERADMIN roles • The difference between the purpose and usage of the USERADMIN/SECURITYADMIN roles and SYSADMIN - Manage Data Governance. • Explain the options available to support column level security including Dynamic Data Masking and External Tokenization • Explain the options available to support row level security using Snowflake Row Access Policies • Use DDL required to manage Dynamic Data Masking and Row Access Policies • Use methods and best practices for creating and applying masking policies on data • Use methods and best practices for Object Tagging	
Data Transformation	- Define User-Defined Functions (UDFs) and outline how to use them. • Secure UDFs • SQL UDFs • JavaScript UDFs • Returning table value compared to scalar value - Define and create External Functions. • Secure external functions - Design, build, and leverage Stored Procedures. • Transaction management - Handle and transform semi-structured data. • Traverse and transform semi-structured data to structured data • Transform structured to semi-structured data	**25-30%**

Section	Objectives	Weight
	- Use Snowpark for data transformation. • Query and filter data using the Snowpark library • Perform data transformations using Snowpark (ie., aggregations) • Join Snowpark dataframes	

Snowflake DEA-C01 Sample Questions:

Question: 1

While loading a table from an internal stage, you received the below error. What two options can you choose from below to fix this?

a) Use FIELD_OPTIONALLY_ENCLOSED_BY to enclose strings
b) Add the required column in the TABLE
c) Use SKIP_FILE to skip the file and continue
d) Use error_on_column_count_mismatch=false

Answer: a, b

Question: 2

What are the different locations where you can define the file format options?

a) COPY INTO TABLE statement.
b) Stage definition.
c) Table definition.
d) Database definition

Answer: a, b, c

Question: 3

You get 2MB files per minute everyday. If you consider only snowflake cost, what is the best option to process such files. You can choose any file format that you want.

a) CSV with snowpipe
b) AVRO with snowpipe
c) Create a dedicated virtual warehouse and use JSON file format
d) Create a dedicated virtual warehouse and use AVRO file format

Answer: a

Question: 4

When queried, a stream accesses and returns the historic data in the same shape as the source table (i.e. the same column names and ordering) with additional columns.

What are those columns?

a) METADATA$ACTION
b) METADATA$ROW_NUMBER
c) METADATA$ISUPDATE
d) METADATA$ROW_ID

Answer: a, c, d

Question: 5

What are the two techniques available to query hierarchical data?

a) RECURSIVE CTEs
b) CONNECT WITH
c) CONNECT BY
d) RECURSION

Answer: a, c

Question: 6

Which of the below are benefits of micro partitioning?

a) Micro partitions are derived automatically
b) Micro partitions need to be maintained by users
c) Micro partitions enables extremely efficient DML and fine-grained pruning for faster queries
d) Columns are stored independently within micro-partitions
e) Columns are compressed individually within micro-partitions

Answer: a, c, d, e

Question: 7

In which of the below use cases does Snowflake applies data egress charge?

a) Unloading data from Snowflake
b) Database replication
c) External functions
d) Loading data into Snowflake

Answer: a, b, c

Question: 8

Which information can be obtained from system$clustering_information?

a) max_depth
b) average_depth
c) average_overlaps
d) total_partition_count

Answer: b, c, d

Question: 9

A high churn table has active data size of only 700 GB, however the allocated storage is more than 1 TB. What may be the potential reasons?

a) The table is truncated and loaded everyday. The table has time travel and fail safe enabled.
b) Large loads happen everyday, hence existing micropartitions are deleted and new micro partitions get created
c) The size of the virtual warehouse is too small to handle the table data
d) Incorrect sizing of the table is done

Answer: a, b

Question: 10

Which of the below tools can be used to evaluate the network connection to Snowflake at any time to verify the required configuration settings are correct?

a) SNOWSQL
b) SNOWPIPE
c) SNOWSIGHTS
d) SNOWCD
e) SNOWPARK

Answer: d

Study Guide to Crack Snowflake SnowPro Advanced - Data Engineer DEA-C01 Exam:

- Getting details of the DEA-C01 syllabus, is the first step of a study plan. This pdf is going to be of ultimate help. Completion of the syllabus is must to pass the DEA-C01 exam.
- Making a schedule is vital. A structured method of preparation leads to success. A candidate must plan his schedule and follow it rigorously to attain success.

- Joining the Snowflake provided training for DEA-C01 exam could be of much help. If there is specific training for the exam, you can discover it from the link above.
- Read from the DEA-C01 sample questions to gain your idea about the actual exam questions. In this PDF useful sample questions are provided to make your exam preparation easy.
- Practicing on DEA-C01 practice tests is must. Continuous practice will make you an expert in all syllabus areas.