



**Valid
Questions**

Snowflake

SnowPro-Advanced-Data-Engineer Exam

SnowPro Certification

Questions & Answers

(Demo Version - Limited Content)

Thank you for Downloading SnowPro-Advanced-Data-Engineer exam PDF Demo

Get Full File:

<https://validquestions.com/exam/snowpro-advanced-data-engineer/>

WWW.VALIDQUESTIONS.COM

Version: 6.0

Question: 1

Which are the valid options for the validation_mode parameter in the COPY command

- A. RETURN_ROWSA
- B. RETURN_ERRORB
- C. RETURN_ERRORSC
- D. RETURN_ALL_ERRORS

Answer: A,C,D

Explanation:

VALIDATION_MODE = RETURN_n_ROWS | RETURN_ERRORS | RETURN_ALL_ERRORS

String (constant) that instructs the COPY command to validate the data files instead of loading them into the specified table; i.e. the COPY command tests the files for errors but does not load them. The command validates the data to be loaded and returns results based on the validation option specified:

Supported Values

Notes

RETURN_n_ROWS (e.g. RETURN_10_ROWS)

Validates the specified number of rows, if no errors are encountered; otherwise, fails at the first error encountered in the rows.

RETURN_ERRORS

Returns all errors (parsing, conversion, etc.) across all files specified in the COPY statement.

RETURN_ALL_ERRORS

Returns all errors across all files specified in the COPY statement, including files with errors that were partially loaded during an earlier load because the ON_ERROR copy option was set to CONTINUE during the load.

<https://docs.snowflake.com/en/sql-reference/sql/copy-into-table.html#optional-parameters>

Question: 2

Which of the below functions are recommended to be used to understand the clustering ratio of a table?

- A. SYSTEM\$CLUSTERING_RATIOA
- B. SYSTEM\$CLUSTERING_DEPTHB
- C. SYSTEM\$CLUSTERING_INFORMATION

Answer: B,C

Explanation:

https://docs.snowflake.com/en/sql-reference/functions/system_clustering_ratio.html

Question: 3

You have many files which are loaded onto the cloud storage. Most of them are less than 200 MB in size, but a few are 1GB or more. You need to process them using SNOWPIPE. Which of the below options is recommended

- A. Split the 1 GB files into smaller file sizesA
- B. Use a dedicated large warehouse for the 1 GB files.B
- C. Merger the smaller 200 MB files into a single file.

Answer: A

Explanation:

<https://docs.snowflake.com/en/user-guide/data-load-considerations-prepare.html#general-file-sizing-recommendations>

The number of load operations that run in parallel cannot exceed the number of data files to be loaded. To optimize the number of parallel operations for a load, we recommend aiming to produce data files roughly 100-250 MB (or larger) in size compressed.

Note

Loading very large files (e.g. 100 GB or larger) is not recommended.

If you must load a large file, carefully consider the ON_ERROR copy option value. Aborting or skipping a file due to a small number of errors could result in delays and wasted credits. In addition, if a data loading operation continues beyond the maximum allowed duration of 24 hours, it could be aborted without any portion of the file being committed.

Aggregate smaller files to minimize the processing overhead for each file. Split larger files into a greater number of smaller files to distribute the load among the compute resources in an active warehouse. The number of data files that are processed in parallel is determined by the amount of compute resources in a warehouse. We recommend splitting large files by line to avoid records that span chunks.

If your source database does not allow you to export data files in smaller chunks, you can use a third-party utility to split large CSV files

Question: 4

The employee project details has the project names as array against each employee as shown below

Row	EMP_ID	EMP_NAME	PROJECT_NAMES
1	1	john	['it', 'prod*']

- A. Which of the query below will convert the array into individual rows?
 select emp_id,
 emp_name,
 p.value::string as project_names

- ```

from employee_project_details,table(flatten(employee_project_details.project_names)) p
;
B. select emp_id,
emp_name,
p.value::string as project_names
from employee_project_details,lateral(flatten(employee_project_details.project_names)) p
;
C. select emp_id,
emp_name,
p.value::string as project_names
from employee_project_details,lateral flatten(employee_project_details.project_names)) p

```

---

**Answer: A**

---

**Explanation:**

Try this out in your snowflake instance

Step 1 - Create the table

```
create or replace table employee_project_details(emp_id varchar, emp_name varchar, project_names
array);
```

Step 2 - Insert values

```
insert into employee_project_details
select '1','john',array_cat(to_array('it'),to_array('prod'));
```

Step 3 - Convert to rows

```
select emp_id,
emp_name,
p.value::string as project_names
from employee_project_details,table(flatten(employee_project_details.project_names)) p
```

---

**Question: 5**


---

For using Snowflake with Spark, which of the below privileges are required?

- A. USAGE on the schema that contains the table that you will read from or write toA.
- B. CREATE STAGE on the schema that contains the table that you will read from or write toB.
- C. AccountadminC.
- D. Sysadmin

---

**Answer: A,B**

---

**Explanation:**

<https://docs.snowflake.com/en/user-guide/spark-connector-install.html#requirements>

Requirements

To install and use Snowflake with Spark, you need the following:

A supported operating system. For a list of supported operating systems, see Operating System Support.

Snowflake Connector for Spark.

Snowflake JDBC Driver (the version compatible with the version of the connector).

Apache Spark environment, either self-hosted or hosted in any of the following:

Qubole Data Service.

Databricks.

Amazon EMR.

In addition, you can use a dedicated Amazon S3 bucket or Azure Blob storage container as a staging zone between the two systems; however, this is not required with version 2.2.0 (and higher) of the connector, which uses a temporary Snowflake internal stage (by default) for all data exchange.

The role used in the connection needs USAGE and CREATE STAGE privileges on the schema that contains the table that you will read from or write to.

**Thank You for trying SnowPro-Advanced-Data-Engineer PDF Demo**

<https://validquestions.com/exam/snowpro-advanced-data-engineer/>

**Start Your SnowPro-Advanced-Data-Engineer Preparation**

*[Limited Time Offer]* Use Coupon "**SAVE20**" for extra 20%  
discount the purchase of PDF file. Test your  
SnowPro-Advanced-Data-Engineer preparation with actual exam questions