Eliminates confusion while taking the Snowflake ARA-C01 exam

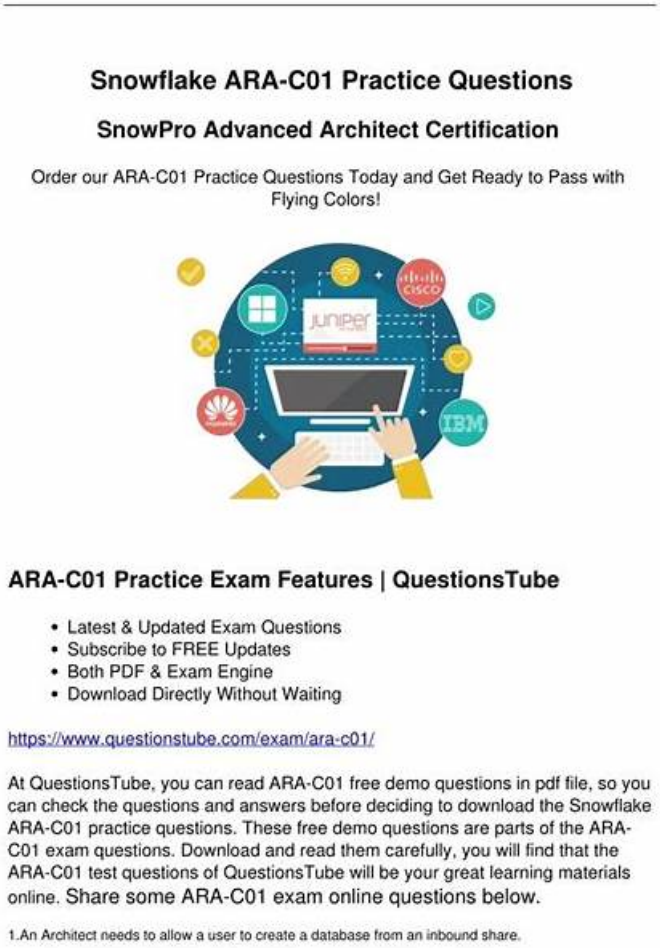

BTW, DOWNLOAD part of ActualVCE ARA-C01 dumps from Cloud Storage: https://drive.google.com/open?id=1HKhzCTO8Nd4iuDgduva6aiAdpxw2vDr7

Our ARA-C01 learning guide is for the world and users are very extensive. In order to give users a better experience, we have been constantly improving. The high quality and efficiency of ARA-C01 test guide has been recognized by users. The high passing rate of ARA-C01 Exam Training is its biggest feature. As long as you use ARA-C01 test guide, you can certainly harvest what you want thing.

Having a Snowflake ARA-C01 certification can enhance your employment prospects,and then you can have a lot of good jobs. ActualVCE is a website very suitable to candidates who participate in the Snowflake certification ARA-C01 exam. ActualVCE can not only provide all the information related to the Snowflake Certification ARA-C01 Exam for the candidates, but also provide a good learning opportunity for them. ActualVCE be able to help you pass Snowflake certification ARA-C01 exam successfully.

>> ARA-C01 Valid Test Testking <<

Snowflake ARA-C01 Valid Test Registration & ARA-C01 Valid Test Vce

The companies do not want to lose them and they offer a good package to convince the candidate to become a part of their organization. So, to fit in the game, you must go for the ActualVCE Snowflake ARA-C01 Practice Exam that will show you where

you stand and how hard you need to work to get the SnowPro Advanced Architect Certification (ARA-C01) certification exam.

In order to prepare for the SnowPro Advanced Architect Certification exam, candidates can take advantage of various resources provided by Snowflake, such as training courses, study guides, and practice exams. Additionally, candidates can also benefit from hands-on experience working with Snowflake's cloud data platform, as well as collaborating with other Snowflake experts and architects.

Snowflake ARA-C01 exam is a challenging test that requires a thorough understanding of Snowflake's architecture and best practices. ARA-C01 exam consists of 60 multiple-choice questions, and candidates have 120 minutes to complete the test. To pass the exam, candidates must score at least 80% or higher. The Snowflake ARA-C01 Certification is valid for two years, after which candidates must retake the exam to maintain their certification.

Snowflake ARA-C01 (SnowPro Advanced Architect Certification) Certification Exam is an essential credential for professionals who want to demonstrate their expertise in Snowflake architecture. SnowPro Advanced Architect Certification certification program is designed to help professionals stay up-to-date with the latest industry trends and technologies, and it can help them advance their careers by demonstrating their skills and knowledge to potential employers.

Snowflake SnowPro Advanced Architect Certification Sample Questions (Q19-Q24):

NEW QUESTION # 19
A table contains five columns and it has millions of records. The cardinality distribution of the columns is shown below:

Column	Number of Distinct Values
C1	10,790
C2	108
C3	302,605
C4	1.117,736
C5	2.205,400

Column C4 and C5 are mostly used by SELECT queries in the GROUP BY and ORDER BY clauses. Whereas columns C1, C2 and C3 are heavily used in filter and join conditions of SELECT queries.
The Architect must design a clustering key for this table to improve the query performance.
Based on Snowflake recommendations, how should the clustering key columns be ordered while defining the multi-column clustering key?

- A. C1, C3, C2
- B. C3, C4, C5
- C. C5, C4, C2
- D. C2, C1, C3

Answer: A

Explanation:
According to the Snowflake documentation, the following are some considerations for choosing clustering for a table1:
Clustering is optimal when either:
You require the fastest possible response times, regardless of cost.
Your improved query performance offsets the credits required to cluster and maintain the table.
Clustering is most effective when the clustering key is used in the following types of query predicates:
Filter predicates (e.g. WHERE clauses)
Join predicates (e.g. ON clauses)
Grouping predicates (e.g. GROUP BY clauses)
Sorting predicates (e.g. ORDER BY clauses)
Clustering is less effective when the clustering key is not used in any of the above query predicates, or when the clustering key is used in a predicate that requires a function or expression to be applied to the key (e.g. DATE_TRUNC, TO_CHAR, etc.).
For most tables, Snowflake recommends a maximum of 3 or 4 columns (or expressions) per key. Adding more than 3-4 columns tends to increase costs more than benefits.
Based on these considerations, the best option for the clustering key columns is C. C1, C3, C2, because:
These columns are heavily used in filter and join conditions of SELECT queries, which are the most effective types of predicates for clustering.
These columns have high cardinality, which means they have many distinct values and can help reduce the clustering skew and

improve the compression ratio.
These columns are likely to be correlated with each other, which means they can help co-locate similar rows in the same micro-partitions and improve the scan efficiency.
These columns do not require any functions or expressions to be applied to them, which means they can be directly used in the predicates without affecting the clustering.

NEW QUESTION # 20
How can the Snowpipe REST API be used to keep a log of data load history?

- A. Call loadHistoryScan every minute for the maximum time range.
- B. Call loadHistoryScan every 10 minutes for a 15-minutes range.
- C. Call insertReport every 20 minutes, fetching the last 10,000 entries.
- D. Call insertReport every 8 minutes for a 10-minute time range.

Answer: B

Explanation:
The Snowpipe REST API provides two endpoints for retrieving the data load history: insertReport and loadHistoryScan. The insertReport endpoint returns the status of the files that were submitted to the insertFiles endpoint, while the loadHistoryScan endpoint returns the history of the files that were actually loaded into the table by Snowpipe. To keep a log of data load history, it is recommended to use the loadHistoryScan endpoint, which provides more accurate and complete information about the data ingestion process. The loadHistoryScan endpoint accepts a start time and an end time as parameters, and returns the files that were loaded within that time range. The maximum time range that can be specified is 15 minutes, and the maximum number of files that can be returned is 10,000. Therefore, to keep a log of data load history, the best option is to call the loadHistoryScan endpoint every 10 minutes for a 15-minute time range, and store the results in a log file or a table. This way, the log will capture all the files that were loaded by Snowpipe, and avoid any gaps or overlaps in the time range. The other options are incorrect because:
Calling insertReport every 20 minutes, fetching the last 10,000 entries, will not provide a complete log of data load history, as some files may be missed or duplicated due to the asynchronous nature of Snowpipe. Moreover, insertReport only returns the status of the files that were submitted, not the files that were loaded.
Calling loadHistoryScan every minute for the maximum time range will result in too many API calls and unnecessary overhead, as the same files will be returned multiple times. Moreover, the maximum time range is 15 minutes, not 1 minute.
Calling insertReport every 8 minutes for a 10-minute time range will suffer from the same problems as option A, and also create gaps or overlaps in the time range.
Reference:
Snowpipe REST API
Option 1: Loading Data Using the Snowpipe REST API
PIPE_USAGE_HISTORY

NEW QUESTION # 21
Consider the following COPY command which is loading data with CSV format into a Snowflake table from an internal stage through a data transformation query.

```
copy into home_sales(city, zip, sale_date, price)
from (select t.$1, t.$2, t.$6, t.$7 from @mystage/sales.csv.qz t)
file_format =
(
format_name = mycsvformat
empty_field_as_null = true
field_optionally_enclosed_by = '"'
)
validation_mode = return_all_errors
;
```


This command results in the following error:
SQL compilation error: invalid parameter 'validation_mode'
Assuming the syntax is correct, what is the cause of this error?

- A. The VALIDATION_MODE parameter supports COPY statements that load data from external stages only.
- B. The value return_all_errors of the option VALIDATION_MODE is causing a compilation error.

- C. The VALIDATION_MODE parameter does not support COPY statements with CSV file formats.
- D. The VALIDATION_MODE parameter does not support COPY statements that transform data during a load.

Answer: D

Explanation:
The VALIDATION_MODE parameter is used to specify the behavior of the COPY statement when loading data into a table. It is used to specify whether the COPY statement should return an error if any of the rows in the file are invalid or if it should continue loading the valid rows. The VALIDATION_MODE parameter is only supported for COPY statements that load data from external stages1.
The query in the question uses a data transformation query to load data from an internal stage. A data transformation query is a query that transforms the data during the load process, such as parsing JSON or XML data, applying functions, or joining with other tables2.
According to the documentation, VALIDATION_MODE does not support COPY statements that transform data during a load. If the parameter is specified, the COPY statement returns an error1. Therefore, option C is the correct answer.

NEW QUESTION # 22
A user needs access to create materialized view on a shema mydb.myschema.
What is the appropriate command to provide the access?

- A. GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO USER1;
- B. GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO USER USER1;
- C. GRANT ROLE MYROLE TO USER USER1; GRANT CREATE MATERIALIZED VIEW ON SCHEMA MYDB.MYSCHEMA TO MYROLE;

Answer: C

NEW QUESTION # 23
When using the COPY INTO
command with the CSV file format, how does the MATCH_BY_COLUMN_NAME parameter behave?

- A. It expects a header to be present in the CSV file, which is matched to a case-sensitive table column name.
- B. The command will return a warning stating that the file has unmatched columns.
- C. The parameter will be ignored.
- D. The command will return an error.

Answer: D

Explanation:
Comprehensive and Detailed Explanation From Exact Extract:
The MATCH_BY_COLUMN_NAME parameter in the COPY INTO
command is used to load semi-structured or structured data, such as CSV, into columns of the target table by matching column names in the data file with those in the table. For CSV files, this parameter requires specific conditions to be met, particularly the presence of a header row in the file, which is used to map columns to the target table.
According to the official Snowflake documentation, when the MATCH_BY_COLUMN_NAME parameter is used with CSV files, it is only supported in specific scenarios and requires the PARSE_HEADER file format option to be set to TRUE. This option indicates that the first row of the CSV file contains column headers, which Snowflake uses to match with the target table's column names. The matching behavior can be configured as CASE_SENSITIVE or CASE_INSENSITIVE, but the default behavior is case-sensitive unless specified otherwise.
However, there is a critical limitation when using MATCH_BY_COLUMN_NAME with CSV files: as of the latest Snowflake documentation, this feature is in Open Private Preview for CSV files and is not generally available for all accounts. When the MATCH_BY_COLUMN_NAME parameter is specified for a CSV file in an environment where this feature is not enabled, or if the PARSE_HEADER option is not set to TRUE, the COPY INTO command will return an error. This is because Snowflake cannot process the column name matching without the header parsing capability, which is not fully supported for CSV files in general availability.
The exact extract from the Snowflake documentation states:
"For loading CSV files, the MATCH_BY_COLUMN_NAME copy option is available in preview. It requires the use of the above-mentioned CSV file format option PARSE_HEADER = TRUE." Additionally, the documentation clarifies:
"Boolean that specifies whether to use the first row headers in the data files to determine column names. This file format option is

applied to the following actions only: Automatically detecting column definitions by using the INFER_SCHEMA function. Loading CSV data into separate columns by using the INFER_SCHEMA function and MATCH_BY_COLUMN_NAME copy option."
Furthermore, a known issue is noted:
"For CSV only, there is a known issue when the INCLUDE_METADATA copy option is used with MATCH_BY_COLUMN_NAME. Do not use this copy option when loading CSV files until the known issue is resolved." Given that the MATCH_BY_COLUMN_NAME parameter is not fully supported for CSV files in general availability and requires specific preview conditions, attempting to use it without meeting those conditions, such as PARSE_HEADER = TRUE or enabling the preview feature, results in an error. Therefore, option C is correct: The command will return an error.
Option A is incorrect because, while MATCH_BY_COLUMN_NAME expects a header in the CSV file for matching when the feature is enabled, the case-sensitive matching is only true when explicitly set to CASE_SENSITIVE. Additionally, the feature's limited availability means it is not guaranteed to work without causing an error. Option B is incorrect because the parameter is not simply ignored; it triggers an error if the conditions are not met. Option D is incorrect because Snowflake does not issue a warning for unmatched columns in this context; it fails with an error when the parameter is unsupported or misconfigured.
References:
Snowflake Documentation: COPY INTO
Snowflake Documentation: Transforming Data During a Load
Stack Overflow: COPY INTO Snowflake Table with Extra Columns

NEW QUESTION # 24
.....

We can say that the Snowflake ARA-C01 exam practice questions are real, valid, and updated SnowPro Advanced Architect Certification (ARA-C01) exam questions that will provide you with everything that you need to learn to prepare and pass the ARA-C01 exam. The Snowflake ARA-C01 Exam Questions will not only assist you in SnowPro Advanced Architect Certification (ARA-C01) exam preparation but also give you sight knowledge about the SnowPro Advanced Architect Certification (ARA-C01) exam topics that will help you in your professional career.

ARA-C01 Valid Test Registration: https://www.actualvce.com/Snowflake/ARA-C01-valid-vce-dumps.html

- Snowflake ARA-C01 Valid Test Testking: SnowPro Advanced Architect Certification - www.prepawaypdf.com 100% Safe Shopping Experience □ Search for ➥ ARA-C01 □ and download it for free immediately on □ www.prepawaypdf.com □ □ARA-C01 Labs
- ARA-C01 - SnowPro Advanced Architect Certification Authoritative Valid Test Testking □ Easily obtain free download of ⇛ ARA-C01 ⇚ by searching on ➡ www.pdfvce.com □□□ □Related ARA-C01 Exams
- 100% Pass Quiz Snowflake - ARA-C01 The Best Valid Test Testking □ Go to website ➤ www.prepawaypdf.com □ open and search for ➤ ARA-C01 □ to download for free □ARA-C01 Dumps Download
- Sample ARA-C01 Exam □ ARA-C01 Real Question □ ARA-C01 Real Question □ Search for ⇛ ARA-C01 ⇚ and download exam materials for free through ▸ www.pdfvce.com ◂ □Sample ARA-C01 Exam
- 100% Free ARA-C01 – 100% Free Valid Test Testking | High Pass-Rate SnowPro Advanced Architect Certification Valid Test Registration □ Go to website 《 www.prepawaypdf.com 》 open and search for " ARA-C01 " to download for free □New ARA-C01 Exam Simulator
- Fantastic ARA-C01 Valid Test Testking – Find Shortcut to Pass ARA-C01 Exam □ Search for ➽ ARA-C01 □ on ☀ www.pdfvce.com □☀□ immediately to obtain a free download □Related ARA-C01 Exams
- Authorized ARA-C01 Pdf □ Sample ARA-C01 Exam □ Exam ARA-C01 Review □ Easily obtain free download of ➽ ARA-C01 □ by searching on ▹ www.prepawayete.com ◃ □Related ARA-C01 Exams
- Pass Guaranteed Snowflake - ARA-C01 - Pass-Sure SnowPro Advanced Architect Certification Valid Test Testking □ Simply search for ➤ ARA-C01 □ for free download on ➥ www.pdfvce.com □ □ARA-C01 Knowledge Points
- ARA-C01 Reliable Dumps Sheet □ ARA-C01 Reliable Dumps Sheet □ Popular ARA-C01 Exams □ Simply search for ⇛ ARA-C01 ⇚ for free download on 【 www.practicevce.com 】 □Popular ARA-C01 Exams
- Authorized ARA-C01 Pdf □ ARA-C01 Labs □ ARA-C01 Reliable Dumps Sheet □ Search for ➥ ARA-C01 □ and download it for free on □ www.pdfvce.com □ website □ARA-C01 Labs
- Reliable ARA-C01 Test Camp □ ARA-C01 Test Dates □ New ARA-C01 Test Papers □ ➽ www.exam4labs.com □ is best website to obtain " ARA-C01 " for free download ♣New ARA-C01 Test Papers
- www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, www.stes.tyc.edu.tw, lizellehartley.com.au, www.dmb-pla.com, msadvisory.co.zw, www.stes.tyc.edu.tw, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, Disposable vapes

P.S. Free & New ARA-C01 dumps are available on Google Drive shared by ActualVCE: https://drive.google.com/open?

id=1HKhzCTO8Nd4iuDgduva6aiAdpxw2vDr7