

Snowflake ARA-C01 Reliable Test Experience - New ARA-C01 Test Pdf

Questions & Answers PDF

Page 1



Thank you for Downloading ARA-C01 exam PDF Demo

You can Buy Latest ARA-C01 Full Version Download

<https://www.certkillers.net/Exam/ARA-C01>

<https://www.certkillers.net>

P.S. Free 2025 Snowflake ARA-C01 dumps are available on Google Drive shared by Actual4Cert: https://drive.google.com/open?id=1tKU-3G2cUy_ygW5uN8lazYP9jVKcaL1m

Consistent practice with it relieves exam stress and boosts self-confidence. The web-based ARA-C01 practice exam does not require additional software installation. All operating systems also support this SnowPro Advanced Architect Certification (ARA-C01) practice test. We update our SnowPro Advanced Architect Certification (ARA-C01) pdf format regularly so keep calm because you will always get updated SnowPro Advanced Architect Certification (ARA-C01) questions.

The Snowflake ARA-C01 exam consists of 90 multiple-choice questions that must be completed within two hours. The questions are designed to test an individual's knowledge of Snowflake's architecture, including multi-cluster warehouses, virtual warehouses, and resource management. It also covers topics such as data modeling, security, performance optimization, and data integration.

Snowflake ARA-C01 certification exam covers a wide range of topics, including Snowflake architecture, data warehousing, data modeling, data integration, security, and performance optimization. ARA-C01 exam consists of 70 multiple-choice questions, and the candidate is given 120 minutes to complete the exam. ARA-C01 Exam is conducted online, and the candidates are provided with a secure browser that is used to access the exam. ARA-C01 exam is designed to test the candidate's knowledge and skills in a variety of areas related to Snowflake architecture and data warehousing. Successful completion of the SnowPro Advanced Architect Certification Exam demonstrates that the candidate possesses the advanced skills and knowledge required to design and implement complex Snowflake solutions using best practices.

>> Snowflake ARA-C01 Reliable Test Experience <<

New ARA-C01 Test Pdf & Detailed ARA-C01 Answers

We hold on to inflexible will power to offer help both providing the high-rank ARA-C01 exam guide as well as considerate after-seals services. With our ARA-C01 study tools' help, passing the exam will be a matter of course. It is our abiding belief to support your preparation of the ARA-C01 study tools with enthusiastic attitude towards our jobs. And all efforts are paid off. Our ARA-C01 Exam Torrent is highly regarded in the market of this field and come with high recommendation. Choosing our ARA-C01 exam guide will be a very promising start for you to begin your exam preparation because our ARA-C01 practice materials with high repute.

Snowflake SnowPro Advanced Architect Certification Sample Questions (Q117-Q122):

NEW QUESTION # 117

A Developer is having a performance issue with a Snowflake query. The query receives up to 10 different values for one parameter and then performs an aggregation over the majority of a fact table. It then joins against a smaller dimension table. This parameter value is selected by the different query users when they execute it during business hours. Both the fact and dimension tables are loaded with new data in an overnight import process.

On a Small or Medium-sized virtual warehouse, the query performs slowly. Performance is acceptable on a size Large or bigger warehouse. However, there is no budget to increase costs. The Developer needs a recommendation that does not increase compute costs to run this query.

What should the Architect recommend?

- A. Create a dedicated size Large warehouse for this particular set of queries. Create a new role that has USAGE permission on this warehouse and has the appropriate read permissions over the fact and dimension tables. Have users switch to this role and use this warehouse when they want to access this data.
- B. Create a task that will run the 10 different variations of the query corresponding to the 10 different parameters before the users come in to work. The query results will then be cached and ready to respond quickly when the users re-issue the query.
- C. Enable the search optimization service on the table. When the users execute the query, the search optimization service will automatically adjust the query execution plan based on the frequently-used parameters.
- D. Create a task that will run the 10 different variations of the query corresponding to the 10 different parameters before the users come in to work. The task will be scheduled to align with the users' working hours in order to allow the warehouse cache to be used.

Answer: C

Explanation:

Enabling the search optimization service on the table can improve the performance of queries that have selective filtering criteria, which seems to be the case here. This service optimizes the execution of queries by creating a persistent data structure called a search access path, which allows some micro-partitions to be skipped during the scanning process. This can significantly speed up query performance without increasing compute costs¹.

Reference

* Snowflake Documentation on Search Optimization Service¹.

NEW QUESTION # 118

Which system functions does Snowflake provide to monitor clustering information within a table (Choose two.)

- A. SYSTEM\$CLUSTERING_DEPTH
- B. SYSTEM\$CLUSTERING_USAGE
- C. SYSTEM\$CLUSTERING_PERCENT
- D. SYSTEM\$CLUSTERING_KEYS
- E. SYSTEM\$CLUSTERING_INFORMATION

Answer: A,E

Explanation:

According to the Snowflake documentation, these two system functions are provided by Snowflake to monitor clustering information within a table. A system function is a type of function that allows executing actions or returning information about the system. A clustering key is a feature that allows organizing data across micro-partitions based on one or more columns in the table. Clustering can improve query performance by reducing the number of files to scan.

* `SYSTEM$CLUSTERING_INFORMATION` is a system function that returns clustering information, including average clustering depth, for a table based on one or more columns in the table. The function takes a table name and an optional column name or expression as arguments, and returns a JSON string with the clustering information. The clustering information includes the cluster by keys, the total partition count, the total constant partition count, the average overlaps, and the average depth¹.

* `SYSTEM$CLUSTERING_DEPTH` is a system function that returns the clustering depth for a table based on one or more columns in the table. The function takes a table name and an optional column name or expression as arguments, and returns an integer value with the clustering depth. The clustering depth is the maximum number of overlapping micro-partitions for any micro-partition in the table. A lower clustering depth indicates a better clustering².

`SYSTEM$CLUSTERING_INFORMATION` | Snowflake Documentation

`SYSTEM$CLUSTERING_DEPTH` | Snowflake Documentation

NEW QUESTION # 119

A retail company has over 3000 stores all using the same Point of Sale (POS) system. The company wants to deliver near real-time sales results to category managers. The stores operate in a variety of time zones and exhibit a dynamic range of transactions each minute, with some stores having higher sales volumes than others.

Sales results are provided in a uniform fashion using data engineered fields that will be calculated in a complex data pipeline.

Calculations include exceptions, aggregations, and scoring using external functions interfaced to scoring algorithms. The source data for aggregations has over 100M rows.

Every minute, the POS sends all sales transactions files to a cloud storage location with a naming convention that includes store numbers and timestamps to identify the set of transactions contained in the files. The files are typically less than 10MB in size.

How can the near real-time results be provided to the category managers? (Select TWO).

- A. An external scheduler should examine the contents of the cloud storage location and issue SnowSQL commands to process the data at a frequency that matches the real-time analytics needs.
- B. All files should be concatenated before ingestion into Snowflake to avoid micro-ingestion.
- C. The copy into command with a task scheduled to run every second should be used to achieve the near-real time requirement.
- **D. A stream should be created to accumulate the near real-time data and a task should be created that runs at a frequency that matches the real-time analytics needs.**
- **E. A Snowpipe should be created and configured with `AUTO_INGEST = true`. A stream should be created to process INSERTS into a single target table using the stream metadata to inform the store number and timestamps.**

Answer: D,E

Explanation:

To provide near real-time sales results to category managers, the Architect can use the following steps:

Create an external stage that references the cloud storage location where the POS sends the sales transactions files. The external stage should use the file format and encryption settings that match the source files² Create a Snowpipe that loads the files from the external stage into a target table in Snowflake. The Snowpipe should be configured with `AUTO_INGEST = true`, which means that it will automatically detect and ingest new files as they arrive in the external stage. The Snowpipe should also use a copy option to purge the files from the external stage after loading, to avoid duplicate ingestion³ Create a stream on the target table that captures the INSERTS made by the Snowpipe. The stream should include the metadata columns that provide information about the file name, path, size, and last modified time. The stream should also have a retention period that matches the real-time analytics needs⁴ Create a task that runs a query on the stream to process the near real-time data. The query should use the stream metadata to extract the store number and timestamps from the file name and path, and perform the calculations for exceptions, aggregations, and scoring using external functions. The query should also output the results to another table or view that can be accessed by the category managers. The task should be scheduled to run at a frequency that matches the real-time analytics needs, such as every minute or every 5 minutes.

The other options are not optimal or feasible for providing near real-time results:

All files should be concatenated before ingestion into Snowflake to avoid micro-ingestion. This option is not recommended because it would introduce additional latency and complexity in the data pipeline. Concatenating files would require an external process or service that monitors the cloud storage location and performs the file merging operation. This would delay the ingestion of new files into Snowflake and increase the risk of data loss or corruption. Moreover, concatenating files would not avoid micro-ingestion, as Snowpipe would still ingest each concatenated file as a separate load.

An external scheduler should examine the contents of the cloud storage location and issue SnowSQL commands to process the data at a frequency that matches the real-time analytics needs. This option is not necessary because Snowpipe can automatically ingest new files from the external stage without requiring an external trigger or scheduler. Using an external scheduler would add more overhead and dependency to the data pipeline, and it would not guarantee near real-time ingestion, as it would depend on the polling interval and the availability of the external scheduler.

The copy into command with a task scheduled to run every second should be used to achieve the near-real time requirement. This

option is not feasible because tasks cannot be scheduled to run every second in Snowflake. The minimum interval for tasks is one minute, and even that is not guaranteed, as tasks are subject to scheduling delays and concurrency limits. Moreover, using the copy into command with a task would not leverage the benefits of Snowpipe, such as automatic file detection, load balancing, and micro-partition optimization. Reference:

1: SnowPro Advanced: Architect | Study Guide
2: Snowflake Documentation | Creating Stages
3: Snowflake Documentation | Loading Data Using Snowpipe
4: Snowflake Documentation | Using Streams and Tasks for ELT
: Snowflake Documentation | Creating Tasks
: Snowflake Documentation | Best Practices for Loading Data
: Snowflake Documentation | Using the Snowpipe REST API
: Snowflake Documentation | Scheduling Tasks
: SnowPro Advanced: Architect | Study Guide
: Creating Stages
: Loading Data Using Snowpipe
: Using Streams and Tasks for ELT
: [Creating Tasks]
: [Best Practices for Loading Data]
: [Using the Snowpipe REST API]
: [Scheduling Tasks]

NEW QUESTION # 120

What transformations are supported in the below SQL statement? (Select THREE).

CREATE PIPE ... AS COPY ... FROM (...)

- A. The ON ERROR - ABORT statement command can be used.
- B. Data can be filtered by an optional where clause.
- C. Columns can be reordered.
- D. Columns can be omitted.
- E. Incoming data can be joined with other tables.
- F. Type casts are supported.

Answer: B,C,D

Explanation:

The SQL statement is a command for creating a pipe in Snowflake, which is an object that defines the COPY INTO <table> statement used by Snowpipe to load data from an ingestion queue into tables1. The statement uses a subquery in the FROM clause to transform the data from the staged files before loading it into the table2.

The transformations supported in the subquery are as follows2:

Data can be filtered by an optional WHERE clause, which specifies a condition that must be satisfied by the rows returned by the subquery. For example:

SQLAI-generated code. Review and use carefully. More info on FAQ.

```
create pipe mypipe as
copy into mytable
from (
select * from @mystage
where col1 = 'A' and col2 > 10
);
```

Columns can be reordered, which means changing the order of the columns in the subquery to match the order of the columns in the target table. For example:

SQLAI-generated code. Review and use carefully. More info on FAQ.

```
create pipe mypipe as
copy into mytable (col1, col2, col3)
from (
select col3, col1, col2 from @mystage
);
```

Columns can be omitted, which means excluding some columns from the subquery that are not needed in the target table. For example:

SQLAI-generated code. Review and use carefully. More info on FAQ.

```
create pipe mypipe as
```

```
copy into mytable (col1, col2)
from (
select col1, col2 from @mystage
);
```

The other options are not supported in the subquery because2:

Type casts are not supported, which means changing the data type of a column in the subquery. For example, the following statement will cause an error:

SQLAI-generated code. Review and use carefully. More info on FAQ.

```
create pipe mypipe as
copy into mytable (col1, col2)
from (
select col1::date, col2 from @mystage
);
```

Incoming data can not be joined with other tables, which means combining the data from the staged files with the data from another table in the subquery. For example, the following statement will cause an error:

SQLAI-generated code. Review and use carefully. More info on FAQ.

```
create pipe mypipe as
copy into mytable (col1, col2, col3)
from (
select s.col1, s.col2, t.col3 from @mystage s
join othertable t on s.col1 = t.col1
);
```

The ON ERROR - ABORT statement command can not be used, which means aborting the entire load operation if any error occurs. This command can only be used in the COPY INTO <table> statement, not in the subquery. For example, the following statement will cause an error:

SQLAI-generated code. Review and use carefully. More info on FAQ.

```
create pipe mypipe as
copy into mytable
from (
select * from @mystage
on error abort
);
```

Reference:

1: CREATE PIPE | Snowflake Documentation

2: Transforming Data During a Load | Snowflake Documentation

NEW QUESTION # 121

It is recommended to assign ACCOUNTADMIN role to atleast two user

- A. FALSE
- B. TRUE

Answer: B

NEW QUESTION # 122

.....

The PDF version of our ARA-C01 guide exam is prepared for you to print it and read it everywhere. It is convenient for you to see the answers to the questions and remember them. After you buy the PDF version of our study material, you will get an E-mail form us in 5 to 10 minutes after payment. Then you can click the link in the E-mail and download your ARA-C01 study engine. You can download it as many times as you need.

New ARA-C01 Test Pdf: <https://www.actual4cert.com/ARA-C01-real-questions.html>

- www.vce4dumps.com Snowflake ARA-C01 Practice Test ↘ Go to website { www.vce4dumps.com } open and search for ☐ ARA-C01 ☐ to download for free ♣ Latest Braindumps ARA-C01 Ppt
- ARA-C01 Reliable Test Online ☐ Certification ARA-C01 Sample Questions ☐ Exam ARA-C01 Papers ☐ Download ➤ ARA-C01 ☐ for free by simply entering (www.pdfvce.com) website ☐ ARA-C01 PDF VCE
- Pass Guaranteed 2026 Snowflake ARA-C01: Valid SnowPro Advanced Architect Certification Reliable Test Experience ☐

P.S. Free 2025 Snowflake ARA-C01 dumps are available on Google Drive shared by Actual4Cert: https://drive.google.com/open?id=1tKU-3G2cUy_ygW5uN8lazYP9jVKcaL1m