

# Valid Test Snowflake ARA-C01 Bootcamp - Latest ARA-C01 Dumps Sheet



BONUS!!! Download part of TorrentExam ARA-C01 dumps for free: <https://drive.google.com/open?id=1bBIVAr-UG7W-8HaqScDhLcxSlf5UnSEB>

It is known to us that our ARA-C01 study materials are enjoying a good reputation all over the world. Our study materials have been approved by thousands of candidates. You may have some doubts about our product or you may suspect the pass rate of it, but we will tell you clearly, it is totally unnecessary. If you still do not trust us, you can choose to download demo of our ARA-C01 Test Torrent. The high quality and the perfect service system after sale of our ARA-C01 exam questions have been apporobated by our local and international customers. So you can rest assured to buy.

Snowflake ARA-C01: SnowPro Advanced Architect Certification is a prestigious and highly sought-after certification that validates the skills and knowledge of experienced architects and consultants working with Snowflake's cloud data platform. ARA-C01 Exam is designed to test the advanced level of expertise in designing, deploying, managing, and optimizing Snowflake solutions.

>> Valid Test Snowflake ARA-C01 Bootcamp <<

## Latest ARA-C01 Dumps Sheet, New ARA-C01 Exam Camp

simulation tests of our ARA-C01 learning materials have the functions of timing and mocking exams, which will allow you to adapt to the exam environment in advance and it will be of great benefit for subsequent exams. After you complete the learning task, the system of our ARA-C01 test prep will generate statistical reports based on your performance so that you can identify your weaknesses and conduct targeted training and develop your own learning plan. For the complex part of our ARA-C01 Exam Question, you may be too cumbersome, but our system has explained and analyzed this according to the actual situation to eliminate your doubts and make you learn better.

## Snowflake SnowPro Advanced Architect Certification Sample Questions (Q119-Q124):

### NEW QUESTION # 119

An Architect is troubleshooting a query with poor performance using the QUERY function. The Architect observes that the

COMPILE\_TIME Is greater than the EXECUTION\_TIME.

What is the reason for this?

- A. The query Is reading from remote storage
- B. The query Is queued for execution.
- C. The query is processing a very large dataset.
- D. The query has overly complex logic.

**Answer: D**

Explanation:

- \* The correct answer is B because the compilation time is the time it takes for the optimizer to create an optimal query plan for the efficient execution of the query. The compilation time depends on the complexity of the query, such as the number of tables, columns, joins, filters, aggregations, subqueries, etc. The more complex the query, the longer it takes to compile.
- \* Option A is incorrect because the query processing time is not affected by the size of the dataset, but by the size of the virtual warehouse. Snowflake automatically scales the compute resources to match the data volume and parallelizes the query execution. The size of the dataset may affect the execution time, but not the compilation time.
- \* Option C is incorrect because the query queue time is not part of the compilation time or the execution time. It is a separate metric that indicates how long the query waits for a warehouse slot before it starts running. The query queue time depends on the warehouse load, concurrency, and priority settings.
- \* Option D is incorrect because the query remote IO time is not part of the compilation time or the execution time. It is a separate metric that indicates how long the query spends reading data from remote storage, such as S3 or Azure Blob Storage. The query remote IO time depends on the network latency, bandwidth, and caching efficiency. References:
  - \* Understanding Why Compilation Time in Snowflake Can Be Higher than Execution Time: This article explains why the total duration (compilation + execution) time is an essential metric to measure query performance in Snowflake. It discusses the reasons for the long compilation time, including query complexity and the number of tables and columns.
  - \* Exploring Execution Times: This document explains how to examine the past performance of queries and tasks using Snowsight or by writing queries against views in the ACCOUNT\_USAGE schema. It also describes the different metrics and dimensions that affect query performance, such as duration, compilation, execution, queue, and remote IO time.
  - \* What is the "compilation time" and how to optimize it?: This community post provides some tips and best practices on how to reduce the compilation time, such as simplifying the query logic, using views or common table expressions, and avoiding unnecessary columns or joins.

## NEW QUESTION # 120

An Architect needs to automate the daily Import of two files from an external stage into Snowflake. One file has Parquet-formatted data, the other has CSV-formatted data.

How should the data be joined and aggregated to produce a final result set?

- A. Use Snowpipe to ingest the two files, then create a materialized view to produce the final result set.
- B. Create a JavaScript stored procedure to read, join, and aggregate the data directly from the external stage, and then store the results in a table.
- C. Create a materialized view to read, Join, and aggregate the data directly from the external stage, and use the view to produce the final result set
- D. Create a task using Snowflake scripting that will import the files, and then call a User-Defined Function (UDF) to produce the final result set.

**Answer: D**

Explanation:

According to the Snowflake documentation, tasks are objects that enable scheduling and execution of SQL statements or JavaScript user-defined functions (UDFs) in Snowflake. Tasks can be used to automate data loading, transformation, and maintenance operations. Snowflake scripting is a feature that allows writing procedural logic using SQL statements and JavaScript UDFs. Snowflake scripting can be used to create complex workflows and orchestrate tasks. Therefore, the best option to automate the daily import of two files from an external stage into Snowflake, join and aggregate the data, and produce a final result set is to create a task using Snowflake scripting that will import the files using the COPY INTO command, and then call a UDF to perform the join and aggregation logic. The UDF can return a table or a variant value as the final result set. References:

- \* Tasks
- \* Snowflake Scripting
- \* User-Defined Functions

### NEW QUESTION # 121

Arrange in order of performance(least to high)

1. External Table
2. External Table with partitioning
3. Materialized View Vs External table
4. Directly querying a file from the stage

- A. 4,3,2,1
- B. 3,2,4,1
- C. 1,2,3,4
- D. 1,2,4,3

**Answer: C**

### NEW QUESTION # 122

What built-in Snowflake features make use of the change tracking metadata for a table? (Choose two.)

- A. A STREAM object
- B. The UPSERT command
- C. The MERGE command
- D. Thee CHANGE DATA\_CAPTURE command
- E. The CHANGES clause

**Answer: A,E**

### NEW QUESTION # 123

A table, EMP\_TBL has three records as shown:

The following variables are set for the session:

Which SELECT statements will retrieve all three records? (Select TWO).

- A. Select \* FROM \$tbl\_ref WHERE \$col\_ref IN ('Name1','Name2','Name3');
- B. SELECT \* FROM identifier<\$tbl\_ref> WHERE NAME IN (\$var1, \$var2, \$var3);
- C. **SELECT \* FROM EMP\_TBL WHERE identifier(\$col\_ref) IN ('Name1','Name2','Name3');**
- D. **SELECT \* FROM \$tbl\_ref WHERE \$col\_ref IN (\$var1, Svar2, Svar3);**
- E. SELECT \* FROM identifier(\$tbl\_ref) WHERE ID IN Cvar1,'var2','var3');

**Answer: C,D**

Explanation:

- \* The correct answer is B and E because they use the correct syntax and values for the identifier function and the session variables.
- \* The identifier function allows you to use a variable or expression as an identifier (such as a table name or column name) in a SQL statement. It takes a single argument and returns it as an identifier. For example, identifier(\$tbl\_ref) returns EMP\_TBL as an identifier.
- \* The session variables are set using the SET command and can be referenced using the \$ sign. For example, \$var1 returns Name1 as a value.
- \* Option A is incorrect because it uses \$tbl\_ref and \$col\_ref, which are not valid session variables or identifiers. They should be \$tbl\_ref and \$col\_ref instead.
- \* Option C is incorrect because it uses identifier<\$tbl\_ref>, which is not a valid syntax for the identifier function. It should be identifier(\$tbl\_ref) instead.
- \* Option D is incorrect because it uses Cvar1, var2, and var3, which are not valid session variables or values. They should be \$var1, \$var2, and \$var3 instead. References:
  - \* Snowflake Documentation: Identifier Function
  - \* Snowflake Documentation: Session Variables
  - \* Snowflake Learning: SnowPro Advanced: Architect Exam Study Guide

### NEW QUESTION # 124

If you still have no confidence for passing test, here we will recommend you an excellent reference material. Our valid ARA-C01 exam collection pdf will help you pass exam and go to success, you will approach to IT field top. You can just spend short time in preparing for real test with our latest ARA-C01 Exam Collection Pdf. You can download free demo in our website for your reference to verify the reliability of our dumps before purchasing.

**Latest ARA-C01 Dumps Sheet:** <https://www.torrentexam.com/ARA-C01-exam-latest-torrent.html>

BTW, DOWNLOAD part of TorrentExam ARA-C01 dumps from Cloud Storage: <https://drive.google.com/open?id=1bBIVAr-UG7W-8HaqScDhLcxSlf5UnSEB>