

# ARA-C01 Latest Mock Exam, Exam ARA-C01 Braindumps



BTW, DOWNLOAD part of ExamBoosts ARA-C01 dumps from Cloud Storage: <https://drive.google.com/open?id=1JizoFBWa9T3Ql9Qoigffidja1vCunUzTr>

As you know that the number of the questions and answers in the real ARA-C01 exam is fixed. So accordingly the information should be collected for you. Our ARA-C01 study materials have done the right thing for you. However, we will never display all the information in order to make the content appear more. Our ARA-C01 learning guide just want to give you the most important information. This is why ARA-C01 actual exam allow you to take the exam in the shortest possible time.

From the moment you decide to contact with us for the ARA-C01 exam braindumps, you are enjoying our fast and professional service. Some of our customers may worry that we are working on certain time about our ARA-C01 study guide. In fact, you don't need to worry at all. You can contact us at any time. The reason why our staff is online 24 hours is to be able to help you solve problems about our ARA-C01 simulating exam at any time. We know that your time is very urgent, so we do not want you to be delayed by some unnecessary trouble.

>> ARA-C01 Latest Mock Exam <<

## Exam ARA-C01 Braindumps & ARA-C01 Examcollection Vce

It is very necessary for candidates to get valid ARA-C01 dumps collection because it can save your time and help you get succeed in IT filed by clearing ARA-C01 actual test. Passing real exam is not easy task so many people need to take professional suggestions to prepare ARA-C01 Practice Exam. The reason that we get good reputation among dump vendors is the most reliable ARA-C01 pdf vce and the best-quality service.

## Snowflake SnowPro Advanced Architect Certification Sample Questions (Q105-Q110):

### NEW QUESTION # 105

An Architect runs the following SQL query:

□

How can this query be interpreted?

- A. FILEROWS is a file. FILE\_ROW\_NUMBER is the file format location.
- B. FILERONS is the file format location. FILE\_ROW\_NUMBER is a stage.
- C. FILEROWS is the table. FILE\_ROW\_NUMBER is the line number in the table.
- D. **FILEROWS is a stage. FILE\_ROW\_NUMBER is line number in file.**

**Answer: D**

### NEW QUESTION # 106

Data is being imported and stored as JSON in a VARIANT column. Query performance was fine, but most recently, poor query performance has been reported.

What could be causing this?

- A. There were variations in string lengths for the JSON values in the recent data imports.
- B. There were JSON nulls in the recent data imports.
- C. **The order of the keys in the JSON was changed.**
- D. The recent data imports contained fewer fields than usual.

**Answer: C**

Explanation:

Data is being imported and stored as JSON in a VARIANT column. Query performance was fine, but most recently, poor query performance has been reported. This could be caused by the following factors:

\* The order of the keys in the JSON was changed. Snowflake stores semi-structured data internally in a column-like structure for the most common elements, and the remainder in a leftovers-like column. The order of the keys in the JSON affects how Snowflake determines the common elements and how it optimizes the query performance. If the order of the keys in the JSON was changed, Snowflake might have to re-parse the data and re-organize the internal storage, which could result in slower query performance.

\* There were variations in string lengths for the JSON values in the recent data imports. Non-native values, such as dates and timestamps, are stored as strings when loaded into a VARIANT column.

Operations on these values could be slower and also consume more space than when stored in a relational column with the corresponding data type. If there were variations in string lengths for the JSON values in the recent data imports, Snowflake might have to allocate more space and perform more conversions, which could also result in slower query performance.

The other options are not valid causes for poor query performance:

\* There were JSON nulls in the recent data imports. Snowflake supports two types of null values in semi-structured data: SQL NULL and JSON null. SQL NULL means the value is missing or unknown, while JSON null means the value is explicitly set to null. Snowflake can distinguish between these two types of null values and handle them accordingly. Having JSON nulls in the recent data imports should not affect the query performance significantly.

\* The recent data imports contained fewer fields than usual. Snowflake can handle semi-structured data with varying schemas and fields. Having fewer fields than usual in the recent data imports should not affect the query performance significantly, as Snowflake can still optimize the data ingestion and query execution based on the existing fields.

References:

- \* Considerations for Semi-structured Data Stored in VARIANT
- \* Snowflake Architect Training
- \* Snowflake query performance on unique element in variant column
- \* Snowflake variant performance

### NEW QUESTION # 107

What are purposes for creating a storage integration? (Choose three.)

- A. Control access to Snowflake data using a master encryption key that is maintained in the cloud provider's key management service.
- B. **Store a generated identity and access management (IAM) entity for an external cloud provider regardless of the cloud provider that hosts the Snowflake account.**
- C. **Avoid supplying credentials when creating a stage or when loading or unloading data.**
- D. **Support multiple external stages using one single Snowflake object.**
- E. Manage credentials from multiple cloud providers in one single Snowflake object.
- F. Create private VPC endpoints that allow direct, secure connectivity between VPCs without traversing the public internet.

**Answer: B,C,D**

Explanation:

A storage integration is a Snowflake object that stores a generated identity and access management (IAM) entity for an external cloud provider, such as Amazon S3, Google Cloud Storage, or Microsoft Azure Blob Storage. This integration allows Snowflake to read data from and write data to an external storage location referenced in an external stage1.

One purpose of creating a storage integration is to support multiple external stages using one single Snowflake object. An integration can list buckets (and optional paths) that limit the locations users can specify when creating external stages that use the integration. Note that many external stage objects can reference different buckets and paths and use the same storage integration for authentication1. Therefore, option C is correct.

Another purpose of creating a storage integration is to avoid supplying credentials when creating a stage or when loading or unloading data. Integrations are named, first-class Snowflake objects that avoid the need for passing explicit cloud provider credentials such as secret keys or access tokens. Integration objects store an IAM user ID, and an administrator in your organization grants the IAM user permissions in the cloud provider account1. Therefore, option D is correct.

A third purpose of creating a storage integration is to store a generated IAM entity for an external cloud provider regardless of the cloud provider that hosts the Snowflake account. For example, you can create a storage integration for Amazon S3 even if your Snowflake account is hosted on Azure or Google Cloud Platform. This allows you to access data across different cloud platforms using Snowflake1. Therefore, option B is correct.

Option A is incorrect, because creating a storage integration does not control access to Snowflake data using a master encryption key. Snowflake encrypts all data using a hierarchical key model, and the master encryption key is managed by Snowflake or by the customer using a cloud provider's key management service. This is independent of the storage integration feature2.

Option E is incorrect, because creating a storage integration does not create private VPC endpoints. Private VPC endpoints are a network configuration option that allow direct, secure connectivity between VPCs without traversing the public internet. This is also independent of the storage integration feature3.

Option F is incorrect, because creating a storage integration does not manage credentials from multiple cloud providers in one single Snowflake object. A storage integration is specific to one cloud provider, and you need to create separate integrations for each cloud provider you want to access4.

**NEW QUESTION # 108**

A company has a Snowflake account named ACCOUNTA in AWS us-east-1 region. The company stores its marketing data in a Snowflake database named MARKET\_DB. One of the company's business partners has an account named PARTNERB in Azure East US 2 region. For marketing purposes the company has agreed to share the database MARKET\_DB with the partner account. Which of the following steps MUST be performed for the account PARTNERB to consume data from the MARKET\_DB database?

- A. Create a new account (called AZABC123) in Azure East US 2 region. From account ACCOUNTA replicate the database MARKET\_DB to AZABC123 and from this account set up the data sharing to the PARTNERB account.
- B. Create a new account (called AZABC123) in Azure East US 2 region. From account ACCOUNTA create a share of database MARKET\_DB, create a new database out of this share locally in AWS us-east-1 region, and replicate this new database to AZABC123 account. Then set up data sharing to the PARTNERB account.
- C. Create a share of database MARKET\_DB, and create a new database out of this share locally in AWS us-east-1 region. Then replicate this database to the partner's account PARTNERB.
- D. From account ACCOUNTA create a share of database MARKET\_DB, and create a new database out of this share locally in AWS us-east-1 region. Then make this database the provider and share it with the PARTNERB account.

**Answer: A**

Explanation:

Explanation

\* Snowflake supports data sharing across regions and cloud platforms using account replication and share replication features.

Account replication enables the replication of objects from a source account to one or more target accounts in the same organization. Share replication enables the replication of shares from a source account to one or more target accounts in the same organization1.

\* To share data from the MARKET\_DB database in the ACCOUNTA account in AWS us-east-1 region with the PARTNERB account in Azure East US 2 region, the following steps must be performed:

\* Create a new account (called AZABC123) in Azure East US 2 region. This account will act as a bridge between the source and the target accounts. The new account must be linked to the ACCOUNTA account using an organization2.

\* From the ACCOUNTA account, replicate the MARKET\_DB database to the AZABC123 account using the account replication feature. This will create a secondary database in the AZABC123 account that is a replica of the primary database in the ACCOUNTA account3.

\* From the AZABC123 account, set up the data sharing to the PARTNERB account using the share replication feature. This will create a share of the secondary database in the AZABC123 account and grant access to the PARTNERB account. The PARTNERB account can then create a database from the share and query the data4.

\* Therefore, option C is the correct answer.

References: : Replicating Shares Across Regions and Cloud Platforms : Working with Organizations and Accounts : Replicating Databases Across Multiple Accounts : Replicating Shares Across Multiple Accounts

### NEW QUESTION # 109

An event table has 150B rows and 1.5M micro-partitions, with the following statistics:

Column NDV\*

A\_ID 11K

C\_DATE 110

NAME 300K

EVENT\_ACT\_0 1.1G

EVENT\_ACT\_4 2.2G

\*NDV = Number of Distinct Values

What three clustering keys should be used, in order?

- A. C\_DATE, A\_ID, EVENT\_ACT\_0
- B. C\_DATE, A\_ID, NAME
- C. A\_ID, NAME, C\_DATE
- D. C\_DATE, A\_ID, EVENT\_ACT\_4

**Answer: A**

Explanation:

Comprehensive and Detailed 150 to 250 words of Explanation From Snowflake SnowPro Architect exam scope and all publicly documented material:

Clustering keys are most beneficial when they improve micro-partition pruning for common filter patterns and when the chosen columns provide a useful ordering that co-locates data. A common heuristic is to place lower-cardinality columns earlier (to quickly narrow partitions) and then add a higher-cardinality column that further reduces scanned partitions for selective access paths. Here, C\_DATE has very low NDV (110), making it an excellent leading key to organize data by date and enable strong pruning for time-bound queries typical of event tables. Next, A\_ID (11K) is moderate cardinality and can further segment data within a date range, helping point lookups or narrow scans by identifier. For the third key, the options force choosing between very high-cardinality event activity columns; selecting EVENT\_ACT\_0 (1.1G) is preferable to EVENT\_ACT\_4 (2.2G) because it is comparatively less distinct while still supporting additional pruning when queries filter by that attribute. This ordering aligns with Snowflake guidance: keep keys few, ordered to match common predicates, and avoid excessively high-cardinality keys unless they directly match frequent selective filters.

---

---

### NEW QUESTION # 110

.....

It is never too late to learn. You still have the chance to obtain the ARA-C01 certificate as long as you want. What is more, many people have harvested happiness and success after passing the ARA-C01 exam. Then you are available for various high salary jobs. You also can become lucky as long as you never give up hopes. Let us make it together. We will be your best friend on your way to get the ARA-C01 Certification with our excellent learning braindumps.

**Exam ARA-C01 Braindumps:** <https://www.examboosts.com/Snowflake/ARA-C01-practice-exam-dumps.html>

Snowflake ARA-C01 Latest Mock Exam It's just like you are the king, then countless people support and work for you, The 98%-99% pass rate has helped many candidates passed the actual test and got the ARA-C01 certification successfully, Why we let you try our ARA-C01 updated dumps free demo before you purchase, Have you ever dreamed about passing the most important exam such as Snowflake ARA-C01 in your field with great ease?

Edit the drummer performance, Streaming protocol database systems process the ARA-C01 query and send results until there are no more results to send, It's just like you are the king, then countless people support and work for you.

**2026 High Hit-Rate ARA-C01: SnowPro Advanced Architect Certification**

## Latest Mock Exam

The 98%-99% pass rate has helped many candidates passed the actual test and got the ARA-C01 Certification successfully, Why we let you try our ARA-C01 updated dumps free demo before you purchase?

Have you ever dreamed about passing the most important exam such as Snowflake ARA-C01 in your field with great ease, With the help of the ARA-C01 valid exam cram, you can not only enlarge your sight, but also master the professional skills.

What's more, part of that ExamBoosts ARA-C01 dumps now are free: <https://drive.google.com/open?id=1JizoFBWa9T3Q19Qoigffdia1vCunUzTr>