

# Useful ARA-C01 Download Demo by ValidTorrent



P.S. Free & New ARA-C01 dumps are available on Google Drive shared by ValidTorrent: [https://drive.google.com/open?id=1wnKPVckZAUlt\\_cB3GDQ0E-u0eNsI5USN](https://drive.google.com/open?id=1wnKPVckZAUlt_cB3GDQ0E-u0eNsI5USN)

The accuracy rate of ARA-C01 test training materials of ValidTorrent is high with wide coverage. It will be the most suitable ARA-C01 test training materials and the one you need most to pass ARA-C01 exam. We promise that we will provide renewal service freely as long as one year after you purchase our ARA-C01 Dumps; if you fail ARA-C01 test or there are any quality problem of our ARA-C01 exam dumps and training materials, we will give a full refund immediately.

The SnowPro Advanced Architect Certification is a valuable credential for data architects and engineers who want to demonstrate their expertise in designing and implementing complex Snowflake solutions. SnowPro Advanced Architect Certification certification provides a competitive edge to professionals by showcasing their skills and knowledge in the field of data management, warehousing, and analytics. SnowPro Advanced Architect Certification certification is recognized globally by companies that use Snowflake for their data management needs.

Snowflake ARA-C01 exam consists of multiple-choice questions that cover a wide range of topics related to data warehousing and cloud computing. ARA-C01 exam is timed, with a total duration of 120 minutes, and candidates must achieve a minimum score of 80% to pass. ARA-C01 Exam is available in multiple languages including English, Japanese, and Spanish.

Snowflake ARA-C01 is an advanced certification exam designed for those who want to demonstrate their expertise in Snowflake architecture, design, and implementation. SnowPro Advanced Architect Certification certification is targeted at experienced architects, engineers, and consultants who have a deep understanding of Snowflake's technology and can provide solutions to complex business problems. Passing the ARA-C01 exam demonstrates that an individual has the necessary skills and knowledge to

design and implement advanced Snowflake solutions.

>> ARA-C01 Download Demo <<

## Exam Dumps ARA-C01 Demo - Reliable ARA-C01 Exam Pdf

Are you tired of feeling overwhelmed and unsure about how to prepare for your SnowPro Advanced Architect Certification (ARA-C01) exam? Are you ready to take control of your future and achieve the scores you want to get in the SnowPro Advanced Architect Certification (ARA-C01) certification exam? If so, it's time to buy real Snowflake ARA-C01 Dumps of ValidTorrent our team of experts has designed the product that has already helped thousands of students just like you pass the exam.

### Snowflake SnowPro Advanced Architect Certification Sample Questions (Q166-Q171):

#### NEW QUESTION # 166

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

**Answer: D**

Explanation:

\* Snowpipe is a service that automates and optimizes the loading of data from external stages into Snowflake tables. Snowpipe uses a queue to ingest files as they become available in the stage. Snowpipe also provides REST endpoints to load data and retrieve load history reports<sup>1</sup>.

\* The loadHistoryScan endpoint returns the history of files that have been ingested by Snowpipe within a specified time range. The endpoint accepts the following parameters<sup>2</sup>:

\* pipe: The fully-qualified name of the pipe to query.

\* startTimeInclusive: The start of the time range to query, in ISO 8601 format. The value must be within the past 14 days.

\* endTimeExclusive: The end of the time range to query, in ISO 8601 format. The value must be later than the start time and within the past 14 days.

\* recentFirst: A boolean flag that indicates whether to return the most recent files first or last. The default value is false, which means the oldest files are returned first.

\* showSkippedFiles: A boolean flag that indicates whether to include files that were skipped by Snowpipe in the response. The default value is false, which means only files that were loaded are returned.

\* The loadHistoryScan endpoint can be used to keep a log of data load history by calling it periodically with a suitable time range.

The best option among the choices is D, which is to call loadHistoryScan every 10 minutes for a 15-minute time range. This option ensures that the endpoint is called frequently enough to capture the latest files that have been ingested, and that the time range is wide enough to avoid missing any files that may have been delayed or retried by Snowpipe. The other options are either too infrequent, too narrow, or use the wrong endpoint<sup>3</sup>.

References:

\* 1: Introduction to Snowpipe | Snowflake Documentation

\* 2: loadHistoryScan | Snowflake Documentation

\* 3: Monitoring Snowpipe Load History | Snowflake Documentation

#### NEW QUESTION # 167

An Architect uses COPY INTO with the ON\_ERROR=SKIP\_FILE option to bulk load CSV files into a table called TABLEA, using its table stage. One file named file5.csv fails to load. The Architect fixes the file and re-loads it to the stage with the exact same file name it had previously.

Which commands should the Architect use to load only file5.csv file from the stage? (Choose two.)

- A. COPY INTO tablea FROM @%tablea NEW\_FILES\_ONLY = TRUE;
- **B. COPY INTO tablea FROM @%tablea;**
- C. COPY INTO tablea FROM @%tablea RETURN\_FAILED\_ONLY = TRUE;
- **D. COPY INTO tablea FROM @%tablea FILES = ('file5.csv');**

- E. COPY INTO tablea FROM @%tablea MERGE = TRUE;
- F. COPY INTO tablea FROM @%tablea FORCE = TRUE;

**Answer: B,D**

Explanation:

\* Option A (RETURN\_FAILED\_ONLY) will only load files that previously failed to load. Since file5.csv already exists in the stage with the same name, it will not be considered a new file and will not be loaded.

\* Option D (FORCE) will overwrite any existing data in the table. This is not desired as we only want to load the data from file5.csv.

\* Option E (NEW\_FILES\_ONLY) will only load files that have been added to the stage since the last COPY command. This will not work because file5.csv was already in the stage before it was fixed.

\* Option F (MERGE) is used to merge data from a stage into an existing table, creating new rows for any data not already present. This is not needed in this case as we simply want to load the data from file5.csv.

Therefore, the architect can use either COPY INTO tablea FROM @%tablea or COPY INTO tablea FROM @%tablea FILES = ('file5.csv') to load only file5.csv from the stage. Both options will load the data from the specified file without overwriting any existing data or requiring additional configuration

### NEW QUESTION # 168

A Snowflake Architect is designing an application and tenancy strategy for an organization where strong legal isolation rules as well as multi-tenancy are requirements.

Which approach will meet these requirements if Role-Based Access Policies (RBAC) is a viable option for isolating tenants?

- A. Create an object for each tenant strategy if row level security is viable for isolating tenants.
- B. Create an object for each tenant strategy if row level security is not viable for isolating tenants.
- C. Create a multi-tenant table strategy if row level security is not viable for isolating tenants.
- **D. Create accounts for each tenant in the Snowflake organization.**

**Answer: D**

Explanation:

In a scenario where strong legal isolation is required alongside the need for multi-tenancy, the most effective approach is to create separate accounts for each tenant within the Snowflake organization. This approach ensures complete isolation of data, resources, and management, adhering to strict legal and compliance requirements. Role-Based Access Control (RBAC) further enhances security by allowing granular control over who can access what resources within each account. This solution leverages Snowflake's capabilities for managing multiple accounts under a single organization umbrella, ensuring that each tenant's data and operations are isolated from others.

Reference: Snowflake documentation on multi-tenancy and account management, part of the SnowPro Advanced: Architect learning path.

### NEW QUESTION # 169

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. The recent data imports contained fewer fields than usual.
- **C. The order of the keys in the JSON was changed.**
- D. There were JSON nulls in the recent data imports.

**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.

Reference:

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 # 170

A healthcare company is deploying a Snowflake account that may include Personal Health Information (PHI). The company must ensure compliance with all relevant privacy standards.

Which best practice recommendations will meet data protection and compliance requirements? (Choose three.)

- A. Use the External Tokenization feature to obfuscate sensitive data.
- B. Avoid sharing data with partner organizations.
- C. Create Dynamic Data Masking policies and apply them to columns that contain PHI.
- D. Use the Internal Tokenization feature to obfuscate sensitive data.
- E. Use, at minimum, the Business Critical edition of Snowflake.
- F. Rewrite SQL queries to eliminate projections of PHI data based on `current_role()`.

**Answer: A,C,E**

Explanation:

A healthcare company that handles PHI data must ensure compliance with relevant privacy standards, such as HIPAA, HITRUST, and GDPR. Snowflake provides several features and best practices to help customers meet their data protection and compliance requirements<sup>1</sup>.

One best practice recommendation is to use, at minimum, the Business Critical edition of Snowflake. This edition provides the highest level of data protection and security, including end-to-end encryption with customer-managed keys, enhanced object-level security, and HIPAA and HITRUST compliance<sup>2</sup>. Therefore, option A is correct.

Another best practice recommendation is to create Dynamic Data Masking policies and apply them to columns that contain PHI. Dynamic Data Masking is a feature that allows masking or redacting sensitive data based on the current user's role. This way, only authorized users can view the unmasked data, while others will see masked values, such as NULL, asterisks, or random characters<sup>3</sup>. Therefore, option B is correct.

A third best practice recommendation is to use the External Tokenization feature to obfuscate sensitive data. External Tokenization is a feature that allows replacing sensitive data with tokens that are generated and stored by an external service, such as Protegrity. This way, the original data is never stored or processed by Snowflake, and only authorized users can access the tokenized data through the external service<sup>4</sup>. Therefore, option D is correct.

Option C is incorrect, because the Internal Tokenization feature is not available in Snowflake. Snowflake does not provide any native tokenization functionality, but only supports integration with external tokenization services<sup>4</sup>.

Option E is incorrect, because rewriting SQL queries to eliminate projections of PHI data based on `current_role()` is not a best practice. This approach is error-prone, inefficient, and hard to maintain. A better alternative is to use Dynamic Data Masking policies, which can automatically mask data based on the user's role without modifying the queries<sup>3</sup>.

Option F is incorrect, because avoiding sharing data with partner organizations is not a best practice. Snowflake enables secure and governed data sharing with internal and external consumers, such as business units, customers, or partners. Data sharing does not involve copying or moving data, but only granting access privileges to the shared objects. Data sharing can also leverage Dynamic Data Masking and External Tokenization features to protect sensitive data<sup>5</sup>.

