ARA-C01 Ausbildungsressourcen, ARA-C01 Lerntipps

Übrigens, Sie können die vollständige Version der ZertSoft ARA-C01 Prüfungsfragen aus dem Cloud-Speicher herunterladen: https://drive.google.com/open?id=1VqVTG9-i4xcD8qWByfqf2GQsksTmGVUq

Alle wünschen sich Erfolg. Die im IT-Bereich arbeitende Leute wissen sicherlich die Wichtigkeit der Zertifizierung der Snowflake ARA-C01 für die Karriere. Immer mehr Leute nehmen an der Snowflake ARA-C01 Prüfung teil. Wie kann man beim immer schweren Wettbewerb noch siegen? Den richtigen Hilfspartner auszuwählen ist am wichtigsten. ZertSoft hat die Snowflake ARA-C01 Prüfung schon mehrere Jahre geforscht. Wir haben gute Kenntnisse in dieser Prüfung. Mit Hilfe der Snowflake ARA-C01 Prüfungssoftware von uns wird Ihr Sieg bei der Prüfung gesichert.

Die Snowpro Advanced Architect Certification Exam Exam ist ein umfassender Test, der verschiedene Themen abdeckt, einschließlich Datenmodellierung, Sicherheits-, Leistungsoptimierung und Datenintegration. Die Prüfung wurde entwickelt, um die Fähigkeit einer Person zu bewerten, erweiterte Schneeflockenlösungen zu entwerfen, zu implementieren und zu optimieren, die den komplexen Geschäftsanforderungen entsprechen.

Die Snowflake ARA-C01 Prüfung ist für erfahrene Architekten gedacht, die ein tiefes Verständnis für Snowflakes Cloud-Datenplattform und ihre Fähigkeiten haben. Die Prüfung testet die Fähigkeit des Kandidaten, komplexe Data-Warehousing-Lösungen zu entwerfen und umzusetzen, einschließlich Datenmodellierung, Datenintegration, Daten transformation und Datenanalyse.

>> ARA-C01 Ausbildungsressourcen <<

ARA-C01 Fragen & Antworten & ARA-C01 Studienführer & ARA-C01

Prüfungsvorbereitung

Wollen Sie Ihre IT-Fähigkeiten beweisen? Möchten Sie mehr Anerkennung und Berufschancen bekommen? Die Prüfungszertifizierung der Snowflake ARA-C01 ist ein bedeutendester Ausweis für Sie. Die Wichtigkeit der Zertifizierung der Snowflake ARA-C01 wissen fast alle Angestellte aus IT-Branche. Die Tatkraft von Menschen ist limitiert. Wenn Sie in einer kurzen Zeit diese wichtige Snowflake ARA-C01 Prüfung bestehen möchten, brauchen Sie unsere die Prüfungssoftware von uns ZertSoft als Ihr bester Helfer für die Prüfungsvorbereitung. Umfassende Prüfungsaufgaben enthaltende und Mnemotechnik entsprechende Software kann Ihnen beim Erfolg der Snowflake ARA-C01 gut helfen!

Snowflake SnowPro Advanced Architect Certification ARA-C01 Prüfungsfragen mit Lösungen (Q133-Q138):

133. Frage
A DevOps team has a requirement for recovery of staging tables used in a complex set of data pipelines. The staging tables are all located in the same staging schema. One of the requirements is to have online recovery of data on a rolling 7-day basis.
After setting up the DATA_RETENTION_TIME_IN_DAYS at the database level, certain tables remain unrecoverable past 1 day.
What would cause this to occur? (Choose two.)

- A. The tables exceed the 1 TB limit for data recovery.
- B. The staging schema has not been setup for MANAGED ACCESS.
- C. The staging tables are of the TRANSIENT type.
- D. The DATA_RETENTION_TIME_IN_DAYS for the staging schema has been set to 1 day.
- E. The DevOps role should be granted ALLOW_RECOVERY privilege on the staging schema.

Antwort: C,D

Begründung:
* The DATA_RETENTION_TIME_IN_DAYS parameter controls the Time Travel retention period for an object (database, schema, or table) in Snowflake. This parameter specifies the number of days for which historical data is preserved and can be accessed using Time Travel operations (SELECT, CREATE ... CLONE, UNDROP)1.
* The requirement for recovery of staging tables on a rolling 7-day basis means that the DATA_RETENTION_TIME_IN_DAYS parameter should be set to 7 at the database level. However, this parameter can be overridden at the lower levels (schema or table) if they have a different value1.
* Therefore, one possible cause for certain tables to remain unrecoverable past 1 day is that the DATA_RETENTION_TIME_IN_DAYS for the staging schema has been set to 1 day. This would override the database level setting and limit the Time Travel retention period for all the tables in the schema to 1 day. To fix this, the parameter should be unset or set to 7 at the schema level1. Therefore, option B is correct.
* Another possible cause for certain tables to remain unrecoverable past 1 day is that the staging tables are of the TRANSIENT type. Transient tables are tables that do not have a Fail-safe period and can have a Time Travel retention period of either 0 or 1 day. Transient tables are suitable for temporary or intermediate data that can be easily reproduced or replicated2. Tofix this, the tables should be created as permanent tables, which can have a Time Travel retention period of up to 90 days1. Therefore, option D is correct.
* Option A is incorrect because the MANAGED ACCESS feature is not related to the data recovery requirement. MANAGED ACCESS is a feature that allows granting access privileges to objects without explicitly granting the privileges to roles. It does not affect the Time Travel retention period or the data availability3.
* Option C is incorrect because there is no 1 TB limit for data recovery in Snowflake. The data storage size does not affect the Time Travel retention period or the data availability4.
* Option E is incorrect because there is no ALLOW_RECOVERY privilege in Snowflake. The privilege required to perform Time Travel operations is SELECT, which allows querying historical data in tables5.
Understanding & Using Time Travel : Transient Tables : Managed Access : Understanding Storage Cost : Table Privileges

134. Frage
A new table and streams are created with the following commands:
CREATE OR REPLACE TABLE LETTERS (ID INT, LETTER STRING) ;
CREATE OR REPLACE STREAM STREAM_1 ON TABLE LETTERS;
CREATE OR REPLACE STREAM STREAM_2 ON TABLE LETTERS APPEND_ONLY = TRUE;
The following operations are processed on the newly created table:
INSERT INTO LETTERS VALUES (1, 'A');
INSERT INTO LETTERS VALUES (2, 'B');

INSERT INTO LETTERS VALUES (3, 'C');
TRUNCATE TABLE LETTERS;
INSERT INTO LETTERS VALUES (4, 'D');
INSERT INTO LETTERS VALUES (5, 'E');
INSERT INTO LETTERS VALUES (6, 'F');
DELETE FROM LETTERS WHERE ID = 6;
What would be the output of the following SQL commands, in order?
SELECT COUNT (*) FROM STREAM_1;
SELECT COUNT (*) FROM STREAM_2;

- A. 2 & 3
- B. 4 & 3
- C. 4 & 6
- D. 2 & 6

Antwort: B

Begründung:
In Snowflake, a stream records data manipulation language (DML) changes to its base table since the stream was created or last consumed. STREAM_1 will show all changes including the TRUNCATE operation, while STREAM_2, being APPEND_ONLY, will not show deletions like TRUNCATE. Therefore, STREAM_1 will count the three inserts, the TRUNCATE (counted as a single operation), and the subsequent two inserts before the delete, totaling 4. STREAM_2 will only count the three initial inserts and the two after the TRUNCATE, totaling 3, as it does not count the TRUNCATE or the delete operation.

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

- A. Create private VPC endpoints that allow direct, secure connectivity between VPCs without traversing the public internet.
- 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. Manage credentials from multiple cloud providers in one single Snowflake object.
- E. Control access to Snowflake data using a master encryption key that is maintained in the cloud provider's key management service.
- F. Support multiple external stages using one single Snowflake object.

Antwort: B,C,F

Begründung:
* 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.
References: : Encryption and Decryption : Private Link for Snowflake : CREATE STORAGE INTEGRATION : Option 1: Configuring a Snowflake Storage Integration to Access Amazon S3

136. Frage
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 value return_all_errors of the option VALIDATION_MODE is causing a compilation error.
- B. The VALIDATION_MODE parameter does not support COPY statements that transform data during a load.
- C. The VALIDATION_MODE parameter does not support COPY statements with CSV file formats.
- D. The VALIDATION_MODE parameter supports COPY statements that load data from external stages only.

Antwort: B

Begründung:
* 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.
References: : COPY INTO <table> : Transforming Data During a Load\

137. Frage
An Architect on a new project has been asked to design an architecture that meets Snowflake security, compliance, and governance requirements as follows:
1) Use Tri-Secret Secure in Snowflake
2) Share some information stored in a view with another Snowflake customer
3) Hide portions of sensitive information from some columns
4) Use zero-copy cloning to refresh the non-production environment from the production environment To meet these requirements, which design elements must be implemented? (Choose three.)

- A. Define row access policies.
- B. Use Dynamic Data Masking.
- C. Create a materialized view.

- D. Use the Enterprise edition of Snowflake.
- E. Create a secure view.
- F. Use the Business-Critical edition of Snowflake.

Antwort: B,E,F

Begründung:
These three design elements are required to meet the security, compliance, and governance requirements for the project.
To use Tri-Secret Secure in Snowflake, the Business Critical edition of Snowflake is required. This edition provides enhanced data protection features, such as customer-managed encryption keys, that are not available in lower editions. Tri-Secret Secure is a feature that combines a Snowflake-maintained key and a customer- managed key to create a composite master key to encrypt the data in Snowflake1.
To share some information stored in a view with another Snowflake customer, a secure view is recommended.
A secure view is a view that hides the underlying data and the view definition from unauthorized users. Only the owner of the view and the users who are granted the owner's role can see the view definition and the data in the base tables of the view2. A secure view can be shared with another Snowflake account using a data share3.
To hide portions of sensitive information from some columns, Dynamic Data Masking can be used. Dynamic Data Masking is a feature that allows applying masking policies to columns to selectively mask plain-text data at query time. Depending on the masking policy conditions and the user's role, the data can be fully or partially masked, or shown as plain-text4.

138. Frage
.....

Laden Sie die neuesten ZertSoft ARA-C01 PDF-Versionen von Prüfungsfragen kostenlos von Google Drive herunter: https://drive.google.com/open?id=1VqVTG9-i4xcD8qWByfqf2GQsksTmGVUq