

# SPS-C01 German, SPS-C01 Lernhilfe



BONUS!!! Laden Sie die vollständige Version der ZertSoft SPS-C01 Prüfungsfragen kostenlos herunter:  
[https://drive.google.com/open?id=1ZQb1-2AWnRBJB\\_hjX3EkI\\_97u8VAnJWN](https://drive.google.com/open?id=1ZQb1-2AWnRBJB_hjX3EkI_97u8VAnJWN)

Unsere Snowflake SPS-C01 Prüfungsunterlagen (Snowflake Certified SnowPro Specialty - Snowpark) enthalten alle echten, originalen und richtigen Fragen und Antworten. Die Abdeckungsrate unserer Snowflake SPS-C01 Unterlagen (Fragen und Antworten) (Snowflake Certified SnowPro Specialty - Snowpark) ist normalerweise mehr als 98%.

ZertSoft ist eine erstklassige Website für die Snowflake SPS-C01 Zertifizierungsprüfung. Im ZertSoft können Sie Tipps und Prüfungsmaterialien finden. Sie können auch die Examensfragen- und antworten teilweise als Probe kostenlos herunterladen. ZertSoft kann Ihnen umsonst die Updaets der Prüfungsmaterialien für die Snowflake SPS-C01 Prüfung bieten. Alle unseren Zertifizierungsprüfungen enthalten Antworten. Unser Eliteteam von IT-Fachleuten wird die neuesten und richtigen Examenübungen nach ihren fachlichen Erfahrungen bearbeiten, um Ihnen bei der Prüfung zu helfen. Alles in allem, wir werden Ihnen alle einschlägigen Materialien in Bezug auf die Snowflake SPS-C01 Zertifizierungsprüfung bieten.

>> SPS-C01 German <<

## SPS-C01 Lernhilfe, SPS-C01 Deutsche

Sind Sie mit Ihrer Arbeit zufrieden? Sind Sie damit Zufrieden, was Sie jetzt machen? Wollen Sie Ihre Arbeitsfähigkeit erhöhen? Dann müssen Sie zuerst mehr nützliche Fähigkeiten für Ihre Arbeit beherrschen. Und das wichtigste ist, dass Arbeitsgeber wissen, Sie mehr Arbeitsfähigkeiten beherrschen. Dann legen Sie Snowflake SPS-C01 Prüfung ab. SPS-C01 Prüfung kann Ihren Wunsch erreichen. Und es macht nichts, wenn Sie die Prüfungsfragen nicht genug kennen, weil Sie die Hilfe und die Vorbereitungswerkzeuge an ZertSoft finden können. Die Prüfungsfragen und-antworten können Ihnen helfen, Snowflake SPS-C01 Zertifikat zu bekommen.

## Snowflake Certified SnowPro Specialty - Snowpark SPS-C01 Prüfungsfragen mit Lösungen (Q140-Q145):

### 140. Frage

You have a Snowpark DataFrame containing customer transaction data'. Your goal is to save this DataFrame as a set of Parquet files in an existing Snowflake stage named , partitioned by the 'transaction\_date' column. You want to ensure that the files are automatically compressed using the Zstandard codec and that existing files with the same name are overwritten. Which of the

following Snowpark code snippet achieves this with the most optimal approach and respects best practices?

- A. Option B
- **B. Option A**
- C. Option D
- D. Option C
- E. Option E

**Antwort: B**

Begründung:

Option A correctly uses the 'parquet' method directly for writing Parquet files to a stage. It specifies partitioning by 'transaction\_date', overwrites existing files using , and sets the compression codec to 'zstd' using the 'option' method. The 'saveAsTable' method, used in option B & E, is intended for creating or overwriting tables, not writing files to a stage. Option D uses a fully qualified Snowflake URL to save the DataFrame, but using saveAsTable is not for writing files into stage . The 'option('fileFormat', 'parquet')' in option C is not the most direct way to specify the format; using .parquet() is more concise and idiomatic.

#### 141. Frage

You have two Snowpark DataFrames, 'df1' and 'df2', representing customer data'. 'df1' contains columns 'CUSTOMER ID', 'NAME, and 'EMAIL', while 'df2' contains 'CUSTOMER ID' and 'PURCHASE AMOUNT'. You need to create a new DataFrame that combines the information from both DataFrames but only includes customers who exist in BOTH 'df1' and 'df2' and the resulting DataFrame should have columns from both. Which of the following Snowpark DataFrame operations should you use, and what is the correct way to call it?

- A.
- B.
- **C.**
- **D.**
- E.

**Antwort: C,D**

Begründung:

To include only customers present in BOTH DataFrames and include columns from both, you need to perform an INNER JOIN. Options B, C and D are incorrect: intersect, union and subtract operations work at a row level. Also, intersect , union and subtract operations expects the number of columns and datatypes to match and is not relevant to the described scenario. Option A and E are valid way to use the join operation: (A) uses the explicit condition 'df1.CUSTOMER\_ID df2.CUSTOMER\_ID while (E) is the short form which specifies the column name directly. Both achieves the same inner join behavior. You can choose E as a cleaner option when only joining on column name.

#### 142. Frage

You have a Snowflake table named 'raw events' with a VARIANT column named 'event data'. The 'event data' column contains JSON objects with a field 'timestamp' that is sometimes represented as a string and sometimes as a number (Unix epoch). You need to create a Snowpark DataFrame that extracts the 'timestamp' as a timestamp object, handling both string and numeric representations. Which of the following code snippets correctly accomplishes this, avoiding errors when encountering incompatible types?

- A.
- B.
- C.
- **D.**
- E.

**Antwort: D**

Begründung:

Option D correctly uses the 'is\_number' function to check if the timestamp is numeric. If it is, it divides by 1000 (assuming milliseconds) and converts to a timestamp. If it's not numeric, it converts directly to a timestamp (assuming it is a string

representation). Options A, B, C, and E will fail when encountering mixed data types because 'to\_timestamp' expects either a number or a string, not both interchangeably. Casting numeric value as String then passing to 'to\_timestamp' would raise issues, it needs to be divided.

### 143. Frage

You are building a Snowpark application that requires you to connect to Snowflake from an environment where directly specifying credentials in the code is not permitted for security reasons. Which of the following are valid and recommended ways to securely pass authentication information to the Snowpark Session?

- A. Hardcoding the credentials in the Snowpark Python script and obfuscating them using Base64 encoding. This provides security by obscurity, making it a reasonably secure approach.
- B. Storing credentials in a Snowflake stage and retrieving them from there at runtime. This is an acceptable, though more complex, solution.
- C. Using environment variables and retrieving them using 'os.environ' to build the connection parameters. This is a secure and recommended approach.
- D. Using the Snowflake CLI's 'snowflake configure' command and relying on the '.snowflake/config' file. This is suitable for development but not recommended for production due to local file dependency.
- E. Storing credentials in a dedicated secret management service (e.g., HashiCorp Vault, AWS Secrets Manager) and retrieving them using an appropriate API. This is the most secure and recommended approach for production environments.

**Antwort: C,D,E**

Begründung:

Options A, C, and D represent valid ways to handle credentials securely. Environment variables (A) are a standard practice for configurations. Using a secret management service (C) provides the best security posture for production environments. Using the Snowflake CLI (D) is acceptable for development. Storing credentials in a Snowflake stage (B) adds unnecessary complexity and doesn't inherently improve security over other options. Base64 encoding (E) is not a secure method; it's easily decoded and provides a false sense of security. Hardcoding and obfuscating credentials is not recommended.

### 144. Frage

You have a Snowpark Python application that performs several data transformations on a DataFrame representing customer transactions. The application is experiencing performance issues, and you suspect that some transformations are unnecessarily expensive. Which of the following techniques can MOST effectively optimize the performance of your Snowpark application, specifically focusing on minimizing data movement and leveraging Snowflake's query optimization capabilities?

- A. Explicitly call .cache() on the DataFrame after each transformation to materialize intermediate results in memory.
- B. Use User-Defined Functions (UDFs) written in Python for all transformations, regardless of their complexity.
- C. Always use the largest available Snowflake warehouse size to ensure sufficient compute resources.
- D. Leverage Snowpark's built-in DataFrame transformations (e.g., .groupBy()) to allow Snowflake to optimize the query execution plan. Avoid pulling large amounts of data into the client application for simple operations. Only call 'collect()' as the last and final option, as this is the most costly activity of all.
- E. Take the dataframe to Pandas dataframe as soon as possible in between transformations, since Pandas dataframes will be faster.

**Antwort: D**

Begründung:

Snowpark is designed to push down computations to Snowflake, allowing Snowflake's query optimizer to handle the execution. Using Snowpark's built-in DataFrame transformations allows Snowflake to understand the intent and optimize the query accordingly. Materializing intermediate results using .cache() (A) can lead to unnecessary data movement. Python UDFs (B) can be useful for complex logic but should be avoided for simple transformations as they bypass Snowflake's optimization capabilities and are generally slower than native SQL functions. Warehouse size (E) is a factor, but optimizing the query logic is more crucial. Using Pandas dataframe is also costly and performance heavy.

### 145. Frage

.....



