

Snowflake DEA-C02 Online Prüfung, DEA-C02 Examsfragen

Snowflake DSA-C02 Certification Details	
Exam Code	DSA-C02
Full Exam Name	Snowflake SnowPro Advanced - Data Scientist
No. of Questions	65
Online Practice Exam	Snowflake Certified SnowPro Advanced - Data Scientist Practice Test
Sample Questions	Snowflake DSA-C02 Sample Questions
Passing Score	750 + Scaled Scoring from 0 - 1000
Time Limit	115 minutes
Exam Fees	\$375 USD
Become successful with VMEExam.com	

P.S. Kostenlose 2026 Snowflake DEA-C02 Prüfungsfragen sind auf Google Drive freigegeben von It-Pruefung verfügbar:
https://drive.google.com/open?id=1AMAIRa4P_iJBVY6afa93VFuczkNrUEd

Jede Version der Snowflake DEA-C02 Prüfungsunterlagen von uns hat ihre eigene Überlegenheit. PDF Version hat keine Beschränkung für Anlage, deshalb können Sie irgendwo die Unterlagen lesen. Wenn Sie Internet benutzen können, die Online Test Engine der Snowflake DEA-C02 können Sie sowohl mit Windows, Mac als auch Android, iOS benutzen. Mit Simulations-Software können Sie die Prüfungsumwelt der Snowflake DEA-C02 erfahren und bessere Kenntnisse darüber erwerben. Übrigens, Sie dürfen die Prüfungssoftware irgendwie viele Male installieren.

Die Snowflake DEA-C02 Zertifizierung ist den IT-Fachleute eine unentbehrliche Prüfung, weil sie ihres Schicksal bestimmt. Die Fragenkataloge zur Snowflake DEA-C02 Prüfung brauchen alle Kandidaten. Mit ihr kann der Kandidat sich gut auf die DEA-C02 Prüfung vorbereiten und nicht so sehr unter Druck stehen. Und die Fragenkataloge in It-Pruefung sind einzigartig. Mit ihr können Sie die Snowflake DEA-C02 Prüfung ganz mühlos bestehen.

>> Snowflake DEA-C02 Online Prüfung <<

DEA-C02 Examsfragen - DEA-C02 Lerntipps

Das erfahrungsreiche Expertenteam von It-Pruefung hat den effizienten Prüfungsfragen und Antworten zur Snowflake DEA-C02 Zertifizierungsprüfung entwickelt, die geeignet für die Kandidaten ist. Die Produkte von It-Pruefung sind von guter Qualität. Sie können sie als Simulationsprüfung vor der Snowflake DEA-C02 Zertifizierungsprüfung benutzen und sich gut auf die Prüfung vorbereiten.

Snowflake SnowPro Advanced: Data Engineer (DEA-C02) DEA-C02 Prüfungsfragen mit Lösungen (Q285-Q290):

285. Frage

You have a Snowpark DataFrame 'df_products' with columns 'product id', 'category', and 'price'. You need to perform the following transformations in a single, optimized query using Snowpark Python: 1. Filter for products in the 'Electronics' or 'Clothing' categories. 2. Group the filtered data by category. 3. Calculate the average price for each category. 4. Rename the aggregated column to 'average_price'. Which of the following code snippets demonstrates the most efficient way to achieve this?

```

○ df_products.filter((df_products['category'] == 'Electronics') | (df_products['category'] == 'Clothing')).groupBy('category').agg(avg(df_products['price']).alias('average_price')).show()

○ from snowflake.snowpark.functions import col, avg df_products.filter(col('category').isin(['Electronics', 'Clothing'])).groupBy(col('category')).agg(avg(col('price')).as_('average_price')).show()

○ df_products.where(df_products.category.isin(['Electronics', 'Clothing'])).groupBy(df_products.category).agg(mean(df_products.price).name('average_price')).show()

○ from snowflake.snowpark.functions import col, avg df_products.filter(col('category').isin(['Electronics', 'Clothing'])).groupBy('category').agg(avg('price').alias('average_price')).show()

○ from snowflake.snowpark.functions import col, avg df_products.where(col('category').isin(['Electronics', 'Clothing'])).groupBy('category').agg(avg('price').alias('average_price')).to_pandas()

```

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

Antwort: B

Begründung:

Option B is the most efficient and correct. It uses 'col()' from 'snowflake.snowpark.functions' to properly reference the 'category' and 'price' columns, uses 'isin()' for a more concise and efficient filtering of multiple category values, groups by the category using 'groupBy()' and calculates the average price with 'avg(col('price')).as_('average_price')'. Option A, C, and D are syntactically incorrect or less efficient ways to accomplish the same task within Snowpark. Option E is incorrect because it utilizes 'to_pandas()' which returns the result as a Pandas DataFrame rather than a Snowpark DataFrame, failing to adhere to the Snowpark environment. While Option D is very similar, it lacks the proper syntax for specifying column references with 'col('category')' in the groupBy and 'col('price')' in the avg function.

286. Frage

You have a table named 'sales_data' with columns 'region', 'product_category', and 'revenue'. You want to create an aggregation policy to prevent users without the 'FINANCE ADMIN' role from seeing revenue values aggregated across all regions. Instead, these users should only see revenue aggregated at the region level. The policy should return NULL for the 'revenue' column when aggregated across all regions by non-admin users. Which of the following SQL snippets correctly implements this aggregation policy?

```

'''sql CREATE OR REPLACE AGGREGATION POLICY revenue_policy RETURNS NUMBER AS ( revenue NUMBER, region VARCHAR, product_category VARCHAR ) -> CASE WHEN CURRENT_ROLE() = 'FINANCE_ADMIN' THEN revenue WHEN GROUPING(region) = 0 THEN revenue ELSE NULL END; ALTER TABLE sales_data ADD AGGREGATION POLICY revenue_policy ON revenue; '''

'''sql CREATE OR REPLACE AGGREGATION POLICY revenue_policy RETURNS NUMBER AS ( revenue NUMBER, region VARCHAR, product_category VARCHAR ) -> CASE WHEN CURRENT_ROLE() = 'FINANCE_ADMIN' THEN revenue WHEN GROUPING(region) = 1 THEN revenue ELSE NULL END; ALTER TABLE sales_data ADD AGGREGATION POLICY revenue_policy ON revenue; '''

'''sql CREATE OR REPLACE AGGREGATION POLICY revenue_policy RETURNS NUMBER AS ( revenue NUMBER ) -> CASE WHEN CURRENT_ROLE() = 'FINANCE_ADMIN' THEN revenue ELSE NULL END; ALTER TABLE sales_data ADD AGGREGATION POLICY revenue_policy ON revenue; '''

'''sql CREATE OR REPLACE AGGREGATION POLICY revenue_policy RETURNS NUMBER AS ( revenue NUMBER, region VARCHAR ) -> CASE WHEN CURRENT_ROLE() = 'FINANCE_ADMIN' THEN revenue WHEN GROUPING(region) = 1 THEN revenue ELSE NULL END; ALTER TABLE sales_data ADD AGGREGATION POLICY revenue_policy ON revenue; '''

```

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

Antwort: B

Begründung:

Option B correctly uses the GROUPING() function to identify when the aggregation is being performed across all regions (GROUPING(region) = 1). It then returns NULL for non-FINANCE_ADMIN users in this scenario. Option A is incorrect because grouping(region) = 0 means region-level aggregation. Option C does not consider regions in the policy, so non-admin users will always see null revenue. Option D's syntax is incorrect, must be combined with GROUPING as in E to work correctly.

Option E can also work but the CASE statement is clearer and easier to understand.

287. Frage

You are tasked with sharing a subset of a customer table (CUSTOMER_DATA) residing in your organization's Snowflake account with a partner organization. You need to mask personally identifiable information (PII) while providing near real-time updates. You decide to use a secure view. Which of the following SQL statements is the MOST efficient and secure way to accomplish this, assuming the partner only needs 'customer id', 'masked_email', 'city', and 'state'? The email should be masked using SHA256.

```
☐ ""sql CREATE OR REPLACE SECURE VIEW PARTNER_CUSTOMER_VIEW AS SELECT customer_id, SHA2(email, 256) as masked_email, city, state FROM CUSTOMER_DATA; ""
```

```
☐ ""sql CREATE VIEW PARTNER_CUSTOMER_VIEW AS SELECT customer_id, SHA2(email, 256) as masked_email, city, state FROM CUSTOMER_DATA; GRANT SELECT ON VIEW PARTNER_CUSTOMER_VIEW TO SHARE; ""
```

```
☐ ""sql CREATE OR REPLACE SECURE VIEW PARTNER_CUSTOMER_VIEW AS SELECT customer_id, SYSTEM$MASK_EMAIL(email) as masked_email, city, state FROM CUSTOMER_DATA; ""
```

```
☐ ""sql CREATE OR REPLACE VIEW PARTNER_CUSTOMER_VIEW AS SELECT customer_id, SHA2(email, 256) as masked_email, city, state FROM CUSTOMER_DATA; GRANT SELECT ON VIEW PARTNER_CUSTOMER_VIEW TO SHARE; ""
```

```
☐ ""sql CREATE OR REPLACE SECURE VIEW PARTNER_CUSTOMER_VIEW AS SELECT customer_id, SHA2(email, 256) as masked_email, city, state FROM CUSTOMER_DATA; GRANT SELECT ON VIEW PARTNER_CUSTOMER_VIEW TO SHARE; ""
```

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

Antwort: C

Begründung:

Option A correctly creates a SECURE VIEW, which is essential for data sharing as it prevents the partner from seeing the underlying table definition. It directly masks the email using SHA256. While Option C uses a hypothetical 'SYSTEM\$MASK_EMAIL' function, SHA256 is a readily available and standard masking method. Options B, D, and E do not create a secure view or properly grant access for data sharing. The correct method to share is to create secure view and then share using the share object.

288. Frage

A data engineer is tasked with optimizing a Snowflake data pipeline that ingests data from multiple external sources, transforms it, and loads it into a reporting table. The pipeline uses a series of Snowflake tasks orchestrated with a root task and child tasks. Performance monitoring shows inconsistent execution times for the transformation tasks. Which of the following strategies would provide the MOST granular insights into the performance bottlenecks within the pipeline and allow for targeted optimization?

- A. Enable query profiling for all queries executed within the transformation tasks using 'ALTER SESSION SET QUERY PROFILE = 'ON'', then analyze the query profiles for performance bottlenecks after each task run.
- B. Use Snowflake's Resource Monitors to track overall warehouse consumption and assume that high consumption during transformation tasks indicates a bottleneck within those tasks.
- C. Implement a custom logging mechanism within the transformation tasks to record execution times for each stage of the transformation process, and store these logs in a Snowflake table for analysis.
- **D. Leverage Snowflake's event tables like QUERY HISTORY and TASK HISTORY in the ACCOUNT USAGE schema joined with custom metadata tags to correlate specific transformation steps to execution times and resource usage. Also set up alerting based on defined performance thresholds.**
- E. Rely solely on the Snowflake web UI's Task History view to identify slow-running tasks.

Antwort: D

Begründung:

Option E is the most effective because it combines the structured data from Snowflake's account usage views (QUERY HISTORY, TASK_HISTORY) with custom metadata tags that provide context about each task's role in the data pipeline. This allows for granular analysis and correlation of transformation steps to performance metrics. Setting up alerts further proactively notifies the data engineer of performance deviations. Option B is good but harder to scale and maintain. Option D is less scalable and can be harder to interpret. The other options lack sufficient detail or provide only a high-level overview.

289. Frage

You are developing a data pipeline in Snowflake that uses SQL UDFs for data transformation. You need to define a UDF that calculates the Haversine distance between two geographical points (latitude and longitude). Performance is critical. Which of the following approaches would result in the most efficient UDF implementation, considering Snowflake's execution model?

- A. Create an External Function (using AWS Lambda or Azure Functions) to calculate the Haversine distance. This allows for offloading the computation to a separate compute environment.
- **B. Create a SQL UDF that directly calculates the Haversine distance using Snowflake's built-in mathematical functions (SIN, COS, ACOS, RADIANS). This is straightforward and easy to implement.**
- C. Create a Java UDF that calculates the Haversine distance, leveraging optimized mathematical libraries. This allows for potentially faster execution due to lower-level optimizations.
- D. Create a SQL UDF leveraging Snowflake's VECTORIZED keyword, hoping to automatically leverage SIMD instructions, without any code changes to mathematical calculation inside the UDF
- E. Create a SQL UDF that pre-calculates the RADIANS for latitude and longitude only once and stores them in a temporary table, using those values for subsequent distance calculations within the same session.

Antwort: B

Begründung:

SQL UDFs are generally the most efficient for simple calculations within Snowflake because they are executed within the Snowflake engine, minimizing data movement and overhead. While Java UDFs (option B) can offer optimizations, the overhead of invoking the Java environment often outweighs the benefits for this type of calculation. External Functions (option C) introduce significant latency due to network communication. Option D provides temporary performance improvements for the specific session, but is not the most efficient general solution. Vectorized keyword doesn't exist in snowflake to create UDFs, Hence it won't allow compilation. This question's emphasis is on understanding the trade-offs between different UDF types and their performance implications within the Snowflake architecture.

290. Frage

.....

Die von It-Prüfung gebotenen Prüfungsfragen enthalten wertvolle Prüfungserfahrungen und relevante Prüfungsmaterialien von IT-Experten und auch die Prüfungsfragen und Antworten für Snowflake DEA-C02 Zertifizierungsprüfung. Mit unserem guten Ruf in der IT-Branche geben wir Ihnen 100% Garantie. Sie können versuchsweise die Examensübungen und Antworten für die Snowflake DEA-C02 Zertifizierungsprüfung teilweise als Probe umsonst herunterladen. Dann können Sie ganz beruhigt unsere Schulungsunterlagen kaufen.

DEA-C02 Examsfragen: <https://www.it-pruefung.com/DEA-C02.html>

Snowflake DEA-C02 Online Prüfung Wir sind seit vielen Jahren der beste und gültige Anbieter für die Prüfung, Snowflake DEA-C02 Online Prüfung Sie können die Demo auf unserer Website direkt kostenlos downloaden, Wenn die neueste Version von DEA-C02 Examsfragen - SnowPro Advanced: Data Engineer (DEA-C02) erscheint, geben wir Ihnen sofort per E-Mail Bescheid, Snowflake DEA-C02 Online Prüfung Das Zertifikat so schnell wie möglich erwerben.

So kommt ihr ja sein Gesicht nicht sehen, Daß sie dich lieben, Wir sind DEA-C02 seit vielen Jahren der beste und gültige Anbieter für die Prüfung, Sie können die Demo auf unserer Website direkt kostenlos downloaden.

DEA-C02 Pass Dumps & PassGuide DEA-C02 Prüfung & DEA-C02 Guide

Wenn die neueste Version von SnowPro Advanced: Data Engineer (DEA-C02) erscheint, DEA-C02 Testing Engine geben wir Ihnen sofort per E-Mail Bescheid, Das Zertifikat so schnell wie möglich erwerben, Wir sind sehr stolz auf unsere DEA-C02 Prüfungsguide: SnowPro Advanced: Data Engineer (DEA-C02) und hoffen, dass wir damit Ihnen am besten helfen.

- DEA-C02 Fragen Antworten ☐ DEA-C02 Übungsmaterialien ☐ DEA-C02 Vorbereitung ☐ Sie müssen nur zu ► www.examfragen.de ◀ gehen um nach kostenloser Download von ► DEA-C02 ◀ zu suchen ☐ DEA-C02 Prüfungen
- DEA-C02 neuester Studienführer - DEA-C02 Training Torrent prep ☐ Öffnen Sie (www.itzert.com) geben Sie ☐ DEA-C02 ☐ ein und erhalten Sie den kostenlosen Download ☐ DEA-C02 Zertifizierungsfragen
- DEA-C02 neuester Studienführer - DEA-C02 Training Torrent prep ☐ Öffnen Sie die Website ☐ www.zertpruefung.ch ☐ Suchen Sie ☐ DEA-C02 ☐ Kostenloser Download ☐ DEA-C02 Tests
- DEA-C02 Übungsmaterialien - DEA-C02 Lernressourcen - DEA-C02 Prüfungsfragen ☐ Suchen Sie jetzt auf 《 www.itzert.com 》 nach “ DEA-C02 ” und laden Sie es kostenlos herunter ☐ DEA-C02 Prüfungen

- [illegible]

Übrigens, Sie können die vollständige Version der It-Prüfung DEA-C02 Prüfungsfragen aus dem Cloud-Speicher herunterladen:
https://drive.google.com/open?id=1AMAIRaP_iJBVY6afa93VFuclzkNrUeD