

DEA-C02 German & DEA-C02 Prüfungs-Guide



SnowPro Advanced: Data Engineer DEA-C02 Exam Questions

Außerdem sind jetzt einige Teile dieser Zertpruefung DEA-C02 Prüfungsfragen kostenlos erhältlich: <https://drive.google.com/open?id=1rYKLykADrx2ipLf3Jyg3OEF2Kn9L8MVS>

Die Snowflake DEA-C02 Zertifizierungsprüfung ist sehr schwierig. Obwohl die Snowflake DEA-C02 Prüfung sehr schwer ist, gibt es noch viele Prüfungsanmelder. Warum? Weil Snowflake DEA-C02 Prüfung eine sehr wichtige Prüfung. Für ITer gibt es negativen Einfluss für die Arbeit. Diese Zertifizierung können Ihnen viele Vorteile bringen. Und es kann auch Ihnen erhöhte Chance schaffen. Zusammenfassend nehmen Snowflake DEA-C02 Prüfung großen Einfluss auf Ihre Karriere. Wollen Sie diese so wichtige Prüfung ablegen?

Snowflake DEA-C02 Dumps von Zertpruefung sind ganz gleich wie die richtigen Zertifizierungsprüfungen. Die beinhalten alle Prüfungsfragen und Testantworten in aktueller Prüfung. Und die Software-Version simuliert die gleiche Atmosphäre der aktuellen Prüfungen. Bei der Nutzung der Zertpruefung Dumps, können Sie ganz sorglos die Snowflake DEA-C02 Prüfung ablegen und sehr gute Note bekommen.

>> DEA-C02 German <<

DEA-C02 Prüfungs-Guide & DEA-C02 Zertifikatsdemo

Wenn Sie die richtige Methode benutzen, haben Sie schon halben Erfolg erhalten. Wir Zertpruefung bieten Ihnen die effizienteste Methode für Snowflake DEA-C02 Prüfung, die von unseren erfahrenen Forschungs-und Entwicklungsstellen hergestellt wird. Auf unserer offiziellen Webseite können Sie durch Paypal die Snowflake DEA-C02 Prüfungsunterlagen gesichert kaufen. Wir werden Ihre Persönliche Informationen und Zahlungsinformationen gut bewahren und bieten Ihnen nach dem Kauf der Snowflake DEA-C02 Unterlagen immer weiter hochwertigen Dienst.

Snowflake SnowPro Advanced: Data Engineer (DEA-C02) DEA-C02 Prüfungsfragen mit Lösungen (Q235-Q240):

235. Frage

You are designing a data governance strategy for a Snowflake data warehouse. One of the key requirements is to track data lineage for sensitive data, specifically Personally Identifiable Information (PII). You need to understand how PII data flows through various transformations and tables. Which Snowflake feature, when combined with appropriate tagging and metadata management practices, can BEST help you achieve this?

- A. Snowflake Secure Data Sharing
- **B. Snowflake INFORMATION_SCHEMA views and Account Usage views**
- C. Snowflake Cloning
- D. Snowflake's Resource Monitors
- E. Snowflake Data Masking policies

Antwort: B

Begründung:

Snowflake's INFORMATION_SCHEMA views (e.g., COLUMNS, TABLES, VIEWS) and Account usage views (e.g., QUERY_HISTORY, ACCESS_HISTORY) provide metadata about database objects and query execution history. By tagging PII columns and tracking queries that access these columns, you can trace data lineage. Data Masking policies protect data at rest and in transit. Secure Data Sharing allows sharing of data without physical copying. Cloning creates a copy of data. Resource monitors help manage costs.

236. Frage

A company is using Snowflake's web app interface to manage its data. A data engineer needs to create a new table, load data into it from a CSV file stored in an internal stage, and then grant SELECT privileges on the table to a specific role using the web app. Which sequence of actions within the Snowflake web app represents the most efficient and secure way to accomplish this task?

- A. 1. Use the Database Tables interface to create the new table using the table editor. 2. Use the Data Load Data wizard to load the CSV file. 3. Use the Database -> Tables interface, select the table, and use the 'Privileges' tab to grant SELECT privilege to the role.
- B. 1. Use the Database Tables interface to create the new table using the table editor. 2. Upload the CSV file directly to the table using the 'Load Data' option. 3. Use the SQL worksheet to execute GRANT SELECT ON TABLE statement.
- C. 1. Use the Database Tables interface to create the new table using the table editor. 2. Use the Data Load Data wizard to load the CSV file. 3. Use the SQL worksheet to execute GRANT SELECT ON TABLE statement.
- D. 1. Use the SQL worksheet to execute CREATE TABLE statement. 2. Use the Database -> Tables interface, select the table, and use the 'Load Data' option to load the CSV file. 3. Use the Database -> Tables interface, select the table, and use the 'Privileges' tab to grant SELECT privilege to the role.
- E. 1. Use the SQL worksheet to execute CREATE TABLE statement. 2. Use the Data Load Data wizard to load the CSV file. 3. Use the SQL worksheet to execute GRANT SELECT ON TABLE statement.

Antwort: C

Begründung:

Option E is the most efficient and secure option, as it leverages both the graphical interface and SQL for specific tasks. Using the Table interface in option A is clunky. Option B is inefficient as Data -> Load Data Wizard is preferable to uploading via the table interface itself. Option C uses the SQL Worksheet and lacks security. Option D is redundant. Option E provides a good balance. Creating the table via the GUI, loading the data via the load data wizard and then granting privileges by the SQL Worksheet is most efficient.

237. Frage

A critical database, 'PRODUCTION_DB', in your Snowflake account was accidentally dropped. You need to restore it as quickly as possible, but you're unsure if Time Travel retention is sufficient. Which method guarantees restoration of the database even if it falls outside the Time Travel window?

- A. Fail-safe cannot be directly accessed by the user for restoration purposes; it is only used by Snowflake Support in extreme disaster recovery scenarios.
- B. Restore from a Snowflake-managed backup using the 'CREATE DATABASE ... FROM BACKUP' command. Specify the timestamp before the drop occurred.
- C. Contact Snowflake Support and request restoration from Fail-safe.
- D. Utilize the data cloning feature: 'CREATE DATABASE CLONE PRODUCTION_DB BEFORE (STATEMENT 'DROP DATABASE PRODUCTION_DB');'
- E. Use the 'UNDROP DATABASE PRODUCTION' command.

Antwort: A

Begründung:

Fail-safe is a last resort for data recovery managed entirely by Snowflake. Users cannot directly access or restore data from Fail-safe. Options B and C are valid for Time Travel, but fail if the data falls outside of that window. Option A is partially correct; you contact Snowflake support, who then might use fail-safe if appropriate, but option E is the most accurate answer. Option D uses Time Travel, which may not work.

238. Frage

You are building a data pipeline that involves ingesting data from AWS S3 into Snowflake using Snowpipe. The data arrives in small files frequently, and you are experiencing performance issues with delayed data availability. You need to optimize the pipeline for near real-time ingestion. Which combination of strategies will MOST effectively address this scenario?

- A. Manually refresh the Snowpipe using the 'ALTER PIPE ... REFRESH' command every few minutes to force ingestion of new data.
- B. Implement a micro-batching process using a third-party tool (like Apache Spark) to aggregate the small files into larger batches before loading them into S3, then configure Snowpipe to ingest the larger files.
- C. Configure S3 event notifications to trigger Snowpipe only when a sufficient number of files have arrived in the S3 bucket, using a serverless function (like AWS Lambda) to manage the file count threshold.
- **D. Enable auto-ingest on the Snowpipe and increase the frequency of S3 event notifications to Snowflake. Combine this with clustering the target table on a relevant column to optimize query performance after loading.**
- E. Increase the warehouse size used for the Snowpipe load process and adjust the 'MAX FILE SIZE' parameter on the pipe definition to match the size of the incoming files.

Antwort: D

Begründung:

Enabling auto-ingest on the Snowpipe and increasing the frequency of S3 event notifications is the most effective way to handle frequent small files. Auto-ingest automatically loads files as soon as they are available, and frequent notifications ensure minimal delay. Clustering the target table improves query performance by organizing data based on a relevant column, leading to faster data retrieval. The option A related to MAX FILE SIZE is invalid since it does not exist in the context of Snowpipe. Options involving third-party tools add complexity and latency. Manually refreshing the pipe is not a scalable solution. Option C is not ideal because you want low latency, not delaying trigger. Option D provides the lowest latency at scale, while A, B, C and E would likely have a negative impact on latency, add extra steps, or are plain incorrect.

239. Frage

You have a Snowflake table 'ORDERS' with columns 'ORDER ID', 'CUSTOMER ID', 'ORDER DATE', and 'TOTAL AMOUNT'. You notice that many queries filtering by 'ORDER DATE' are slow, even after enabling query acceleration. You decide to implement a caching strategy to improve performance. Which of the following approaches will be most effective in leveraging Snowflake's caching capabilities and improving the performance of date-filtered queries, especially when the data volume for each date is large and varied? Assume virtual warehouse is medium size.

- **A. Create a clustered table on 'ORDER_DATE'. This will physically organize the data on disk, allowing Snowflake to quickly retrieve the relevant data for date-filtered queries.**
- B. Create a materialized view that pre-aggregates the data by 'ORDER_DATE', such as calculating the sum of 'TOTAL_AMOUNT' for each date. This will allow Snowflake to serve the results directly from the materialized view for queries that require aggregation.
- C. Apply a WHERE clause with a date range in all the SELECT statements. This forces the metadata caching.
- D. Increase the data retention period for the 'ORDERS' table. A longer retention period will ensure that more data is available in the Snowflake cache.
- E. Use after running a query filtered by 'ORDER_DATE'. This will cache the result of the query in the current session for subsequent queries with the same filter.

Antwort: A

Begründung:

Clustering on 'ORDER DATE' (A) is the most effective approach. Clustering physically organizes the data, allowing Snowflake to efficiently prune partitions and retrieve only the necessary data for date-filtered queries. Materialized views (B) are suitable for pre-aggregation, but less suitable for ad-hoc queries on different date ranges where only filtering is required. Increasing the data retention period (C) does not impact caching performance and may increase storage costs. 'RESULT_SCAN' (D) is session-specific and not persistent. Adding a WHERE clause (E) helps, but is not as impactful as Clustering.

240. Frage

.....

Das Ziel der Snowflake DEA-C02 Prüfungssoftware ist: Bei Ihrer Vorbereitung der Snowflake DEA-C02 Prüfung Ihnen die effektivste Hilfe zu bieten, um Ihre Geld nicht zu verschwenden und Ihre Zeit zu sparen. Unsere Software hat schon zahllose

2026 Die neuesten Zertprüfung DEA-C02 PDF-Versionen Prüfungsfragen und DEA-C02 Fragen und Antworten sind kostenlos verfügbar: <https://drive.google.com/open?id=1rYKLykADrx2ipLf3Jvg3OEF2K9L8MVS>