ARA-C01試験資料 & ARA-C01キャリアパス

P.S. JapancertがGoogle Driveで共有している無料かつ新しいARA-C01ダンプ：https://drive.google.com/open?id=1H6w2DZ9tpXizJP1WrfVR0P_bif-oqD7g

あらゆる人にとって、時間は非常に大切です。ARA-C01試験に対して、いろいろな資料があります。そのような資料を勉強するには、長い時間がかかります。でも、ARA-C01問題集を利用すれば、短い時間でARA-C01試験に合格できます。そして、ARA-C01問題集は安くて、便利です。誰でも、ARA-C01問題集を選択すれば、試験に合格する可能性が大きいです。もし、ARA-C01問題集を勉強すれば、もし、将来にITエリートになります。

Snowflake ARA-C01認定試験は、気弱な人向けではありません。これは、スノーフレークアーキテクチャ、データモデリング、パフォーマンスの最適化、セキュリティ、および管理を深く理解する必要がある厳密で挑戦的な試験です。試験は、120分以内に完了する必要がある60の複数選択質問で構成されています。 ARA-C01試験の合格スコアは80％であり、試験に合格した候補者にはSnowpro Advanced Architect認定が授与されます。

>> ARA-C01試験資料 <<

試験の準備方法-100％合格率のARA-C01試験資料試験-正確的なARA-C01キャリアパス

当社の設立以来、私たちはARA-C01試験資料に大規模な人材、資料、および財源を投入してきましたが、これまで、私たちは間違いなく研究資料を全世界に紹介し、幸運を求めるすべての人々を作るという大胆な考えを持っています より良い機会は、彼らの人生の価値を実現するためのアクセス権を持っています。 したがって、当社のARA-C01練習問題は、試験に合格し、より良い未来を勝ち取るのに役立ちます。 また、常に先駆的な精神を持ち続け、あなたの道を歩むプロジェクトに積極的に取り組みます。

Snowflake ARA-C01（SnowPro Advanced Architect Certification）認定試験は、高度なSnowflakeアーキテクチャに精通した専門家が自分の専門知識を証明するために設計された包括的な認定プログラムです。この認定プログラムは、Snowflakeのアーキテクチャと設計原則に深い理解を持ち、ビジネス要件をスケーラブルで安全で高性能なSnowflakeソリューションに変換できる個人のスキルと知識を検証することを目的としています。

Snowflake SnowPro Advanced Architect Certification 認定 ARA-C01 試験問題 (Q140-Q145):

質問 # 140
A company is using Snowflake in Azure in the Netherlands. The company analyst team also has data in JSON format that is stored in an Amazon S3 bucket in the AWS Singapore region that the team wants to analyze.
The Architect has been given the following requirements:
1. Provide access to frequently changing data
2. Keep egress costs to a minimum
3. Maintain low latency
How can these requirements be met with the LEAST amount of operational overhead?

- A. Use AWS Transfer Family to replicate data between the S3 bucket in AWS Singapore and an Azure Netherlands Blob storage, then use an external table against the Blob storage.
- B. Use a materialized view on top of an external table against the S3 bucket in AWS Singapore.
- C. Use an external table against the S3 bucket in AWS Singapore and copy the data into transient tables.
- D. Copy the data between providers from S3 to Azure Blob storage to collocate, then use Snowpipe for data ingestion.

正解：B

解説：
Option A is the best design to meet the requirements because it uses a materialized view on top of an external table against the S3 bucket in AWS Singapore. A materialized view is a database object that contains the results of a query and can be refreshed periodically to reflect changes in the underlying data1. An external table is a table that references data files stored in a cloud storage service, such as Amazon S32. By using a materialized view on top of an external table, the company can provide access to frequently changing data, keep egress costs to a minimum, and maintain low latency. This is because the materialized view will cache the query results in Snowflake, reducing the need to access the external data files and incur network charges.
The materialized view will also improve the query performance by avoiding scanning the external data files every time. The materialized view can be refreshed on a schedule or on demand to capture the changes in the external data files1.
Option B is not the best design because it uses an external table against the S3 bucket in AWS Singapore and copies the data into transient tables. A transient table is a tablethat is not subject to the Time Travel and Fail- safe features of Snowflake, and is automatically purged after a period of time3. By using an external table and copying the data into transient tables, the company will incur more egress costs and operational overhead than using a materialized view. This is because the external table will access the external data files every time a query is executed, and the copy operation will also transfer data from S3 to Snowflake. The transient tables will also consume more storage space in Snowflake and require manual maintenance to ensure they are up to date.
Option C is not the best design because it copies the data between providers from S3 to Azure Blob storage to collocate, then uses Snowpipe for data ingestion. Snowpipe is a service that automates the loading of data from external sources into Snowflake tables4. By copying the data between providers, the company will incur high egress costs and latency, as well as operational complexity and maintenance of the infrastructure.
Snowpipe will also add another layer of processing and storage in Snowflake, which may not be necessary if the external data files are already in a queryable format.
Option D is not the best design because it uses AWS Transfer Family to replicate data between the S3 bucket in AWS Singapore and an Azure Netherlands Blob storage, then uses an external table against the Blob storage. AWS Transfer Family is a service that enables secure and seamless transfer of files over SFTP, FTPS, and FTP to and from Amazon S3 or Amazon EFS5. By using AWS Transfer Family, the company will incur high egress costs and latency, as well as operational complexity and maintenance of the infrastructure.
The external table will also access the external data files every time a query is executed, which may affect the query performance.
References: 1: Materialized Views 2: External Tables 3: Transient Tables 4: Snowpipe Overview 5: AWS Transfer Family

質問 # 141
How can the Snowpipe REST API be used to keep a log of data load history?

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

正解：A

解説：
* 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 reports1.
* The loadHistoryScan endpoint returns the history of files that have been ingested by Snowpipe within a specified time range. The endpoint accepts the following parameters2:
* 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 endpoint3.
1: Introduction to Snowpipe | Snowflake Documentation
2: loadHistoryScan | Snowflake Documentation
3: Monitoring Snowpipe Load History | Snowflake Documentation

質問 # 142
A company has several sites in different regions from which the company wants to ingest data.
Which of the following will enable this type of data ingestion?

- A. The company must have a Snowflake account in each cloud region to be able to ingest data to that account.
- B. The company should use a storage integration for the external stage.
- C. The company must replicate data between Snowflake accounts.
- D. The company should provision a reader account to each site and ingest the data through the reader accounts.

正解： B

解説：
This is the correct answer because it allows the company to ingest data from different regions using a storage integration for the external stage. A storage integration is a feature that enables secure and easy access to files in external cloud storage from Snowflake. A storage integration can be used to create an external stage, which is a named location that references the files in the external storage. An external stage can be used to load data into Snowflake tables using the COPY INTO command, or to unload data from Snowflake tables using the COPY INTO LOCATION command. A storage integration can support multiple regions and cloud platforms, as long as the external storage service is compatible with Snowflake12.
Reference:
Snowflake Documentation: Storage Integrations
Snowflake Documentation: External Stages

質問 # 143
Which data models can be used when modeling tables in a Snowflake environment? (Select THREE).

- A. Inmon/3NF
- B. Graph model
- C. Bayesian hierarchical model
- D. Data vault
- E. Data lake
- F. Dimensional/Kimball

正解： A、D、F

質問 # 144
The following table exists in the production database:
A regulatory requirement states that the company must mask the username for events that are older than six months based on the current date when the data is queried.
How can the requirement be met without duplicating the event data and making sure it is applied when creating views using the table or cloning the table?

- A. Use a row level policy on the user_events table using a entitlement table with valid dates.
- B. Use a masking policy on the username column with event_timestamp as a conditional column.
- C. Use a masking policy on the username column using a entitlement table with valid dates.
- D. Use a secure view on the user_events table using a case statement on the username column.

正解： B

質問 # 145
.....

ARA-C01キャリアパス: https://www.japancert.com/ARA-C01.html

- 実用的なARA-C01試験資料試験-試験の準備方法-効率的なARA-C01キャリアパス □ 今すぐ➠ www.jpshiken.com □を開き、➤ ARA-C01 □を検索して無料でダウンロードしてくださいARA-C01試験過去問
- ハイパスレートのARA-C01試験資料試験-試験の準備方法-最高のARA-C01キャリアパス □ URL ➥ www.goshiken.com □をコピーして開き、“ARA-C01 ”を検索して無料でダウンロードしてくださいARA-C01日本語版と英語版
- ARA-C01的中合格問題集 □ ARA-C01最新対策問題 □ ARA-C01的中合格問題集 □ ▸ ARA-C01 ◂を無料でダウンロード➽ jp.fast2test.com □ウェブサイトを入力するだけARA-C01最新関連参考書
- ARA-C01最新関連参考書 □ ARA-C01最新関連参考書 □ ARA-C01日本語的中対策 □ 「 www.goshiken.com 」を開き、「 ARA-C01 」を入力して、無料でダウンロードしてくださいARA-C01資格関連題
- ARA-C01試験の準備方法 | 便利なARA-C01試験資料試験 | 真実的なSnowPro Advanced Architect Certificationキャリアパス □ 今すぐ□ www.xhs1991.com □で➠ ARA-C01 □を検索し、無料でダウンロードしてくださいARA-C01難易度受験料
- 試験の準備方法-正確的なARA-C01試験資料試験-素敵なARA-C01キャリアパス ☀ ➽ www.goshiken.com □から簡単に“ARA-C01 ”を無料でダウンロードできますARA-C01試験番号
- ARA-C01的中合格問題集 □ ARA-C01資格参考書 □ ARA-C01日本語版と英語版 □ ➡ www.passtest.jp □□□サイトにて最新➥ ARA-C01 □問題集をダウンロードARA-C01日本語版と英語版
- ハイパスレートのARA-C01試験資料試験-試験の準備方法-最高のARA-C01キャリアパス □ 今すぐ【 www.goshiken.com 】で➽ ARA-C01 □を検索し、無料でダウンロードしてくださいARA-C01日本語的中対策
- ARA-C01模擬対策 □ ARA-C01トレーニング □ ARA-C01資格参考書 ↘ ➽ www.passtest.jp □サイトにて ▸ ARA-C01 ◂問題集を無料で使おうARA-C01資格練習
- ARA-C01的中合格問題集 □ ARA-C01最新対策問題 □ ARA-C01合格率書籍 □ 《 www.goshiken.com 》から“ARA-C01 ”を検索して、試験資料を無料でダウンロードしてくださいARA-C01模擬対策
- 試験の準備方法-正確的なARA-C01試験資料試験-素敵なARA-C01キャリアパス □ ☀ ARA-C01 □☀□を無料でダウンロード【 jp.fast2test.com 】ウェブサイトを入力するだけARA-C01資格練習
- myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.stes.tyc.edu.tw, bookmarkforest.com, robertoybw559741.tkzblog.com, oisiyuub056575.blogsvirals.com, www.stes.tyc.edu.tw, jesseuqry357366.dreamyblogs.com, joshmblv720905.blogsumer.com, nanniefzgb128096.blogsvirals.com, hannaumus050985.bloggip.com, Disposable vapes

ちなみに、Japancert ARA-C01の一部をクラウドストレージからダウンロードできます：https://drive.google.com/open?id=1H6w2DZ9tpXizJP1WrfVR0P_bif-oqD7g