

ARA-C01日本語講座、ARA-C01復習攻略問題



無料でクラウドストレージから最新のJPTestKing ARA-C01 PDFダンプをダウンロードする: <https://drive.google.com/open?id=1tx7x5UdH1umb-Mg3bB7-LkdOIK3sIVop>

ARA-C01学習教材を世界中に確実に紹介し、幸運とより良い機会を求めるすべての人々が自分の人生の価値を実現できるようにするという大胆な考えを持っています。したがって、ARA-C01練習問題は、ARA-C01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。ARA-C01トレーニング資料は、その素晴らしい品質のためにあなたを決して失望させません。

Snowflake ARA-C01試験は、スノーフレークソリューションの設計と実装の専門知識を実証したい上級スノーフレークアーキテクト向けに設計された認定試験です。この試験は、スノーフレークアーキテクチャ、実装、およびベストプラクティスに関する候補者の知識をテストするように設計されています。この試験に合格することは、スノーフレークの高レベルの技術的専門知識を実証する素晴らしい成果であり、新しいキャリアの機会にドアを開けることができます。

Snowflake ARA-C01 試験は、Snowflake ソリューションを実装および運用する経験豊富な高度なアーキテクトを対象とした認定試験です。この試験は、Snowflake が提供する最高レベルの認定であり、複雑な Snowflake 環境を設計、構築、管理する責任を持つ専門家のスキルと専門知識を確認することを意図しています。

Snowflake ARA-C01 (SnowPro Advanced Architect Certification) は、複雑な Snowflake クラウドデータプラットフォームを設計および実装することに特化したデータアーキテクトおよびエンジニアのスキルと知識を認定するためのプロフェッショナル認定試験です。この認定試験は、データ管理、データウェアハウジング、および分析ニーズに Snowflake を使用する企業によって、世界的に認識されています。

>> ARA-C01日本語講座 <<

最新ARA-C01 | 信頼的なARA-C01日本語講座試験 | 試験の準備方法 SnowPro Advanced Architect Certification復習攻略問題

JPTestKingは、他の学習教材と比較した場合、ARA-C01トレーニング教材の品質が高いことを約束できます。10年以上のビジネス経験を持つARA-C01調査ツールは、顧客の購入権をずっと重視してきました。当社のウェブサイトのARA-C01学習資料は、ユーザーの通常の作業と学習に影響を与えず、時間の利用率を大幅に向上させ、1石で2羽の鳥を殺します。弊社の学習教材が、最短でARA-C01試験に合格するのに役立つことは間違いありません。

Snowflake SnowPro Advanced Architect Certification 認定 ARA-C01 試験問題 (Q205-Q210):

質問 # 205

Data is being imported and stored as JSON in a VARIANT column. Query performance was fine, but most recently, poor query performance has been reported.

What could be causing this?

- A. There were variations in string lengths for the JSON values in the recent data imports.
- B. There were JSON nulls in the recent data imports.
- C. The order of the keys in the JSON was changed.
- D. The recent data imports contained fewer fields than usual.

正解: A、C

解説:

Data is being imported and stored as JSON in a VARIANT column. Query performance was fine, but most recently, poor query performance has been reported. This could be caused by the following factors:

The order of the keys in the JSON was changed. Snowflake stores semi-structured data internally in a column-like structure for the most common elements, and the remainder in a leftovers-like column. The order of the keys in the JSON affects how Snowflake determines the common elements and how it optimizes the query performance. If the order of the keys in the JSON was changed, Snowflake might have to re-parse the data and re-organize the internal storage, which could result in slower query performance. There were variations in string lengths for the JSON values in the recent data imports. Non-native values, such as dates and timestamps, are stored as strings when loaded into a VARIANT column. Operations on these values could be slower and also consume more space than when stored in a relational column with the corresponding data type. If there were variations in string lengths for the JSON values in the recent data imports, Snowflake might have to allocate more space and perform more conversions, which could also result in slower query performance.

The other options are not valid causes for poor query performance:

There were JSON nulls in the recent data imports. Snowflake supports two types of null values in semi-structured data: SQL NULL and JSON null. SQL NULL means the value is missing or unknown, while JSON null means the value is explicitly set to null.

Snowflake can distinguish between these two types of null values and handle them accordingly. Having JSON nulls in the recent data imports should not affect the query performance significantly.

The recent data imports contained fewer fields than usual. Snowflake can handle semi-structured data with varying schemas and fields. Having fewer fields than usual in the recent data imports should not affect the query performance significantly, as Snowflake can still optimize the data ingestion and query execution based on the existing fields.

Reference:

Considerations for Semi-structured Data Stored in VARIANT

Snowflake Architect Training

Snowflake query performance on unique element in variant column

Snowflake variant performance

質問 # 206

An Architect has chosen to separate their Snowflake Production and QA environments using two separate Snowflake accounts.

The QA account is intended to run and test changes on data and database objects before pushing those changes to the Production account. It is a requirement that all database objects and data in the QA account need to be an exact copy of the database objects, including privileges and data in the Production account on at least a nightly basis.

Which is the LEAST complex approach to use to populate the QA account with the Production account's data and database objects on a nightly basis?

- A. 1) Enable replication for each database in the Production account
2) Create replica databases in the QA account
3) Create clones of the replica databases on a nightly basis
4) Run tests directly on those cloned databases
- B. 1) In the Production account, create an external function that connects into the QA account and returns all the data for one specific table
2) Run the external function as part of a stored procedure that loops through each table in the Production account and populates each table in the QA account
- C. 1) Create a stage in the Production account
2) Create a stage in the QA account that points to the same external object-storage location
3) Create a task that runs nightly to unload each table in the Production account into the stage
4) Use Snowpipe to populate the QA account
- D. 1) Create a share in the Production account for each database
2) Share access to the QA account as a Consumer
3) The QA account creates a database directly from each share

- 4) Create clones of those databases on a nightly basis
- 5) Run tests directly on those cloned databases

正解: A

解説:

Explanation

* This approach is the least complex because it uses Snowflake's built-in replication feature to copy the data and database objects from the Production account to the QA account. Replication is a fast and efficient way to synchronize data across accounts, regions, and cloud platforms. It also preserves the privileges and metadata of the replicated objects. By creating clones of the replica databases, the QA account can run tests on the cloned data without affecting the original data. Clones are also zero-copy, meaning they do not consume any additional storage space unless the data is modified. This approach does not require any external stages, tasks, Snowpipe, or external functions, which can add complexity and overhead to the data transfer process.

References:

- * Introduction to Replication and Failover
- * Replicating Databases Across Multiple Accounts
- * Cloning Considerations

質問 # 207

An Architect entered the following commands in sequence:

USER1 cannot find the table.

Which of the following commands does the Architect need to run for USER1 to find the tables using the Principle of Least Privilege? (Choose two.)

- A. GRANT ALL PRIVILEGES ON DATABASE SANDBOX TO ROLE INTERN;
- B. GRANT USAGE ON SCHEMA SANDBOX.PUBLIC TO ROLE INTERN;
- C. GRANT ROLE PUBLIC TO ROLE INTERN;
- D. GRANT OWNERSHIP ON DATABASE SANDBOX TO USER INTERN;
- E. GRANT USAGE ON DATABASE SANDBOX TO ROLE INTERN;

正解: B、E

解説:

According to the Principle of Least Privilege, the Architect should grant the minimum privileges necessary for the USER1 to find the tables in the SANDBOX database.

The USER1 needs to have USAGE privilege on the SANDBOX database and the SANDBOX.PUBLIC schema to be able to access the tables in the PUBLIC schema. Therefore, the commands B and C are the correct ones to run.

The command A is not correct because the PUBLIC role is automatically granted to every user and role in the account, and it does not have any privileges on the SANDBOX database by default.

The command D is not correct because it would transfer the ownership of the SANDBOX database from the Architect to the USER1, which is not necessary and violates the Principle of Least Privilege.

The command E is not correct because it would grant all the possible privileges on the SANDBOX database to the USER1, which is also not necessary and violates the Principle of Least Privilege.

質問 # 208

How do Snowflake databases that are created from shares differ from standard databases that are not created from shares? (Choose three.)

- A. Shared databases must be refreshed in order for new data to be visible.
- B. Shared databases can also be created as transient databases.
- C. Shared databases are read-only.
- D. Shared databases will have the PUBLIC or INFORMATION_SCHEMA schemas without explicitly granting these schemas to the share.
- E. Shared databases are not supported by Time Travel.
- F. Shared databases cannot be cloned.

正解: C、E、F

解説:

According to the SnowPro Advanced: Architect documents and learning resources, the ways that Snowflake databases that are created from shares differ from standard databases that are not created from shares are:

Shared databases are read-only. This means that the data consumers who access the shared databases cannot modify or delete the data or the objects in the databases. The data providers who share the databases have full control over the data and the objects, and can grant or revoke privileges on them¹.

Shared databases cannot be cloned. This means that the data consumers who access the shared databases cannot create a copy of the databases or the objects in the databases. The data providers who share the databases can clone the databases or the objects, but the clones are not automatically shared².

Shared databases are not supported by Time Travel. This means that the data consumers who access the shared databases cannot use the AS OF clause to query historical data or restore deleted data. The data providers who share the databases can use Time Travel on the databases or the objects, but the historical data is not visible to the data consumers³.

The other options are incorrect because they are not ways that Snowflake databases that are created from shares differ from standard databases that are not created from shares. Option B is incorrect because shared databases do not need to be refreshed in order for new data to be visible. The data consumers who access the shared databases can see the latest data as soon as the data providers update the data¹. Option E is incorrect because shared databases will not have the PUBLIC or INFORMATION_SCHEMA schemas without explicitly granting these schemas to the share. The data consumers who access the shared databases can only see the objects that the data providers grant to the share, and the PUBLIC and INFORMATION_SCHEMA schemas are not granted by default⁴. Option F is incorrect because shared databases cannot be created as transient databases. Transient databases are databases that do not support Time Travel or Fail-safe, and can be dropped without affecting the retention period of the data. Shared databases are always created as permanent databases, regardless of the type of the source database⁵. Reference: Introduction to Secure Data Sharing | Snowflake Documentation, Cloning Objects | Snowflake Documentation, Time Travel | Snowflake Documentation, Working with Shares | Snowflake Documentation, CREATE DATABASE | Snowflake Documentation

質問 # 209

Multi-cluster warehouses are best utilized for

- A. Scaling resources to improve concurrency for users/queries
- B. Improving the performance of data loading
- C. Improving the performance of slow-running queries

正解: A

質問 # 210

.....

ARA-C01の学習教材で20~30時間準備したと主張することができます。ARA-C01試験に簡単に合格して、期待されるスコアを取得できます。またSnowflake、ARA-C01試験問題の無料デモを提供しており、ARA-C01トレーニング資料の有効性と正確性を確認できます。JPTTestKing やって来てみてください! ARA-C01トレーニング資料の高い精度に驚かれることでしょう。そして、ARA-C01練習問題集の高い合格率は99%から100%なので、SnowPro Advanced Architect Certification試験に簡単に合格します。

ARA-C01復習攻略問題: <https://www.jpctestking.com/ARA-C01-exam.html>

- 素敵ARA-C01 | 正確なARA-C01日本語講座試験 | 試験の準備方法SnowPro Advanced Architect Certification 復習攻略問題 □ ➡ www.it-passports.com □ サイトにて“ARA-C01”問題集を無料で使おうARA-C01日本語版復習資料
- 実用的-最新のARA-C01日本語講座試験-試験の準備方法ARA-C01復習攻略問題 □ 今すぐ☀ www.goshiken.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テキスト
- www.mogixexam.comはSnowflake ARA-C01の試験問題集を提供する ♣ ▷ www.mogixexam.com ◁ から簡単に ➡ ARA-C01 □ を無料でダウンロードできますARA-C01資格トレーニング
- GoShikenはSnowflake ARA-C01の試験問題集を提供する □ 今すぐ (www.goshiken.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受験料過去問
- ARA-C01対策学習 □ ARA-C01受験料過去問 □ ARA-C01模擬試験サンプル □ 時間限定無料で使える (ARA-C01) の試験問題は □ www.xhs1991.com □ サイトで検索ARA-C01練習問題
- 素敵ARA-C01 | 正確なARA-C01日本語講座試験 | 試験の準備方法SnowPro Advanced Architect Certification 復習攻略問題 □ (ARA-C01) を無料でダウンロード { www.goshiken.com } で検索するだけARA-C01練習問題

