

ARA-C01높은통과율 시험공부자료100% 시험패스가능한덤프자료



2025 ITDumpsKR 최신 ARA-C01 PDF 버전 시험 문제집과 ARA-C01 시험 문제 및 답변 무료 공유:
<https://drive.google.com/open?id=1v5-7ZE6NxXVwCFO9QgsavakSAK1zZ7sb>

IT업계에 종사하는 분들은 치열한 경쟁을 많이 느낄것입니다. 치열한 경쟁속에서 자신의 위치를 보장하는 길은 더 많이 배우고 더 많이 노력하는것 뿐입니다.국제적으로 인정받은 IT인증자격증을 취득하는것이 제일 중요한 부분이 아닌가 싶기도 합니다. 다른 분이 없는 자격증을 내가 소유하고 있다는 생각만 해도 뭔가 안전감이 느껴지지 않나요? 더는 시간낭비하지 말고ITDumpsKR의Snowflake인증 ARA-C01덤프로Snowflake인증 ARA-C01시험에 도전해보세요.

우리ITDumpsKR에서 여러분은 아주 간단히Snowflake ARA-C01시험을 패스할 수 있습니다. 만약 처음Snowflake ARA-C01시험에 도전한다면 우리의Snowflake ARA-C01시험자료를 선택하여 다운받고 고부를 한다면 생각보다는 아주 쉽게Snowflake ARA-C01시험을 통과할 수 있으며 무엇보다도 시험시의 자신감 충만에 많은 도움이 됩니다. 다른 자료판매사이트도 많겠지만 저희는 저희 자료에 자신이 있습니다. 우리의 시험자료는 모두 하이퀄리티한 문제와 답으로 구성되었습니다, 그리고 우리는 업데이트를 아주 중요시 생각하기에 어느 사이트보다 더 최신버전을 보실 수 있을것입니다. 우리의Snowflake ARA-C01자료로 자신만만한 시험 준비하시기를 바랍니다. 우리를 선택함으로써 자신의 시간을 아끼는 셈이라고 생각하시면 됩니다.Snowflake ARA-C01로 빠른시일내에 자격증 취득하시고 SnowflakeIT업계중에 엘리트한 전문가되시기를 바랍니다.

>> ARA-C01높은 통과율 시험공부자료 <<

ARA-C01최신버전 시험공부자료 - ARA-C01인기자격증 덤프자료

Snowflake ARA-C01 덤프가 고객님의 기대를 가득 채워드릴수 있도록 정말로 노력하고 있는 ITDumpsKR입니다. Snowflake ARA-C01 덤프는 pdf버전과 소프트웨어버전으로만 되어있었는데 최근에는 휴대폰에서가 사용가능한 온라인버전까지 개발하였습니다. 날따라 새로운 시스템을 많이 개발하여 고객님의게 더욱 편하게 다가갈수 있는 ITDumpsKR가 되겠습니다.

Snowflake ARA-C01 시험요강:

주제	소개
주제 1	<ul style="list-style-type: none">• Determine the appropriate data loading or data unloading solution to meet business needs• Design an architecture that meets data security, privacy, compliance, and governance requirements
주제 2	<ul style="list-style-type: none">• Design data sharing solutions, based on different use cases• Determine the appropriate data transformation solution to meet business needs
주제 3	<ul style="list-style-type: none">• Design a Snowflake account and database strategy, based on business requirements• Troubleshoot performance issues with existing architectures

주제 4	<ul style="list-style-type: none"> • Create architecture solutions that support Development Lifecycles as well as workload requirements • Outline Snowflake security principles and identify use cases where they should be applied
주제 5	<ul style="list-style-type: none"> • Outline performance tools, best practices, and appropriate scenarios where they should be applied • Determine the appropriate data recovery solution in Snowflake and how data can be restored

최신 SnowPro Advanced Certification ARA-C01 무료샘플문제 (Q19-Q24):

질문 # 19

A user is executing the following command sequentially within a timeframe of 10 minutes from start to finish:

```
use role sysadmin;
use warehouse compute_wh;
use schema sales.public;
create table t_sales (numeric integer) data_retention_time_in_days=1;
create or replace table t_sales_clone clone t_sales at(offset => -60*30);
```

What would be the output of this query?

- A. Table T_SALES_CLONE successfully created.
- B. Time Travel data is not available for table T_SALES.
- C. The offset -> is not a valid clause in the clone operation.
- D. Syntax error line 1 at position 58 unexpected 'at'.

정답: A

설명:

The query is executing a clone operation on an existing table t_sales with an offset to account for the retention time. The syntax used is correct for cloning a table in Snowflake, and the use of the at(offset=> -60*30) clause is valid. This specifies that the clone should be based on the state of the table 30 minutes prior (60 seconds * 30). Assuming the table t_sales exists and has been modified within the last 30 minutes, and considering the data_retention_time_in_days is set to 1 day (which enables time travel queries for the past 24 hours), the table t_sales_clone would be successfully created based on the state of t_sales 30 minutes before the clone command was issued.

질문 # 20

What is a characteristic of event notifications in Snowpipe?

- A. Notifications identify the cloud storage event and the actual data in the files.
- B. When a pipe is paused, event messages received for the pipe enter a limited retention period.
- C. The load history is stored in the metadata of the target table.
- D. Snowflake can process all older notifications when a paused pipe is resumed.

정답: B

설명:

Event notifications in Snowpipe are messages sent by cloud storage providers to notify Snowflake of new or modified files in a stage. Snowpipe uses these notifications to trigger data loading from the stage to the target table. When a pipe is paused, event messages received for the pipe enter a limited retention period, which varies depending on the cloud storage provider. If the pipe is not resumed within the retention period, the event messages will be discarded and the data will not be loaded automatically. To load the data, the pipe must be resumed and the COPY command must be executed manually. This is a characteristic of event notifications in Snowpipe that distinguishes them from other options. References: Snowflake Documentation: Using Snowpipe, Snowflake Documentation: Pausing and Resuming a Pipe

질문 # 21

An Architect has been asked to clone schema STAGING as it looked one week ago, Tuesday June 1st at 8:00 AM, to recover some objects.

The STAGING schema has 50 days of retention.

The Architect runs the following statement:

CREATE SCHEMA STAGING_CLONE CLONE STAGING at (timestamp => '2021-06-01 08:00:00'); The Architect receives the following error: Time travel data is not available for schema STAGING. The requested time is either beyond the allowed time travel period or before the object creation time.

The Architect then checks the schema history and sees the following:

CREATED_ON|NAME|DROPPED_ON

2021-06-02 23:00:00 | STAGING | NULL

2021-05-01 10:00:00 | STAGING | 2021-06-02 23:00:00

How can cloning the STAGING schema be achieved?

- A. Rename the STAGING schema and perform an UNDROP to retrieve the previous STAGING schema version, then run the CLONE statement.
- B. Modify the statement: CREATE SCHEMA STAGING_CLONE CLONE STAGING at (timestamp => '2021-05-01 10:00:00');
- C. Undrop the STAGING schema and then rerun the CLONE statement.
- D. Cloning cannot be accomplished because the STAGING schema version was not active during the proposed Time Travel time period.

정답: A

설명:

Explanation

* The error message indicates that the schema STAGING does not have time travel data available for the requested timestamp, because the current version of the schema was created on 2021-06-02 23:00:00, which is after the timestamp of 2021-06-01 08:00:00. Therefore, the CLONE statement cannot access the historical data of the schema at that point in time.

* Option A is incorrect, because undropping the STAGING schema will not restore the previous version of the schema that was active on 2021-06-01 08:00:00. Instead, it will create a new version of the schema with the same name and no data or objects.

* Option B is incorrect, because modifying the timestamp to 2021-05-01 10:00:00 will not clone the schema as it looked one week ago, but as it looked when it was first created. This may not reflect the desired state of the schema and its objects.

* Option C is correct, because renaming the STAGING schema and performing an UNDROP to retrieve the previous STAGING schema version will restore the schema that was dropped on 2021-06-02

23:00:00. This schema has time travel data available for the requested timestamp of 2021-06-01

08:00:00, and can be cloned using the CLONE statement.

* Option D is incorrect, because cloning can be accomplished by using the UNDROP command to access the previous version of the schema that was active during the proposed time travel period.

References: : Cloning Considerations : Understanding & Using Time Travel : CREATE <object> ... CLONE

질문 # 22

What is a key consideration when setting up search optimization service for a table?

- A. The table must be clustered with a key having multiple columns for effective search optimization.
- B. Search optimization service works best with a column that has a minimum of 100 K distinct values.
- C. Search optimization service can help to optimize storage usage by compressing the data into a GZIP format.
- D. Search optimization service can significantly improve query performance on partitioned external tables.

정답: B

설명:

Search optimization service is a feature of Snowflake that can significantly improve the performance of certain types of lookup and analytical queries on tables. Search optimization service creates and maintains a persistent data structure called a search access path, which keeps track of which values of the table's columns might be found in each of its micro-partitions, allowing some micro-partitions to be skipped when scanning the table1.

Search optimization service can significantly improve query performance on partitioned external tables, which are tables that store data in external locations such as Amazon S3 or Google Cloud Storage. Partitioned external tables can leverage the search access path to prune the partitions that do not contain the relevant data, reducing the amount of data that needs to be scanned and transferred from the external location2.

The other options are not correct because:

* A. Search optimization service works best with a column that has a high cardinality, which means that the column has many distinct values. However, there is no specific minimum number of distinct values required for search optimization service to work effectively.

The actual performance improvement depends on the selectivity of the queries and the distribution of the data¹.

* C. Search optimization service does not help to optimize storage usage by compressing the data into a GZIP format. Search optimization service does not affect the storage format or compression of the data, which is determined by the file format options of the table. Search optimization service only creates an additional data structure that is stored separately from the table data¹.

* D. The table does not need to be clustered with a key having multiple columns for effective search optimization. Clustering is a feature of Snowflake that allows ordering the data in a table or a partitioned external table based on one or more clustering keys. Clustering can improve the performance of queries that filter on the clustering keys, as it reduces the number of micro-partitions that need to be scanned. However, clustering is not required for search optimization service to work, as search optimization service can skip micro-partitions based on any column that has a search access path, regardless of the clustering key³.

References:

* 1: Search Optimization Service | Snowflake Documentation

* 2: Partitioned External Tables | Snowflake Documentation

* 3: Clustering Keys | Snowflake Documentation

질문 # 23

A company's client application supports multiple authentication methods, and is using Okta.

What is the best practice recommendation for the order of priority when applications authenticate to Snowflake?

- A. 1) External browser, SSO 2) Key Pair Authentication, mostly used for development environment users 3) Okta native authentication 4) OAuth (either Snowflake OAuth or External OAuth) 5) Password
- **B. 1) OAuth (either Snowflake OAuth or External OAuth) 2) External browser 3) Okta native authentication 4) Key Pair Authentication, mostly used for service account users 5) Password**
- C. 1) Okta native authentication 2) Key Pair Authentication, mostly used for production environment users 3) Password 4) OAuth (either Snowflake OAuth or External OAuth) 5) External browser, SSO
- D. 1) Password 2) Key Pair Authentication, mostly used for production environment users 3) Okta native authentication 4) OAuth (either Snowflake OAuth or External OAuth) 5) External browser, SSO

정답: B

설명:

This is the best practice recommendation for the order of priority when applications authenticate to Snowflake, according to the Snowflake documentation and the web search results. Authentication is the process of verifying the identity of a user or application that connects to Snowflake. Snowflake supports multiple authentication methods, each with different advantages and disadvantages. The recommended order of priority is based on the following factors:

* Security: The authentication method should provide a high level of security and protection against unauthorized access or data breaches. The authentication method should also support multi-factor authentication (MFA) or single sign-on (SSO) for additional security.

* Convenience: The authentication method should provide a smooth and easy user experience, without requiring complex or manual steps. The authentication method should also support seamless integration with external identity providers or applications.

* Flexibility: The authentication method should provide a range of options and features to suit different use cases and scenarios. The authentication method should also support customization and configuration to meet specific requirements.

Based on these factors, the recommended order of priority is:

* OAuth (either Snowflake OAuth or External OAuth): OAuth is an open standard for authorization that allows applications to access Snowflake resources on behalf of a user, without exposing the user's credentials. OAuth provides a high level of security, convenience, and flexibility, as it supports MFA, SSO, token-based authentication, and various grant types and scopes. OAuth can be implemented using either Snowflake OAuth or External OAuth, depending on the identity provider and the application^{1,2}.

* External browser: External browser is an authentication method that allows users to log in to Snowflake using a web browser and an external identity provider, such as Okta, Azure AD, or Ping Identity.

External browser provides a high level of security and convenience, as it supports MFA, SSO, and federated authentication.

External browser also provides a consistent user interface and experience across different platforms and devices^{3,4}.

* Okta native authentication: Okta native authentication is an authentication method that allows users to log in to Snowflake using Okta as the identity provider, without using a web browser. Okta native authentication provides a high level of security and convenience, as it supports MFA, SSO, and federated authentication. Okta native authentication also provides a native user interface and experience for Okta users, and supports various Okta features, such as password policies and user management^{5,6}.

* Key Pair Authentication: Key Pair Authentication is an authentication method that allows users to log in to Snowflake using a public-private key pair, without using a password. Key Pair Authentication provides a high level of security, as it relies on asymmetric encryption and digital signatures. Key Pair Authentication also provides a flexible and customizable authentication option, as it supports various key formats, algorithms, and expiration times. Key Pair Authentication is mostly used for service account users, such as applications or scripts that connect to Snowflake programmatically⁷.

* Password: Password is the simplest and most basic authentication method that allows users to log in to Snowflake using a

