

Snowflake ADA-C01 Interactive Practice Exam, Valid ADA-C01 Exam Syllabus



What's more, part of that PracticeTorrent ADA-C01 dumps now are free: <https://drive.google.com/open?id=1L8LLLYvOE6xorWDciliGmutKJqWXT29p>

Passing an SnowPro Advanced Administrator exam on the first attempt can be stressful, but Snowflake ADA-C01 exam questions can help manage stress and allow you to perform at your best. We at PracticeTorrent give you the techniques and resources to make sure you get the most out of your exam study. We provide preparation material for the SnowPro Advanced Administrator exam that will guide you when you sit to study for it. ADA-C01 updated questions give you enough confidence to sit for the Snowflake exam.

Snowflake ADA-C01 Exam Syllabus Topics:

Topic	Details
Topic 1	<ul style="list-style-type: none">• Account Management and Data Governance: This section of the exam measures the skills of Data Governance Managers and Database Administrators and covers account organization, access control, and regulatory data protection. Candidates will learn how to manage organizational accounts, encryption keys, and Tri-Secret Secure implementations. It focuses on applying best practices in ORGADMIN and ACCOUNTADMIN roles, implementing masking and row access policies, and performing data classification and tagging. The domain also emphasizes data auditing, account identifiers, and effective management of tables, views, and query operations to support enterprise-wide governance standards.
Topic 2	<ul style="list-style-type: none">• Disaster Recovery, Backup, and Data Replication: This section of the exam measures the skills of Disaster Recovery Engineers and Cloud Operations Managers and covers Snowflake methods for ensuring business continuity. Candidates must understand how to replicate databases and account-level objects, implement failover strategies, and perform backup and restoration through Time Travel and Fail-safe features. The domain emphasizes replication across accounts, handling data consistency during failover, and applying cost-efficient disaster recovery strategies to maintain availability during outages or regional failures.

Topic 3	<ul style="list-style-type: none"> • Snowflake Security, Role-Based Access Control (RBAC), and User Administration: This section of the exam measures the skills of Snowflake Administrators and Cloud Security Engineers and covers authentication, access control, and network management in Snowflake. Candidates must understand how to configure authentication methods such as SSO, MFA, OAuth, and key-pair authentication, and how to manage network policies and private connectivity. The domain also tests knowledge of user and role management using SCIM, designing access control architecture, and applying the RBAC framework to ensure secure user authorization and data protection within Snowflake environments.
Topic 4	<ul style="list-style-type: none"> • Performance Monitoring and Tuning: This section of the exam measures the skills of Cloud Infrastructure Engineers and Performance Analysts and focuses on optimizing Snowflake compute and storage resources. Candidates will need to understand how to configure and manage virtual warehouses, evaluate query profiles, and apply caching and clustering strategies for performance tuning. It also includes monitoring concurrency, resource utilization, and implementing cost optimization strategies. The ability to interpret, explain plans, apply search optimization, and manage cost controls is key for maintaining efficient Snowflake environments.
Topic 5	<ul style="list-style-type: none"> • Data Sharing, Data Exchange, and Snowflake Marketplace: This section of the exam measures the skills of Data Integration Specialists and Data Platform Administrators and covers managing and implementing data-sharing solutions within Snowflake. It evaluates understanding of data sharing models across regions and clouds, secure data sharing methods, and managing provider-consumer relationships. The domain also includes the use of Snowflake Data Exchange and Marketplace to publish, consume, and manage data listings, ensuring secure collaboration and efficient data monetization.

>> Snowflake ADA-C01 Interactive Practice Exam <<

Evaluate Your Exam Preparation with Online Snowflake ADA-C01 Practice Test Engine

The contents of ADA-C01 exam torrent was all compiled by experts through the refined off textbooks. Hundreds of experts simplified the contents of the textbooks, making the lengthy and complex contents easier and more understandable. With ADA-C01 study tool, you only need 20-30 hours of study before the exam. ADA-C01 Guide Torrent provides you with a brand-new learning method. In the course of doing questions, you can memorize knowledge points. You no longer need to look at the complicated expressions in the textbook.

Snowflake SnowPro Advanced Administrator Sample Questions (Q23-Q28):

NEW QUESTION # 23

When adding secure views to a share in Snowflake, which function is needed to authorize users from another account to access rows in a base table?

- **A. CURRENT_USER**
- B. CURRENT_ACCOUNT
- C. CURRENT_ROLE
- D. CURRENT_CLIENT

Answer: A

Explanation:

According to the Working with Secure Views documentation, secure views are designed to limit access to sensitive data that should not be exposed to all users of the underlying table(s). When sharing secure views with another account, the view definition must include a function that returns the identity of the user who is querying the view, such as CURRENT_USER, CURRENT_ROLE, or CURRENT_ACCOUNT. These functions can be used to filter the rows in the base table based on the user's identity. For example, a secure view can use the CURRENT_USER function to compare the user name with a column in the base table that contains the authorized user names. Only the rows that match the user name will be returned by the view. The CURRENT_CLIENT function is not suitable for this purpose, because it returns the IP address of the client that is connected to Snowflake, which is not related to the user's identity.

NEW QUESTION # 24

A Snowflake Administrator needs to persist all virtual warehouse configurations for auditing and backups.

Given a table already exists with the following schema:

Table Name:VWH_META

Column 1:SNAPSHOT_TIME TIMESTAMP_NTZ

Column 2:CONFIG VARIANT

Which commands should be executed to persist the warehouse data at the time of execution in JSON format in the table VWH META?

- A. 1. SHOW WAREHOUSES;
2. INSERT INTO VWH_META
SELECT CURRENT_TIMESTAMP (),
OBJECT_CONSTRUCT (*)
FROM TABLE (RESULT_SCAN (LAST_QUERY_ID ()));
- B. 1. SHOW WAREHOUSES;
2. INSERT INTO VWH META
SELECT CURRENT_TIMESTAMP (), *
FROM TABLE (RESULT_SCAN (LAST_QUERY_ID ())) ;
- C. 1. SHOW WAREHOUSES;
2. INSERT INTO VWH META
SELECT CURRENT_TIMESTAMP (), *
FROM TABLE (RESULT_SCAN (SELECT
LAST_QUERY_ID(-1)));
- D. 1. SHOW WAREHOUSES;
2. INSERT INTO VWH META
SELECT CURRENT_TIMESTAMP (),
FROM TABLE (RESULT_SCAN (LAST_QUERY_ID(1))) ;

Answer: A

Explanation:

Explanation

According to the Using Persisted Query Results documentation, the RESULT_SCAN function allows you to query the result set of a previous command as if it were a table. The LAST_QUERY_ID function returns the query ID of the most recent statement executed in the current session. Therefore, the combination of these two functions can be used to access the output of the SHOW WAREHOUSES command, which returns the configurations of all the virtual warehouses in the account. However, to persist the warehouse data in JSON format in the table VWH_META, the OBJECT_CONSTRUCT function is needed to convert the output of the SHOW WAREHOUSES command into a VARIANT column. The OBJECT_CONSTRUCT function takes a list of key-value pairs and returns a single JSON object. Therefore, the correct commands to execute are:

1.SHOW WAREHOUSES;

2.INSERT INTO VWH_META SELECT CURRENT_TIMESTAMP (), OBJECT_CONSTRUCT (*) FROM TABLE (RESULT_SCAN (LAST_QUERY_ID ())); The other options are incorrect because:

*A. This option does not use the OBJECT_CONSTRUCT function, so it will not persist the warehouse data in JSON format. Also, it is missing the * symbol in the SELECT clause, so it will not select any columns from the result set of the SHOW WAREHOUSES command.

*B. This option does not use the OBJECT_CONSTRUCT function, so it will not persist the warehouse data in JSON format. It will also try to insert multiple columns into a single VARIANT column, which will cause a type mismatch error.

*D. This option does not use the OBJECT_CONSTRUCT function, so it will not persist the warehouse data in JSON format. It will also try to use the RESULT_SCAN function on a subquery, which is not supported. The RESULT_SCAN function can only be used on a query ID or a table name.

NEW QUESTION # 25

The following commands were executed:

Grant usage on database PROD to role PROD_ANALYST;

Grant usage on database PROD to role PROD_SUPERVISOR;

Grant ALL PRIVILEGES on schema PROD. WORKING to role PROD_ANALYST;

Grant ALL PRIVILEGES on schema PROD. WORKING to role PROD_SUPERVISOR;

Grant role PROD_ANALYST to user A;

Grant role PROD SUPERVISOR to user B;

What authority does each user have on the WORKING schema?

- A. All existing tables in schema PROD. WORKING will be visible to both users.
- B. Only user B can create tables, because all privileges were transferred to PROD_SUPERVISOR.
- **C. Both user A and user B can create tables in the PROD. WORKING schema.**
- D. Tables created by either user A or user B will be visible to both users.

Answer: C

NEW QUESTION # 26

In general, the monthly billing for database replication is proportional to which variables? (Select TWO).

- A. The frequency of the secondary database refreshes from the primary database
- B. The number of times data moves across regions and/or cloud service providers between the primary and secondary database accounts
- C. The number and size of warehouses defined in the primary account
- **D. The amount of table data in the primary database that changes as a result of data loading or DML operations**
- **E. The frequency of changes to the primary database as a result of data loading or DML operations**

Answer: D,E

Explanation:

Explanation

Snowflake charges for database replication based on two categories: data transfer and compute resources¹.

Data transfer costs depend on the amount of data that is transferred from the primary database to the secondary database across regions and/or cloud service providers². Compute resource costs depend on the use of Snowflake-provided compute resources to copy data between accounts across regions¹. Both data transfer and compute resource costs are proportional to the frequency and amount of changes to the primary database as a result of data loading or DML operations³.

Therefore, the answer is A and B. The other options are not directly related to the replication billing, as the frequency of secondary database refreshes does not affect the amount of data transferred or copied⁴, and the number and size of warehouses defined in the primary account do not affect the replication process⁵.

NEW QUESTION # 27

A team is provisioning new lower environments from the production database using cloning. All production objects and references reside in the database, and do not have external references.

What set of object references needs to be re-pointed before granting access for usage?

- **A. Sequences, storage integrations, views, secure views, and materialized views**
- B. Sequences, views, secure views, and materialized views
- C. There are no object references that need to be re-pointed
- D. Sequences, views, and secure views

Answer: A

Explanation:

Explanation

According to the Snowflake documentation¹, when an object in a schema is cloned, any future grants defined for this object type in the schema are applied to the cloned object unless the COPY GRANTS option is specified in the CREATE statement for the clone operation. However, some objects may still reference the source object or external objects after cloning, which may cause issues with access or functionality. These objects include:

*Sequences: If a table column references a sequence that generates default values, the cloned table may reference the source or cloned sequence, depending on where the sequence is defined. To avoid conflicts, the sequence reference should be re-pointed to the desired sequence using the ALTER TABLE command².

*Storage integrations: If a stage or a table references a storage integration, the cloned object may still reference the source storage integration, which may not be accessible or valid in the new environment. To avoid errors, the storage integration reference should be re-pointed to the desired storage integration using the ALTER STAGE or ALTER TABLE command^{3,4}.

*Views, secure views, and materialized views: If a view references another view or table, the cloned view may still reference the source object, which may not be accessible or valid in the new environment. To avoid errors, the view reference should be re-

