Snowflake ADA-C01 Exam | ADA-C01 Examcollection - Test Engine Simulation of Practice ADA-C01 Exam Pdf

BONUS!!! Download part of ExamsLabs ADA-C01 dumps for free: https://drive.google.com/open?id=1XBwWSsP9Q_ruR80zjcPU1zgzkVErZZ4k

ExamsLabs has built customizable Snowflake ADA-C01 practice exams (desktop software & web-based) for our customers. Users can customize the time and ADA-C01 questions of Snowflake ADA-C01 Practice Tests according to their needs. You can give more than one test and track the progress of your previous attempts to improve your marks on the next try.

Snowflake ADA-C01 Exam Syllabus Topics:

Topic	Details
Topic 1	• 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	• 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	• 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	• 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.
Topic 5	• 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.

ADA-C01 Exam Questions - ADA-C01 Guide Torrent & SnowPro Advanced Administrator Test Guide

These real and updated Snowflake ADA-C01 dumps are essential to pass the ADA-C01 exam on the first try. Don't waste further time and money, get real Snowflake ADA-C01 pdf questions and practice test software, and start ADA-C01 Test Preparation today. ExamsLabs will also provide you with up to 365 days of free exam questions updates.

Snowflake SnowPro Advanced Administrator Sample Questions (Q71-Q76):

NEW QUESTION # 71
A resource monitor named MONTHLY_FINANCE_LIMIT has been created and applied to two virtual warehouses (fin_wh1 and fin_wh2) using the following SQL:

```
ALTER RESOURCE MONITOR MONTHLY_FINANCE_LIMIT SET CREDIT_QUOTA = 1000
    FREQUENCY = MONTHLY
    START_TIMESTAMP = '2022-12-01 00:00 PST'
    NOTIFY_USERS_ON_80_PERCENT DO SUSPEND
    NOTIFY_USERS_ON_100_PERCENT DO SUSPEND_IMMEDIATE;

ALTER WAREHOUSE fin_wh1 SET RESOURCE_MONITOR = MONTHLY_FINANCE_LIMIT;
ALTER WAREHOUSE fin_wh2 SET RESOURCE_MONITOR = MONTHLY_FINANCE_LIMIT;
```

Given that the combined total of credits consumed by fin_wh1 and fin_wh2 (including cloud services) has reached 800 credits and both warehouses are suspended, what should the ACCOUNTADMIN execute to allow both warehouses to be resumed? (Select TWO).

- A. ALTER WAREHOUSE fin_wh2 UNSET RESOURCE_MONITOR MONTHLY_FINANCE_LIMIT;
- B. ALTER WAREHOUSE fin_wh2 RESUME;
- C. ALTER RESOURCE MONITOR MONTHLY_FINANCE_LIMIT SET CREDIT_QUOTA = 1500;
- D. ALTER WAREHOUSE fin_wh1 RESUME;
- E. ALTER RESOURCE MONITOR MONTHLY_FINANCE_LIMIT RESET;
- F. ALTER WAREHOUSE fin_wh1 UNSET RESOURCE_MONITOR MONTHLY_FINANCE_LIMIT;
- G. ALTER WAREHOUSE fin_wh1 UNSET RESOURCE_MONITORS;

Answer: C,E

Explanation:
 Scenario:

Resource Monitor MONTHLY_FINANCE_LIMIT has a credit quota of 1000.
800 credits have been used and warehouses are already suspended.
According to monitor configuration:
At 80%, warehouses are suspended.
At 100%, warehouses would be suspended immediately.
Warehouses cannot resume until the monitor is reset or the quota is increased.
 E. SET CREDIT_QUOTA = 1500
Increases the monthly credit limit to 1500.
Since current usage is 800 < 1500, this puts usage below 80%.
This allows resumption of warehouses.
 F. RESET
sql
CopyEdit
ALTER RESOURCE MONITOR MONTHLY_FINANCE_LIMIT RESET;
Resets usage to zero for the current period.
Allows warehouses to resume immediately - same effect as a fresh cycle.
 Why Other Options Are Incorrect:
A . / B. ALTER WAREHOUSE ... RESUME
Won't work while the resource monitor is actively suspending the warehouses due to limits.
C . / D. UNSET RESOURCE_MONITOR
You can't remove a resource monitor from a warehouse while it is currently suspended by that same monitor.
You must first reset or adjust the monitor.
G . UNSET RESOURCE_MONITORS
Invalid syntax - there's no RESOURCE_MONITORS plural keyword.
SnowPro Administrator Reference:
Resource Monitors Overview
ALTER RESOURCE MONITOR
Best Practices for Controlling Warehouse Credit Usage

NEW QUESTION # 72
An organization's sales team leverages this Snowflake query a few times a day:
SELECT CUSTOMER ID, CUSTOMER_NAME, ADDRESS, PHONE NO
FROM CUSTOMERS
WHERE LAST UPDATED BETWEEN TO_DATE (CURRENT_TIMESTAMP) AND (TO_DATE (CURRENT_TIMESTAMP) -7); What can the Snowflake Administrator do to optimize the use of persisted query results whenever possible?

- A. Assign everyone on the sales team to the same virtual warehouse.
- B. Wrap the query in a User-Defined Function (UDF) to match syntax execution.
- C. Leverage the CURRENT_DATE function for date calculations.
- D. Assign everyone on the sales team to the same security role.

Answer: C

Explanation:
According to the web search results from my predefined tool search_web, one of the factors that affects the reuse of persisted query results is the exact match of the query syntax1. If the query contains functions that return different values for successive runs, such as CURRENT_TIMESTAMP, then the query will not match the previous query and will not benefit from the cache. To avoid this, the query should use functions that return consistent values for the same day, such as CURRENT_DATE, which returns the current date without the time component2. Option A is incorrect because wrapping the query in a UDF does not guarantee the syntax match, as the UDF may also contain dynamic functions. Option B is incorrect because the virtual warehouse does not affect the persisted query results, which are stored at the account level1. Option C is incorrect because the security role does not affect the persisted query results, as long as the role has the necessary privileges to access the tables and views used in the query1.
1: Using Persisted Query Results | Snowflake Documentation 2: Date and Time Functions | Snowflake Documentation

NEW QUESTION # 73
Which actions are considered breaking changes to data that is shared with consumers in the Snowflake Marketplace? (Select TWO).

- A. Unpublishing the data listing
- B. Dropping a column from a table
- C. Deleting data from a table
- D. Renaming a table
- E. Adding region availability to the listing

Answer: B,D

Explanation:
Explanation
According to the Snowflake documentation1, breaking changes are changes that affect the schema or structure of the shared data, such as dropping or renaming a column or a table. These changes may cause errors or unexpected results for the consumers who query the shared data. Deleting data from a table, unpublishing the data listing, or adding region availability to the listing are not breaking changes, as they do not alter the schema or structure of the shared data.
1: Managing Data Listings in Snowflake Data Marketplace | Snowflake Documentation

NEW QUESTION # 74
An Administrator needs to create a sample of the table LINEITEM. The sample should not be repeatable and the sampling function should take the data by blocks of rows.
What select command will generate a sample of 20% of the table?

- A. select * from LINEITEM sample system (20);
- B. select * from LINEITEM sample bernoulli (20);
- C. select * from LINEITEM tablesample system (20) seed (1);
- D. select * from LINEITEM tablesample block (20 rows);

Answer: A

Explanation:
Explanation
This command will generate a sample of 20% of the table by using the SYSTEM (or BLOCK) sampling method, which selects each block of rows with a probability of 20/100. This method is suitable for taking data by blocks of rows, as the question requires. According to the Snowflake documentation, "SYSTEM (or BLOCK): Includes each block of rows with a probability of p/100. Similar to flipping a weighted coin for each block of rows. This method does not support fixed-size sampling." The other options are either incorrect or do not meet the requirements of the question. Option A uses the BERNOULLI (or ROW) sampling method, which selects each row with a probability of 20/100, but does not take data by blocks of rows. Option C uses the BLOCK sampling method, but specifies a fixed number of rows (20) instead of a percentage (20%).
Option D uses the SYSTEM sampling method, but specifies a seed value (1), which makes the sampling repeatable, contrary to the question.

NEW QUESTION # 75
A Snowflake customer is experiencing higher costs than anticipated while migrating their data warehouse workloads from on-premises to Snowflake. The migration workloads have been deployed on a single warehouse and are characterized by a large number of small INSERTs rather than bulk loading of large extracts. That single warehouse has been configured as a single cluster, 2XL because there are many parallel INSERTs that are scheduled during nightly loads.
How can the Administrator reduce the costs, while minimizing the overall load times, for migrating data warehouse history?

- A. The INSERTS should be converted to several tables to avoid contention on large tables that slows down query processing.
- B. The 2XL warehouse should be changed to 4XL to increase the number of threads available for parallel load queries.
- C. There should be another 2XL warehouse deployed to handle a portion of the load queries.
- D. The warehouse should be kept as a SMALL or XSMALL and configured as a multi-cluster warehouse to handle the parallel load queries.

Answer: D

Explanation:
Explanation
According to the Snowflake Warehouse Cost Optimization blog post, one of the strategies to reduce the cost of running a

warehouse is to use a multi-cluster warehouse with auto-scaling enabled. This allows the warehouse to automatically adjust the number of clusters based on the concurrency demand and the queue size. A multi-cluster warehouse can also be configured with a minimum and maximum number of clusters, as well as a scaling policy to control the scaling behavior. This way,the warehouse can handle the parallel load queries efficiently without wasting resources or credits. The blog post also suggests using a smaller warehouse size, such as SMALL or XSMALL, for loading data, as it can perform better than a larger warehouse size for small INSERTs. Therefore, the best option to reduce the costs while minimizing the overall load times for migrating data warehouse history is to keep the warehouse as a SMALL or XSMALL and configure it as a multi-cluster warehouse to handle the parallel load queries. The other options are incorrect because:
*A. Deploying another 2XL warehouse to handle a portion of the load queries will not reduce the costs, but increase them. It will also introduce complexity and potential inconsistency in managing the data loading process across multiple warehouses.
*B. Changing the 2XL warehouse to 4XL will not reduce the costs, but increase them. It will also provide more compute resources than needed for small INSERTs, which are not CPU-intensive but I/O-intensive.
*D. Converting the INSERTs to several tables will not reduce the costs, but increase them. It will also create unnecessary data duplication and fragmentation, which will affect the query performance and data quality.

NEW QUESTION # 76
.....

BTW, DOWNLOAD part of ExamsLabs ADA-C01 dumps from Cloud Storage: https://drive.google.com/open?id=1XBwWSsP9Q_ruR80zjcPU1zgzkVErZZ4k