

New ADA-C01 Exam Quizzes | Reliable Snowflake Valid ADA-C01 Exam Fee: SnowPro Advanced Administrator

Solutions Third Edition
Name: _____

Listening

1 Listen to a radio interview and choose the best answers (a-d).



2 What does Kauffman claim in his book?

- It doesn't take as long as people think to become very good at something.
- Twenty hours is enough time to learn any language fluently.
- Everyone can quickly become reasonably competent at anything.
- Learning any new skill takes no effort at all.

3 What does the speaker say about the 10,000-hour figure?

- It's very inaccurate, and not well researched.
- It's only true for some people.
- It discourages people from trying to learn new things.
- It doesn't match her experience of learning.

4 What is the first step in learning something new, according to Kauffman's book?

- Breaking the skill down into smaller parts.
- Discovering shortcuts to learning.
- Learning how to memorise things fast.
- Finding enough time to practise.

5 What does the speaker say is a major barrier to learning?

- A lack of natural ability.
- Inability to memorise or recall new information.
- A lack of time to practise.
- An inflexible attitude to learning.

6 What do Kaufmann and Gladwell both believe about learning?

- The amount of time you spend learning something is totally unimportant.
- There's no such thing as natural talent.
- Talent isn't everything when it comes to learning.
- People give up learning new skills too easily.

Reading

3 Read the texts about vision and match them with questions 1-8. Each text can be matched with more than one question.

Which text ...

- mentions the brain creating something? _____
- talks about qualified people being influenced? _____
- mentions the link between language and perception? _____
- mentions our inability to remember things? _____
- discusses differences in colour perception? _____
- suggests a reason for something by mentioning the natural world? _____
- mentions seeing more when you pay attention? _____
- suggests competitions are not as fair as they seem? _____

A Do you see the same colour as me?

The Candoshi native people of Peru have one word for orange and yellow; when most people in Peru would use separate Spanish words for these two colours. More surprisingly, they have the same colour word (*ayshana*) for colours from purple to blue to green, but another word (*shamshap*) for what most Peruvians would call 'dark green'. It's clear that different groups of people tend to divide up colours differently, but do they actually see them in a different way?

Tests with the Himba people of Namibia suggest that we do see colours differently. In one experiment, a slightly different tone of green in a circle of green dots really stood out for the Himba people. This is interesting, bearing in mind it was almost impossible for Europeans to see the difference. In contrast, the Himba found a test where they had to find a light blue spot amongst light green ones very difficult. This may be owing to the fact that the Himba word for these two colours is the same. It seems that people around the world really do see things differently.



Solutions Third Edition Advanced Tests
Introduction - Unit 1



What's more, part of that Actual4Dumps ADA-C01 dumps now are free: <https://drive.google.com/open?id=1P1xANCct0u-U0XiemMbCrEETLQXVdoK9>

The contents of ADA-C01 study materials are all compiled by industry experts based on the ADA-C01 examination outlines and industry development trends over the years. It does not overlap with the content of the ADA-C01 question banks on the market, and avoids the fatigue caused by repeated exercises. Our ADA-C01 Exam Guide is not simply a patchwork of exam questions, but has its own system and levels of hierarchy, which can make users improve effectively.

Snowflake ADA-C01 Exam Syllabus Topics:

Topic	Details
Topic 1	<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 2	<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 3	<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.
Topic 4	<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 5	<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.

>> ADA-C01 Exam Quizzes <<

Valid ADA-C01 Exam Fee | ADA-C01 Updated CBT

This is an era of high efficiency, and how to prove your competitiveness, perhaps only through the ADA-C01 certificates you get is the most straightforward. But the time is limited for many people since you may be caught with other affairs. With our ADA-C01 study materials, all your problems will be solved easily without doubt. We can provide not only the trustable and valid ADA-C01 Exam Torrent but also the most flexible study methods. And we can confirm that you are bound to pass your ADA-C01 exam just as numerous of our other customers do.

Snowflake SnowPro Advanced Administrator Sample Questions (Q16-Q21):

NEW QUESTION # 16

A user with the proper role issues the following commands when setting up and activating network policies:

```
CREATE OR REPLACE NETWORK POLICY foo_policy
ALLOWED_IP_LIST = ('1.1.1.0/24', '2.2.2.0/24', '3.3.3.0/24')
BLOCKED_IP_LIST = ('1.1.1.1')
COMMENT = 'Account level policy';
ALTER ACCOUNT SET NETWORK_POLICY=FOO_POLICY;
CREATE OR REPLACE NETWORK POLICY bar_policy
ALLOWED_IP_LIST = ('3.3.3.0/24')
BLOCKED_IP_LIST = ('3.3.3.10')
COMMENT = 'user level policy';
ALTER USER user1 SET NETWORK_POLICY=BAR_POLICY;
```

Afterwards, user1 attempts to log in to Snowflake from IP address 3.3.3.10.

Will the login be successful?

- A. Yes, because 3.3.3.10 is found in the ALLOWED_IP_LIST of foo_policy.
- **B. No, because 3.3.3.10 is found in the BLOCKED_IP_LIST of bar_policy.**
- C. Yes, because 3.3.3.10 is found in the ALLOWED_IP_LIST of bar_policy.
- D. No, because 3.3.3.10 is not found in the ALLOWED_IP_LIST of foo_policy.

Answer: B

Explanation:

According to the Snowflake documentation¹, network policies are a feature that allows restricting access to your account based on user IP address. A network policy can be applied to an account, a user, or a security integration, and can specify a list of allowed IP addresses and a list of blocked IP addresses. If there are network policies applied to more than one of these, the most specific network policy overrides more general network policies. In this case, the user1 has a network policy (bar_policy) applied to them, which overrides the account-level network policy (foo_policy). The bar_policy allows access only from the IP range 3.3.3.0/24, and blocks access from the IP address 3.3.3.10. Therefore, the user1 will not be able to log in to Snowflake from IP address 3.3.3.10, as it is found in the BLOCKED_IP_LIST of bar_policy. Option A is incorrect because the ALLOWED_IP_LIST of bar_policy does not override the BLOCKED_IP_LIST of bar_policy. Option C is incorrect because the ALLOWED_IP_LIST of foo_policy does not apply to user1, as it is overridden by the user-level network policy. Option D is incorrect because the ALLOWED_IP_LIST of foo_policy does not matter, as it is overridden by the user-level network policy.

NEW QUESTION # 17

If the query matches the definition, will Snowflake always dynamically rewrite the query to use a materialized view?

- A. No, because joins are not supported by materialized views.
- **B. No, because the optimizer might decide against it.**
- C. Yes, because materialized views are always faster.
- D. No, because the materialized view may not be up-to-date.

Answer: B

Explanation:

Snowflake's query optimizer can automatically rewrite queries against the base table or regular views to use the materialized view instead, if the query matches the definition of the materialized view¹. However, this is not always guaranteed, as the optimizer might decide against using the materialized view based on various factors, such as the freshness of the data, the size of the result set, the complexity of the query, and the availability of the materialized view². Therefore, the answer is no, because the optimizer might decide against it.

NEW QUESTION # 18

Which masking policy will mask a column whenever it is queried through a view owned by a role named MASKED_VIEW_ROLE?

- A. create or replace masking policy maskstring as (val string) returns string -> case when invoker_role() in ('MASKED_VIEW_ROLE') then else val end;
' **
- B. create or replace masking policy maskString as (val string) returns string -> case when current_role() in ('MASKED_VIEW_ROLE') then '*****' else val end;
- C. create or replace masking policy maskString as (val string) returns string -> case when array_contains ('MASKED_VIEW_ROLE' :: variant, parse_json(current_available_roles ())) then
'*
else val
end;
** '
- **D. create or replace masking policy maskstring as (val string) returns string -> case when is_role_in_session ('MASKED_VIEW_ROLE') then ' ** else val end;
*,**

Answer: D

Explanation:

Explanation

A masking policy is a SQL expression that transforms the data in a column based on the role that queries the column¹. The `is_role_in_session` function returns true if the specified role is in the current session². Therefore, the masking policy in option A will mask the column data with asterisks whenever it is queried through a view owned by the `MASKED_VIEW_ROLE`³. The other options use different functions that do not check the ownership of the view, but rather the current role, the invoker role, or the available roles in the session^{4,5}.

These functions may not return the desired result if the role that owns the view is different from the role that queries the view.

NEW QUESTION # 19

A user has enrolled in Multi-factor Authentication (MFA) for connecting to Snowflake. The user informs the Snowflake Administrator that they lost their mobile phone the previous evening.

Which step should the Administrator take to allow the user to log in to the system, without revoking their MFA enrollment?

- A. Alter the user and set `MINS TO BYPASS MFA` to a value that will disable MFA long enough for the user to log in.
- B. Alter the user and set `DISABLE_MFA` to true, which will suspend the MFA requirement for 24 hours.
- C. Instruct the user to append the normal URL with `?mode=mfa_bypass&code=` to log on.
- D. Instruct the user to connect to Snowflake using SnowSQL, which does not support MFA authentication.

Answer: A

Explanation:

Explanation

The `MINS_TO_BYPASS_MFA` property allows the account administrator to temporarily disable MFA for a user who has lost their phone or changed their phone number¹. The user can log in without MFA for the specified number of minutes, and then re-enroll in MFA using their new phone¹. This does not revoke their MFA enrollment, unlike the `DISABLE_MFA` property, which cancels their enrollment and requires them to re-enroll from scratch¹. The other options are not valid ways to bypass MFA, as SnowSQL does support MFA authentication², and there is no such URL parameter as `?mode=mfa_bypass&code=` for Snowflake³.

NEW QUESTION # 20

A Snowflake Administrator has a multi-cluster virtual warehouse and is using the Snowflake Business Critical edition. The minimum number of clusters is set to 2 and the maximum number of clusters is set to 10. This configuration works well for the standard workload, rarely exceeding 5 running clusters. However, once a month the Administrator notes that there are a few complex long-running queries that are causing increased queue time and the warehouse reaches its maximum limit at 10 clusters.

Which solutions will address the issues happening once a month? (Select TWO).

- A. Examine the complex queries and determine if they can be made more efficient using clustering keys or materialized views.
- B. Use a task to increase the cluster size for the time period that the more complex queries are running and another task to reduce the size of the cluster once the complex queries complete.
- C. Increase the multi-cluster maximum to 20 or more clusters.
- D. Have the group running the complex monthly queries use a separate appropriately-sized warehouse to support their workload.
- E. Increase the minimum number of clusters started in the multi-cluster configuration to 5.

Answer: B,D

Explanation:

Explanation

According to the Snowflake documentation¹, a multi-cluster warehouse is a virtual warehouse that consists of multiple clusters of compute resources that can scale up or down automatically to handle the concurrency and performance needs of the queries submitted to the warehouse. A multi-cluster warehouse has a minimum and maximum number of clusters that can be specified by the administrator. Option A is a possible solution to address the issues happening once a month, as it allows the administrator to use a task to increase the cluster size for the time period that the more complex queries are running and another task to reduce the size of the cluster once the complex queries complete. This way, the warehouse can have more resources available to handle the complex queries without reaching the maximum limit of 10 clusters, and then return to the normal cluster size to save costs. Option B is another possible solution to address the issues happening once a month, as it allows the administrator to have the group running the complex monthly queries use a separate appropriately-sized warehouse to support their workload. This way, the warehouse can isolate the complex queries from the standard workload and avoid queue time and resource contention. Option C is not a recommended solution to address the issues happening once a month, as it would increase the costs and complexity of managing the multi-cluster warehouse, and may not solve the underlying problem of inefficient queries. Option D is a good practice to improve the

