



**Free Questions for ADA-C01 by dumpssheet**

**Shared by Roth on 15-04-2024**

**For More Free Questions and Preparation Resources**

**Check the Links on Last Page**

# Question 1

---

## Question Type: MultipleChoice

---

A company has many users in the role ANALYST who routinely query Snowflake through a reporting tool. The Administrator has noticed that the ANALYST users keep two

small clusters busy all of the time, and occasionally they need three or four clusters of that size.

Based on this scenario, how should the Administrator set up a virtual warehouse to MOST efficiently support this group of users?

### Options:

---

- A-** Create a multi-cluster warehouse with MIN\_CLUSTERS set to 1. Give MANAGE privileges to the ANALYST role so this group can start and stop the warehouse, and increase the number of clusters as needed.
- B-** Create a multi-cluster warehouse with MIN\_CLUSTERS set to 2. Set the warehouse to auto-resume and auto-suspend, and give USAGE privileges to the ANALYST role. Allow the warehouse to auto-scale.
- C-** Create a standard X-Large warehouse, which is equivalent to four small clusters. Set the warehouse to auto-resume and auto-suspend, and give USAGE privileges to the ANALYST role.
- D-** Create four virtual warehouses (sized Small through XL) and set them to auto-suspend and auto-resume. Have users in the ANALYST role select the appropriate warehouse based on how many queries are being run.

## Answer:

---

B

## Explanation:

---

According to the Snowflake documentation<sup>1</sup>, 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 B is the most efficient way to support the group of users, as it allows the administrator to create a multi-cluster warehouse with `MIN_CLUSTERS` set to 2, which means that the warehouse will always have two clusters running to handle the standard workload. The warehouse can also auto-scale up to the maximum number of clusters (which can be set according to the peak workload) when there is a spike in demand, and then scale down when the demand decreases. The warehouse can also auto-resume and auto-suspend, which means that the warehouse will automatically start when a query is submitted and automatically stop after a period of inactivity. The administrator can also give `USAGE` privileges to the `ANALYST` role, which means that the users can use the warehouse to execute queries and load data, but not modify or operate the warehouse. Option A is not efficient, as it requires the users to manually start and stop the warehouse, and increase the number of clusters as needed, which can be time-consuming and error-prone. Option C is not efficient, as it creates a standard X-Large warehouse, which is equivalent to four small clusters, which may be more than needed for the standard workload, and may not be enough for the peak workload. Option D is not efficient, as it creates four virtual warehouses of different sizes, which can be confusing and cumbersome for the users to select the appropriate warehouse based on how many queries are being run, and may also result in wasted resources and costs.

## Question 2

---

**Question Type: MultipleChoice**

---

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).

**Options:**

---

- A-** 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.
- B-** Have the group running the complex monthly queries use a separate appropriately-sized warehouse to support their workload.
- C-** Increase the multi-cluster maximum to 20 or more clusters.
- D-** Examine the complex queries and determine if they can be made more efficient using clustering keys or materialized views.
- E-** Increase the minimum number of clusters started in the multi-cluster configuration to 5.

**Answer:**

---

A, B

**Explanation:**

---

According to the Snowflake documentation<sup>1</sup>, 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 performance of the queries, but it is not a direct solution to address the issues happening once a month, as it requires analyzing and optimizing the complex queries using clustering keys or materialized views, which may not be feasible or effective in all cases. Option E is not a recommended solution to address the issues happening once a month, as it would increase the costs and waste resources by starting more clusters than needed for the standard workload.

## Question 3

---

Question Type: MultipleChoice

---

What is required for stages, without credentials, to limit data exfiltration after a storage integration and associated stages are created?

### Options:

---

- A-** ALTER ACCOUNT my\_account SET  
REQUIRE\_STORAGE\_INTEGRATION\_FOR\_STAGE\_CREATION = true;  
ALTER ACCOUNT my\_account SET  
REQUIRE\_STORAGE\_INTEGRATION\_FOR\_STAGE\_OPERATION = true;  
ALTER ACCOUNT my\_account SET  
PREVENT\_UNLOAD\_TO\_INLINE\_URL = false;
- B-** ALTER ACCOUNT my\_account SET  
REQUIRE\_STORAGE\_INTEGRATION\_FOR\_STAGE\_CREATION = false;  
ALTER ACCOUNT my\_account SET  
REQUIRE\_STORAGE\_INTEGRATION\_FOR\_STAGE\_OPERATION = false;  
ALTER ACCOUNT my\_account SET  
PREVENT\_UNLOAD\_TO\_INLINE\_URL = true;
- C-** ALTER ACCOUNT my\_account SET  
REQUIRE\_STORAGE\_INTEGRATION\_FOR\_STAGE\_CREATION = false;  
ALTER ACCOUNT my\_account SET

```
REQUIRE_STORAGE_INTEGRATION_FOR_STAGE_OPERATION = false;  
ALTER ACCOUNT my_account SET  
PREVENT_UNLOAD_TO_INLINE_URL = false;
```

```
D- ALTER ACCOUNT my_account SET  
REQUIRE_STORAGE_INTEGRATION_FOR_STAGE_CREATION = true;  
ALTER ACCOUNT my_account SET  
REQUIRE_STORAGE_INTEGRATION FOR STAGE_OPERATION = true;  
ALTER ACCOUNT my_account SET  
PREVENT_UNLOAD_TO_INLINE_URL = true;
```

### **Answer:**

---

D

### **Explanation:**

---

According to the Snowflake documentation<sup>1</sup>, stages without credentials are a way to create external stages that use storage integrations to access data files in cloud storage without providing any credentials to Snowflake. Storage integrations are objects that define a trust relationship between Snowflake and a cloud provider, allowing Snowflake to authenticate and authorize access to the cloud storage. To limit data exfiltration after a storage integration and associated stages are created, the following account-level parameters can be set:

\* `REQUIRE_STORAGE_INTEGRATION_FOR_STAGE_CREATION`: This parameter enforces that all external stages must be created using a storage integration. This prevents users from creating external stages with inline credentials or URLs that point to unauthorized locations.

\* `REQUIRE_STORAGE_INTEGRATION_FOR_STAGE_OPERATION`: This parameter enforces that all operations on external stages, such as `PUT`, `GET`, `COPY`, and `LIST`, must use a storage integration. This prevents users from performing operations on external stages with inline credentials or URLs that point to unauthorized locations.

\* `PREVENT_UNLOAD_TO_INLINE_URL`: This parameter prevents users from unloading data from Snowflake tables to inline URLs that do not use a storage integration. This prevents users from exporting data to unauthorized locations.

Therefore, the correct answer is option D, which sets all these parameters to true. Option A is incorrect because it sets `PREVENT_UNLOAD_TO_INLINE_URL` to false, which allows users to unload data to inline URLs that do not use a storage integration. Option B is incorrect because it sets both `REQUIRE_STORAGE_INTEGRATION_FOR_STAGE_CREATION` and `REQUIRE_STORAGE_INTEGRATION_FOR_STAGE_OPERATION` to false, which allows users to create and operate on external stages without using a storage integration. Option C is incorrect because it sets all the parameters to false, which does not enforce any restrictions on data exfiltration.

## Question 4

---

**Question Type:** MultipleChoice

---

A virtual warehouse `report_wh` is configured with `AUTO_RESUME=TRUE` and `AUTO_SUSPEND=300`. A user has been granted the role `accountant`.

An application with the `accountant` role should use this warehouse to run financial reports, and should keep track of compute credits used by the warehouse.



What minimal privileges on the warehouse should be granted to the role to meet the requirements for the application? (Select TWO).

**Options:**

---

- A- OPERATE
- B- MODIFY
- C- MONITOR
- D- USAGE
- E- OWNERSHIP

**Answer:**

---

C, D

**Explanation:**

---

According to the Snowflake documentation<sup>1</sup>, the MONITOR privilege on a warehouse grants the ability to view the warehouse usage and performance metrics, such as the number of credits consumed, the average and maximum run time, and the number of queries executed. The USAGE privilege on a warehouse grants the ability to use the warehouse to execute queries and load data. Therefore, the minimal privileges on the warehouse that should be granted to the role to meet the requirements for the application are MONITOR and USAGE. Option A is incorrect because the OPERATE privilege on a warehouse grants the ability to start, stop, resume, and suspend the warehouse, which is not required for the application. Option B is incorrect because the MODIFY privilege on a warehouse

grants the ability to alter the warehouse properties, such as the size, auto-suspend, and auto-resume settings, which is not required for the application. Option E is incorrect because the OWNERSHIP privilege on a warehouse grants the ability to drop the warehouse, grant or revoke privileges on the warehouse, and transfer the ownership to another role, which is not required for the application.

## Question 5

---

**Question Type:** MultipleChoice

---

Which type of listing in the Snowflake Marketplace can be added and queried immediately?

### Options:

---

- A- Monetized listing
- B- Standard listing
- C- Regional listing
- D- Personalized listing

### Answer:

---

B

## Explanation:

---

According to the Snowflake documentation<sup>1</sup>, a standard listing is a type of listing that provides free access to the full data product, with no payment required. A standard listing can be added and queried immediately by the consumer, as long as they accept the terms and conditions of the listing. A monetized listing is a type of listing that charges for access to the data product, using the pricing models offered by Snowflake. A monetized listing requires the consumer to provide payment information and agree to the billing terms before accessing the data product. A regional listing is not a type of listing, but a way to specify the regions where the listing is available. A personalized listing is a type of listing that provides limited trial access to the data product, with unlimited access to the full data product available upon request. A personalized listing requires the consumer to request access from the provider and wait for the provider to grant access before accessing the data product. Therefore, the only type of listing that can be added and queried immediately is the standard listing.

## Question 6

---

### Question Type: MultipleChoice

---

A Snowflake account is configured with SCIM provisioning for user accounts and has bi-directional synchronization for user identities. An Administrator with access to SECURITYADMIN uses the Snowflake UI to create a user by issuing the following commands:

```
use role USERADMIN;
```

```
create or replace role DEVELOPER_ROLE;
```

```
create user PTORRES PASSWORD = 'hello world!' MUST_CHANGE_PASSWORD = FALSE
```

```
default_role = DEVELOPER_ROLE;
```

The new user named PTORRES successfully logs in, but sees a default role of PUBLIC in the web UI. When attempted, the following command fails:

```
use DEVELOPER_ROLE;
```

Why does this command fail?

### Options:

---

- A-** The DEVELOPER\_ROLE needs to be granted to SYSADMIN before user PTORRES will be able to use the role.
- B-** The new role can only take effect after USERADMIN has logged out.
- C-** USERADMIN needs to explicitly grant the DEVELOPER\_ROLE to the new USER.
- D-** The new role will only take effect once the identity provider has synchronized by way of SCIM with the Snowflake account.

### Answer:

---

C

### Explanation:

---

According to the Snowflake documentation<sup>1</sup>, creating a user with a default role does not automatically grant that role to the user. The user must be explicitly granted the role by the role owner or a higher-level role. Therefore, the USERADMIN role, which created the DEVELOPER\_ROLE, needs to explicitly grant the DEVELOPER\_ROLE to the new user PTORRES using the GRANT ROLE command. Otherwise, the user PTORRES will not be able to use the DEVELOPER\_ROLE and will see the default role of PUBLIC in the web UI. Option A is incorrect because the DEVELOPER\_ROLE does not need to be granted to SYSADMIN before user PTORRES can use the role. Option B is incorrect because the new role can take effect immediately after it is created and granted to the user, and does not depend on the USERADMIN role logging out. Option D is incorrect because the new role will not be affected by the identity provider synchronization, as it is created and managed in Snowflake.

## Question 7

---

**Question Type:** MultipleChoice

---

What are benefits of using Snowflake organizations? (Select TWO).

### Options:

---

- A- Administrators can change Snowflake account editions on-demand based on need.
- B- Administrators can monitor and understand usage across all accounts in the organization.

- C-** Administrators can simplify data movement across all accounts within the organization.
- D-** User administration is simplified across all accounts within the organization.
- E-** Administrators have the ability to create accounts in any available cloud provider or region.

### **Answer:**

---

B, E

### **Explanation:**

---

According to the Snowflake documentation<sup>1</sup>, organizations are a feature that allows linking the accounts owned by a business entity, simplifying account management and billing, replication and failover, data sharing, and other account administration tasks. Some of the benefits of using organizations are:

\* Administrators can monitor and understand usage across all accounts in the organization using the ORGANIZATION\_USAGE schema, which provides historical usage data for all accounts in the organization via views in a shared database named SNOWFLAKE2. This can help to optimize costs and performance across the organization.

\* Administrators have the ability to create accounts in any available cloud provider or region using the CREATE ACCOUNT command, which allows specifying the cloud platform and region for the new account<sup>3</sup>. This can help to meet the business needs and compliance requirements of the organization.

Option A is incorrect because administrators cannot change Snowflake account editions on-demand based on need, but rather have to contact Snowflake Support to request an edition change<sup>4</sup>. Option C is incorrect because administrators cannot simplify data movement across all accounts within the organization, but rather have to enable account database replication for both the source and target

accounts, and use the ALTER DATABASE ... ENABLE REPLICATION TO ACCOUNTS command to promote a local database to serve as the primary database and enable replication to the target accounts<sup>5</sup>. Option D is incorrect because user administration is not simplified across all accounts within the organization, but rather requires creating and managing users, roles, and privileges for each account separately, unless using a federated authentication method such as SSO or SCIM.

**To Get Premium Files for ADA-C01 Visit**

**<https://www.p2pexams.com/products/ada-c01>**

**For More Free Questions Visit**

**<https://www.p2pexams.com/snowflake/pdf/ada-c01>**

